


AGING POWER DELIVERY INFRASTRUCTURES

Second Edition



H. Lee Willis
Randall R. Schrieber



CRC Press
Taylor & Francis Group

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2013 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works
Version Date: 20121205

International Standard Book Number-13: 978-1-4398-9691-4 (eBook - PDF)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Contents

<i>Series Introduction</i>	<i>iii</i>
<i>Preface</i>	<i>v</i>
1 Aging Power Delivery Infrastructures	1
1.1 Introduction and Overview	1
1.2 Aging Infrastructure Management Must Become A Mainstream Part of Utility Operations?	4
1.3 What is the Infrastructure Involved?	5
1.4 Characteristics of an Aging Infrastructure	7
1.5 Power Delivery Systems	8
1.6 The Business Environment Has Changed	11
1.7 Five Factors Contribute to Aging Infrastructure Problems	14
1.8 Condition, Criticality, and Consequences	34
1.9 Concluding Remarks	38
References and Bibliography	48
2 Power Delivery Systems	51
2.1 Introduction	51
2.2 T&D System's Mission	52
2.3 The "Laws of T&D"	54
2.4 Levels of the T&D System	57
2.5 Utility Distribution Equipment	68
2.6 T&D Costs	74
2.7 Types of Delivery System Design	83
	xi

2.8	Smart Grid	96
2.9	Conclusion	98
	References and Bibliography	100
3	Customer Demand for Power and Reliability of Service	103
3.1	The Two-Q's: Quantity and Quality of Power	103
3.2	Electric Consumer Need for Quantity of Power	104
3.3	Electric Consumer Need for Quality of Power	113
3.4	Two-Q Analysis: Quantity and Quality of Power Are Both Important Aspects of Consumer Value	129
3.5	Conclusion and Summary	132
	References and Bibliography	133
4	Power System Reliability and Reliability of Service	135
4.1	Introduction	135
4.2	Outages Cause Interruptions	137
4.3	Reliability Indices	141
4.4	Reliability and Contingency Criteria for Planning	146
4.5	Cost Is Usually <i>the</i> Major Consideration	151
4.6	Two-Q Analysis of Power Systems	155
4.7	Conclusion and Summary	157
	References and Bibliography	158
5	Cost and Economic Evaluation	161
5.1	Introduction	161
5.2	Costs	162
5.3	Time Value of Money	165
5.4	Decision Bases and Cost-Effectiveness Evaluation	181
5.5	Budget-Constrained Planning: Marginal Benefit vs. Cost Analysis	182

Contents	xiii
5.6 Asset Management and Pareto Analysis	195
5.7 Conclusion	200
References and Bibliography	201
 6 Equipment Inspection, Testing, and Diagnostics	 203
6.1 Introduction	203
6.2 Inspection, Testing, and Diagnostic Evaluation	203
6.3 Equipment Testing and Diagnosis Methods	213
6.4 Tests and Diagnosis of Insulating Oil	221
6.5 On-Line and Real-Time Monitoring and Testing	229
6.5 Summary and Final Comments	231
References	232
 7 Aging Equipment and Its Impacts	 233
7.1 Introduction	233
7.2 Equipment Aging	234
7.3 Equipment Failure Rate Increases with Age	249
7.4 Impact of Escalating Failure Rates	256
7.5 Summary of Key Points	262
References	263
 8 Obsolete System Structures	 265
8.1 Introduction	265
8.2 Obsolete Systems Layouts	266
8.3 Impacts on the Sub-transmission – Substation Level	272
8.4 Feeder System Impacts	278
8.5 “Fixes” For Outdated System Structures	292
8.6 Summary of Key Points	295
References and Bibliography	296

9	Traditional Reliability Engineering Tools and Their Limitations	297
9.1	Introduction	297
9.2	Contingency-Based Planning Methods	298
9.3	Limitations of N-1 Methodology	304
9.4	Other Planning Related Concerns	317
9.5	Summary and Conclusion	326
	References and Bibliography	330
10	Primary Distribution Planning and Engineering Interactions	331
10.1	Introduction	331
10.2	Distribution Planning and the Perceived Role of Distribution	332
10.3	Flexibility and Effectiveness in Feeder Level Planning	342
10.4	Smart Distribution Systems	355
10.5	Conclusion	333
	References and Bibliography	360
11	Equipment Condition Assessment	361
11.1	Introduction	361
11.2	What Does “Condition” Mean?	362
11.3	Analysis, Modeling, and Prediction	367
11.4	Power Transformers	374
11.5	Switchgear and Circuit Breakers	381
11.6	Underground Equipment and Cables	384
11.7	Overhead Lines and Associated Equipment	387
11.8	Service Transformers and Service Circuits	393
11.9	Evaluating and Prioritizing Equipment Condition	394
11.10	Condition Tracking	403
11.11	Summary and Final Comments	406

References and Bibliography	407
12 Optimization	409
12.1 Introduction	409
12.2 Prioritizing Inspection, Maintenance, and Service	410
12.3 Reliability Centered Maintenance	414
12.4 Basic Reliability-Centered Prioritization	418
12.5 Prioritization of the <i>Type</i> of Maintenance	429
12.6 Practical Aspects for Implementation	437
12.7 Extending Reliability-Centered Prioritization and Optimization to Other Operations Functions	445
12.8 Optimization	450
12.9 Conclusion and Recommendations	454
References and Bibliography	459
13 Planning Methods for Aging T&D Infrastructures	457
13.1 Introduction	457
13.2 Planning: Finding the Best Alternative	460
13.3 Short- and Long-Range Planning	470
13.4 The T&D Planning Process	480
13.5 The Systems Approach	500
13.6 Summary of Planning in an Aging T&D Infrastructure	504
References and Bibliography	506
14 Reliability Can Be Planned and Engineered	507
14.1 Introduction	507
14.2 Reliability Can Be Engineered	510
14.3 Methods for Distribution System Reliability Assessment	514
14.4 Application of Analytical Simulation for Detailed Reliability	

Assessment	518
14.5 Use of a Hybrid Analytical Solution	524
14.6 Conclusion and Key Points	532
References and Bibliography	533
15 Six Aging Infrastructure Management Stories	537
15.1 Introduction	537
15.2 Four Philosophies of Ownership	538
15.3 Comparing the Four Ownership Philosophies	563
15.4 Organizational culture	567
15.5 Six Cases: “Success” from Six Different Perspectives	574
15.5.1 Big States Electric – Asset Management	574
15.5.2 Mid-State Electric – Reliability-Centered Maintenance	579
15.5.3 International Bulk Petroleum – Buy, Burn, and Bash	583
15.5.4 Big State University System – Standards Based	586
15.5.5 Third-World Orphan Rescue Foundation – Uniquely Standards Based	589
15.5.6 Huey Longwaites International Airport – Buy, Burn, and Bash	594
15.6 Conclusion	596
16 Managing an Aging Infrastructure	599
16.1 Introduction	599
16.2 An Inauspicious Beginning	600
16.3 Overview of Approach	603
16.4 The Data Chase	604
16.5 Building the Historical Model of Pole Ownership	612
16.6 Building the Predictive Model	623
16.7 What the Model Says About the Future	632

Contents **xvii**

16.8	Projections of Pole-Related Costs and Performance	638
16.9	Finding the Bad Poles	648
16.10	Pole Replacement Merit and Economics	658
16.11	Options for Pole Replacement Programs	667
16.12	CL&P's Plan for Dealing with Aging Wooden Distribution Poles	678

17 Guidelines and Recommendations **689**

17.1	Introduction	689
17.2	Five Interrelated Factors	689
17.3	Results-Driven Management (RDM)	692
17.4	Steps to Mitigate Aging Infrastructure Effects	695
17.5	The Aging Infrastructure Management Program Itself	705
	Bibliography	714

Appendix A – Notes on a Spreadsheet-Based Model for Aging Trends

A.1	Introduction	1
A.2	The Concept: Incrementing an Age Distribution	3
A.3	Multi-Sheet Structure of a More Practical Model	7
A.4	Fetting and Calibrating a Model	9
A.5	Real World Data Situations: Getting A Solution	14
A.6	Comments on Application	19
	Reference	21

Appendix B – Sustainable Point Analysis

B.1	Introduction	1
B.2	Elements of Aging Infrastructure Analysis	4
B.3	Quantitative Examples	7
B.4	Metrics for Measuring Aging Infrastructures	24
B.5	Effect of Different Failure Rate Curve Shapes	31

B.6 Reading Distributions	36
B.7 Applying Economic Analysis	42
B.8 Summary of Key Points	46

Index

Preface

Like the first, this second edition is both a reference book and a tutorial guide on aging power delivery systems, the problems they cause, and the technical and managerial approaches that power system owners can take to manage them. Electric power is unique among major public utilities in that the quality and usability of its product is entirely determined by the characteristics of the system that delivers it, not the way in which it is manufactured. America's aging power delivery systems mean potentially poor power quality, frequent and lengthy service interruptions, and frustrating challenges for the customers, owners and employees of the electric utilities and large heavy industries who depend on these systems.

The first edition, released in 2001, contained everything we could find on aging power delivery systems at that time. Such is the pace of work in this field that only twelve years later, a majority of this second edition is all new. In early 1994, it first became clear to us that the aging of power delivery systems throughout the U.S. was an important and growing problem which potentially threatened the entire industry. Although the visible symptoms of aging infrastructures – frequent and lengthy interruptions of electric service and rising repair and maintenance costs – were largely incipient at that time – projections of equipment age, system performance, and utility financial performance made it clear that the problem would grow steadily. Events proved this projection to be accurate. During the final five years of the 20th century, aging utility systems throughout the United States began to give their first clear indications of a growing level of operating problems, leading to an uptick in degraded customer service and reduced public trust in electric utilities.

From 1994 through 2001, we devoted considerable effort personally and

professionally to the problem of aging power delivery infrastructures and its possible solutions. We performed original research and testing. We collected a host of innovative, proven, and in some cases unproven techniques from around the world, and studied and vetted them carefully. We developed ways of combining these so they work together, and we tested them until we were confident we knew what worked, what didn't, and why. The first edition (2001) summarized what we had learned.

But, as is so often the case, our major breakthrough came just after publication of that first edition: the *sustainable point* approach. Originally we viewed sustainable point analysis as an interesting and powerful mathematical technique that could tease additional information out of historical data and lead to more accurate and useful planning and decision-support models. But gradually we realized the concept was even more useful as a management approach.

Partly this is because the concept is quite simple: *any set of equipment in which failed units are replaced will eventually trend to a sustainable combination of age, reliability, performance and cost, after which conditions will basically remain the same from year to year, not worsen further.* Aging infrastructures do not deteriorate forever. As long as failures are replaced with new or repaired equipment, the system will eventually deteriorate only to a stable point at which new replacements just balance old failures and the whole mix remains pretty much the same from year to year. The fact that, for nearly every utility in the world, this point is still in the future (so things will continue to get a bit worse every year) and the fact that this sustainable point represents a totally unacceptable combination of poor service and high cost, does not diminish the power of the concept.

The concept is useful in a practical sense because there are many possible sustainable points, each a function of different operating policies. A company's policy could be to install the most robust equipment possible, then neglect it entirely while in service, spending only for a replacement when a unit fails. For most types of equipment, that policy will trend to a high failure rate requiring a lot of replacements. On the other hand, that company could spend lavishly from the beginning on frequent inspections, service "whether needed or not," and very attentive pro-active maintenance and refurbishment. The equipment would last far longer. Failures would be less frequent. But overall cost might be higher, at least initially, when most of the system is new.

In the long term, something between that "deliberate neglect" and the "intense attention" plan will be closer to optimum. The important managerial points are that a power system owner has a choice, that the sustainable point can be moved by making changes in ownership and operating policies, and that the best ownership and operating policies will change as a function of the age and condition of the equipment. There is a wide range of sustainable points available to a utility depending on the operating, maintenance, and loading policies it decides to put in place. What is best for a new system may not be best for an older system. What works for one utility will work for another, but perhaps not in the best possible way.

One consequence of this concept is that it makes clear that the term “aging infrastructure” must include more than just all the equipment. It must also include the loading standards, operating guidelines, and inspection-maintenance-service and refurbishment policies that the owner uses. Those standards and guidelines can become deteriorated in their own way – obsolete with respect to current needs. The aging infrastructure issues facing utilities and industry are about worn out equipment *and* outdated designs *and* constricted sites and rights of way *and* outdated engineering standards and operating procedures. Good aging infrastructure management consists of optimizing the choice of equipment and its refurbishment while also making compatible changes in all those operating and ownership policies, the whole combination aimed at optimizing the business results the power system owner desires.

This book is composed of five major groups of chapters. The first consists of Chapters 1 and 17, which together present an “executive summary” of the problem and its major effects, the sustainable point approach, practical solutions that work, and the limitations and consequences these have in the real world.

Chapters 2 - 6 provide a series of tutorial background discussions on topics necessary to understand some of the subtleties of aging infrastructures and their solutions. The third group of Chapters, 7 – 10, looks at various aspects of aging equipment condition deterioration and its systemic effects, and how all that interrelates with the many functions within a power system and an electric utility company or industrial power system user’s business needs. Chapters 11 – 14 discuss methods and technologies to measure, study, plan, model and decide how to best manage aging equipment and the effects it creates.

Chapters 15 and 16 present case studies. Chapter 15 summarizes aging equipment, its effects, and its solution as done at six companies that depend on power systems for their business success. It shows the range of approaches and methods used, and demonstrates the very real differences that can exist among “optimum policies” depending on the owner’s perspective and needs.

Chapter 16 follows, in considerable detail, a major utility’s aging infrastructure management program from inception to completion as a new optimized process implemented as part of its on-going operations. Two Appendices have been added to this latest addition.

Appendix A discusses the structure and use of an electronic spreadsheet model that fits a sustainable point model to historical equipment and failure data. Appendix B provides a discourse on sustainable point concepts and mathematics, with examples to illustrate key points.

*H. Lee Willis
Randall R. Schrieber*

1

Aging Power Delivery Infrastructures

1.1 INTRODUCTION AND OVERVIEW

This chapter provides an overview of aging power delivery infrastructures, their characteristics, the issues they create, the challenges they present, and the methods and processes utilities and the power industry will use to meet those challenges. It is divided into nine sections that cover different aspects of aging power infrastructures, all from a high-level, big picture perspective. The rest of this book fills in the details. This Introduction and Chapter 17 - Guidelines and Recommendations, taken together constitute an executive summary.

Gradually Growing Older

America's electric utility systems are growing older. In many systems, significant portions of the equipment and facilities in service date from the economic boom following WWII, or from the sustained growth period of the 1950s and 1960s that many American cities and towns experienced. A lot of equipment installed then, and still in service today, is between 50 and almost 70 years old.

As electrical equipment gets older, it deteriorates from the wear and tear of service. At some point, it wears out. Then, it may fail outright, or more often, succumb to forces it would have withstood when new but no longer can - the high winds of a storm or the rare electrical spike caused by switching or lightning. Such failures cause interruptions in service to the utility's customers, and require expensive emergency repairs and system restoration work. Monies will have to be found to pay for that work and the cost of replacement equipment.

All equipment eventually reaches an age where deterioration and increase in failure rate is significant. This is the first fundamental truth of aging power

infrastructures. Electric utility equipment is designed and built to be incredibly robust, but eventually time wins: *sooner, but usually later, everything fails*. A second fundamental truth for most utilities is that it is no longer sooner and that it will soon be later: large portions of old equipment are approaching an age when they are not going to provide dependable service without considerable attention.

A portion of the equipment installed fifty to sixty years ago has already failed in the decades since. Storms, floods and other natural disasters took their toll. Accidents, such as a power pole being hit by a car or an underground cable being dug up by a construction crew, happened. Sometimes equipment just failed for no obvious reason. Regardless, all that equipment was replaced when it failed. As a result, even in the oldest areas of a system, there is a mix of old and new.

But the old predominates. Even in areas built immediately after WWII, rarely has more than 25% of the equipment been replaced in the intervening decades. Well over 70%, and often more than 85%, of the equipment is original in spite of being more than five decades old. That creates a practical operating problem for a utility. Some of the equipment is new. The biggest portion, although having spent decades in service, is still good enough to provide satisfactory service for a decade or more to come. But a tiny fraction of the old is badly deteriorated, perhaps not to the extent that it shows visible signs of imminent failure, but to the extent that it cannot do its original job *dependably*. Those failures will poison the performance of the system overall.

It might seem that this problem is easy to solve. Just have the utility find and replace all the equipment that is going to fail in the next few years, and do that on a continuing basis. Realistically, that is not easy to do. First, the level of failure considered “bad” is remarkably low. Power systems depend on long chains of transmission lines, substations, transformers, breakers, feeder circuits, reclosers, sectionalizers, fuses, cutouts, service transformers and service drops. All must function well if the utility is to deliver reliable power to its customers. If anything fails in this chain, some utility customers go without power until service can be restored.

Thus, an equipment failure rate of only half a percent per year would be disastrous. It would mean service interruption rates that are completely unacceptable to customers, regulators, and stockholders alike, and emergency field work rates that exceed the utility’s ability to dispatch crews to do the work. Under normal circumstances, less than two tenths of one percent of equipment in good condition fails in a year. Thus, finding the future “bad actors” in a power system is a challenge – the proverbial needle in the haystack. A utility might have 8,000 steel lattice transmission towers in service, and it will be a very bad year indeed if more than four give serious problems.

Finding that flawed equipment is not easy. In many cases, badly deteriorated equipment does not give obvious signs of its undependability. Tests can be done to improve the knowledge the utility has of equipment condition. But those require skilled personnel, always in short supply, and expensive specialized

equipment. And many tests are intrusive to customers: utility technicians have to go into backyards and factory property to set up and run their tests. In downtown areas they must cordon off lanes of major streets while they work in equipment vaults under the pavement.

Furthermore, test results are not entirely dependable. A test that is 95% accurate may sound as if it is a good test – and it may be the best available – but its unrestricted use creates problems. Suppose a utility has 10,000 older wood distribution poles, of which 500 (1/2 percent) are truly so deteriorated as to probably fail in the next decade. A 95% accurate test will find 475 of those, leaving 25 bad poles in the field. That alone is a concern, but not the major worry from this test. The real concern is that if it has a 5% false positive rate, too, then it will identify 475 of the 9,500 good poles as bad, for a total of 950 poles it indicates should be replaced. If the utility follows those test results it will replace twice as many poles as necessary – and incur twice the cost. The test may be the best that can be done in the circumstances, but it is not the type of clean, effective solution that inspires confidence among all stakeholders.

Furthermore, replacement of equipment already in place is very costly. Of course, there was a significant cost associated with the original installation – what is often called “greenfield” site work. Money was required to buy a pole or transformer or cable, and skilled labor was needed in every aspect from design to construction. But the cost of *replacing* that equipment many years later is much greater. First, before the new equipment can be installed the old equipment has to be removed and recycled, etc. Then new equipment has to be put back into place and reconnected to the rest of the system. Further, typically this work must be done while the rest of the system is in operation – i.e., energized – so that service to customers is not interrupted. That means the utility’s personnel must work in and around live electric equipment. Such work can be done safely, but will proceed much slower and cost more than the de-energized site labor that was used in the original installation. As a rough rule of thumb, replacement work is two to three times as expensive as an initial greenfield installation.

Thus, replacement of even a small portion of an aged system can be a formidable financial burden to a utility. Going back to the 10,000 wood poles discussed above, the utility would follow up its test results by replacing nearly 1000 poles in order to obtain a failure rate over the next decade only slightly above that for all-new equipment. That may sound like a good bargain. But those replacements will each cost about three times as much as new poles cost when originally installed, and again, it is replacing more than two poles for every failure it will avoid, even though the test procedure is 95% accurate. And those tests, themselves, have a noticeable cost, as does the effort to correlate test results and manage the pro-active replacement in and around all the other activity the utility has keeping the lights on 8,760 hours a year. Altogether, the utility will spend roughly 5% of the cost of the original construction of the entire line for these 10,000 poles (adjusted for inflation, etc.) in order to keep failures and their

consequences under control for another decade: or about half a percent per year.

To put that half a percent capital expenditure in perspective, consider that a typical utility might add about 1% new equipment each year for system expansion – as its customer base and energy sales grow. Traditionally, the rates that it charges its customers, the budgets it puts together and manages to each year, the number and types of field crews and materials and equipment supply infrastructure it has, and its management bandwidth, too, are all geared to that level of capital spending and construction. This additional half a percent cost to pro-actively handle aging equipment is a *50% increase* above that traditional rate of capital spending on system equipment “additions.” To a company that may have only a single-digit profit margin, that finds it difficult to hire skilled field personnel and managers, this can appear to be an untenable challenge.

Against this, however, is the undeniable fact that O&M costs for repairs, restorations, and replacement for failures are and will continue to rise in the foreseeable future as the system continues to age, and that if the utility manages pro-active programs well, they will ultimately reduce cost as compared to that. The utility has no choice but to adapt with new testing, tracking, evaluation, analysis, planning, and management approaches to control its infrastructure aging.

In any practical sense this issue of aging infrastructures is new to the utility industry. During the last half of the 20th century, electric utilities in North America, Europe and in many other places around the globe developed and institutionalized incredibly efficient methods and processes to run a utility system while providing high levels of reliability ($\approx 99.98\%$), all at costs that regulators accepted as being as low as practically possible. Then, equipment that today is 60 years old and beginning to give problems was only 20 to 40 years old and still quite robust.

Since the problem of aging infrastructures was not a priority then, utilities and the power industry did not set themselves up to address it. Nor was it a focus for equipment manufacturers and service providers that support utilities: they go where the market leads them. Even only twenty years ago – half a career ago for senior managers and the most experienced engineers and operators at a utility – this problem did not exist in any measure serious enough to make it a priority. The challenge the industry now faces is to change this situation in time to effectively control aging and its impacts before the problems they cause become serious.

1.2 AGING INFRASTRUCTURE MANAGEMENT MUST BECOME A MAINSTREAM PART OF UTILITY OPERATIONS

In most cases it makes no sense to replace a system that was originally installed seventy years ago in a wholesale fashion. Even among the oldest equipment a majority may still be in serviceable condition and capable of providing many more years of good service. A utility’s best course is to do what it can to find the bad actors and the questionable equipment and replace only that, continuing into

the future each year with a system that is old but in which its *condition is being managed* to a level optimized to balance service reliability and cost. The system will be old, but condition will be “good enough.”

Among other things, this means aging infrastructures are here to stay. Any utility’s system will age, until it at some point it has become “old,” a term the authors will define here as meaning the system has a sufficient portion of equipment deteriorated badly enough to cause the utility concerns about failures, breakdowns, reliability, and cost. Once a utility starts managing the condition of that aging infrastructure well, it will replace or refurbish deteriorated equipment that is in unsatisfactory condition, but retain all the old equipment that is still in satisfactory condition. The system will continue to include a lot of old equipment, and that older equipment will always have the potential to go bad in the future. Further, that equipment kept in service will continue to age. The aging infrastructure condition management program will not necessarily reverse aging, just control condition to some (hopefully optimum economic) point. The system will not be new, but neither will it be *badly* deteriorated.

To do that the utility will need some combination of on-going testing, tracking, mitigation of continued deterioration, and pro-active replacement and refurbishment of deteriorated equipment, will be needed on a continuing basis into the future. Such programs will control equipment condition to an “optimum” level, a stable, *sustainable point* at which equipment in service continues to age, deteriorated equipment is replaced, and the utility continues to test, maintain, and service and replace equipment sparingly but in a targeted manner, with overall cost kept at a minimum.

Utilities that today have no processes and budgets in place to do this type of condition management will have develop those capabilities. Those new processes will be folded into the traditional planning and construction and O&M management processes and functions that the utility has had for decades. And condition management will be mainstream utility practice: It is likely that by 2035, a significant portion of electric utilities in North America, perhaps as many as 40%, will be spending slightly more capital each year on replacing aged and deteriorated equipment than on system expansion due to load growth.

1.3 WHAT IS THE INFRASTRUCTURE INVOLVED?

Even a power system serving only a few thousand customers is made up of tens of thousands of individual components – poles, ducts, underground vaults, substations, transformers, circuit breakers, voltage regulators, monitoring and control systems, service transformers, meters, transmission towers, and electrical conductors (wires). To many people this all-inclusive set of equipment is the infrastructure. But the comprehensive view that leads to the best solution to aging and the problems it causes says there is more. The *methods* of operation of the system are a big part of how the power system performs. For example, three

cornerstones of utility operation are:

- 1) Loading specifications – the guidelines the utility sets on the amount of electrical or mechanical load permitted on a unit – influence how much wear and tear it experiences, how much burden it is expected to carry, and how much its failure affects service reliability.
- 2) System operating guidelines – the procedures and priorities for operating the system have a lot to do with how failures affect cost, customer service, public safety, and the utility’s business performance.
- 3) Inspection, service and maintenance guidelines. Every utility has such guidelines in place along with processes that apply them.

Some if not all these guidelines, and the institutionalized habits built around them, are themselves quite old – “deteriorated” as compared to modern needs. Revision of loading and operating guidelines can reduce stress on older equipment and thus decrease the potential for failures. Modification of O&M rules and guidelines for equipment inspection and service can help find deterioration *before* it causes problems. There are ways to prepare and organize for contingencies so as to minimize customer inconvenience and utility cost, even when unexpected failures cannot be avoided. And, given that a utility will have to perform equipment testing, track condition, and perform pro-active refurbishment and replacement as part of a continuing condition management program, those activities should be merged with traditional maintenance and operating activities and all of it managed in the most efficient manner possible. Finally, the design and technology of the utility’s control systems is part of the infrastructure. Different control schemes can mitigate some of the reliability and cost effects of failures due to aging. New technologies can be retrofitted to detect insipient failure, to respond automatically to failure, or to further reduce reliability and cost impacts.

In many cases, these changes to specifications, to operating guidelines and procedures and to system configuration and design are more cost effective at reducing reliability problems and utility costs than efforts aimed directly at refurbishing or replacing equipment. In many more cases – a majority perhaps – they can affect the efficiency and quality of results from equipment testing, lifetime extension, and replacement programs. They, too, are a vital part of the “infrastructure” that should be managed. Thus:

An aging infrastructure includes all the power system’s equipment and facilities along with all processes, methods and guidelines and control systems used to design, operate, maintain, repair, and restore service and manage resources and system performance.

1.4 CHARACTERISTICS OF AN AGING INFRASTRUCTURE

Situations vary, but most aging infrastructure problem areas have most of the same characteristics, which are:

- The area was originally developed, or last experienced a rebuilding boom, prior to the 1970s. Most of the existing system structure (substation sites, rights of way) dates from before that time.
- The majority of equipment in the area is more than forty years old.
- The system is well engineered and planned, in the sense that a good deal of concern and attention has been devoted to it. It fully meets the utility's minimum engineering criteria.
- The area is seeing steady but perhaps slow growth in load. It may be undergoing, or about to undergo a rebuilding boom.
- The area is plagued by above-average equipment failure rates. Overtime labor may be high due to large amounts of unscheduled repair and restoration.
- SAIFI degradation leads SAIDI degradation by a few years. Typically, frequency of interruption begins rising from five to eight years prior to big increases in SAIDI.
- When a major "event" such as a widespread interruption of service occurs, it is sometimes due to a series of events that are quite unusual or even characterized as a bizarre combination of events.
- Things go bad a lot.

The next to last bullet above deserves special attention. In a badly deteriorated system, breakdowns and minor problems as well as outright failures have all increased in frequency compared to a system in satisfactory condition. This in itself leads to many operating problems. But it also means that the incidence of *simultaneous problems* increases drastically: *double* the chance that any particular relay in a system may not operate as expected, and the likelihood that two relays will not operate together when they should *quadruples*.

In a badly deteriorated system, events begin occurring more frequently due to combinations of breakdowns and failures in which the existing deteriorated condition of surrounding equipment or systems affected the consequences of or contributed to a failure and outage, often in bizarre combinations and ways that were never anticipated and will, at first glance, be dismissed as "once in a lifetime" occurrences that "could never happen again." Utility executives and

managers need to realize that such statements are specifically true but not true overall: *that particular* bizarre combination will very likely never happen again. But other equally un-anticipatable and bizarre events will occur, and more frequently as the system continues in a downward spiral of condition deterioration. This pattern is a sure sign of an aged and badly deteriorating infrastructure.

1.5 POWER DELIVERY SYSTEMS

The authors have deliberately used the term “Power Delivery” rather than “T&D” throughout this book in order to make a point that is both crucial to success in dealing with aging infrastructures, but very often overlooked. The infrastructure involved in the issues being discussed here includes a good deal of the T&D system, but not all of it. In order to understand the problem, and appreciate how to solve it effectively at the lowest possible cost, a person must appreciate that both problems and solutions must be examined from the standpoint of what does the equipment involved do?

Modern power T&D systems can be divided into two parts, or levels, based on the function of that part of the power system under a de-regulated structure. In terms of the new de-regulated industry paradigm these are most easily distinguished as the wholesale level and the retail level. *The point of this section, and the focus of this entire book, is that the problems of aging infrastructure, as they impact reliability of customer service, are almost exclusively limited to the retail level.*

This is not to say that the wholesale power transmission level is unimportant. It is absolutely critical, and problems at that level can cause customer service interruptions. What is most important at that level, however, is *system security*, the continued ability of the system to remain “up and running.” That is slightly different than reliability, and handled well by FERC (Federal Energy Regulatory Commission) and NERC (North American Electric Reliability Corporation) requirements. Programs and methods needed to comply with these requirements are composed of the same testing, tracking monitoring, evaluation analysis, and planning techniques covered in this book. However, the priorities and focus are generated by regulatory requirements and regional grid considerations. Those aspects are not dealt with here.

The authors recognize that the consumer does not care about whether his lights are out because the wholesale grid cannot supply all the power needed regionally, or because a part of the local power delivery chain has failed – he is without electric power. But this book is about understanding and fixing the problems caused by aging power delivery infrastructures and eradicating the customer service quality, operating and business problems they cause to the local delivery utility. In that context, the distinction between wholesale and retail matters a




Figure 1.1 The electric power industry, including generation (now de-regulated), the high-voltage transmission network (open-access), and power delivery, which consists of sub-transmission, substations, feeders and laterals, and the utilization voltage system leading to all customers.

Table 1.1 Seven Characteristics of Distribution (Power Delivery) Systems as Opposed to Wholesale Transmission (FERC)

1.	Local distribution is normally in close proximity to retail customers.
2.	Local distribution system is primarily radial in character.
3.	Power flows into local distribution. It rarely, if ever, flows out.
4.	Power on the local distribution system is not re-consigned or transported to another market.
5.	Power on the local distribution system is consumed in a comparatively restricted geographic area.
6.	Meters at the interface from transmission to distribution measure the flow onto the distribution system.
7.	Local distribution is of reduced voltage compared to the transmission grid.

great deal. The record shows that the vast majority of customer service quality issues seen throughout the industry are created at the local delivery, or retail, level and that the solution to those problems lies there as well. Problems at the wholesale level are generation shortages and transmission congestion. These are very serious in their own way, but an entirely different set of concerns, and are not addressed here.

Delivering Power to Mrs. Watts

Figure 1.1 shows a typical power system from the generation plants to the home of a customer, Mrs. Watts. The power Mrs. Watts uses to heat her home this week came through the wholesale grid from any number of power generation plants and not always from the same ones. But it flows into the same local distribution system all the time. And except during emergencies or contingencies, always takes the same path to her home. This “power delivery system” is owned and operated by the local utility company.

Power Delivery: Getting Power to People

In contrast to the wholesale level, the retail, or power delivery, level exists *solely* to route power to consumers. Table 1.1 lists criteria that FERC (Federal Energy Regulatory Commission) established to identify the local delivery level. Most important to this discussion is a point mentioned earlier: a particular consumer’s power *always* flows through the same part of the delivery system, except on those rare occasions when there is a contingency. The power that Mrs. Watts uses each day is intended to *always* come through the same set of service drops, through the same service transformer, lateral branch, feeder trunk, substation switchgear, substation transformer, high-side switchgear, and substation transmission-level

feed line. Only during an equipment failure will an alternate route be used for some part of this supply chain.

Power delivery infrastructures are those portions of the power system whose purpose is solely to deliver power to energy consumers, and do *not* include those elements associated with the wholesale electric power grid.

Power delivery includes more than just local distribution, although it certainly includes that. It also includes a portion of the transmission-level system in what was traditionally called the “sub-transmission lines – those lines that deliver power to distribution substations. Figure 1.1 indicated the portion of the power system identified as the power delivery system. This is normally the transmission that despite its voltage level, falls under FERC’s distinction of distribution function, as opposed to wholesale transmission (Table 1.1).

1.6 THE BUSINESS ENVIRONMENT HAS CHANGED

If the changes above were not enough, the technical, social and political environment in which utilities must function has changed dramatically in the last thirty years. In the 1970s electronic calculators were still in their infancy; digital computers were large mainframes with far less computing power than today’s average laptop computer, and not yet widely deployed, and automated monitoring and system communications for power systems was used only in major generating plants. Similarly, at home most consumers did not have any appliances or other products that relied heavily on electronics other than perhaps their radio, television or stereo system – all analog devices.

The transistor had been invented in 1947. Discrete component technology found its way into consumer products in the 1960s with the introduction of the “Transistor Radio.” However, the proliferation of electronics in consumer mass-market products did not begin to occur on a large scale until the development of the integrated circuit. Intel introduced the 8008™ microprocessor in 1972. The capability to support computational complexity and the sensitivity of products to disturbances on the electric power supply system have been increasing at accelerating rates ever since. Today one can find integrated circuit technology throughout the home. Televisions, clock radios, home video game systems, desk top computers, programmable microwaves, programmable ovens, programmable dish washing machines, programmable VCRs and an increasing number of other appliances incorporating programmable controllers are becoming commonplace.

Office technology has seen a similar evolution to greater complexity and sensitivity. Mechanical adding machines and calculators were replaced first with electronic calculators and more recently desktop computers. Mechanical typewriters were replaced with electric models, then the stand-alone word processor. Now word processing programs run on the same desktop computers

used for accounting, spreadsheet analysis and computationally intensive engineering calculations or graphic design. Many businesses own or lease their own telephone switching systems, another application of microprocessor technology.

Likewise, production and manufacturing facilities have evolved. Manual control of machinery gave way to analog controls. Analog controls have given way to numerical control, first to punched tape, then to process control computers. Electronics have made possible variable-speed AC motor drives. Efforts to increase productivity and efficiency in the factory have introduced many of the same systems used in the modern office and factory, such as desktop computers and customer-owned telephone systems. Additionally, the internal power distribution systems of the plants have become more demanding as they provided faster protection of valuable equipment. Today's microprocessor controlled relays, motor control systems, and recording and monitoring systems are much more sensitive than the electromechanical relays, analog instrumentation and motor control centers of 50 years ago.

While the microprocessor revolution was taking place in homes, offices and factories, little was changing in the construction of the power delivery system. Granted, solid-state relays were developed and deployed, improving the security of the bulk power system. The computer revolution also improved the ability of the utility engineer to model and analyze the operation of the power delivery system further optimizing technical aspects of system design in accord with the planning and design parameters developed in the 50s and 60s. Advances in measurement technology based on the microprocessor also allowed technicians to monitor the operation of the system in near real time. Digital event recorders capture the transient operation of the system for later analysis. With the right instruments and communication channels, one can monitor various attributes of electric service for many points on the power delivery system.

However, the system itself has not changed in any significant way. Although there have been improvements in the quality of the material and equipment, the system itself remains essentially the same. The power delivery system is composed of poles, wires, insulated cables, switches, fuses and transformers. It is subject to damage inflicted by man and nature in addition to the random and rather infrequent failure of the components. Finally, the way in which the system reacts when subjected to a tree limb falling on an overhead wire or a back hoe digging into a cable is the same today as it was 50 years ago. A fuse or a breaker or other protective device opens to isolate the damaged segment of the system. In the time required for protective equipment to operate, high fault currents flow, depressing the voltage for customers and customers on other circuits served from the same source until the protective device operates. In many systems, fast acting reclosers or circuit breaker relay schemes open and reclose before the protective device can operate in an attempt to clear temporary faults, subjecting all customers on the affected circuit to a series of momentary interruptions.

The end result of the mismatch between the sensitivity of the equipment used by customers and the operation of the delivery system has been a steady decline of the customer's perception of system performance. Residential customers may arrive home and find all of the electronic clocks blinking or with the words "power failure" showing on the display. The initiating event may have been a momentary interruption due to a tree branch brushing against an overhead conductor, or just a voltage sag due to a short circuit on another circuit out of the same substation. Fifty years ago with electric clocks driven by small motors, the event would have gone unnoticed by anyone not at home – at most they would have noticed was a blink or a momentary dimming of the lights.

The situation was similar in the industrial and commercial sectors. In the 1980s it was not unusual to encounter problems associated with the application of new technology that were internal to the facility but nonetheless blamed on the utility. In one instance a manufacturer had installed microprocessor-controlled robots on an arc-welding line. The controls were supplied from the same circuit in the plant that supplied the welding equipment. It was relatively easy to demonstrate that the controls worked fine when the welders did not strike an arc but caused control problems when they did. In another instance a plastics injection molding plant replaced all of the electromechanical voltage and current protection devices with new solid state devices. Because he was able to do so, the plant electrician had set them at maximum sensitivity and zero time delay. The end result was that they would trip the supply to the compressor motors any time there was a switching operation on the supply line, despite being only several hundred yards from the substation. The problem was resolved by reducing the sensitivity of the protective devices to the same level of those they replaced.

But most of today's power quality problems are not as easy to solve. Sensitive plant control equipment may be built into the system that it is controlling, and some equipment cannot tolerate even a momentary power interruption. Sensitivity of equipment throughout a plant may be such that even routine voltage sags are intolerable. A single momentary or voltage sag may cause the loss of hundreds of thousands of dollars. Solutions to these problems have been available at utilization voltages for some time. There are un-interruptible power supplies (UPS), motor generator sets and constant voltage transformers to name just a few. Yet, it may be that there is not enough space available for the manufacturer to install the necessary mitigation equipment. Installation of low voltage mitigation equipment at multiple locations throughout the plant may not be cost effective.

It is in this context that the aging power delivery infrastructure must be viewed. Not only does system performance suffer as equipment condition deteriorates, but today's customers have become much less tolerant of any disturbance in their power system. The potentially poorer service quality in combination with rising customer expectations has put the power delivery companies in the spotlight.

1.7 FIVE FACTORS CONTRIBUTE TO AGING INFRASTRUCTURE PROBLEMS

Throughout the United States, and in many other places around the world, portions of electric utility systems have become aged: in the sense as described above. The average age of equipment in the region or area of the system exceeds the design lifetime for electrical equipment. Figure 1.2 shows the distribution of average ages of substation power transformers throughout the power industry, estimated by the authors. The mean age is about 40 years, meaning that the average transformer is roughly $\frac{3}{4}$ of the way through its normal lifetime. The figure indicates that more than a few transformers have been in service over 60 years.

The fact that a transformer, circuit breaker, or other piece of electrical equipment is old does not mean that can no longer do its job. Old equipment will often still function, and although not as efficiently as modern units, nonetheless get the job done. But what age does mean is that the unit is much more likely to fail in the next year, as compared to a new unit of similar size and design. As a result, a system composed of old equipment has a much higher incidence of equipment failures. Equipment failures often lead to interruptions of service to customers, and that means the system provides relatively poor service – at least poor service as compared to a system that is composed of newer equipment.




Figure 1.2 Estimated percent of distribution substation transformers over a certain age in the US. Roughly half are probably over 40 years in service.

Table 1.2 Five Contributing Factors to Aging Infrastructure Problems

Factor	% of Problem	Description
Aging power equipment	20% - 40%	Older equipment has higher failure rates, leading to higher customer interruption rates. It creates higher inspection maintenance costs and much more unexpected O&M costs due to repair and restoration work it creates. It creates uncertainty and unpredictability in both operating and business results.
Obsolete system layouts	25% - 30%	Older areas of the system often are in bad need of additional room for substation sites and rights of way that cannot be obtained, and thus the utility must “make do” with only that obtained decades ago. The result is a system where peak demand and reliability needs have outgrown the basic designed infrastructure’s ability to support it well.
Outdated engineering	12% - 25%	Traditional tools for power delivery planning and engineering are ineffective in recognizing and correcting some of the problems that result for use of aged equipment, obsolete system layouts, and modern deregulated loading levels.
Non-optimum distribution	5% - 15%	The utility industry’s technology, both in the field for system equipment and maintenance, and in the office for engineering and back-office support, has advanced by decades since many power systems were designed. Yet they continue with the same basic layout and design originally conceived, even though these designs are not constrained by conditions as are the obsolete system layouts (Factor 2, above). Revision of layout and use of smart equipment (see Section 2.8) can be used to provide some “immunity” to the negative effects of aging equipment.
Old cultural values	5% - 50%	Planning, engineering, and operation of systems using concepts and procedures that worked fine in the vertically integrated industry that exacerbate problems when applied under deregulation.

This perspective provides the easiest way to define and identify an aging power delivery infrastructure – simply look at the age of its equipment. But that provides only a partial view of the problem, and given that condition management seeks to keep old equipment but old good old equipment, a view incompatible with the condition management approach. What this perspective that age is the key factor misses is that four other factors common to older parts of the utility infrastructure heavily interact with equipment age to exacerbate its effects and contribute to the overall problems of poor customer reliability and high operating costs. Table 1.2 summarizes these five interrelated, contributing factors, which are all at work in an aging power delivery infrastructure. Together, they work against efficient, reliable operation of the older power system. These are discussed in more detail below, and in much greater detail elsewhere in this book.

Contributing Factor 1: Aging Equipment

As cited above, old equipment is a serious contributing factor to poor reliability and high costs in many utility systems. In fact, there is usually some area in any electric distribution utility's system where old equipment is a worrisome factor in providing good customer service quality.

All power system equipment will eventually fail. At some point, after a period of service, something will wear out or deteriorate to the point that it no longer adequately performs its role. The equipment will cease to provide electrical power delivery to those customers downstream of it. The nature of electrical equipment is that very often failures lead to a fault, requiring protective equipment action that isolates a larger portion of the system than just the customers served by that failed unit.

Regardless, electrical equipment does not have a uniform probability of failure with age. Usually electrical equipment fails only after many years of completely satisfactory service. Thus, old equipment is much more likely to fail than new equipment (roughly three to ten times as likely). As a result, a system or area of a system composed of old equipment will have a much higher failure rate than one composed of new equipment. This has three consequences, all of which are undesirable:

- 1a. Equipment failure* causes lower customer service quality. Customers are more likely to have interruptions of their electrical service due to the unexpected failures of this older equipment. The frequency of interruption rises – SAIFI increases.¹
- 1b. Higher repair and restoration efforts.* More failures mean more and more unpredicted events requiring emergency or contingency

¹ SAIFI, System Average Interruption Frequency Index, the average number of interruptions of service experienced annually per customer. For a discussion of this and other reliability-related indices, see Chapter 4.

mitigation efforts. The impact of this on all but the most efficiently managed and staffed organizations is to increase duration of equipment outages (and hence customer interruptions). There are more numerous outages to repair, taxing the restoration and repair resources to the limit and often requiring repair of one to wait on another. SAIDI increases slightly as a result.

Also quite important for the utility's bottom line, these higher repair and restoration needs increase overall costs.

- 1c. More inspection and testing efforts required.* Equipment more prone to fail often benefits from more frequent and comprehensive inspection and testing, which in many cases can reveal developing problems before they cause failures. Catching trouble early allows what would have become unscheduled emergency repair work to be done in a more efficient, less costly scheduled manner. This reduces the utility's costs, but the inspection and testing increases required increased costs. Compared to newer systems, costs are higher.

Chapters 6 and 7 will focus on various aspects of equipment aging, its management, and the mitigation of its effects in more detail. Usually, a utility cannot afford to replace all of its aging equipment. As a result, the solution to this particular problem source is a combination of partial measures that seek only to reduce failure rate slightly and control its effects:

- a) Careful assessment and tracking of condition,
- b) Good utilization of capacity and expected reliability of equipment,
- c) Artful management of restoration and replacement policies,
- d) Design of the system for failure – planning that assumes that things will go wrong.

Chapters 10 and 11 will discuss these in more detail.

Contributing Factor 2: Obsolete System Layouts

This, not aging equipment *per se*, is perhaps the strongest contributor to poor reliability in a sizeable portion of aging infrastructure cases. This is often not recognized or identified as an aging issue, but the facts speak for themselves.

Usually, the design of the power system in an area of a power system, its layout of substations and rights of way and the configuration of how they are interconnected and supported by one another during contingencies, dates from the time the equipment now in service was installed, or before. For example: the central part of the power system in many urban areas of the U.S. The “downtown power system” in most large cities is based on substation sites and rights of way that were laid out in the 1940s through the 1960s, when the post WWII boom

created a period of sustained economic growth. These designs were executed with great skill to meet both the needs of the moment and reasonable amounts of load growth two or more decades into the future.

But between four and six (not just two) decades have passed since these systems were first laid out and built. In many urban areas throughout the United States, the electric demand has doubled or tripled since the power system last had a major revision in *layout*. By default, these systems had their layout “locked in.” As the city continued to grow, the utility found it increasingly difficult to obtain new substation sites or rights of way in the crowded, growing downtown area. Both economic and political pressures stood in the way of expanding its existing sites or adding new ones. Adding capacity at existing sites instead accommodated load growth. The utility often had to compromise preferred configuration (switchgear type and interconnection) and/or maintainability (quick accessibility) in order to squeeze every last MVA of capacity possible into crowded spaces.

At some point, equipment additions and renewals notwithstanding, the power system reached the limits of the existing sites and rights of way. Every bit of room had been used, every innovative trick within design standards had been exercised, and there was simply no more room. From thereon out, the utility either had to find some way to get the “unattainable” addition sites, or face the only other choice left: load the existing equipment in its sites and rights of way to higher levels than it would prefer.

In many cases, the rather extreme situations lead the utilities to overcome significant economic and political pressures and obtain some additional sites or space. But in all cases, the major manner of accommodation of load growth was by simply allowing loading on existing equipment to rise above traditional limits. A certain amount of this increase is by design, and does not necessarily have a negative impact. Since the mid 1980s, equipment utilization ratios (the ratio of peak loading to equipment capability) has risen through the power industry as electric distribution utilities sought to increase efficiency of equipment usage and cut back on expenses. It was not uncommon to see an electric distribution utility raise its target utilization ratio for transformers from 66%, about typical for the mid-1970s, to 70 – 83% during the last 15 years of the 20th century. This increase was by design.

But the increases in loading forced on aging systems due to a combination of obsolete system layouts and other factors (to be discussed below) went far beyond these limits. In the central core of many cities in the U.S., power delivery equipment such as sub-transmission cables, substation power transformers, and primary distribution circuits are all loaded to their maximum rating at time of peak demand, leaving next to nothing for contingencies or emergencies. In a few cases, loadings at peak average close 100% throughout the downtown area and exceed that at a few places. Several of the utilities that experience serious problems during the summer of 1999 had some major delivery equipment in their central systems that was loaded to over 110% of nameplate rating during peak

periods, not by accident, but by intention.

There is plenty of evidence that this trend is widespread. Figure 1.3 compares 1999 utilization ratios in the central core and the suburbs of five major metropolitan systems, all among the largest power delivery systems in the U.S. Not all urban utilities follow this pattern, but a majority seems to match this profile. This is not to say that the utility’s engineers preferred this higher loading level or that management and operators were entirely comfortable with the high loadings. But it was accepted, not as an emergency situation, but as the “normal” plan for serving peak demand.

Distribution Designs That Are Not Obsolete, But Are Not Entirely Modern

At the distribution feeder and service level there are similar issues of design and configuration that interact with aging infrastructure impacts. Traditional distribution systems are nearly all radial in operation and serial in their sectionalization. As a result individual equipment failures cause noticeable customer power interruptions, and events involving multiple failures (such as during storms or bizarre events) tend to cascade to large numbers of customers out of service.

Considerable improvement in a distribution systems SAIDI and SAIFI, even in the presence of a slowly rising tide of equipment breakdowns and failures, can be made by using a combination of smart equipment and changes in the circuit layouts and operating configuration. The margin cost of reliability for such changes to a system is often far less than a the cost of dealing directly with the aging equipment failures via a replacement program, and is often a first step in mitigating the continued effects of system and equipment deterioration.




Figure 1.3 Substation transformer utilization ratios in the central core and the suburbs of five major metropolitan areas. Dotted line shows the utility’s own design guidelines.

High Loadings: A Complicated Interaction with Design

These higher loading levels were not accepted without a good deal of preparation and often only with changes to the system to keep it, at least ostensibly, within engineering criteria. In all cases the authors are aware of the utility planning and engineering departments' applied intense engineering methods to make arrangements so that the system could, at least based on traditional analysis, tolerate equipment outages in spite of the potential reliability problems such high loadings could cause.

However, the full implications of higher loadings are difficult to evaluate fully with traditional power system engineering tools. As a result, some of the potential problems of accommodating load growth in aging system areas by using higher loading levels were not fully appreciated. (This is part of contributing cause number 3 and will be discussed later in this chapter.) While the systems in question met traditional reliability criteria, the systems failed to perform to the level the criteria were supposed to assure.

Configuration limitations

In addition, it is important to realize that the obsolete system presented the utility with challenges beyond just those caused by the higher loadings that had to result from the limited amounts of substation and rights of way space. The utility had to distribute a growing amount of power out of the same limited number of substation sites. This had detrimental effects on the performance and reliability of the primary feeder system. Beyond a certain point, a primary feeder system simply cannot be reinforced to deliver more power out of a limited number of sites, without accepting a certain amount of degradation in reliability and/or cost as a consequence of that limitation. This is due to a host of secondary, but important effects related to primary feeder configuration and performance that will be discussed in Chapters 7, 8, 11, and 12.

Symptoms of the obsolete system layout

The net result of obsolete infrastructure is that higher loading levels have three undesirable consequences:

- 2a. *Accelerated aging.* Time to failure for most electrical equipment is often to a large degree a function of its loading. The higher utilization rates forced on utilities in these areas of the system cause electrical equipment there, particularly cables and transformers, to "age" or deteriorate faster and fail sooner. Thus, in those areas where age and failure rate is already a serious concern, the obsolete system layout forces the utility to exacerbate those problems just to get by. Over time, this tends to increase the frequency of service interruptions (i.e., SAIFI goes up).

- 2b. *Limited contingency capability.* High loading levels make it very difficult to provide service to all customers when a unit has failed. To begin with, a system with high utilization has “a lot of eggs in each basket.” When a failure does occur, if service to consumers is to be maintained, a lot of load has to be re-distributed to other, still-functioning equipment nearby. But those remaining nearby units are already at high loading levels, meaning they have little available capacity to pick up additional load. The result is that interruptions, when they do occur, are difficult to restore, involving more load transfers, more emergency measures, and taking more time. While this does not increase the number of consumer interruptions to any large degree, it tends to increase the average duration of interruptions that do occur. At some point, SAIDI increases – the number of interruptions has gone up (due to age and the accelerated aging cited above) *and* now the length of interruptions goes up due to this reason. CAIDI (essentially their product) increases considerably.
- 2c. *Higher delivery operating costs.* The distribution systems in question operate at high loading levels. One consequence of this is the need in places for the utility to replace existing lines with a larger conductor simply to handle the MW-miles of delivery burden. Regardless, such systems typically have relatively high electrical losses. In addition, voltage drops are extreme. Tap changers and voltage regulators operate frequently, both fail more often and both require more preventive maintenance. These reasons and a host of similar secondary effects increase operating cost at the distribution level by 10%–20%.

In Table 1.2 and the discussion above, the authors have indicated that this particular cause among the four presented here typically contributes as much or more to a utility’s reliability and cost problems than any of the other three. This is because the obsolete layout interacts so much with the other aging issues. It exacerbates equipment aging. Solutions require an analysis that is very sensitive to the weaknesses of traditional engineering methods. And, on its own, the obsolete system layout creates operating and reliability challenges.

Both the problem and its solutions are complicated because they interact so strongly with both the other three causes of the aging infrastructure problem and with many other aspects of the power delivery utility’s operation. Generally, the solution is some combination of:

- a) Somehow obtaining a few new sites and ROW to improve the system layout,
- b) Use of innovative means of squeezing more capacity into existing sites, and most important,

- c) Optimized configuration of substation switching and buswork, and primary feeder-level switching, based on sound reliability-based engineering (to be discussed below).

Chapter 8 focuses on obsolete system layout and its effects in more detail.

Contributing Cause Number 3: Old Engineering Methods

Many T&D utilities are using engineering methods that worked well in the 1970s, but cannot fully guarantee reliability operation in a world where substation and line loadings are pushing beyond traditional levels. As a result, engineering planning functions at many utilities cannot identify precisely what the problems are or prescribe the most effective and economical remedies. This particular subsection will provide a particularly comprehensive summary of this subject. Engineering tools are in an area where utilities can make immediate improvements and one where improvement must be made if the aging infrastructure issue is to be properly addressed. Chapters 9, 12, and 13 provide more detail on this subject.

N-1 criterion

Traditionally, reliability was designed into a power system by making certain that there was a “backup” of sufficient capability for every major element of the power system. If that element failed, negative consequences of that failure were covered. This is called “N–1” planning, or use of the N–1 criterion. Any one element of the power system, composed of N elements, can be removed, and the system will still perform its job, even under the most stressful (peak) conditions. Simply stated, one bought reliability with capacity

N–1 engineering is done by running N separate studies of a power system, one in which each of the N elements is removed (outaged) and the system switched as appropriate to have other equipment pick up the load it was serving. Computer programs that use a load flow algorithm (the basic power system simulator which estimates voltage, current, and loading throughout a power system for a specified set of conditions) have been written to do this automatically. The power system engineer inputs a description of the system and its loads, and the program both solves the “normal” (everything working) load flow case and all N “single contingency” (one element out) cases, and reports all with any problems such as overloads, low voltages, poor power factor, etc.

More generally, this approach is called the “N–X” criterion – a power system of N elements can be designed to tolerate the loss of any X elements. This requires N^X cases to be run by the automatic analysis program.

The use of this N–1 approach is dogma throughout the power industry. Developed in the 1960s, it proved very successful: systems designed with it worked well. Every major utility developed the design of its high voltage transmission grid, sub-transmission and substation systems, and other critical

elements of its system, using this method. Nearly every utility modified the method in numerous small respects to fit its needs or preferences, but overall, this approach was universally applied. The reasons were:

- 1) The basic concept makes tremendous sense – design the system so that if anything goes wrong, nothing bad happens as a result.
- 2) Experience showed that this method worked: traditional power systems designed with this approach provided good reliability.
- 3) There were no alternatives, at least originally.

N-1's limitations

The N-1 approach is a sound engineering method, in fact one that the authors recommend as a part of any power system planning and engineering procedure. But like all engineering methods it makes assumptions, uses approximations, and takes short cuts to gain efficiency – in other words, it has its limits. There is no doubt (based on publications from the 1960s and 1970s) that the engineers who developed and first applied N-1 recognized these limitations. However, because the traditional applications of N-1 seldom encountered these limitations, and because institutional memories are short, many utility planners and engineers working today never considered that as conditions change, their systems might run into those limitations.

By any practical standpoint, N-1 capability is a *necessary requirement* for power system reliability. To be reliable, a power system must have the ability to tolerate the loss of any one element of the system and still do its job. However, that does not make it a *necessary and sufficient* criterion. One that, if met, assures adequate reliability of service, which is how the industry often has applied the criterion and methods based upon it: design a system to meet N-1 and it was by definition “reliable enough.” The fact that a power system meets N-1 criterion is not sufficient to assure that it will provide satisfactory reliability, or even fairly good reliability. As proof, it is only necessary to look at the system outages that occurred in summer of 1999. All of the systems that experienced severe reliability problems fully met N-1 and many cases met N-2, and even N-3 design criteria in critical places.

N-1's “ragged edge” as a dependable planning tool revolves around its lack of probabilistic expectation in its computations. N-1 methods and the N-1 criteria assure power system planners and engineers that there is *some feasible way* to back up every unit in the system, should it fail. However, they make no assessment of how likely the system is to get into situations that go beyond that. N-1 methods do not provide any analysis of:

- How likely is it that backup will be needed?
- How reasonable is the feasible plan for each contingency situation?

- Is the planner is actually building a “house of cards” by expecting “too many things to go right” once one thing has gone wrong?
- How much stress might the system be under during such contingency situations, and what are the long-term implications for both equipment life and operating feasibility of the system?
- How often do conditions occur which cannot be backed up (e.g., multiple failures) and how bad could the situation become when that is the case?

As a result, systems that meet the N–1 criteria can be less reliable than needed engineers might expect, even though the N–1 criteria “guarantees” there is a way to back up every unit in the system. This is much more likely to happen in modern power systems than it was in traditional, regulated power systems, due to changes in utilization and design made in the period 1990–2000, and the more volatile operating environment of de-regulation.

The reason is that modern power systems are much “leaner” than those of thirty or forty years ago, when N–1 was first developed and applied. In the 1960s through the early 1980s, electric utilities typically loaded key equipment, such as substation power transformers and downtown sub-transmission cables, to only about 2/3 or a little more (typically about 66%) of their capacity, even during peak periods. The remaining capacity was kept as “operating reserve” or “contingency margin.” Engineers and planners at distribution utilities designed their power systems using the N–1 and other criteria, while counting on this margin. In this environment, N–1 was both necessary and sufficient to assure good reliability.

Beginning in the 1980s and increasingly throughout the 1990s, utilities pushed equipment utilization upwards to where, in some systems, the average substation transformer was loaded to 83% of its rating during peak periods. As will be discussed in Chapter 8, in aging areas of these systems, which suffered from “obsolete system layout,” utilization rates often averaged close to 100% under “normal” peak conditions.

Whenever a power delivery utility increased its planned equipment utilization rates, changes were carefully engineered into the system, so that it could fully meet the N–1 criteria everywhere, and even attain N–2 capability in critical places. This despite the greatly reduced contingency margin levels now available: *there were feasible ways to back up any and all failures, despite the higher utilization rates.*

However, as will be discussed in Chapter 9, this fell far short of providing traditional levels of reliability. N–1 is not sufficient to assure traditional levels of reliability, when the system is operated at high utilization rates. Basically the reasons are that such systems are much more likely to experience “high stress events” to the point that situations that go beyond the N–1 backup capability and are far more likely to occur. Such systems may have from three to five times the

potential for customer service quality problems as traditional systems. Coupled with the higher failure rates caused by aging equipment, this creates a serious service quality handicap for the utility.

Other Engineering and Planning Related Problems

Power systems that operate with either an obsolete system layout and/or at high utilization rates are sensitive to more than just probability-related increases in interruption rates. The sensitivity of the system service quality to other problems increases exponentially. In particular, problems are encountered with partial failures, configuration complexity, and load forecast sensitivity, which are discussed in further detail in the following paragraphs.

Partial failures

Contingency-based methods such as N-1 are basically “zero-one” models of failure: a unit of equipment is modeled as completely in, or completely out of service. But modern power systems often encounter partial failures or operating constraints of some components:

- A transformer may be in service but its tap changer has been diagnosed as problematic and is locked in one position.
- An oil-filled UG cable’s pumps are disabled and the cable has been de-rated.
- A bus ground that failed tests has dictated opening a tiebreaker to balance fault duties.

These and other similar problems impact an area of the system with obsolete system layout or an area of the system operating at high utilization ratios, far more than they impact a traditional (66% – 70% utilization) power system. The obsolete layout is stressed with respect to voltages, flows, etc., and limitations in its full equipment usage due to such partial failures add stress and often “uses up” contingency capability. A power transformer loaded to 85% at peak, whose tap changer is locked into one position, is subject to voltage regulation problems that can easily reduce its ability to handle load by more than 15%. The contingency margin (85%-100%) that the typical N-1 method assumes is there, is in fact gone.

Configuration complexity

At high utilization levels a system needs additional switching flexibility to accommodate its need to back up equipment operating at high loading levels with other equipment that is also already operating at high loading levels. Traditional N-1 engineering methods are capable of analyzing the additional complexity this additional switching adds to the system design, at least from the standpoints of

both electrical flows and feasibility evaluation of contingency operation. However, configuration decisions at substations (i.e., ring-bus vs. breaker-and-a-half vs. double bus design) and feeder layout and branching/switching design (see Chapter 10) affect reliability – how likely and how often – failures or outages will lead to interruptions. When configuration becomes more complex, as it does in most cases where adequate levels of reliability are being sought when utilization rates are high, the interaction of configuration and reliability is non-trivial and not intuitively obvious. Candidate plans that look like they will provide an improvement can, in fact be less reliable overall.

N-1 methods cannot assess, nor provide good feedback to engineers on the interaction of configuration with actual customer-level reliability. For this reason they leave a gap in the engineer-planner's understanding of the system being designed. Generally, systems laid out only with N-1 methods have configurations, which will, providing feasible contingency capability, offer far less operating flexibility than really necessary for reliable operation.

Load Forecast Sensitivity

A system that is under a good deal of stress, either because it has a system layout that is less than ideal for its load levels, or because equipment is being operated at high utilization rates, is particularly sensitive to errors in the load forecast used in its planning. Poor load forecasts used in the planning or operation of a power delivery system effectively “use up” its contingency capability [Willis and Powell, 1985]. This will be discussed in more detail in Section 9.2. Poor normalization of weather data for forecasting or poor spatial forecasting (poor correlation of loads with areas and equipment) result in deterioration of a systems contingency to withstand capability. This greatly exacerbates the reliability-of-service problems discussed up to this point in this chapter.

Contributed Cause 4: Non-Optimum Use Of The Distribution Level

Averaged over the entire power industry, the primary distribution system is without a doubt both the most important and least well-used part of the power delivery chain, for several reasons. First, it accomplishes more of the power delivery function than any other level of the system, in the sense that power is routed into a few locations in fairly large chunks. Then the distribution system routes and sub-divides that power into very small allotments delivered to many different locations. No other layer (See Figure 2.3) of the power system accomplishes so much dispersion over area or so much sub-division of power.

Secondly, the primary distribution level has more impact on customer power quality than any other layer. In the vast majority of power systems, it is both immediately downstream of the lowest voltage regulation equipment (substation voltage regulators or load-tap changers) in the system, and by the nature of its

topography, the source of the majority of voltage drop from that regulation point to the customer. That topography also guarantees that the primary distribution system sees the majority of outages that lead to customer service interruptions. Thus, in most systems, the distribution system is the source of the majority of both the voltage-related and availability-related power quality problems seen by energy consumers.

Modern Capabilities Are Not Well Used

At many power distribution utilities, the primary distribution system is planned and engineered using principles and layout concepts that were developed in the 1920s and 1930s and that have been handed down through generations of planners and engineers without substantial revision or re-examination. Foremost among these long-time “habits” are the overall layout style of the primary feeder system: whether the large trunk, multi-branch of the hierarchical subdivision, often-called “feathered” template for radial feeder layout, guides that. Additional factors include switching zone and contingency backup policy, voltage regulation and “reach” design, and conductor set selection and tapering policy.

Any reasonable layout style, conductor size set, switching zone, and conductor sizing rule base, will result in a distribution system that does its basic job as interpreted from the traditional perspective: route power from substations to customers without violation of voltage, loading, or power factor standards. However, this interpretation of the distribution system’s job is obsolete, and incomplete. It looks at the distribution function only from the perspective of moving power and neglects its role in and capability for improvement of customer-service reliability.

Very few primary distribution systems are planned and engineered to make maximum use of the distribution system as a resource for customer reliability. As mentioned above, the guiding rules for primary distribution planning and design at most utilities were developed in the 1930s, and last seriously re-examined and revised in the late 1960s. Additionally, they are applied in a “cookie cutter” manner, using rules-of-thumb, informal guidelines for the larger issues of layout, switch planning (along with copious formal guidelines about the details), and rote rules for overall design. In a majority of cases, these rules are often too conservative, not allowing the full use of distribution as a resource.

“Hidden” Capability May Thus Be An Untapped Resource

As a result, distribution systems in most utilities are underutilized in two ways:

1. Economics – when viewed just from their standpoint of their traditional function moving and dispersing power – they cost more than is necessary.

- 2. Reliability – the potential of the primary system to maximally contribute to improved reliability is partially untapped in the vast majority of cases.

Table 1.3 gives statistics on the results of “distribution optimization” studies from nine utilities, with which the authors have recently worked. The first six cases involve only changes in configuration, etc. Cases 7 - 9 include changes to circuit configuration and layout guidelines and the use of smart equipment used in combination. These cases are representative of all the utilities the authors have worked with over the past twenty years. They are an average set, not a set picked to show particularly good results. They demonstrate both the margin of improvement that is possible by updating and modernizing primary distribution planning and design paradigms, and the variation that is seen among different utilities.

Table 1.3 lists four statistics for each system. The first, “Electrical - \$,” is the reduction in overall (lifetime) cost of distribution that could be wrought with respect to the traditional primary distribution paradigm of moving power. This is the savings in *new* capital additions that result from improving how the distribution system is designed to do its traditional job – moving power, without the use of appropriate “smart” equipment and using smart equipment.

The second statistic, “Existing - \$” shows how much the useful utilization of an *existing* distribution system can be improved with respect to that traditional paradigm. Existing systems encompass both the good and bad points of their past engineering and design, and those systems cannot be thrown away and re-designed based on new and improved rules. However, as shown, some improvement can be wrought in the MW/\$ capability of these systems. However,

Table 1.3 Improvements in Distribution System Cost and Performance Derived from Revision of Traditional Design Guidelines and Engineering Methods

Case	Utility System Type	Electrical \$	Existing \$	Reliability \$	System Reli.
1	Urban UG, Atlantic coast	8%	2%	79%	3.4:1
2	Rural, small towns, Central US	12%	0%	25%	.8:1
3	Metropolitan, Midwest US	6%	3%	40%	1.8:1
4	Entire system, IOU, Central US	2%	3%	36%	1.5:1
5	Rural, small towns, Southern US	12%	11%	25%	1.6:1
6	Entire system, Northeast US	22%	16%	33%	.92:1
7	Metro suburban system, US	3%	5%	62%	1.8:1
8	Metropolitan system, Central US	7%	15%	42%	1.5:1
9	Statewide system, Western US	4%	9%	34%	1.4:1

improvement is only useful where needed (i.e., if one improves a 5 MVA feeder so it can carry 5.4 MVA, that is useless unless it is in an area where one *needs* .4 MVA more capability). These values reflect both what is possible, and what was found useful, on the respective systems.

The third statistic, “Reliability - \$,” represents the reduction in the cost of improving reliability that was attained by revising design guidelines and applying the most effective reliability-based planning and engineering methods to these systems. These improvements while rather dramatic, are typical. As stated earlier, distribution systems are simply not designed from a reliability standpoint. When this is done and the engineering is optimized, the result is a very considerable improvement.

The final statistic, “System Reli.,” represents the margin in “bang for the buck” that the revised distribution reliability improvement made over the existing cost of improving reliability that was in effect before the revision. This statistic goes to the heart of the use of distribution as a resource in aging infrastructure areas, and of the authors’ overall theme of using all of the power system chain optimally to obtain the greatest “bang for the buck.” In four of the six cases listed, reliability improvement was between 20% and 40% less expensive to buy at the primary level than at other levels of the system. In these cases, spending on the distribution system could deliver reliability improvement for less cost than spending on other levels of the system.

In two other cases, this was not the case, and the value shown is negative, which means that their optimum business case was to spend more money for more results, to the extent marginal cost dropped. Of note here is that the utilities were unaware of this entire issue, because the traditional tools they were using for the design of all three levels were not capable of directing engineering reliability or of even measuring such things as reliability improvement per dollar. Thus, in all cases, improvements of between 9% (1/ 92 – see final column) and 70% were possible by adopting reliability-based engineering methods and using them at the distribution system and for coordination of design among the distribution and sub-transmission-substation levels.

Table 1.3 shows that revision of distribution system guidelines and adoption of reliability-based engineering methods should be viewed as a valuable improvement in terms of the traditional (electrical) perspective on distribution. The average improvements in cost effectiveness of around 8% are noticeable and valuable.

However, from what might be termed a “modern perspective” – one that views the system’s role as one of both performing the electrical duties and providing reliability of power availability, the improvement represents a quantum leap in cost effectiveness. By looking at the distribution system as a reliability resource, and optimizing its use for that purpose, an average 40% increase in cost effectiveness of reliability improvement is obtained at the distribution level. Coordination of that with the planning and design of other levels of the system

achieve a 70% increase in overall effectiveness of money spent on reliability improvement.

The authors are aware that many readers, particularly some with long experience as planners at power delivery utilities, will dispute that these improvements shown are theoretical and were not implementable (they were), or that that their system is different (not likely), or that “it won’t work here” (it will). Chapter 10 will go into more detail on these improvements.

It is rocket science

The rather significant improvement in effectiveness that is possible by revising guidelines and adopting newer engineering methods at the primary distribution level brings up a key point about why the industry clung to what are essentially outdated engineering practices for so long. Distribution planning and engineering at most utilities is kept at a rather low level.

Engineering of a distribution system so that it performs the basic electrical functions that were traditionally required of it is rather straightforward. It can be done and was done for decades with nothing more than rough calculations (slide rules at best), nomographs and tables, rote-memorized rules of thumb, and institutionalized memory informally handed down from one generation of engineers to another. When digital computers are used, that level of engineering becomes quite simple. For this reason, one often hears terms such as “it isn’t rocket science” or “distribution planning and design are quite simple.” This cultural paradigm is exacerbated by the following facts: 1) Distribution is of relatively low voltage (compared to other parts of the power system; 2) the equipment involved is rather small and inexpensive (compared to other levels) and; 3) the system is normally laid out in radial form (which makes power flow easy to calculate).

In fact, the power distribution system is by far the most difficult of all layers in the power system to plan and engineer to a truly high standard. It is rocket science, because distribution systems are like piranha. Yes, all the equipment is low voltage, all of it is small, and each unit is some type of very common, “commodity” design.

But a distribution system consists of tens or hundreds or thousands of thousands of these components. In aggregate it is large, expensive, and most importantly, *very* complicated. A distribution system represents a tremendous challenge in *combinatorial decision-making* – in deciding how to select and arrange equipment, line routes and capacity tapering, interconnection, normal and contingency flow patterns, and a myriad of other factors to best achieve the required goals.

It is possible to select and arrange this equipment based on standardized layout rules, traditional (and rather straightforward) guidelines on equipment and engineering methods, so that it works well enough from an electrical standpoint.

Beyond this, distribution systems present a number of other challenges which ultimately create a large gap between the quality of a plan that works, and that of a really good plan. First, the distribution system has to handle the vast majority (roughly 95%) of the load coincidence (diversity of peak) dynamics on the system. Secondly, distribution is closer to the consumer than any other part of the power system: it has more than its share of constraints on routing, equipment placement, and design. Third, distribution must deal with a good deal of 'fuzziness' in terms of the demand levels it serves. Unlike at the transmission level, accurate load measurements are few and far between. Usually loads are estimated from metered, kWh data, and assumed values of diversity.

For all of these reasons, getting the most value possible out of an existing distribution system, or from new investment about to be made in it, is more difficult than in any other level of a power systems. Distribution planning, done to the most effective manner possible, is rocket science.

Contributing Cause 5: Old Cultural Identities and Ideas

Section 1.1 stressed the changes that aging power delivery infrastructures will require of utilities, their organizations, and their personnel. Simply put, the problem did not exist to any noticeable degree in past decades, and methods to address it were not part of the utility's operations or culture.

Old habits die hard. At some utilities, the managers, planners, engineers and operators still tend to think in terms of the culture and goals that were appropriate for a purely regulated, vertically integrated industry. New goals may have been articulated, but for various reasons a change in thinking, in priorities, and in method has not followed, one which makes the planning, engineering, and operating compatible with those new goals. The use of the out-dated engineering methods cited earlier could be characterized as part of this problem. Particularly, the considerable resistance to new ideas that prevents many utilities from fully adopting new reliability approaches (Contributed Cause 3, above) and distribution planning practices and guidelines (Contributing Cause 4, above) is mostly due to cultural resistance to a change in institutionalized values and methods. However, the authors have identified those as separate issues – methodology, rather than culture -- in order to focus attention on broader issues. In a number of cases, perhaps half of the aging infrastructure problems the authors have analyzed – retention of outmoded priorities, corporate values, and concepts - has been a contributing factor in the problems the utilities experience with aging infrastructures. This is often not a result of resisting change as much as one of recognizing when, how and where to change.

Executive management's perspective: business-driven

Since the power industry was de-regulated in the mid 1990s, the 21st century power companies and distribution utilities that fell out of the dis-aggregation of

traditional vertically integrated utilities, began taking a more “business driven” approach to their decision-making. Investments and operating policies were justified on the basis of their business case: if and how much they helped the bottom line of the business. This perspective is a huge departure from the values traditionally used at many utilities. Executives at many modern utilities drive all priorities into their organization starting from their financial model. “What will it take for us to achieve our financial goals?” Budgets and priorities are set by that requirement.

While electric utilities always recognized they were businesses, many of the values they used in decision-making were based on a “public stewardship” concept that does the right thing, regardless of cost. This attitude made sense within the traditionally regulated utility business model. Any expenditure justified within that framework could be put into the rate base. Ultimately the utility would recover its expenditures and earn a reasonable return on invested capital.

Deregulation, and shifts in attitudes within the industry and regulatory agencies, has changed that perspective. No longer are capital expenditures and operating expenses acceptable to a utility’s management just because projected revenues will cover all projected future costs. Executives insist that spending be contained for a variety of business-related reasons. Regulators are concerned about price as much as quality. The result is a system that is much more focused on the business implications of each project, not just the assurance that a project is the minimum revenue requirements solution to a given “problem,” as defined by the planning guides.

In this new paradigm, traditional values such as good customer service, high utilization of physical assets, and efficient operation are all-important priorities, because they contribute to that end. It is interesting that there seems to be more focus on these “virtues” of good utility operation under deregulation than there was under a traditionally regulated operation. Certainly utilization rates are higher, attention paid to reliability is greater than ever, and more effort is being directed at improving efficiency than ever before. The reason is that when one has less of something, more attention is given and more effort is put into getting the most out of it. Modern utilities understand that all of the virtues they desire boil down to one resource: money, and that they will never have enough to meet all possible needs based on traditional planning standards.

Additionally, in aging infrastructure areas, a power distribution utility will have nowhere near enough money to take the “obvious” approach to the aging infrastructure: replace all the old equipment and upgrade the obsolete layout. There simply is not enough money. This makes the whole planning situation budget-constrained. The overall, executive management perspective is thus,

“The power delivery system is a physical asset that we purchase and maintain to provide good customer service and stockholder

profits. It is doing its job when it supports our business model.”

Changes in cultural “standards” may be needed

Many of today’s power delivery planning, engineering and operating concepts were developed by or for vertically integrated utilities in a regulated, non-competitive industry. Some are quite inappropriate for a business-driven utility environment. They use procedures that are rule- rather than results-based procedures, and thus not easily adaptable to an environment, in which decisions are ultimately driven by financial, rather than standards or policy, considerations. Required changes are scattered throughout this book, but in general, the utility needs to re-examine:

- *Standards and guidelines* – one utility that encountered severe problems in 1999 had not re-examined and revised its distribution planning and engineering guidelines and standards since 1978, forgetting entirely the issues of de-regulation and competition, load characteristics, capital vs. operating cost considerations, and a host of other changes in system requirements have occurred in the intervening 28 years.
- *Budget base*. Many utilities operate from year-to-year by revising the current year budget incrementally for the following year. Budget allocations and functions among departments are effectively institutionalized along traditional lines.
- *Protection and restoration practices* at many utilities follow practices developed in the late 60s and early 70s. These are often incompatible with modern consumer needs of fewer interruptions and do not make use of the most modern equipment and techniques for elimination of interruption propagation through the network whenever an outage occurs.
- *Prioritization of projects* is still based on models and concepts developed while utilities were regulated, vertically integrated companies. Most of these prioritization methods are entirely incapable of allocating a limited (constrained) budget among needed projects in the best possible manner [Willis, 1997, Chapter 18].
- *Reliability engineering* at many utilities is mostly reactive. Reliability analysts review operating records and experience in order to identify poor performing areas or equipment. Once identified, these areas are turned over to Engineering and Planning for rehabilitation, using what are basically traditional contingency-based tools.
- *System design* at many utilities throughout the United States still follows specific styles and guidelines developed decades ago. Nearly every

utility in the industry has slightly different preferences in how it lays out and specifies equipment for substations, how it routes, sizes, and switches feeders, how it designs its protection schemes and contingency backup plans, and how it operates its system at times of high stress. Every utility insists that its methods are best. But they can't all be correct – in fact few have it right, and all can learn by reviewing what other utilities do best.

In these cases, and a host of others, utilities can make improvements by re-examining what are often long cherished practices and unique characteristics or policies, which they firmly believed in. This fresh approach should both look at the changes in customer loads and needs since the last time practices were examined, and should make certain that the new standards, guidelines, methods and practices are maximally supportive of business-driven focus of the utility.

Inappropriate paradigms reduce the effectiveness of efforts the utility makes to improve or maintain performance, particularly in areas where reliability is relatively costly to maintain (aging areas). The impact is raised costs, which translates to increases in SAIFI or SAIDI, whichever the utility needs, because money is used inefficiently, which would otherwise go to make improvements as needed.

1.8 CONDITION, CRITICALITY, AND CONSEQUENCES

Condition is a classification, scoring, or categorization of equipment that identifies which equipment is most in need of attention based on its characteristics and features. It is determined for all information known about an item, including negative information (explicit recognition of what one does not know).

To be most useful, a condition score or categorization must be *single-valued*. For example, a condition evaluation system that lists a transformer as [A, D, C] on an A-F scale, where A, D, and C respectively indicate the shell and outside systems are in first class condition, the bushings test very weak, and the core is in average shape for cores of this type, will still have to be boiled down to a single valued score that can help answer a fundamental question: “Do we do something about its condition or leave it in service as is?” Besides, the utility has all the base data and information that led to each of those three categorizations. It can use them to come up with one score (D: the bushings are bad so go out and replace them and check everything else while there.).

A high level of quantization is not necessarily an advantage in a condition score or classification and may be a hindrance. A score that ranges from 0 to 100 in units of 1 gives the impression of great accuracy, but in practice leads to no better results than a 1-10 scale or an A-B-C-D-F scale. The utility's goal is to find equipment that is in bad condition. Frankly, a digital scale of [bad, not bad] can get a big portion of the job done.

Conditions Have Consequences

Conditions have consequences – they mean something with respect to the utility’s goals. The goal is not to develop a scoring system that puts equipment into “bins” on the basis of available data. The purpose is to find equipment that helps ferret out equipment that is counterproductive to corporate goals.

As an example, a condition scoring system based on DGA (dissolved gas analysis) data could be devised that would group the 900 power transformers a utility had into five bins of 180 each. The 15 transformers (1.7% of the total) in most need of attention are in bin F – but so are 165 other units. A three-tiered condition method that would perhaps put 850 units in class A (good to go as is, no special attention needed), 35 of them in class D (inspect very frequently so that developing problems can probably be caught in time) and 15 in class F (has a high probability of failure or breakdown or other problems).

Criticality

Condition measures should be based on the status and characteristics of the equipment itself: its DGA readings, its impulse test results, the evaluation of inspectors training in visual classification of defects, and its age, model type, and installation details. A separate measure is *criticality*, which recognizes that a utility might care about one potential failure or problem due to deterioration more than another of equal condition and likelihood, because it was more critical to achieving the utility’s goals.

For example, in Chapter 16’s case study, among a utility’s wood distribution poles is one that would cause 67,800 customer minutes of service interruption if it failed: among all 723,000 poles in the system it is the most critical to service, in the sense that it would create the most customer disruption and SAIDI increase if it failed. At the other end of the criticality scale are thousands of poles supporting single-phase laterals and service lines that lead to only one or a few customers and thus will have far less impact on customer service and reliability, creating perhaps fewer than 400 customer minutes of outage if they fail. Clearly, the utility wants to put more focus on tracking and controlling the condition of the most critical poles and prioritize these poles of lesser potential impact on operations to a lower level of attention.

Condition versus Criticality

Generally, a first step in the study on any particular class of equipment (e.g., medium voltage breakers) in an aging infrastructure is to classify the equipment in that class in two dimensions -- by both condition and criticality. Equipment in the top ranking (that of most concern) in both categories will get the most attention. Figure 1.2 is a matrix for the 894 medium voltage breakers in a particular power system, showing how many are in each combination of four categories of condition and four categories of criticality.

Fuzzy Evaluations and Staged Prioritization

A reality of aging infrastructure management is that nearly everything that must be analyzed and evaluated in development of a condition estimate is “fuzzy.” Data and models might provide good results and significant statistical fits when applied to an entire class of equipment (e.g., 900 breakers over twenty years of operation, all 723,000 poles in the system). But the data and evaluations for any one breaker’s condition or one pole’s level of deterioration and losses of strength are subject to a good deal of uncertainty. Experience has led the authors to strongly favor aging-asset management approaches that acknowledge the uncertainty and that cannot be practically reduced beyond a certain amount. Such approaches are therefore designed to deal with it efficiently.

From a practical standpoint the uncertainty or fuzziness means that the units assigned to each box in Figure 1.4 have some likelihood of belonging in a nearby box instead, particularly as regards the vertical (condition evaluation) ranking. There is little doubt that the three units in the upper right most box are in need of very focused attention: anything scoring that high in the initial screening will definitely “make the playoffs.” But the twelve units in the box right below are an example of the conundrum the utility faces due to uncertainty. A few of these units probably belong be in the box above and should receive attention now. Several probably should be in the box below and are not really a priority this year. For this reason, condition versus criticality matrices are only a first step in staged

Med Voltage Breakers		Criticality				Row Total
		1	2	3	4	
Condition	D	0	6	4	3	13
	C	7	10	11	12	40
	B	65	150	134	37	386
	A	24	51	145	235	455
	Col Total	96	217	294	287	894

Figure 1.4 Condition versus Criticality Matrix for 894 medium voltage (feeder) breakers at a medium size utility in the central US identifies three transformers as having both a poor condition and high criticality of failure/breakdown consequences, and 16 others of high concern.




Figure 1.5 Condition and Criticality leads to actions that recognize the underlying uncertainty in condition evaluation and seeks to clarify condition in all but the most urgent situations before prescribing further action.

or hierarchical process as depicted in Figure 1.3. As shown, the sixteen breakers in the two categories closest to the upper right are scheduled for immediately inspection and evaluation. Based on those results several may be moved into the upper category and would be assigned for refurbishment or replacement. In this way, every one of the 894 medium voltage breakers in the system is assigned to an appropriate category of work to support the aging asset management program.

1.9 CONCLUDING REMARKS

Aging power delivery infrastructures include utility system equipment and facilities that are old and of concern due to deterioration of condition, which leads to high failure rates, repair and replacement costs, and customer service interruptions (Table 1.4). The infrastructure also includes the operating and management processes and methods that the utility uses to determine and control its use of that system and its equipment. In many cases these processes and methods are as old as that system’s equipment and in need of revision and augmentation just as the system is. Regardless, making changes and additions to

its operating processes and methods and perhaps the power delivery system, too, a utility can control condition deterioration and impacts on customer and system performance so that it can get the best business results possible.

Manage Both Effects and Causes

Managing an aging infrastructure is about *both* reducing the effects of aging equipment as much as possible (if you can’t control the disease, control the symptoms) and controlling the deterioration of the system as much as possible (but keep the disease under control as much as you can). Some combination of effort and budget split between these two will be optimum. That is the primary focal point for executive management. For example, a distributed secondary network design renders a power system almost completely immune to a high incidence of small random failures through the system. Compared to the traditional radial central trunk feeder system, they are an order of magnitude more tolerant of failures. Should a utility with aging equipment shift to such a design so the system and the customer base “just doesn’t care” about increasing failure rates? Of course not. Making that move would be very expensive, and distribution networks carry their own operating characteristics that would be difficult to live with en masse over an entire system.

But a move in that direction might be a good idea. There is a huge range of distribution configurations between the distributed secondary network approach and the traditional central trunk radial feeder approach, such as multi-branch feeder/highly sectionalizable feeder layouts, open and closed loop systems of various types. Particularly in combination with “smart” equipment and control

Table 1.4 Characteristics of the Typical Aging Infrastructure Area

1.	The system layout and design was first put in place more than forty years ago.
2.	The majority of equipment in the area is more than forty years old.
3.	The system is well engineered and fully meets minimum engineering criteria.
4.	The area is seeing steady, if perhaps low, load growth.
5.	The area is plagued by above average equipment failure rates. Overtime is high due to large amounts of unscheduled repair and restoration.
6.	SAIFI began rising some years ago; SAIDI is beginning to rise now.
7.	Major interruption events always occur due to a bizarre series of outages.
8.	Things go bad a lot.
9.	No comprehensive plan exists.

systems, these can provide a certain amount of immunity from increasing rates of equipment failure. The failures still occur, but they lead to far lower SAIDI and SAIFI increases.

But ultimately, the utility has to address, or at least control, the root cause, which is deterioration of equipment due to the wear and tear of service. This correlates with age more than anything else, but age is not the direct cause. There is also a noticeable element of random luck involved: a brand new transformer can be struck by lightning and greatly weakened, too. As much as justifiable based on cost and results, the utility needs to find a way to address this.

Condition Data, Modeling, and Tracking is the Key to Success

Regardless of the concepts and approach taken in managing an aging infrastructure, condition estimation is the foundation of any aging infrastructure management program. The extent to which the utility can, in an affordable and practicable manner, separate the bad equipment (poor condition) from the good (otherwise) determines the potential it has to manage its infrastructure and achieve economical success. Ultimately, any path to success requires the utility to separate the equipment that is in a condition to give satisfactory service, regardless of age, from equipment, old or new, that is deteriorated or flawed to the point that it will not provide satisfactory service.

The Future System Will Not Necessarily Be Good, Just “Good Enough”

Power system engineers and operators dream of a system in which all of the equipment is in first class shape: in top operating condition, clean, nicely painted, shiny, and new. That is only a dream. It is not the practical future of any power system (even one that is brand new today). Any well-managed aging power delivery infrastructure will have high portions of old equipment in it. Quality and condition will not be uniformly first-class throughout, but only “barely good enough” in places.

The service transformers shown in Figure 1.6 are old, have a patina of rust and mildew, and do not live up to that dream. They are a bit noisy, too, indicating their cores are a tad loose. If tested, it is likely they would return mediocre results – within specifications, but not near the “new” end of the test range. Yet for all that, these transformers are in place and doing their job, and their condition and criticality does justify replacement. Replacing them makes little economic sense,



Figure 1.6 A bank of service transformers does not look new, because they aren't, but they are good to go in a system that is trying to get the most from its past investments in equipment. The utility's challenge is to separate these units, which can be expected to give satisfactory service for many years to come, from other units that won't.

particularly when money is needed to replace and repair equipment that is in much more need of attention.

An Optimum Path Toward an Optimum Sustainable Point

In Section 1.1, the authors introduced the term *sustainable point* to mean an operating status at which the infrastructure no longer deteriorates and performance no longer worsens each year. Ideally, a utility would find an affordable and effective way to manage aging equipment to achieve a good sustainable point. It would put the resources, tools and management processes needed in place to affect that program. The system and its performance would then be controlled. The system would be an infrastructure with *controlled aging*, managed to a set of cost and results goals.

A good deal of this book is devoted to the sustainable point concept and its application to the management of the aging power delivery infrastructures. The final chapters of this book deal almost exclusively with determining what is the best sustainable point and then managing to it. The Appendices describe how to model it from operating and historical data and give details about the characteristics of using it in practical application. Key features of sustainable points are summarized here. They will be presented in more detail, proven, discussed and then used later in this book.

A sustainable point exists from every power delivery infrastructure. The sustainable point for a particular infrastructure is a function of the power system's layout and design, the equipment it is composed of and the ages and conditions of that equipment, and the ownership strategy, operating rules, guidelines and standards being used to care for and operate the system made up of that equipment.

A different sustainable point can exist, if changes are made. Change anything within this infrastructure – change the system layout and the way circuits are configured, or the equipment or its condition, or how it is loaded or how it is inspected and maintained, or the ownership strategy and operating rules for refurbishment and replacement, etc., and the sustainable point will change.

A power delivery infrastructure trends toward its sustainable point – slowly. As a mental exercise, suppose one installs 100,000 wood distribution poles, all new, and puts all of them into service on January 1 of particular a year. Ignore infant mortality failure rates in this example and operate them as most utilities do (basically, do nothing until they fall over or a period of five to ten years, when they are inspected and treated with an anti-fungal spray).

Very few will be replaced in the following 365 days of operation because they are all new and in good condition and few will fail. As a result, nearly all the poles will make it into year two. The “system” will age almost exactly 1.0 year in its first year of operation, and enter the following year with average age just slightly

less than one year.² In that following, second year of operation, the poles will again exhibit a good failure rate, etc.

Over time the average age continues to climb at about one year per year, but as the age increases, the average failure rate begins to increase – very slowly at first because deterioration rates are very gradual initially. More and more failures will occur each year, meaning that more and more new poles are installed each year, until at some point the new poles counteract the continued aging, at least on average. To make the sustainable point and this trend understandable for the moment, suppose that this balance occurs when average age reaches 50 years, at which failure rate is 2%. This means in this year, 2% of poles will be replaced with new (because they fail) and 98% will not fail and will be a year older going into the next year. Average age will then be $98\% \times 51 \text{ years} + 2\% \times 0 \text{ years} =$ very nearly exactly 50 years again. The system will not have aged: it is at its sustainable point.

The reader should note that this is a *very* simplified example, far too simple for useful application or generalization. It is used here only to quickly communicate the concept that *eventually* a balance point – the sustainable point – is reached. The fact that this point exists for any and all power infrastructures can be proven mathematically and demonstrated by example based on just about any power system. But in the real world, the sustainable point is more complicated than this example makes out. In particular, in the real world failure rate and age are never reciprocal (in this example, the age, is equal to the inverse of the failure rate of 2%, i.e., $50 = 1/.02$).

Furthermore, due to non-linearities in just about every relationship and factor involved, and non-uniformities in nearly all the statistical distributions involved, it is much more common for a sustainable point for a set of wood utility poles to be something like 64 years at .61%. In addition, the sustainable point is reached when not just average, but distributions about the mean are stable. This is often quite important to a utility. All this will be discussed as needed in later chapters. It is covered in detail in Appendix B's tutorials.

Practical Implications: the Good and the Bad

From a practical standpoint, what is important to electric utilities and the people who manage and operate them about the sustainable point is:

Good: Most power systems are not at their sustainable point – yet. With no significant exceptions, all utility systems in the US are not yet at their sustainable point. Despite age and aging, including perhaps pockets of seriously old equipment and facilities (80+ years old) even the oldest systems in the US are still getting older ever year.

² Perhaps the failure rate for new poles is .25%, meaning that 250 poles failed. Treating the new poles as age = 0 and the originals as age =1, at the beginning of the next year the average age is .9975 years.

Bad: The sustainable point is much worse than now. In almost every power system throughout North America, the sustainable point for the power delivery infrastructure (power system in combination with current operating and management rules) would be completely unacceptable to customers, utility, and regulators alike. It represents a combination of failure rate, reliability, and cost that is far worse than present.

Good: The system will take time to get there: one can count on the slow, steady trend. Power systems do not fall “off the edge” with regard to aging and its effects, nor does aging typically lead to catastrophic “system meltdowns” when failure rates and costs skyrocket over a short period of several years. Aging power delivery infrastructures instead are very similar to quicksand: the situation slowly but inexorably worsens, perhaps until it is very bad. This is a particularly apt analogy in other ways, too. First, often measures taken to get out of the mess are counter-productive, and only worsen the situation in the long run. Second, only a sustained steady pull will extract one from the situation.

Bad: all this means it's going to get worse every year. Things move slowly, and storms and year-to-year randomness mean volatility will be seen in annual results, but the long term trend is inevitable: things will gradually get a lot worse.

Good: the sustainable point can be moved. The utility can change its strategy, its ownership and operating policies, its maintenance and service guidelines, etc., and move the sustainable point to something more to its liking.

Bad: it takes a good deal of time for systemic changes to have a significant impact. The “sluggish” response of the system to aging and its effects works the other way, too. Perhaps a utility is at an uncomfortable point with regard to particular equipment (i.e., failure rates are 3X that of new equipment due to significant deterioration of condition) in a system on its way to a sustainable point that is much worse than that, perhaps 10X). A change in inspection, maintenance, repair and refurbishment policies that results in a sustainable point that has a failure rate “as good as new” will not make a huge impact overnight. Even targeted as effectively as possible, these changes might take five years or more to drive failure rates back down to only 2X.

Effective programs cost a lot of money. As discussed in Section 1.1, it costs more money. The industry will ultimately have no choice but to either accept worsening reliability or somewhat higher costs.

Good: changes in design and configuration along with the addition of “smart equipment” reduce the system’s vulnerability and sensitivity to its aging and deterioration-related effects. Basically they mitigate the consequences of failure. Old, deteriorated equipment still fails as frequently as before, it just does not matter as much or have as much effect.

More good: these measures lower the cost all around. Actually, these changes move the sustainable point themselves, but the way many people chose to look at them is that they make changes to operating guidelines, maintenance and inspection programs, etc., more cost effective (less is needed).

The Optimum Plan Does Not *Quickly* Reach Its Sustainable Point

The sustainable point is a useful concept because it makes the basic behaviors and trends of an aging power delivery infrastructure (see above) more understandable and intuitive. It also can help in some statistical analytics applications and numerical modeling steps (knowledge of the characteristics of an answer can help solve for it with minimal error). However, most often the optimum business strategy is to not quickly jump to the sustainable point. Examples later in this book will demonstrate this, but generally, given the “sluggish” response of aging systems and the effectiveness of methods like circuit configuration revision and the addition of smart equipment in reducing system sensitivity to equipment failures, the optimum business strategy is to start now, on a path that optimizes results but might takes as long as twenty years to reach a sustainable point: as the system continues to age and deterioration continues to take its toll each year, the utility makes staged changes that keep the effects of aging just barely at bay: the program just barely erases the increases in reliability problems and operating costs that would have occurred.

Thus, every year, as the system’s failure rates increase just a small amount, the utility spends a little more on smart equipment and targeted circuit configuration changes to make the system a little less susceptible to failures than it was the year before. It begins to remove the most worrisome equipment and replace it with new or refurbished equipment. It applies life extension in areas and to equipment where it makes economic sense. It starts to apply inspection and maintenance programs that will have an effect over the long run. All of this is done in a coordinated manner, with the most effective measures being “used up” first, and generally only to an extent needed each year.

Thus, when all is said and done, a lot is changing each year, but what is not changing are the key performance indicators. SAIFI, SAIDI, safety, system performance, utility profit margin, and business results basically stay where they have been and were traditionally. The utility is on a long-term, *gradual* path that will take years but will gradually change the sustainable point to an acceptable status. Perhaps the original sustainable point represented reliability and cost problems that are three times what last year’s were. And perhaps it would take twenty years for the quicksand-like pace of deterioration to take its toll, (these values are representative of many real systems). Then the roll out of the program can be viewed as taking about that long. The utility gradually moves the sustainable point to that operating point it wants.

This works: performance stays where it was. It is often optimum or close to it from just about any and all standpoints. But it is necessary for one practical

reason. Any solution to the management of an aging power delivery infrastructure will involve increasing spending. This can be done optimally but it will still represent a significant increase – the discussion in Section 1.1 mentioned a 50% increase and that is representative of what most utilities will face. Rolling that out gradually, a few percent a year, is really the only option that will be viable to the utility, to regulators, and to the utility's customers. That said, there are two final points executive management will need to consider:

1. *The best sustainable point may not have today's performance results.* When all is considered, a gradual change to less reliability, or different reliability (higher SAIFI, lower SAIDI, or vice versa) may be best. Given historical trends, the authors expect our society will expect better reliability, particularly with respect to storm damage. Regardless, executives and planners should not assume that the only answer lies with keeping results where they are now.

It is never too early to start. The sustainable point might be 20 years in the future. The point where the continuing gradual trend becomes untenably unreliable and costly serious might be eight to ten years in the future. Regardless, the time to begin putting a solution in place is now: a 50% increase in spending spread over 20 years is much more palatable and manageable than a similar increase spread over only 15 years.

What Can and Should Be Done

The most easily identifiable characteristic of an aging delivery infrastructure is that a lot the equipment is old. However, "old" alone is not enough to cause concern and age is not the real issue. A good deal of that old equipment is very likely in satisfactory condition – good enough to remain in service for at least a few and perhaps many more years. The problem in an aging infrastructure system or area is that an unacceptably high portion of the old equipment is not dependable enough for continued service. It is deteriorated to the point that the failure rate or capability is seriously degraded, and from a practical and economic standpoint, it is incapable of providing satisfactory service in a dependable manner in the future.

Major Challenges Need to Be Overcome

The first major challenge to a utility is *finding* that equipment with poor condition amidst the greater mass of merely old equipment, determining what it faces overall (when it looks at all old equipment, its entire system, and its options) and then doing something *affordable* and *effective* about the situation. The second major challenge a utility faces is to determine what its viable alternatives are and which to pick. This usually involves folding aging infrastructure management into its long-term business strategy and mainstream processes.

Table 1.5 Summary of Suggested Fixes for Aging Infrastructure Problems

Priority	Action	Comments
	1. <i>Use a results-focused approach.</i>	The two results desired are customer service quality and lowest possible cost. Thus, the guiding principle above all others is to maximize them. A focus on them rather than old equipment and deteriorating condition is not sweeping the problem under the rug. It is a management and planning method that leads to optimum results.
	2. <i>Build the strategy, policies, and technology plans first.</i>	This is a long-term problem that affects nearly every aspect of utility operations and business results. Create a sound long-term solution that matches it in scale.
	3. <i>Zero-based all decisions and budgets.</i>	Nothing is sacred, and no level of expense or traditional function should be considered untouchable as the company considers revising and changing its processes and approaches. Planning should consider going both ways: increasing as well as decreasing everything.
	4. <i>Marginal benefit/cost ratio should be used as the basis for all evaluation of “bang for the buck.”</i>	This provides far superior results to prioritization done on the basis of benefit/cost ratio or other means. This concept is discussed and developed in detail in Chapters 5 (Section 5.3), 11, and 14.
	5. <i>Macro-comparison of intra-project alternatives using that marginal benefit/cost ratio should be used across departmental and functional boundaries to prioritize spending decisions among candidate projects and budgets.</i>	This rather lengthy, “technical title” actually describes a rather simple but critical procedure, in which the keys to making results-centered management work, is covered elsewhere (Willis and Brown, 2012)

Priority	Action	Comments
6.	<i>Reliability-based planning, engineering, operating and maintenance methods should be employed throughout the organization and a customer reliability-focused approach should guide all decisions on IT.</i>	
7.	<i>Planning for failure.</i>	The best results come by recognizing that equipment failures cannot be avoided but can be managed. Failures are a natural consequence of equipment use, and they will happen. However, their likelihood and consequences can be engineered and managed to the best business case.
8.	<i>Make inspection and tracking a large part of any aging power delivery system infrastructure management effort.</i>	As stated earlier, significant amounts of old equipment will be a reality for the foreseeable future. Old equipment that is still in satisfactory shape can go from “good enough” to “bad” in a relatively short period of time. The only protection is to inspect and track condition well and frequently.
9.	<i>Gather all the data practically possible.</i>	Historical ownership and operating records are often useful in ways not immediately obvious. When data sources or old records are found, they should at least be preserved so that planners have the option to use them in the future.
10.	<i>Be patient but skeptical.</i>	Aging infrastructures are a problem that won’t go away. But neither will it defy solution. That said, all aspects are fuzzy. Data are never perfect or sufficient. Models are approximate. Staff are skilled but probably not as confident as they would like. Expect reversal of aging trends to take time and to initially move slowly, but look for signs it is not working according to plan and be prepared to take appropriate measures.

The final challenge for a utility is in finding the money to implement its plan. Even at its best, any successful program to keep an aging infrastructure at a stable, sustainable and acceptable operating point will cost more than the utility was spending prior to aging having become a problem. A big part of the effort will be communicating this to other stakeholders and arranging cost escalation to be slow enough that it is palatable to others.

Table 1.5 gives ten key recommendations for how to proceed. Much this book will be about these recommendations and the details of their application. The central theme is a condition-based, sustainable-point focused, multi-attribute asset management approach. This complicated-sounding approach is straight-forward in concept even if rather intricate to apply well in many of its technical steps:

- Condition based. It is condition, not age, that is the major focus. Old equipment actually preferable: it means the utility is getting its value from what it spent – long ago.
- Sustainable point means that the utility does focus on the long-term aging trends and where they are going, and in determining how to manage them best in the long run.
- Multi-attribute means the utility is considering all its important goals and targets, such as customer reliability, system efficiency and cost, safety, public image, capital spending limits, and overall business risk.
- Asset management is a method of performing such work in a way that prioritize work and planning alternatives so that spending is as cost effective as possible (Willis and Brown, 2012).

REFERENCES AND BIBLIOGRAPHY

- R. E. Brown, *Electric Power Distribution Reliability – Second Edition*, CRC Press, Boca Raton, 2009.
- R. E. Brown, *Business Essentials for Utility Engineers*, CRC Press, Boca Raton, 2009.
- R E Brown, et al., “Spatial Load Forecasting Using Non-Uniform Areas” in *Proceedings of the IEEE T&D Conference*, April, 1999.
- R. E. Brown, et al., “Reliability and Capacity: A Spatial Load Forecasting Method for a Performance-Based Regulatory Environment,” in *Proceedings of the IEEE PICA Conference*, May, 1999, San Jose.
- P. Carrier (Chairman), *Interim Report of the U.S. Department of Energy’s Power Outage Study Team: Findings from the summer of 1999*, United States Department of Energy, Office of Policy, Washington, January 2000.

- M. V. Engel et al., editors, *Tutorial on Distribution Planning*, IEEE Course Text EHO 361-6-PWR, Institute of Electrical and Electronics Engineers, Hoes Lane, NJ, 1992.
- H. N. Tram et al., "Load Forecasting Data and Database Development for Distribution Planning," *IEEE Trans. on Power Apparatus and Systems*, November 1983, p. 3660.
- H. L. Willis, "Load Forecasting for Distribution Planning, Error and Impact on Design," *IEEE Transactions on Power Apparatus and Systems*, March 1983, p. 675.
- H. L. Willis, *Spatial Electric Load Forecasting – Second Edition*, Marcel Dekker, New York, 2002.
- H. L. Willis, *Power Distribution Planning Reference Book – Second Edition*, Marcel Dekker, New York, 2004.
- H. L. Willis and R. E. Brown, "Substation Asset Management," Chapter 19 in *Electric Power Substations Engineering- Third Edition*, J. D. McDonald, editor, CRC Press, Boca Raton, 2012.
- H. L. Willis and R. W. Powell, "Load Forecasting for Transmission Planning," *IEEE Transactions on Power Apparatus and Systems*, August 1985, p. 2550.
- H. L. Willis and J. E. D. Northcote-Green, "Spatial Load Forecasting -- A Tutorial Review," *Proceedings of the IEEE*, Feb. 1983, p. 232.
- H. L. Willis and H. N. Tram, "Load Forecasting for Transmission Planning," *IEEE Trans. on Power Apparatus and Systems*, paper 83 SM 379-5, March 1984.

2

Power Delivery Systems

2.1 INTRODUCTION

Retail sale of electric energy involves the delivery of power in ready to use form to the final consumers. Whether marketed by a local utility, load aggregator, or direct power retailer, this electric power must flow through the local power delivery system on its way from power production to consumer. This transmission and distribution (T&D) system consists of thousands of transmission and distribution lines, substations, transformers, and other equipment scattered over a wide geographical area and interconnected so that all function in concert to deliver power as needed to the utility's consumers.

This chapter is a quick tutorial review on T&D systems, their mission, characteristics, and design constraints. It examines the natural phenomena that shape T&D systems and explains the key physical relationships and their impact on design and performance. For this reason experienced planners are advised to scan this chapter, or at least its conclusions, so they understand the perspective upon which the rest of the book builds.

In a traditional electric system, power production is concentrated at only a few large, usually isolated, power stations. The T&D system moves the power from those often-distant generating plants to the many consumers who consume the power. In some cases, cost can be lowered and reliability enhanced through the use of distributed generation (DG): numerous smaller generators placed at strategically selected points throughout the power system in proximity to the

consumers.¹ This and other distributed resources - so named because they are distributed throughout the system in close proximity to consumers - including storage systems and demand-side management, often provide great benefit.

But regardless of the use of distributed generation or demand-side management, the T&D system is the ultimate distributed resource, consisting of thousands, perhaps millions, of units of equipment scattered throughout the service territory, interconnected and operating in concert to achieve uninterrupted delivery of power to the electric consumers. These systems represent an investment of billions of dollars, require care and precision in their operation, and provide widely available, economical, and reliable energy.

This chapter begins with an examination of the role and mission of a T&D system: why it exists and what it is expected to do, as shown in Section 2.2. Section 2.3 looks at several fundamental physical “laws” that constrain T&D systems design. The typical hierarchical system structure that results and the costs of its equipment are summarized in sections 2.4, 2.5 and 2.6. In section 2.7, a number of different ways to lay out a distribution system are covered, along with their advantages and disadvantages. Section 2.8 covers the “Smart Grid” approach, perhaps the most important modern innovation in retail electric delivery systems. Section 2.9 summarizes key points and provides some final comments on power system performance and aging.

2.2 T&D SYSTEM’S MISSION

A T&D system’s primary mission is to *deliver* power to electrical consumers at their *place* of consumption and in *ready-to-use* form. This means it must be dispersed throughout the utility service territory, at every location with a capacity in rough proportion to consumer demand in that neighborhood (Figure 2.1). This is the primary requirement for a T&D system -- *the system must cover ground* -- reaching every consumer with an electrical path of sufficient load capability and voltage and voltage regulation strength to satisfy that consumer's demand for power, power quality, and reliability of service.

That last point has become more critical still since the first edition of this book was published. The electrical path must provide an uninterrupted flow of stable power to the utility’s consumers. *Reliable power delivery means delivering all of the power demanded not just some of the power needed, and doing so all of the time.* Anything less than near perfection in meeting this goal is considered unacceptable. While 99.9% reliability of service may sound impressive, it means nearly nine hours of electric service interruption each year, an amount that would be unacceptable in nearly any first-world country.

¹ See *Distributed Power Generation – Planning and Evaluation*, by H. L. Willis and W. G. Scott, Marcel Dekker, 2000.




Figure 2.1 Map of electrical demand for a major U.S. city shows where the total demand of more than 2,000 MW peak is located. Degree of shading indicates electric load distribution. The T&D system must cover the region with sufficient capacity at every location to meet the consumer needs there.

Beyond the need to deliver power to the consumer, the utility's T&D system must also deliver it in ready-to-use form - at the utilization voltage required for electrical appliances and equipment, and free of large voltage fluctuations, high levels of harmonics or transient electrical disturbances (Engel et al., 1992).

Most electrical equipment in the United States is designed to operate properly when supplied with 60 cycle alternating current at between 114 and 126 volts, a plus or minus five percent range centered on the nominal *utilization voltage* of 120 volts (RMS average of the alternating voltage). In many other countries, utilization standards vary from 230 to slightly over 250 volts, at either 50 or 60 cycles AC.² But regardless of the utilization voltage, a utility must maintain the voltage provided to each consumer within a narrow range centered within the voltages that electric equipment is designed to tolerate.

² Power is provided to consumers in the United States by reversed alternating current legs (+120 volts and -120 volts wire to ground). This scheme provides 240 volts of power to any appliance that needs it, but for purposes of distribution engineering and performance, acts like only 120 volt power.

A ten- percent range of delivery voltage throughout a utility's service area may be acceptable, but a ten- percent range of fluctuation in the voltage supplied to any one consumer is not. An instantaneous shift of even three percent in voltage causes a perceptible (and to some people, disturbing) flicker in electric lighting. More importantly, voltage fluctuations can cause erratic and undesirable behavior of some electrical equipment.

Thus, whether high or low within the allowed range, the delivery voltage of any one consumer must be maintained at about the same level all the time - normally within a range of three to six percent - and any fluctuation must occur slowly. Such *stable voltage* can be difficult to obtain, because the voltage at the consumer end of a T&D system varies inversely with electric demand, falling as the demand increases, rising as it decreases. If this range of load fluctuation is too great, or if it happens too often, the consumers may consider it poor service.

Thus, a T&D system's mission is to:

1. Cover the service territory, reaching all consumers
2. Have sufficient capacity to meet the peak demands of its consumers
3. Provide highly reliable delivery to its consumers
4. Provide stable voltage quality to its consumers

And of course, above everything else, achieve these four goals at the *lowest cost* possible.

2.3 THE LAWS OF T&D

The complex interaction of a T&D system is governed by a number of physical laws relating to the natural phenomena that have been harnessed to produce and move electric power. These interactions have created a number of "truths" that dominate the design of T&D systems:

1. *It is more economical to move power at high voltage.* The higher the voltage, the lower the costs per kilowatt to move power any distance.
2. *The higher the voltage, the greater the capacity and the greater the cost of otherwise similar equipment.* Thus, high voltage lines, while potentially economical, cost a great deal more than low voltage lines, but have a much greater capacity. They are only economical in practice if they can be used to move a lot of power in one block - they are the giant economy size, but while always giant, they are only economical if one truly needs the giant size.

3. *Utilization voltage is useless for the transmission of power.* The 120/240 volt single-phase utilization voltage used in the United States, or even the 250 volt/416 volt three-phase used in European systems is not equal to the task of economically moving power more than a few hundred yards. The application of these lower voltages for anything more than very local distribution at the neighborhood level results in unacceptably high electrical losses and high costs.
4. *It is costly to change voltage level* - not prohibitively so, for it is done throughout a power system (that's what transformers do) - but voltage transformation is a major expense, which does nothing to move the power any distance in and of itself.
5. *Power is more economical to produce in very large amounts.* There is a significant economy of scale in generation, notwithstanding the claims by advocates of modern distributed generators. Large generators produce power more economically than small ones. Thus, it is most efficient to produce power at a few locations utilizing large generators.³
6. *Power must be delivered in relatively small quantities at low (120 to 250 volt) voltage level.* The average consumer has a total demand equal to only 1/10,000th or 1/100,000th of the output of a large generator.

An economical T&D system builds upon these concepts. It must “pick up” power at a few, large sites (generating plants), and deliver it to many, many more small sites (consumers). It must somehow achieve economy by using high voltage, but only when power flow can be arranged so that large quantities are moved simultaneously along a common path (line). Ultimately, power must be subdivided into “house- or business-sized” quantities, reduced to utilization voltage, and routed into each business and home via equipment whose compatibility with individual consumer needs means it will be relatively inefficient compared to the system as a whole. Regardless, this is the mission assigned to the retail power delivery system, what traditionally was called the “distribution system.”

³ The issue is more complicated than just a comparison of the cost of big versus small generation, as will be addressed later in this book. In some cases, distributed generation provides the lowest cost overall, regardless of the economy of scale, due to constraints imposed by the T&D system. Being close to the consumers, distributed generation does not carry with it the costs of adding T&D facilities to move the power from generation site to consumer. Often this is the margin of difference, as will be discussed later in this book.

Hierarchical Voltage Levels

The overall concept of a power delivery system layout that has evolved to best handle these needs and “truths” is one of hierarchical voltage levels as shown in Figures 2.2 and 2.3. As power is moved from generation (large bulk sources) to consumer (small demand amounts) it is first moved in bulk quantity at high voltage - this makes particular sense since there is usually a large bulk amount of power to be moved out of a large generating plant. As power is dispersed throughout the service territory, it is gradually moved down to lower voltage levels, where it is moved in ever smaller amounts (along more separate paths) on lower capacity equipment until it reaches the consumers. The key element is a “lower voltage and split” concept.

Thus, in Figure 2.3, the 3 kW used by a particular consumer -- Mrs. Watts at 412 Oak Street in Metropolis City - might be produced at a 800 MW power plant more than three hundred miles to the north. Her power is moved as part of a 750 MW block of wholesale from plant to city on a 345 kV transmission line, to a switching substation. Here, the voltage is *lowered* to 138 kV through a 345 to 138 kV transformer, and immediately after that, this large block is *split* into five separate smaller portions in the switching substation’s bus work. Each of these five parts, roughly 150 MW each, is routed onto a transmission line leaving that substation.




Figure 2.2 A power system is structured in a hierarchical manner with various voltage levels. A key concept is “lower voltage and split” which is done from three to five times during the course of power flow from generation to consumer.

Now part of one of these smaller blocks of power, Mrs. Watts's electricity is routed to her side of Metropolis on a 138 kV transmission line that snakes 20 miles through the northern part of the city. Ultimately this line connects to another switching substation on the south side of the city. Along the way, it feeds power to several distribution substations along its route,⁴ among which it feeds 40 MW into the particular substation that serves Mrs. Watts's neighborhood as well as others around that. Here, her power is run through a 138-kV to 12.47kV distribution transformer. At the low side of that substation transformer, at 12.47 kV (the primary distribution voltage) the 40 MW is split into six parts, each about 7 MW, with each 7 MVA part routed onto a different distribution feeder. Mrs. Watts's power flows along one particular feeder for two miles, until it gets to within a few hundred feet of her home. Here, a much smaller amount of power, 50 kVA (sufficient for perhaps ten homes), is routed to a service transformer, one of several hundred scattered up and down the length of the feeder.

Once on the service transformer Mrs. Watts's power and that of her neighbors is lowered to utilization voltage, the voltage at which the utility's customers will use the power. In this case that is 120/240 volts (208 V phase to phase). As it emerges, from the service transformer, her power is routed onto the secondary system, operating at that 120/240 volt level (250/416 volts in Europe and many other countries). The secondary wiring splits the 50 kVA into small blocks of power, each about 5 kVA, and routes one of these to Mrs. Watts's home along a secondary conductor to her service drops - the wires leading directly to her house.

Over the past one hundred years, this hierarchical system structure has proven a most effective way to move and distribute power from a few large generating plants to a widely dispersed consumer base. Neglecting light variations in design and operating practice from one utility to another, all power delivery systems are like this. The key element in this structure is the "reduce voltage and split" function - a splitting of the power flow being done simultaneously with a reduction in voltage. Usually, this happens between three and five times as power makes its way from generator to consumers.

2.4 LEVELS OF THE T&D SYSTEM

As a consequence of this hierarchical structure, a power delivery system can be thought of very conveniently to be composed of several distinct levels of equipment, as illustrated in Figure 2.3. Each level consists of many units of fundamentally similar equipment, doing roughly the same job, but located in different parts of the system. For example, all of the distribution substations are planned and designed in the same manner and do roughly the same job. Likewise

⁴ Transmission lines whose sole or major function is to feed power to distribution substations are often referred to as "sub-transmission" lines.




Figure 2.3 A T&D system consists of several levels of power delivery equipment, each feeding the one below it in an unbroken chain that delivers power to Mrs. Watts' home.

all feeders are similar in equipment type, layout, and mission, and all service transformers have the same basic mission and are designed with similar planning goals and to similar engineering standards.

Power can be thought of as flowing “down” through these levels, on its way from power production to consumer. As it moves from the generation plants (system level) to the consumer, the power travels through the transmission level, to the sub-transmission level, to the substation level, through the primary feeder level, and onto the secondary service level, where it finally reaches the consumer.

Each level takes power from the next higher level in the system and delivers it to the next lower level in the system. While each level varies in the types of equipment it has, its characteristics, mission, and manner of design and planning, all share several common characteristics:

- Each level is fed power by the one above it, in the sense that the next higher level is electrically closer to the generation.
- Both the nominal voltage level and the average capacity of equipment drops from level to level, as one moves from generation to consumer. Transmission lines operate at voltages of between 69 kV and 1,100 kV and have capacities between 50 and 2,000 MW. By contrast, distribution feeders operate between 2.2 kV and 34.5 kV and have capacities somewhere between 2 and 35 MW.
- Each level has many more pieces of equipment in it than the one above. A system with several hundred thousand consumers might have fifty transmission lines, one hundred substations, six hundred feeders, and forty thousand service transformers.
- As a result, the net capacity of each level (number of units times average size) increases as one moves toward the consumer. A power system might have 4,500 MVA of substation capacity but 6,200 MVA of feeder capacity and 9,000 MVA of service transformer capacity installed.⁵
- Reliability drops as one moves closer to the consumer. A majority of service interruptions are a result of failure (either due to aging or to damage from severe weather) of transformers, connectors, or conductors very close to the consumer, as shown in Figure 2.4.

⁵ This greater-capacity-at-every-level is deliberate and required both for reliability reasons and to accommodate coincidence of load, which will be discussed in Chapter 3.




Figure 2.4 Ten years of consumer interruptions for a large electric system, grouped by level of cause. Interruptions due to generation and transmission often receive the most attention because they usually involve a large number of consumers simultaneously. However, such events are rare whereas failures and interruptions at the distribution level create a constant background level of interruptions.

Table 2.1 gives statistics for a typical system. The net effect of the changes in average size and number of units is that each level contains a greater total capacity than the level above it. The service transformer level in any utility system has considerably more installed capacity (number of units times average capacity) than the feeder system or the substation system. Total capacity increases as one heads toward the consumer because of non-coincidence of peak load (which will be discussed in Chapter 3) and for reliability purposes.

Table 2.1 Equipment Level Statistics for a Medium-Sized Electric System

Level of System	Voltage of Level - kV	Number of Units	Avg. Cap. MVA	Total Cap MVA
Transmission	345, 138	12	150	1,400
Sub-transmission	138, 69	25	65	1,525
Substations	138/23.9 & 69/13.8	45	44	1,980
Feeders	23.9, 13.8	227	11	2,497
Service Trans.	.12, .24	60,000	.05	3,000
Secondary/Service	.12, .24	250,000	.014	3,500
Consumer	.12	250,000	.005	1,250

The (Wholesale) Transmission Level

The transmission system is a network of three-phase lines operating at voltages generally between 115 kV and 765 kV. Capacity of each line is between 50 MVA and 2,000 MVA. The term “network” means that there is more than one electrical path between any two points in the system (Figure 2.5). Networks are laid out in this manner for reasons of reliability and operating flow--if any one element (line) fails, there is an alternate route and power flow is (hopefully) not interrupted.

In addition to their function in moving power, portions of the transmission system, the largest elements, namely its major power delivery lines, are designed, at least in part, for stability needs. The transmission grid provides a strong electrical tie between generators, so that each can stay synchronized with the system and with the other generators. This arrangement allows the system to operate and to function evenly as the load fluctuates and to pick up load smoothly if any generator fails - what is called stability of operation. A good deal of the equipment put into transmission system design, and much of its cost, is for these stability reasons, not solely or even mainly, for moving power.




Figure 2.5 A network is an electrical system with more than one path between any two points, meaning that (if properly designed) it can provide electrical service even if any one element fails.

The Sub-transmission Level

The sub-transmission lines in a system take power from the transmission switching stations or generation plants and deliver it to substations along their routes. A typical sub-transmission line may feed power to three or more substations. Often, portions of the transmission system, bulk power delivery lines, lines designed at least in part for stability as well as power delivery needs to do this too. The distinction between transmission and sub-transmission lines becomes rather blurred.

Normally, sub-transmission lines are in the range of capacity of 30 MVA up to perhaps 250 MVA, operating at voltages from 34.5 kV to as high as 230 kV. With occasional exceptions, sub-transmission lines are part of a network grid - they are part of a system in which there is more than one route between any two points. Usually, at least two sub-transmission routes flow into any one substation, so that feed can be maintained if one fails.⁶

The Substation Level

Substations, the meeting point between the transmission grid and the distribution feeder system, are where a fundamental change takes place within most T&D systems. The transmission and sub-transmission systems above the substation level usually form a network, as discussed above, with more than one power flow path between any two parts. But from the substation on to the consumer, arranging a network configuration would simply be prohibitively expensive. Thus, most distribution systems are radial - there is only one path through the other levels of the system.

Typically, a substation occupies an acre or more of land on which the necessary substation equipment is located. Substation equipment consists of high and low voltage racks and busses for the power flow, circuit breakers for both the transmission and distribution level, metering equipment, and the “control house,” where the relaying, measurement, and control equipment is located. But the most important equipment, what gives this substation its capacity rating, are the substation *transformers*. These transformers convert the incoming power from transmission voltage levels to the lower primary voltage for distribution.

Individual substation transformers vary in capacity, from less than 10 MVA to as much as 150 MVA. They are often equipped with tap-changing mechanisms and control equipment to vary their winding ratio so that they maintain the distribution voltage within a very narrow range, regardless of larger fluctuations

⁶ Radial feed - only one line - is used in isolated, expensive, or difficult transmission situations, but for reliability reasons is not recommended.

on the transmission side. The transmission voltage can swing by as much as 5%, but the distribution voltage provided on the low side of the transformer stays within a narrow band, perhaps only $\pm .5\%$.

Very often, a substation will have more than one transformer. Two is a common number, four is not uncommon, and occasionally six or more are located at one site. Having more than one transformer increases reliability. In an emergency, a transformer can handle a load much over its rated load for a brief period (e.g., perhaps up to 140% of rating for up to four hours). Thus, the T&D system can pick up the load of the outaged portions during brief repairs and in emergencies.

Equipped with one to six transformers, substations range in “size” or capacity from as little as five MVA for a small, single-transformer substation, serving a sparsely populated rural area, to more than 400 MVA for a truly large six-transformer station, serving a very dense area within a large city.

Often T&D planners will speak of a *transformer unit*, which includes the transformer and all the equipment necessary to support its use - “one fourth of the equipment in a four-transformer substation.” This is a much better way of thinking about and estimating cost for equipment in T&D plans. While a transformer itself is expensive (between \$50,000 and \$1,000,000); the buswork, control, breakers, and other equipment required to support its use can double or triple that cost. Since that equipment is needed in direct proportion to the transformer’s capacity and voltage, and since it is needed *only* because a transformer is being added, it is normal to associate it with the transformer as a single planning unit. Add the transformer and add the other equipment along with it.

Substations consist of more equipment and involve more costs than just the electrical equipment. The site has to be purchased and prepared. Preparation is not trivial. The site must be excavated, a grounding mat - wires running under the substation to protect against an inadvertent flow during emergencies - laid down, and foundations and control ducting for equipment must be installed. Transmission towers to terminate incoming transmission must be built. Feeder getaways - ducts or lines to bring power out to the distribution system - must be added.

The Feeder Level

Feeders, typically either overhead distribution lines mounted on wooden poles or underground buried or ducted cable sets, route the power from the substation throughout its service area. Feeders operate at the primary distribution voltage. The most common primary distribution voltage in use throughout North America is 12.47 kV, although anywhere from 4.2 kV to 34.5 kV is widely used. Worldwide, there are primary distribution voltages as low as 1.1 kV and as high

as 66 kV. Some distribution systems use several primary voltages - for example 23.9 kV, 13.8 kV and 4.16 kV.

A feeder is a small transmission system in its own right, distributing between 2 MVA to more than 30 MVA, depending on the conductor size and the distribution voltage level. Normally between two and 12 feeders emanate from any one substation, in what has been called a *dendrillic* configuration - repeated branching into smaller branches as the feeder moves out from the substation toward the consumers. In combination, all the feeders in a power system constitute the *feeder system* (Figure 2.6). An average substation has between two and eight feeders, and can vary between one and forty feeders.




Figure 2.6 Distribution feeders route power away from the substation, as shown (in idealized form - configuration is never so evenly symmetric in the real world) for two substations. Positions of switches make the system electrically radial, while parts of it are physically a network. Shown here are two substations, each with four feeders.

The main, three-phase trunk of a feeder is called the *primary trunk* and may branch into several main routes, as shown in the diagram on the next page. These main branches end at open points where the feeder meets the ends of other feeders - points at which a *normally open switch* serves as an emergency tie between two feeders.

Additionally, normally closed switches in several switchable elements will divide each feeder. During emergencies, segments can be re-switched to isolate

damaged sections and route power around outaged equipment to consumers who would otherwise have to remain out of service until repairs were made.

By definition, the feeder consists of all primary voltage level segments between the substations and an open point (switch). Any part of the distribution level voltage lines - three-phase, two-phase, or single-phase - that is switchable is considered part of the primary feeder. The primary trunks and switchable segments are usually built using three phases. The largest size of distribution conductor (typically this is about 500-600 MCM conductor, but a conductor over 1,000 MCM is not uncommon, and one of the authors designed and built a feeder with a 2,000 MCM conductors) being used for reasons other than just maximum capacity (e.g., contingency switching needs). Often a feeder has excess capacity because it needs to provide back up for other feeders during emergencies.

The vast majority of distribution feeders used worldwide and within the United States are overhead construction, wooden pole with wooden crossarm or post insulators. Only in dense urban areas, or in situations where esthetics are particularly important, can the higher cost of underground construction be justified. In this case, the primary feeder is built from insulated cable, which is pulled through concrete ducts that are first buried in the ground. Underground feeder costs are three to ten times that of overhead costs.

Many times, however, the first several hundred yards of an overhead primary feeder are built underground even if the system is overhead. This underground portion is used as the *feeder get-away*. Particularly at large substations, the underground get-away is dictated by practical necessity, as well as by reliability and esthetics. At a large substation, ten or twelve 3-phase, overhead feeders leaving the substation mean from 40 to 48 wires hanging in mid-air around the substation site, with each feeder needing the proper spacing for electrical insulation, safety, and maintenance.

One significant aging infrastructure issue is that at many large-capacity substations in a tight location, there is simply not enough overhead space for enough feeders to be routed out of the substation. Even if there is, the resulting tangle of wires looks unsightly, and perhaps most importantly, is potentially unreliable. One broken wire falling in the wrong place can disable a lot of power delivery capability. The solution to this dilemma is the underground feeder get-away, usually consisting of several hundred yards of buried, ducted cable that takes the feeder out to a riser pole, where it is routed above ground and connected to overhead wires. Very often, this initial underground link sets the capacity limit for the entire feeder. The underground cable ampacity is the limiting factor for the feeder's power transmission.

The Lateral Level

Laterals are short stubs or line-segments that branch off the primary feeder and represent the final primary voltage part of the power's journey from the substation

to the consumer. A lateral is directly connected to the primary trunk and operates at the same nominal voltage. A series of laterals tap off the primary feeder as it passes through a community, each lateral routing power to a few dozen homes.

Normally, laterals do not have branches, and many laterals are only one or two-phase. All three phases are used only if a relatively substantial amount of power is required, or if three-phase service must be provided to some of the consumers. Normally, single and two-phase laterals are arranged to tap alternately into different phases on the primary feeder. An attempt by the Distribution Planning Engineer to balance the loads as closely as possible is shown below.

Typically, laterals deliver from as little as 10 kVA for a small single-phase lateral, to as much as 2 MVA. In general, even the largest laterals use small conductors (relative to the primary size). When a lateral needs to deliver a great deal of power, the planner will normally use all three phases, with a relatively small conductor for each, rather than employ a single-phase and use a large conductor. This approach avoids creating a significant imbalance in loading at the point where the lateral taps into the primary feeder. Power flow, loadings and voltage are maintained in a more balanced state if the power demands of a “large lateral” are distributed over all three phases.

Laterals (wooden poles) are built overhead or underground. Unlike primary feeders and transmission lines, single-phase laterals are sometimes buried directly. In this case, the cable is placed inside a plastic sheath (that looks and feels much like a vacuum cleaner hose). A trench is dug, and the sheathed cable is unrolled into the trench and buried. Directly buried laterals are no more expensive than underground construction in many cases.

The Service Transformers

Service transformers lower voltage from the primary voltage to the utilization or consumer voltage, normally 120/240-volt two-leg service in most power systems throughout North America. In overhead construction, service transformers are pole mounted and single-phase, between 5-kVA and 166 -kVA capacity. There may be several hundred scattered along the trunk and laterals of any given feeder; since power can travel efficiently only up to about 200 feet at utilization voltage, there must be at least one service transformer located reasonably close to every consumer. In underground areas, service transformers may be located in underground vaults, or at ground level (pad-mounted).

Passing through these transformers, power is lowered in voltage once again, to the final utilization voltage (120/240 volts in the United States) and routed to the secondary system or directly to the consumers. In cases where the system is supplying power to large commercial or industrial consumers, or the consumer requires three-phase power, between two and three transformers may be located together in a transformer bank, and be interconnected in such a way as to provide

multi-phase power. Several different connection schemes are possible for varying situations.

Padmount and vault-type service transformers provide underground service, as opposed to overhead pole-mounted service. The concept is identical to overhead construction, with the transformer and its associated equipment changed to accommodate incoming and outgoing lines that are underground.

The Secondary and Service Level

Secondary circuits fed by the service transformers route power at utilization voltage within very close proximity to the consumer. Usually this is an arrangement in which each transformer serves a small radial network of utilization voltage secondary and service lines. These lead directly to the meters of consumers in the immediate vicinity.

At most utilities, the layout and design of the secondary level is handled through a set of standardized guidelines and tables. Engineering technicians use these and clerks to produce work orders for the utilization voltage level equipment. In the United States, the vast majority of this system is single-phase. In European systems, much of the secondary is 3-phase, particularly in urban and suburban areas.

What is Transmission and What is Distribution?

Definitions and nomenclature defining “transmission” and “distribution” vary greatly among different countries, companies, and power systems. Generally, four types of distinction between the two are made, only three of which are dealt with in this book:

By economic societal function: As noted in Chapter 1, in the United States, the wholesale transmission, or grid level, is defined and operated differently from local power delivery, by Federal mandate. This distinction is not addressed in this book.

By voltage class: transmission is anything above 34.5 kV; primary distribution is anything below that but above the utilization voltage.

By function: distribution includes all utilization voltage equipment, plus all lines that feed power to service transformers.

By configuration: transmission is nearly always a network or heavily looped circuits. Distribution is nearly always radial or a simple loop scheme.

Generally, the last three definitions apply simultaneously since in most utility systems any transmission above 34.5 kV is configured as a network and does not feed service transformers directly. On the other hand, all distribution is radial, built of only 34.5 kV or below, and does feed service transformers. Substations – the meeting places of transmission lines (incoming) and distribution lines (outgoing) – are often included in one or the other category, and sometimes are considered as separate entities.

2.5 UTILITY DISTRIBUTION EQUIPMENT

The preceding section made it clear that a power delivery system is a very complex entity, composed of thousands, perhaps even millions, of components, which function together as a *T&D system*. Each unit of equipment has only a small part to play in the system, and is only a small part of the cost, yet each is critical for satisfactory service to at least one or more consumers or it would not be included in the system.

T&D system planning is complex because each unit of equipment influences the electrical behavior of its neighbors. It must be designed to function well in conjunction with the rest of the system, under a variety of different conditions, regardless of shifts in the normal pattern of loads or the status of equipment nearby. While the modeling and analysis of a T&D system can present a significant challenge, individually its components are relatively simple to understand, engineer, and plan. In essence, there are only two major types of equipment that perform the power delivery function:

1. Transmission and distribution lines move power from one location to another
2. Transformers change the voltage level of the power

Added to these three basic equipment types are two categories of equipment used for a very good reason:

1. Protective equipment which provides safety and “fail safe” operation.
2. Voltage regulation equipment that is used to maintain voltage within an acceptable range as the load changes.
3. Monitoring and control equipment that measures equipment and system status feeds this information to control systems so that the utility knows what the system is doing and can control it.

Transmission and Distribution Lines

By far the most omnipresent part of the power distribution system is the portion devoted to actually moving the power flow from one point to another. Transmission lines, sub-transmission lines, feeders, laterals, secondary and service drops, all consist of electrical conductors, suitably protected by isolation (transmission towers, insulator strings, and insulated wrappings) from voltage leakage and ground contact. It is this conductor that carries the power from one location to another.

Electrical conductors are available in various capacity ranges, with capacity generally corresponding to the metal cross section (other things being equal, thicker wire carries more power). Conductors can be all steel (rare, but used in some locations where winter ice and wind loadings are quite severe), all aluminum, or copper, or a mixture of aluminum and steel. Underground transmission can use various types of high-voltage cable. Line capacity depends on the current-carrying capacity of the conductor or the cable, the voltage, the number of phases, and constraints imposed by the line's location in the system.

The most economical method of handling a conductor is to place it overhead, supported by insulators on wooden poles or metal towers suitably clear of interference or contact with persons or property. However, underground construction, while generally more costly, avoids esthetic intrusion of the line and provides some measure of protection from weather. It also tends to reduce the capacity of a line slightly due to the differences between underground cable and overhead conductor. Suitably wrapped with insulating material in the form of underground cable, the cable is placed inside concrete or metal ducts or surrounded in a plastic sheath.

Transmission/sub-transmission lines are always 3-phase. There are three separate conductors for the alternating current, sometimes with a fourth neutral (un-energized) wire. Voltage is measured between phases. A 12.47 kV distribution feeder has an alternating current voltage (RMS) of 12,470 volts as measured between any two phases. Voltage between any phase and ground is 7,200 volts (12.47 divided by the square root of three). Major portions of a distribution system - trunk feeders - are as a rule built as three-phase lines, but lower-capacity portions may be built as either two-phase, or single-phase.⁷

Regardless of type or capacity, every electrical conductor has impedance (a resistance to electrical flow through it) that causes voltage drop and electrical losses whenever it is carrying electric power. Voltage drop is a reduction in the voltage between the sending and receiving ends of the power flow. Losses are a reduction in the net power, and are proportional to the square of the power. Double the load and the losses increase by four. Thus, 100 kilowatts at 120 volts

⁷ In most cases, a single-phase feeder or lateral has two conductors: the phase conductor and the neutral.

might go in one end of a conductor, only to emerge at the other as 90 kilowatts at 114 volts at the other end. Both voltage drop and losses vary in direct relation to load - within very fine limits if there is no load, there are no losses or voltage drop. Voltage drop is proportional to load. Double the load and voltage drop doubles. Losses are *quadratic*, however. Double the load and losses quadruple.

Transformers

At the heart of any alternating power system are transformers. They change the voltage and current levels of the power flow, maintaining (except for a very small portion of electrical losses), the same overall power flow. If voltage is reduced by a factor of ten from the high to low side, then Current is multiplied by ten, so that their overall product (Voltage times Current equals Power) is constant in and out.

Transformers are available in a diverse range of types, sizes, and capacities. They are used within power systems in four major areas:

1. At power plants, where power which is minimally generated at about 20,000 volts is raised to transmission voltage (100,000 volts or higher),
2. At switching stations, where transmission voltage is changed (e.g., from 345,000 volts to 138,000 volts before splitting onto lower voltage transmission lines),
3. At distribution substations, where incoming transmission-level voltage is reduced to distribution voltage for distribution (e.g., 138 kV to 12.47 kV),
4. And at service transformers, where power is reduced in voltage from the primary feeder voltage to utilization level (12.47 kV to 120/240 volts) for routing into consumers' homes and businesses.

Larger transformers are generally built as three-phase units, in which they simultaneously transform all three phases. Often these larger units are built to custom or special specifications, and can be quite expensive - over \$3,000,000 per unit in some cases. Smaller transformers, particularly most service transformers, are single-phase - it takes three installed side by side to handle a full three-phase line's power flow. They are generally built to standard specifications and bought in quantity.

Transformers experience two types of electrical losses - no-load losses (often called core, or iron, losses) and load-related losses.

No-load losses are electrical losses inherent in operating the transformer - due to its creation of a magnetic field inside its core. They occur simply

because the transformer is connected to an electrical power source. They are constant, regardless of whether the power flowing through the transformer is small or large. No-load losses are typically less than one percent of the nameplate rating. Only when the transformer is seriously overloaded, to a point well past its design range, will the core losses change (due to magnetic saturation of the core).

Load-related losses are due to the current flow through the transformer's impedance and correspond very directly with the level of power flow. Like those of conductors and cables they are proportional to current squared or quadrupling whenever power flow doubles. The result of both types of losses is that a transformer's losses vary as the power transmitted through it varies, but always at or above a minimum level set by the no-load losses.

Switches

Occasionally, it is desirable to be able to vary the connection of line segments within a power delivery system, particularly in the distribution feeders. Switches are placed at strategic locations so that the connection between two segments can be opened or closed. Switches are planned to be normally closed (NC) or normally open (NO), as was shown in Figure 2.6.

Switches vary in their rating (how much current they can carry) and their load break capacity (how much current they can interrupt or switch off), with larger switches being capable of opening a larger current. They can be manually, automatically, or remotely controlled in their operation.

Protection

When electrical equipment fails, for example if a line is knocked down during a storm, the normal function of the electrical equipment is interrupted. Protective equipment is designed to detect these conditions and isolate the damaged equipment, even if this means interrupting the flow of power to some consumers. Circuit breakers, sectionalizers, and fused disconnects, along with control relays and sensing equipment, are used to detect unusual conditions and interrupt the power flow whenever a failure, fault, or other unwanted condition occurs on the system.

These devices and the protection engineering required to apply them properly to the power system are not the domain of the utility planners and will not be discussed here. Protection *is* vitally important, but the planner is sufficiently involved with protection if he or she produces a system design that *can* be protected within standards, and if the cost of that protection has been taken into account in the budgeting and planning process. Both of these considerations are very important.

Protection puts certain constraints on equipment size and layout - for example, in some cases a very large conductor is too large (because it would permit too high a short circuit current) to be protected safely by available equipment and cannot be used. In other cases, long feeders are too long to be protected (because they have too low a short circuit current at the far end). A good deal of protective equipment is quite complex, containing sensitive electro-mechanical parts (many of which move at high speeds and in a split-second manner), and depending on precise calibration and assembly for proper function. As a result, the cost of protective equipment and control and the cost of its maintenance are often significant. Differences in protection cost *can* make the deciding difference between two plans.

Voltage Regulation

Voltage regulation equipment includes line regulators and line drop compensators, as well as tap changing transformers. These devices vary their turns-ratio (ratio of voltage in to voltage out) to react to variations in voltage drop. If voltage drops, they raise the voltage; if voltage rises, they reduce it to compensate. Properly used, they can help maintain voltage fluctuation on the system within acceptable limits, but they can only reduce the range of fluctuation, not eliminate it altogether.

Capacitors

Capacitors are a type of voltage regulation equipment. By correcting power factor they can improve voltage under many heavy loads (hence large voltage drop) cases. Power factor is a measure of how well voltage and current in an alternating system are in step with one another. In a perfect system, voltage and current would alternately cycle in conjunction with one another - reaching a peak, then reaching a minimum, at precisely the same times. But on distribution systems, particularly under heavy load conditions, current and voltage fall out of phase. Both continue to alternate 60 times a second, but during each cycle voltage may reach its peak slightly ahead of current - there is a slight lag of current versus voltage, as shown in Figure 2.7.

It is the precise, simultaneous peaking of both voltage and current that delivers maximum power. If out of phase, even by a slight amount, effective power drops, as does power factor - the ratio of real (effective) power to the maximum possible power, if voltage and current were locked in step.

Power engineers refer to a quantity called VARs (Volt-Amp Reactive) that is caused by this condition. Basically, as power factors worsen (as voltage and current fall farther apart in terms of phase angle) a larger percentage of the electrical flow is VARs, and a smaller part is real power. The frustrating thing is that the voltage is still there, and the current is still there, but because of the shift

in their timing, they produce only VARS, not power. The worse the power factor, the higher the VAR content. Poor power factor creates considerable cost and performance consequences for the power system: a large conductor is still required to carry the full level of current even though power delivery has dropped. And because current is high, the voltage drop is high, too, further degrading quality of service.

Unless one has worked for some time with the complex variable mathematics associated with AC power flow analysis, VARS are difficult to picture. A useful analogy is to think of VARS as “electrical foam.” If one tried to pump a highly carbonated soft drink through a system of pipes, turbulence in the pipes, particularly in times of high demand (high flow) would create foam. The foam would take up room in the pipes, but contribute little value to the net flow - the equivalent of VARS in an electrical system.

Poor power factor has several causes. Certain types of loads create VARS - loads that cause a delay in the current with respect to voltage as it flows through them. Among the worst offenders are induction motors - particularly small ones as almost universally used for blowers, air conditioning compressors, and the powering of conveyor belts and similar machinery. Under heavy load conditions, voltage and current can get out of phase to the point that the power factor can drop below 70%. Additionally, transmission equipment itself can often create this lag and “generate” a poor power factor.




Figure 2.7 Current and voltage in phase deliver maximum power (left). If current and voltage fall out of phase (right), actual power delivered drops by very noticeable amounts -- the *power factor* falls.

Capacitors correct the poor power factor. They “inject” VARs into a T&D line to bring power factor close to 1.0, transforming VAR flow back into real power flow, regaining the portion of capacity lost to poor power factor. Capacitors can involve considerable cost depending on location and type. They tend to do the most good if put on the distribution system, near the consumers, but they cost a great deal more in those locations than if installed at substations.

2.6 T&D COSTS

A T&D system can be expensive to design, build, and operate. Equipment at every level incurs two types of costs. Capital costs include the equipment and land, labor for site preparation, construction, assembly and installation, and any other costs associated with building and putting the equipment into operation. Operating costs include labor and equipment for operation, maintenance and service, taxes and fees, as well as the value of the power lost to electrical losses. Usually, capital costs are a one-time cost (once it's built, the money's been spent). Operating costs are continuous or periodic.

Electrical losses vary depending on load and conditions. While these losses are small by comparison to the overall power being distributed (seldom more than 8%), they constitute a very real cost to the utility. The present worth of the lifetime losses through a major system component such as a feeder or transformer can be a significant factor impacting its design and specification, often more than the original capital cost of the unit. Frequently, a more costly type of transformer will be selected for a certain application because its design leads to overall savings due to lower losses. Or a larger capacity line (larger conductor) will be used when really needed due to capacity requirements, purely because the larger conductor will incur lower cost losses.

Cumulatively, the T&D system represents a considerable expense. While a few transmission lines and switching stations are composed of large, expensive, and purpose-designed equipment, the great portion of the sub-transmission-substation-distribution system is built from “small stuff” - commodity equipment bought mostly “off the shelf” to standard designs. Individually inexpensive, they amount to a significant cost when added together.

Transmission Costs

Transmission line costs are based on a per-mile cost and a termination cost at either end of the line associated with the substation at which it is terminated. Costs can run from as low as \$50,000/mile for a 46 kV wooden pole sub-transmission line with perhaps 50 MVA capacity (\$1 per kVA-mile) to over \$1,000,000 per mile for a 500 kV double circuit construction with 2,000 MVA capacity (\$.5/kVA-mile).

Substation Costs

Substation costs include all the equipment and labor required to build a substation, including the cost of land and easements/ROW. For planning purposes, substations can be thought of as having four costs:

1. *Site cost* - the cost of buying the site and preparing it for a substation, including the cost of all legal work required to gain permissions and licenses (this can sometimes exceed the cost of the land).
2. *Transmission cost* - the cost of terminating the incoming sub-transmission lines at the site.
3. *Transformer cost* - the transformer and all metering, control, oil spill containment, fire prevention, cooling, noise abatement, and other related equipment, along with buswork, switches, metering, relaying, and breakers associated with this type of transformer, and their installation.
4. *Feeder buswork/Get-away costs* - the cost of beginning distribution at the substation, includes getting feeders out of the substation.

To expedite planning, estimated costs of feeder buswork and get-aways are often folded into the transformer costs. The feeders to route power out of the substation are needed in conjunction with each transformer and in direct proportion to the transformer capacity installed, so that their cost is sometimes considered together with the transformer as a single unit. Regardless, the transmission, transformer, and feeder costs can be estimated fairly accurately for planning purposes.

Cost of land is another matter entirely. Site and easements or ROW into a site have a cost that is a function of local land prices, which vary greatly depending on location and real estate markets. Site preparation includes the cost of preparing the site which includes grading, grounding mat, foundations, buried ductwork, control building, lighting, fence, landscaping, and an access road.

Substation costs vary greatly depending on type, capacity, local land prices and other variable circumstances. In rural settings where load density is quite low and minimal capacity is required, a substation may involve a site of only several thousand square feet of fenced area. This single area includes a single incoming transmission line (69 kV), one 5 MVA transformer; fusing for all fault protection; and all “buswork” built with wood poles and conductor, for a total cost of perhaps no more than \$90,000. The substation would be applied to serve a load of perhaps 4 MW, for a cost of \$23/kW. This substation in conjunction with the system around it would probably provide service with about ten hours of service interruptions per year under average conditions.

However, a typical substation built in most suburban and urban settings would be fed by two incoming 138 kV lines feeding two 40 MVA, 138 kV to 12.47 kV transformers. These transformers would each feed a separate low side (12.47 kV) bus, each bus with four outgoing distribution feeders of 9 MVA peak capacity each, and a total cost of perhaps \$2,000,000. Such a substation's cost could vary from between about \$1.5 million and \$6 million, depending on land costs, labor costs, the utility equipment and installation standards, and other special circumstances. In most traditional vertically integrated, publicly regulated electric utilities, this substation would have been used to serve a peak load of about 60 MVA (75% utilization of capacity), meaning that at its nominal \$2,000,000 cost works out to \$33/kW. In a competitive industry, with tighter design margins and proper engineering measures taken beforehand, this could be pushed to a peak loading of 80 MVA (100% utilization, \$25/kW). This substation and the system around it would probably provide service with about two to three hours of service interruptions per year under "normal," average conditions.

Feeder System Costs

The feeder system consists of all the primary distribution lines, including three-phase trunks and their lateral extensions. These lines operate at the primary distribution voltage - 23.9 kV, 13.8 kV, 12.47 kV, 4.16kV or whatever – and may be three-, two-, or single-phase construction as required. Typically, the feeder system is also considered to include voltage regulators, capacitors, voltage boosters, sectionalizers, switches, cutouts, fuses, and any intertie transformers (required to connect feeders of different voltage at tie points - for example: 23.9 and 12.47 kV) that are installed on the feeders (not at the substations or at consumer facilities).

As a rule of thumb, construction of three-phase overhead, wooden pole crossarm type feeders using a normal, large conductor (about 600 MCM per phase) at a medium distribution primary voltage (e.g., 12.47 kV) costs about \$150,000/mile. However, cost can vary greatly due to variations in labor, filing and permit costs among utilities, as well as differences in design standards, and terrain. Where a thick base of topsoil is present, a pole can be installed by simply auguring a hole for the pole. In areas where there is rock close under the surface, holes have to be jack-hammered or blasted, and costs go up accordingly. It is generally less expensive to build feeders in rural areas than in suburban or urban areas. Thus, while \$150,000 is a good average cost, a mile of new feeder construction could cost as little as \$55,000 in some situations and as much as \$500,000 in others.

A typical distribution feeder (three-phase, 12.47 kV, 600 MCM/phase) would be rated at a thermal (maximum) capacity of about 15 MVA and a recommended economic (design) peak loading of about 8.5 MVA peak, depending on losses and




Figure 2.8 Here, a service transformer, fed from a distribution primary-voltage lateral, feeds in turn ten homes through a secondary circuit operating at utilization voltage.

other costs. At \$150,000/mile, this capacity rating gives somewhere between \$10 to \$15 per kW-mile as the cost for basic distribution line. Underground construction of a three-phase primary feeder is more expensive, requiring buried ductwork and cable, and usually works out to a range of \$30 to \$50 per kW-mile.

Lateral lines, short primary-voltage lines working off the main three-phase circuit, are often single- or two-phase and consequently have lower costs but lower capacities. Generally, they are about \$5 to \$15 per kW-mile overhead, with underground costs of between \$5 to \$15 per kW-mile (direct buried) to \$30 to \$100 per kW-mile (ducted).

Cost of other distribution equipment, including regulators, capacitor banks and their switches, sectionalizers, line switches, etc., varies greatly depending on the specifics of each application. In general, the cost of the distribution system will vary from between \$10 and \$30 per kW-mile.

Service Level Costs

The service or secondary system consists of the service transformers that convert primary voltage to utilization voltage, the secondary circuits that operate at utilization voltage, and the service drops that feed power directly to each consumer. Without exception these are very local facilities, meant to move power no more than a few hundred feet at the very most and deliver it to the consumer “ready to use.”

Many electric utilities develop cost estimates for this equipment on a per-consumer basis. A typical service configuration might involve a 50 MVA pole-mounted service transformer feeding ten homes, as shown in Figure 2.8. Costs for this equipment might include:

Heavier pole & hardware for transformer application	\$250
50 kW transformer, mounting equipment, and installation	\$850
500 feet secondary (120/240 volt) single-phase @ \$2/ft.	\$1,100
10 service drops including installation at \$100	<u>\$1,100</u>
	\$3,300

For a cost of about \$330 per consumer, or about \$66/kW of coincident load.

Maintenance and Operating Costs

Once put into service, T&D equipment must be maintained in sound, operating condition, hopefully in the manner intended and recommended by the manufacturer. This will require periodic inspection and service, and may require repair due to damage from storms or other contingencies. In addition, many utilities must pay taxes or fees for equipment (T&D facilities are like any other business property). Operating, maintenance, and taxes are a continuing annual expense.

It is very difficult to give any generalization of O&M&T costs, partly because they vary so greatly from one utility to another, but mostly because utilities account for and report them in very different ways. Frankly, the authors have never been able to gather a large number of comparable data sets from which to produce even a qualitative estimate of average O&M&T costs.⁸ With that caveat, a general rule of thumb: O&M&T costs for a power delivery system probably run between 1/8 and 1/30 of the capital cost, annually.

The Cost to Upgrade Exceeds the Cost to Build

One of the fundamental factors affecting design of T&D systems and management of an electric utility is that it costs more to upgrade equipment to a higher capacity than to build to that capacity in the original construction. For example, a 12.47 kV overhead, three-phase feeder with a 9 MW peak load capability (336 MCM phase conductor) might cost \$240,000/mile to build (\$26.66 per kW-mile). Building it with a 600 MCM conductor instead, for a capacity of 15 MVA, would cost in the neighborhood of \$300,000 (\$20/kW-mile).

⁸ For example, some utilities include part of O&M expenses in overhead costs, others do not. A few report repairs (including storm damage) as part of O&M, others accumulate major repair work separately. Others report parts of routine service (periodic rebuilding of breakers) as a type of capital cost because it extends equipment life or augments capacity. Others report all such work as O&M. Such differences are common among utilities.

However, upgrading an *existing* 336 MCM, 9 MW capacity line to 600 MCM and 15 MVA capacity could probably cost the utility at least \$500,000/mile - over \$83 per kW-mile for the additional capacity. It is more expensive because replacement entails removing the old conductor and installing a new conductor along with brackets, cross-arms, and other hardware required supporting the heavier new conductor. Another factor is that typically this work has to be done “hot” (i.e., with the feeder energized and feeding power to customers) which means the work must be undertaken with extreme care and following a number of safety-related procedure that require additional equipment and labor. By contrast the original construction is done cold – with the power turned off and no customers yet served by the circuit.

Thus, T&D planners have an incentive to look at their long-term needs carefully, and to “overbuild” against initial requirements if growth trends show future demand will be higher. The cost of doing so must be weighed against long-term savings, but often T&D facilities are built with considerable margin (50%) above existing load to allow for future load growth.

The very high cost per kW for upgrading a T&D system in place creates one of the best-perceived opportunities for DSM and DG reduction. Note that the capital cost/kW for the upgrade capacity in the example above (\$33/kW) is nearly three times the cost of similar new capacity. Thus, planners often look at areas of the system where slow, continuing load growth has increased load to the point that local delivery facilities are considerably taxed, as areas where DSM and DG can deliver significant savings.

In some cases, distributed resources *can* reduce or defer significantly the need for T&D upgrades of the type described above. However, this does not assure significant savings for the situation is more complicated than an analysis of capital costs to upgrade may indicate. If the existing system (e.g., the 9 MW feeder) needs to be upgraded, then it is without a doubt highly loaded, which means its losses may be high, even off-peak. The upgrade to a 600 MCM conductor will cut losses 8,760 hours per year. Financial losses may drop by a significant amount, *enough in many cases to justify the cost of the upgrade alone*. The higher the annual load factor in an area, the more likely this is to occur, but it is often the case even when load factor is only 40%. However, DSM and in some cases DG also lowers losses, making the comparison quite involved, as will be discussed later in this book.

Electrical Losses Costs

Movement of power through any electrical device, be it a conductor, transformer, regulator or whatever, incurs a certain amount of electrical loss due to the impedance (resistance to the flow of electricity) of the device. These losses are a result of inviolable laws of nature. They can be measured, assessed, and minimized through proper engineering, but never eliminated completely.

Losses are an operating cost

Although losses do create a cost (sometimes a considerable one) it is not always desirable to reduce them as much as possible. Perhaps the best way to put them in proper perspective is to think of T&D equipment as *powered by electricity* - the system that moves power from one location to another runs on electric energy itself. Seen in this light, losses are revealed as what they are - a necessary operating expense to be controlled and balanced against other costs.

Consider a municipal water department, which uses electric energy to power the pumps that drive the water through the pipes to its consumers. Electricity is an acknowledged operating cost, one accounted for in planning, and weighed carefully in designing the system and estimating its costs. The water department could choose to buy highly efficient pump motors. Motors that command a premium price over standard designs but provide a savings in reduced electric power costs, use piping that is coated with a friction-reducing lining to promote rapid flow of water (thus carrying more water with less pump power), all toward reducing its electric energy cost. Alternatively, after weighing the cost of this premium equipment against the energy cost savings it provides, the water department may decide to use inexpensive motors and piping and simply pay more over the long run. The point is that the electric power required to move the water is viewed merely as one more cost that had to be included in determining the lowest "overall" cost. *It takes power to move power.*

Since its own delivery product powers electric delivery equipment, this point often is lost. However, in order to do its job of delivering electricity, a T&D system must be provided with power itself, just like the water distribution system. Energy must be expended to move the product. Thus, a transformer consumes a small portion of the power fed into it. In order to move power 50 miles, a 138 kV transmission line similarly consumes a small part of the power given to it.

Initial cost of equipment can always be traded against long-term electrical losses and higher costs. Highly efficient transformers can be purchased that use considerably less power to perform their function than standard designs. Larger conductors can be used in any transmission or distribution line, which will lower impedance, and thus losses for any level of power delivery. But both will cost more money initially - the efficient transformer may cost three times what a standard design does and the larger conductor might entail a need for not only larger and more expensive wire, but heavier hardware to hold it in place and stronger towers and poles to keep it in the air. In addition, these changes may produce other costs - for example, use of a larger conductor not only lowers losses, but a higher fault duty (short circuit current), which increases the required rating and cost for circuit breakers. Regardless, initial equipment costs can be balanced against long-term financial losses through careful study of needs, performance, and costs, to establish a minimum overall (present worth) worth.

Load-related losses

Flow of electric power through any device is accompanied by what are called load-related losses, which increase as the power flow (load) increases. These are due to the impedance of the conductor or device. Losses increase as the *square* of the load - doubling the power flowing through a device quadruples the losses. Tripling power flow increases the losses by nine.

With very few exceptions, larger electrical equipment always has lower impedance, and thus lower load-related losses, for any given level of power delivery. Hence, if the losses inherent in delivering 5 MW using a 600 MCM conductor are unacceptably large, the use of a 900 MCM conductor will reduce them considerably. The cost of the larger conductor can be weighed against the savings in reduced losses to decide if it is a sound economic decision.

No-load losses

“Wound” T&D equipment - transformers and regulators - have load-related losses as do transmission lines and feeders but in addition, they have a type of electrical loss that is constant, regardless of loading. No-load losses constitute the electric power required to establish a magnetic field inside these units, without which they would not function. Regardless of whether a transformer has any load - any power passing through it at all - it will consume a small amount of power, generally less than 1% of its rated full power, simply because it is energized and “ready to work.” No-load losses are constant, and occur 8,760 hours per year.

Given similar designs, a transformer will have no load losses proportional to its capacity. A 10 MVA substation transformer will have twice the no-load losses of a 5 MVA transformer of similar voltage class and design type. Therefore, unlike the situation with a conductor, selection of a larger transformer does not always reduce net transformer losses, because while the larger transformer will always have lower load-related losses, it will have higher no-load losses, and this increase might outweigh the reduction in load-related losses. Again, low-loss transformers are available, but cost more than standard types. Lower-cost-than-standard but higher-loss transformers are also available (often a good investment for backup and non-continuous use applications). Overall, minimizing lifetime cost is a balanced act of now versus then and one factor against another.

The costs of losses

The electric power required to operate the T&D system - the electrical losses - is typically viewed as having two costs, demand and energy. Demand cost is the cost of providing the peak capacity to generate and deliver power to the T&D equipment. A T&D system that delivers 1,250 MW at peak might have losses during this peak of 100 MW. This means the utility must have generation, or buy

power at peak to satisfy this demand, whose cost is calculated using the utility's power production cost at time of peak load. This is usually considerably above its average power production cost.

Demand cost also ought to include a considerable T&D portion of expense. Every service transformer in the system (and there are many) is consuming a small amount of power in doing its job at peak. Cumulatively, this might equal 25 MW of power - up to 1/4 of all losses in the system. That power must not only be generated by the utility but also transmitted over its transmission system, through its substations, and along its feeders to *reach* the service transformers. Similarly, the power for electrical losses in the secondary and service drops (while small, are numerous and low voltage, so that their cumulative contribution to losses is noticeable) has to be moved even farther, through the service transformers and down to the secondary level.

Demand cost of losses is the total cost of the capacity to provide the losses and move them to their points of consumption.

Energy losses occur whenever the power system is in operation, which generally means 8,760 hours per year. While losses vary as the square of load, so they drop by a considerable margin off-peak. Their steady requirement every hour of the year imposes a considerable energy demand over the course of a year. This cost is the cost of the energy to power the losses.

Example: Consider a typical 12.47 kV, three-phase, OH feeder, with 15 MW capacity (600 MCM phase conductor), serving a load of 10 MW at peak with 4.5% primary-level losses at peak (450 kW losses at peak), and having a load factor of 64% annually. Given a levelized capacity cost of power delivered to the low side bus of a substation of \$10/kW, the demand cost of these losses is \$4,500/year. Annual energy cost, at 3.5¢ /kWhr, can be estimated as:

$$450 \text{ kW losses at peak} \times 8,760 \text{ hr} \times (64\% \text{ load factor})^2 \times 3.5\text{¢} = \$56,500 \quad (2.1)$$

Thus, the financial losses (demand plus energy costs) for this feeder are over \$55,000 annually. At a present worth discount factor of around 11%, this means losses have an estimated present worth of about \$500,000. This computation used a simplification - squaring the load factor to estimate load factor impact on losses - which tends to underestimate losses slightly. Actual losses probably would be more in the neighborhood of \$565,000 PW. If the peak load on this feeder were run up to its maximum rating (about 15 MW instead of 10 MW) with a similar load factor of 64%, annual losses and their cost would increase to $(15/10)^2$ or \$1,250,000.

This feeder, in its entirety, might include four miles of primary trunk (at \$150,000/mile) and thirty miles of laterals (at \$50,000/mile), for a total capital cost of about \$2,100,000. Thus, *total losses are on the order and magnitude of original cost of the feeder itself*, and in cases where loading is high, can approach that cost. Similar loss-capital relations exist for all other levels of the T&D

system, with the ratio of losses capital cost increasing as one nears the consumer level (lower voltage equipment has higher losses/kW).

Total T&D Costs

What is the total cost to deliver electric power? Of course this varies from system to system as well as from one part of a system to another - some consumers are easier to reach with electric service than others (Figure 2.9). Table 2.2 shows the cost of providing service to a “typical” residential consumer. Generally, a utility amortizes the capital cost of 30 or more years and tries to minimize the total present worth lifetime cost. More detail in how costs are accounted and minimized in planning and operations can be found in Brown, 2009.

2.7 TYPES OF DELIVERY SYSTEM DESIGN

There are three fundamentally different ways to lay out a power distribution system used by electric utilities, each of which has variations in its own design. As shown in Figure 2.10, radial, loop, and network systems differ in how the distribution feeders are arranged and interconnected about a substation.

Most power distribution systems are designed as *radial distribution systems*. The radial system is characterized by having only one path between each consumer and a substation. The electrical power flows exclusively away from the substation and out to the consumer along a *single path*, which, if interrupted, results in complete loss of power to the consumer. Radial design is by far the most widely used form of distribution design, accounting for over ninety-nine percent of all distribution construction in North America. Its predominance is due to two overwhelming advantages: it is much less costly than the other two alternatives and it is much simpler in planning, design, and operation.

In most radial plans, both the feeder and the secondary systems are designed and operated radially. Each radial primary distribution feeder serves a definite service area (all consumers in that area are provided power by only that feeder) of somewhere around 1,000 to 4,000 customers. Most radial feeder systems are actually built as networks, but operate radially by opening switches at certain points throughout the physical network (shown earlier in Figure 2.6), so that the resulting configuration is electrically radial. The planner determines the layout of the network and the size of each feeder segment in that network and decides where the open points should be for proper operation as a set of radial feeders.

Table 2.2 Cost of Providing Service to Typical Residential Consumers

Level	Cost Components	Cost
Transmission	4 kW x 100 miles x \$.75/kW mile	\$300
Substation	4 kW x \$60/kW	\$240
Feeder	4 kW x 1.5 miles x \$10/kW-mile	\$60
Service	1/10th of 50 kVA local service system	\$300
Smart Meter		<u>\$200</u>
Total Initial Cost (Capital)		\$1,100
Operations, Maintenance, and Taxes (PW next 30 years) =		\$700
Cost of Electrical Losses (PW next 30 years) =		<u>\$700</u>
Estimated Cost of Power Delivery, 30 years, PW		\$2,500




Figure 2.9 Cost of power delivery varies depending on location. Shown here are the annual capacity costs of delivery evaluated on a ten-acre basis throughout a coastal city of population 250,000. Cost varies from a low of \$85/kW to a high of \$270/kW.




Figure 2.10 Simplified illustration of the concepts behind three types of power distribution configuration. Radial systems have only one electrical path from the substation to the consumer, loop systems have two, and networks have several. Arrows show most likely direction of electric flows.

A further attribute of most radial feeder system designs, but not an absolutely critical element, is the use of single-phase laterals. Throughout North America, most utilities use single- and two-phase laterals to deliver power over short distances, tapping off only one or two phases of the primary feeder. This minimizes the amount of wire that needs to be strung for the short segment required to get the power in the general vicinity of a dozen or so consumers. These laterals are also radial, but seldom, if ever, end in a switch (they just end). There are many utilities, particularly urban systems in Europe, Africa, and Asia that build every part of the radial distribution system, including laterals, with all three phases.

Each service transformer in these systems feeds power into a small radial system around it, basically a single electrical path from each service transformer to the consumers nearby.

Regardless of whether it uses single-phase laterals or not, the biggest advantages of the radial system configuration, in addition to its lower cost, is the simplicity of analysis and predictability of performance. Because there is only one path between each consumer and the substation, the direction of power flow is absolutely certain. Equally important, the load on any element of the system can be determined in the most straightforward manner - by simply adding up the entire consumer loads “downstream” from that piece of equipment.

Before the advent of economical and widely available computer analysis, this alone was an overwhelming advantage, for it allowed simple, straightforward, “back of the envelope” design procedures to be applied to the distribution system with confidence that the resulting system would work well. The simplicity of

analysis, and confidence that operating behavior is strictly predictable are still great advantages.

Because load and power flow direction are easy to establish, voltage profiles can be determined with a good degree of accuracy without resorting to exotic calculation methods. Capacity requirements can be ascertained exactly. Fault levels can be predicted with accuracy and protective devices - breaker-relays and fuses - can be coordinated in an absolutely assured manner, without resorting to network methods of analysis. Regulators and capacitors can be sized, located, and set using relatively simple procedures (simple compared to those required for similar applications to non-radial designs in which the power flow direction is not a given).

On the debit side, radial feeder systems are less reliable than loop or network systems because there is only one path between the substation and the consumer. Thus, if any element along this path fails, a loss of power delivery results. Generally, when such a failure occurs, a repair crew is dispatched to re-switch temporarily the radial pattern network, transferring the interrupted consumers onto another feeder, until the damaged element can be repaired. This minimizes the period of outage, but an outage still occurred because of the failure.

Despite this apparent flaw, radial distribution systems, if well designed and constructed, generally provide very high levels of reliability. For all but the most densely populated areas, or absolutely critical loads (hospitals, important municipal facilities, the utility's own control center) the additional cost of an inherently more reliable configuration (loop or network) cannot possibly be justified for the slight improvement that is gained over a well-designed radial system.

An alternative to purely radial feeder design is a *loop system* consisting of a distribution design with two paths between the power sources (substations, service transformers) and every consumer. Such systems are often called "European," because this configuration is the preferred design of many European utilities, as well as European electrical contractors when called upon to lay out a power distribution system anywhere in the world. Equipment is sized and each loop is designed so that service can be maintained regardless of where an open point might be on the loop. Because of this requirement, whether operated radially (with one open point in each loop), or with closed loops, the basic equipment capacity requirements of the loop feeder design do not change.

Some urban areas in Europe and Asia are fed by multiple hierarchical loop systems. A 100+kV sub-transmission loop routes power to several substations, out of which several loop feeders distribute power to service transformers, which each route powers through a long loop secondary.

In terms of complexity, a loop feeder system is only slightly more complicated than a radial system. Power usually flows out from both sides toward the middle,

and in all cases can take only one of two routes. Voltage drop, sizing, and protection engineering are only slightly more complicated than for radial systems.

But if designed thus, and if the protection (relay-breakers and sectionalizers) is also built to proper design standards, the loop system is more reliable than radial systems. Service will not be interrupted to the majority of consumers whenever a segment is outaged, because there is no "downstream" portion of any loop.

The major disadvantage of loop systems is capacity and cost. A loop must be able to meet all power and voltage drop requirements when fed from only one end, not both. It needs extra capacity on each end, and the conductor must be large enough to handle the power and voltage drop needs of the entire feeder if fed from either end. This makes the loop system inherently more reliable than a radial system, but the larger conductor and extra capacity increase costs.

Distribution networks are the most complicated, the most reliable, and in very rare cases also the most economical method of distributing electric power. A network involves multiple paths between all points in the network. Power flow between any two points is usually split among several paths, and if a failure occurs it instantly and automatically re-routes itself.

Rarely does a distribution network involve primary voltage-level network design, in which all or most of the switches between feeders are closed so that the feeder system is connected between substations. This is seldom done because it proves very expensive and often will not work well.⁹ Instead, a "distribution network" almost always involves "interlaced" radial feeders and a network secondary system - a grid of electrically strong (i.e., larger than needed to feed consumers in the immediate area) conductor connecting all the consumers together at utilization voltage. Most distribution networks are underground simply because they are employed only in high-density areas, where overhead space is not available.

In this type of design, the secondary grid is fed from radial feeders through service transformers, basically the same way the secondary is fed in radial or loop systems. The feeders are radial, but laid out in an interlaced manner - none has a sole service area, but instead they overlap. The interlaced configuration means that alternate service transformers along any street or circuit path are fed from alternate feeders, as shown in Figure 2.11.

While segments from two feeders always run parallel in any part of the system, the same two feeders never overlap for all of their routing. The essence of the interlaced system (and a design difficulty in any practical plan) is to mix up feeders so that each feeder partially parallels a few other feeders. Thus, if the

⁹ The major reason is this puts feeder network paths in parallel with transmission between substations, which results in unacceptable loop and circular flows as well as large dynamic shifts in load on the distribution system.




Figure 2.11 To obtain an interlacing factor of 2, two feeders are routed down each street, with alternating network transformers fed from each.




Figure 2.12 Top: a non-interlaced feeder system experiences the loss of one feeder, and all transformers in the lower right part of the system are lost - service is certain to be interrupted. Bottom: the same system interlaced. Loss of the feeder is a serious contingency, but can be withstood because the feeder losses are distributed in such a way that transformers still in service surround each transformer out of service.

feeder fails, it spreads its load out over a few other feeders, overloading none severely (Figure 2.12).

At a minimum, distribution networks use an interlacing factor of two, meaning that two feeders overlap in any one region, each feeding every other service transformer. But such a system will fail when any two feeders are out of service. Interlacing factors as high as five (four overlapping feeders, each feeding every fourth

Networks have proven more expensive but more reliable than radial distribution systems. In urban applications where the load density is high, the electric facilities must be placed underground, and repairs and maintenance are difficult because of traffic and congestion, networks may cost only a little more to build and maintain than other types of underground systems. The best network systems can tolerate the loss of any three feeders (the other two in any area picking up the remaining load, although often very overloaded) without any interruption of consumer service. Networks require only a little more conductor capacity than a loop system. The loop configuration required “double capacity” everywhere to provide increased reliability.

Networks have two disadvantages. First, they are *much* more complicated, than other forms of distribution, and thus much more difficult to analyze and more expensive to operate. There is no “downstream” side to each unit as in a radial or loop system, so that the load seen by any unit of equipment cannot be obtained by merely adding up consumers on one side of it, nor can the direction of power flow through it be assumed. Loadings and power flow, fault currents and protection must be determined by network techniques such as those used by transmission planners. However, even more sophisticated calculation methods than those applied to transmission may be required, because a large distribution network can consist of 50,000 nodes or more - the size of the very largest transmission-level power pool. Distribution network load flows are often more difficult to solve than transmission systems because the range of impedance in the modeled circuits is an order of magnitude wider.

The second disadvantage distributed and distribution networks can have in comparison to other types of distribution designs is in system protection and in the field force safety procedures. Networks can – in fact usually do – have *extremely* high fault duties in the vaults and manholes where electrical equipment is located. The authors have seen fault levels as high as 160,000 amps on the low side of distributed network transformers in large downtown networks. Proper use of protective suits and gear, and good work practices, can mitigate any hazard to workers in such vaults, but these prudent precautions cost money and slow the pace of maintenance and service compared to what can be done on systems with lower fault levels. This cost issue has caused more than one utility to decide to convert older distributed networks to another type of distribution system, such as primary select radial switching.




Figure 2.13 Two ways to route a radial feeder to 108 service transformers. Left, a "multi-branch" configuration. Right, a "large trunk" design.




Figure 2.14 A power system is divided by substation service boundaries into a set of substation service areas, as shown.

Large-Trunk vs. Multi-Branch Feeder Layout

Figure 2.13 illustrates the basic concept behind two different ways to approach the layout of a radial distribution system. Each has advantages and disadvantages with respect to the other in certain situations, and neither is superior to the other in terms of reliability, cost, ease of protection, and service quality. Either can be engineered to work in nearly any situation. Most planning utilities have an institutionalized preference for one or the other. Beyond showing that there are significantly different ways to lay out a distribution system, this brings to light an important point about distribution design: major differences in standards exist among electric utilities, as a result of which comparison of statistics or practice from one to the other is often not valid. Feeder layout types and practices are discussed in much greater detail in the *Power Distribution Planning Reference Book*.

Service Areas

As mentioned earlier, in most power systems, each substation is usually the sole provider of electrical service to the region around it - its service area. Similarly, feeders and distribution networks also have distinct service areas. Usually, the service area for a substation, feeder, or other unit of equipment is the immediate area surrounding it, and usually these service areas are contiguous (i.e. not broken into several parts) and exclusive - no other similar distribution unit serves any of the load in an area. As an example, Figure 2.14 shows a map of substation service areas for a rectangular portion of a power system. Each distribution substation exclusively serves all consumers in the area containing it.

Cumulatively, the simultaneous peak demand of all customers in its service area defines the maximum power the substation must serve. Within a power system, each individual part, such as a substation or service transformer, will see its peak load at whatever time and in whatever season the consumers in its service area generate their cumulative peak demand. One result of this is that the peak loads for different substations often occur at different seasons of the year or hours of the day. But whenever the peak occurs, it defines the maximum power the unit is required to deliver. Peak demand is one of the most important criteria in designing and planning distribution systems. Usually it defines the required equipment capacity.

Dynamic Service Area Planning

By making switching changes in the distribution system, it is possible to expand or shrink a substation or feeder's service area significantly, increasing or decreasing its net load. This is an important element of T&D planning, as shown in Figure 2.15, which illustrates a very typical T&D expansion situation. Two neighboring substations, A and B, each have a peak load near the upper limit of

their reliable load-handling range. Load is growing slowly throughout the system, so that in each substation annual peak load is increasing at about 1 MW per year. Under present conditions, both will need to be upgraded soon. Approved transformer types, required to add capacity to each, are available only in 25 MVA increments, costing \$500,000 each. Both substations do not need to be reinforced. A new 25 MVA transformer and associated equipment are added to substation A, increasing its ability to handle a peak load by about 20 MVA. Ten MW of substation B's service area is then transferred to substation A. The result is that each substation has 10 MW of margin for continued load growth - further additions are not needed for 10 years.




Figure 2.15 Load in both substations is growing at about 1 MW per year. Each substation has sufficient capacity to handle present load within contingency criteria (a 25% margin above peak) but nothing more. By transferring load as shown above, only one substation has to be reinforced with an additional (25 MVA) capacity, yet both end up with sufficient margin for another ten years' growth. Service area shifts keep expansion costs down in spite of the fact that equipment like transformers are available only in large, discrete sizes.

This type of planned variation in service areas is a major tool used to keep T&D expansion costs low, a key element in building a reliable, economical schedule of expansion. Optimization of this particular aspect of planning can be a challenge, not because of any inherent difficulty in analysis, but simply because there are so many parts of the system and constraints to track at one time. This is one reason for the high degree of computerization in distribution planning at many utilities. Balancing the myriad requirements of many substations and their design and operating constraints is an ideal problem for numerical optimization techniques.

The Systems Approach

One complication in determining the most economical equipment for a power system is that its various levels - transmission, substation, and distribution - are interconnected, with the distribution in turn, connected to the consumers and hence a function of the local load density. The best size and equipment type at each level of the system is a function of the types of equipment selected for the other levels of the system and the load density. In general, the equipment is so interconnected that it is impossible to evaluate any one aspect of design without taking all others into account.

Consider the question of substation spacing - determining how far apart substations should be, on average for best utilization and economy. Within any service territory, if substations are located farther apart, there will be fewer of them, saving the utility the cost of buying and preparing some substation sites, as well as reducing the cost of building a large number of substations.

Thus, overall substation cost is a function of spacing, but how it varies depends on land and equipment costs, which must be balanced carefully. With fewer substations however, each substation must serve a larger area of the system. Hence, each substation will have a larger load and thus require a larger capacity, meaning it must have more or larger transformers.

The larger substations will also require a larger amount of power to be brought to each one, which generally calls for a higher sub-transmission voltage. Yet there will be fewer sub-transmission lines required because there are fewer substations to which power must be delivered. All three aspects of layout are related - greater substation spacing calls for larger substations with bigger transformers, and a higher transmission voltage, but fewer lines are needed - and all three create better economies of scale as spacing is increased. Thus, transmission costs generally drop as substation spacing is increased.

Nevertheless, there is a limit to this economy. The distribution feeder system is required to distribute each substation's power through its service area, moving power out to the boundary between each substation's service area and that of its neighbors. Moving substations farther apart means that the distribution system must move power, on average, a greater distance. Distributing power over these

longer distances requires longer and more heavily loaded feeders. This situation increases voltage drop and can produce higher losses, all of which can increase cost considerably. Employing a higher distribution voltage (such as 23.9 kV instead of 13.8 kV) improves performance and economy, but it costs more to distribute power from a few larger substations farther apart, than to distribute it from many smaller substations close together. Feeder costs go up rapidly as substation spacing is increased.

The major point of this section is that all four of the above aspects of system design are interconnected - 1) substation spacing in the system, 2) size and number of substations, 3) transmission voltage and design and, 4) distribution feeder voltage and design. One of these factors cannot be optimized without close evaluation of its interrelationship with the other three. Therefore, determining the most cost-effective design guideline involves evaluating the transmission-substation-feeder system design as a whole against the load pattern, and selecting the best combination of transmission voltage, substation transformer sizes, substation spacing, and feeder system voltage and layout.




Figure 2.16 Overall cost of a T&D system depends on the balanced design of the sub-transmission, substation, and feeder level, as described in the text. Cost can vary by significant margins depending on how well performance and economy at various levels of the system are coordinated.

This economic equipment sizing and layout determination is based on achieving a balance between two conflicting cost relationships:

1. Higher voltage equipment is nearly always more economical on a per-MW basis.
2. Higher voltage equipment is available only in large sizes (lots of MW).

In cases where the local area demands are modest, higher voltage equipment may be more expensive simply because the minimum size is far above what is required - the utility has to buy more than it needs. How these two cost relationships play against one another depends on the load and the distances over which power must be delivered. Other factors to be considered are those unique to each power system, such as the voltages at which power is delivered from the regional power pool, and whether the system is underground or overhead.

Figure 2.16 illustrates the difference that careful coordination of system design between levels of the power system can have in lowering overall cost. Shown are the overall costs from various combinations of a T&D system layout for a large metropolitan utility in the eastern United States. Each line connects a set of cost computations for a system built with the same transmission and distribution voltages (e.g., 161 kV transmission and 13.8 kV distribution) but varying in substation sizes (and hence, implicitly, their spacing).

In all cases, the utility had determined it would build each substation with two equally sized transformers (for reliability), with none over 75 MVA (larger transformers are too difficult to move along normal roads and streets, even on special trailers). Either 161 kV or 69 kV could be used as sub-transmission, either 23.9 kV or 13.8 kV could be used as distribution voltage. Any size transformer, from 15 MVA to 75 MVA could be used, meaning the substation could vary from 30 MVA to 150 MVA in size. (Peak load of such substations can normally be up to 75% of capacity, for a peak load of from 23 to 100 MW.) Substation spacing itself is implicit and not shown. Given the requirement to cover the system, determining transmission voltage, distribution, and substation size defines the system design guidelines entirely.

Overall, the ultimate and lowest cost T&D system guidelines are to build 120 MVA substations (two 60 MVA transformers) fed by 161 kV sub-transmission and distributing power at 23.9 kV. This has a levelized cost (as computed for this utility) of about \$179/kW. In this particular case, a high distribution voltage is perhaps the most important key to good economy - if 13.8 kV is used instead of 23.9 kV as the primary voltage, minimum achievable cost rises to \$193/kW.

The very worst design choices plotted in Figure 2.16, from an economic standpoint, would be to build 25 MVA substations fed by 161 kV

sub-transmission and feeding power to 23.9 kV feeders (\$292/kW). This would require many small substations, each below the effective size of both the transmission and distribution voltages, and lead to high costs. Overall, 161 kV and 23.9 kV are the correct choices for economy, but only if used in conjunction with a few, large substations. If substations are to be 25 MVA, then 69 kV and 13.8 kV do a much more economical job (\$228/kW), but still don't achieve anything like the optimum value. The point: Achieving economy in power delivery involves coordinating the interactions, performance, and economies of the multiple system levels. Chapters 11 and 12 discuss the issues of and techniques for such coordinated multi-level planning.

2.8 SMART GRID

Smart Grid refers to power systems that use sensors, digital and analog electronic control, and data communications among equipment and between the system and utility control centers, to enhance the capabilities and performance of the power system. Although the term “grid” typically refers to only the high-voltage (wholesale transmission) portion of a power system, the term “smart grid” is widely used to refer to any and all parts of a power system that employs such technology.

Beyond the inclusion of a great deal of “high technology” and data communication into the equipment and system, a key element of a Smart Grid is that the design of the power system itself, and way it is operated, are also changed in recognition and to take full advantage of the capabilities that the smart equipment provides.

For example, the traditional design for primary distribution in North America is to lay out circuits in a radial configuration and to arrange sectionalization (where and how fuses opened if there were problems) in a cascading sequence up and down the circuit. As was discussed in Section 2.7, the chief advantage of this approach is simplicity, of design, protection coordination, and operation, to a degree that fit the needs of the period during which the industry settled on this approach (first quarter of the 20th century). Circuits can be successfully designed and equipment properly specified with a minimum of information and, quite literally, on the back of the envelope calculations.¹⁰ The system is

¹⁰ This is not to say one cannot do a somewhat better job at assuring good performance and lowered cost using the greater data gathered by modern IT systems and the computerized tools in wide use today. However, with only a pencil and paper, and estimates of number of customers and peak load/customers, a competent engineer can layout a radial distribution system so it will work satisfactorily when all equipment is functioning as intended and fail safe when something goes wrong. In a period when the most advanced tools available to power engineers were a slide rule and a book of logarithms, this was a necessity.

straightforward to operate, dispatch, and refer in a safe manner with only maps as databases and common sense procedures used to assure safety. The industry stuck with radial distribution systems throughout the 20th century and into the 21st because they provided satisfactory performance for the time and no change was needed.

Considerable *improvement in every and all aspects* of a power distribution system's performance can be achieved if a more complicated basic design approach is used as compared to simple radial configurations. System cost, customer service reliability, voltage stability, power quality, electric losses, and safety can all be improved at the same time. The price paid for the simplicity of early radial distribution systems is a substantial margin in equipment and system capability that is often there but unused. In the early 20th century, this margin covered all the unknowns due to engineers and operators not having comprehensive load and system metering data at their fingertips, and for the approximations used in back-of-the-envelope calculations, so that circuits designed with linear models of non-linear processes would still work well.

A system designed without that margin can be cheaper – perhaps up to 15% lower in overall cost – and more reliable with better quality of service. But engineers will need much more data in order to plan and engineer it. Modern smart data acquisition systems can gather this data, and modern Enterprise IT systems can provide it to the engineers in a way they can efficiently use it. Further, system operators will find the system a challenge to operate compared to the traditional radial system. It will require that they constantly adjust equipment settings in response to changes in load and system conditions, and that from time to time they recognize the system is in a different operating mode and that they have to re-orient some of its basic scheduling rules and priorities in order to assure that it “fails safe” should something go wrong.

All of this can be accommodated, usually in a completely automatic way, by the proper use of smart grid equipment and designs. The distribution system is no longer entirely radial in configuration or operation. It is deliberately designed so that it will have to change its configuration (switching) and equipment settings (voltage tap levels and capacitor switching) in order to function well throughout the course of a year or even perhaps a single day. Automatic equipment can do that. Protection and fail safe capability require that one unit of equipment “know” the status and loading of others nearby: data communication can provide this.

The net result is that a smart distribution system has better performance in every category than a traditional system, including the performance category cost. There is a price to pay, however. Smart systems are *much* more complex from nearly every standpoint, and require advanced design tools and planners, engineers and operators of exceptional skill and training to plan, design, and operate them.

2.9 CONCLUSION

A transmission and distribution system moves power from a utility's power production and purchase points to its consumers. The T&D system's mission of delivering power to the consumers means that it must be composed of equipment spread throughout the service territory and arranged locally so that capacity is always in proportion to local electrical demand. The facilities need to be in each neighborhood, sized and configured both to fit well into the whole and to serve the local needs. In the authors' opinion, engineering a system to meet this challenge is seldom easy, but engineering a system to do so at the minimum possible cost is always extremely challenging.

A T&D system is composed of several interconnected, hierarchical levels, each of which is required for completion of the power delivery task. These levels are:

1. Transmission
2. Sub-transmission
3. Substation
4. Feeder
5. Secondary
6. Consumer

To a certain extent, the burden of power delivery, and costs, can be shifted from one level to another through changes in the specifications of equipment, layout standards, and design of the various levels. For example, costs can be pared at the substation level by using fewer, larger substations. This means, feeders in each substation area will have to carry power farther, and perhaps carry more power per feeder as well, increasing feeder costs. Low overall design cost is achieved by balancing these factors.

The performance and economy and thus the design of a power system are dominated by a number of constraints due to physical laws, and further shaped by a number of practical considerations with regard to equipment, layout, and operating requirements. The more important of these are:

1. *The T&D System Must Cover Ground.* This is *the* rule about T&D - ultimately the electric system must "run a wire" to every consumer. A significant portion of the cost of a distribution system is due to this requirement alone, independent of the amount of peak load or energy supplied.

2. *Interconnectivity of Levels.* The sub-transmission/substation-feeder triad comprises a highly interconnected system with the electrical and economic performance at one level *heavily* dependent on design, siting, and other decisions at another level. To a certain extent the T&D system must be planned and designed as a whole - its parts cannot be viewed in isolation.
3. *Discrete Equipment Sizes.* In many cases equipment is available only in certain discrete sizes. For example, 69 kV/12.47kV transformers may be available only in 5, 10, 20, and 22 MVA sizes. Usually, there is a large economy of scale - the installed cost of a 20 MVA unit being considerably less than two times that for a 10 MVA unit.
4. *Dynamic Service Area Assignment is a Powerful Planning Tool.* Maximization of equipment utilization and minimization of cost per unit of capacity can be obtained despite the fact that capacity upgrades can be built only in one of several discrete sizes. To accomplish this the planner changes the boundaries of the service area during the planning process. When reinforced with new capacity additions, a substation might pick up load from surrounding substations, effectively spreading the capacity addition among several substations.
5. *Losses Cost Can Be Significant.* In some cases, the present worth of future losses on a line or heavily loaded transformer can exceed its total capital cost. In most cases, this does not happen. But in most cases, the present worth of losses is greater than the difference in cost between most of the choices available to the planner. This means loss is a major factor to be considered in achieving overall least-cost design.
6. *Cost to Upgrade is Greater than the Cost to Build.* For example, one mile of a 12.47 kV, three-phase feeder using a 600 MCM conductor (15 MW thermal capacity) might cost \$150,000 to build, and one mile of 9 MW feeder (336 MCM conductor) might cost only \$110,000. But the cost to upgrade the 336 MCM feeder to 600 MCM wire size at a later date would be about \$130,000, for a cumulative total of \$240,000. Therefore, one aspect of minimizing cost is to determine size for equipment not on the basis of immediate need, but by assessing the eventual need and determining whether the present worth of the eventual savings warrants the investment in larger size now.
7. *Coincidence of Peaks.* Not all consumer loads occur at the same time. This has a number of effects. First, peak load in different parts of the system may occur at different times. Second, the system peak load will

always be less than the sum of the peak loads at any one level of the system. For example, in more power systems the sum of all substation peak loads usually exceeds system total peak by 3% to 8%. Diversity of peak loads means that considerable attention must be paid to the pattern and timing of electric load if equipment needs (and consequently costs) are to be minimized.

8. *Reliability is Obtained through Contingency Margin.* Traditionally, T&D reliability in a static system (one without automation, and hence incapable of recognizing and reacting to equipment outages and other contingencies by re-configuring and/or re-calibrating equipment) is assured by adding an emergency, or contingency margin throughout the system. At the T&D level this means that an average of 20% to 50% additional capacity is included in most equipment, so that it can pick up additional load if neighboring equipment fails or is overloaded.

A Complicated System

In conclusion, a power T&D system provides the delivery of an electric utility's product. It must deliver that product to every consumer's site, with high reliability and in ready-to-use form. While composed of equipment that is individually straightforward, most T&D systems are quite complex due to the interactions of thousands, even millions, of interconnected elements. Achieving economy and reliability means carefully balancing a myriad of mutual interactions with conflicting cost tradeoffs.

REFERENCES AND BIBLIOGRAPHY

- R.E. Brown, *Electric Power Distribution Reliability – Second Edition*, CRC Press, Boca Raton, 2009.
- R.E. Brown, *Business Essentials for Utility Engineers*, CRC Press, Boca Raton, 2010.
- J. J. Burke, *Power Distribution Engineering – Fundamentals and Applications*, Marcel Dekker, New York, 1994.
- M. V. Engel et al., editors, *Tutorial on Distribution Planning*, IEEE Course Text EHO 361-6-PWR, Institute of Electrical and Electronics Engineers, Hoes Lane, NJ, 1992.
- A. C. Monteith and C. F. Wagner, *Electrical Transmission and Distribution Reference Book*, Westinghouse Electric Company, Pittsburgh, PA, 1964.

H. L. Willis, *Power Distribution Planning Reference Book – Second Edition*, Marcel Dekker, New York, 1997.

H. L. Willis and W. G. Scott, *Distributed Power Generation – Planning and Evaluation*, Marcel Dekker, New York, 2000.

3

Customer Demand for Power and Reliability of Service

3.1 THE TWO Qs: QUANTITY AND QUALITY OF POWER

Electric consumers require power, whether delivered from the utility grid or generated locally by distributed sources, in order to help accomplish the uses for which they need energy. Their need for electric power, and the value they place upon its delivery to them, has two interrelated but fundamentally separate dimensions. These are the two Qs: Quantity, the amount of power needed, and Quality, the most important aspect of which is usually dependability of supply (reliability of power supply, or availability as it is often called). The relative importance of these two features varies from one consumer to another depending on their individual needs, but each consumer finds value in both the amount of power he obtains, and its availability as a constant, steady source that it will be there whenever needed.

This chapter discusses demand and use of electric power as seen from the customer's standpoint: the utility's job is to satisfy consumer needs as fully as possible within reasonable cost constraints. Cost is very much an important aspect to consumers too, so both the utility and the consumer must temper their plans and desires with respect to power and reliability based on real world economics. Consumers do not get everything they want; only what they are willing to pay for.

Utilities should not aim to provide flawless service, which would be prohibitively expensive, but instead aim to provide the highest level possible within economic constraints of the customers' willingness to pay.

This chapter begins, in Section 3.2, with a discussion of consumer use of electricity and includes the quantity of electric demand as seen from an "end-use" perspective. It continues with how demand varies as a function of consumer type and end-use, and how power demand is represented in electric system studies using load curves and load duration curves. Section 3.3 then discusses reliability and availability as seen by the customers and ways this can be characterized and studied. Section 3.4 briefly reviews Two-Q analysis and planning concepts and their application to customer load analysis. Finally, Section 3.5 provides a summary of key points.

3.2 ELECTRIC CONSUMER NEED FOR QUANTITY OF POWER

People and businesses buy power because they want the products that electrical use can provide

No consumer actually wants the electric energy itself. Consumers want the *products* it can provide – a cool home in summer, hot water on demand, compressed air for manufacturing, electronic robotic factory control, cold beer in the 'fridge and football on color TV. Electricity is only an intermediate means to some end-use.

These different goals and needs are called *end-uses*, and they span a wide range of applications. Some end-uses are unique to electric power (the authors are not aware of any manufacturer of natural gas powered TVs, stereos, or computers). For many other end-uses, electricity is only one of several possible energy sources (water heating, home heating, cooking, or clothes drying). In many other end-uses, electricity is so convenient that it enjoys a virtual monopoly, even though there are alternatives, e.g., gasoline-powered refrigerators, and natural gas for interior lighting and for air conditioning.

Each end-use is satisfied through the application of appliances

Lighting was the first and is still the largest single category or end use of power in many utility systems. A wide range of illumination devices are used, including incandescent bulbs, fluorescent tubes, sodium vapor, high-pressure monochromatic gas-discharge tubes, and in special cases, lasers. Each type of lighting device has differences from the others that give it an appeal to some consumers or for certain types of applications. Regardless, each requires electric power to function, creating an electric load when it is activated.




Figure 3.1 Electric peak demand of a utility in the southeastern United States broken down by consumer class and within the residential class, by contribution to peak for the major uses for which electricity is purchased at time of peak by the residential class.

Similarly, for other end-uses, such as space heating, there are various types of appliances, each with advantages or disadvantages in initial cost, operating efficiency, reliability and maintenance, noise and vibration, or other aspects. Buildings can be heated by heat pumps of various types and efficiencies, circulating hot water, or simple resistive heat. Each produces an electric load when used to produce heat for a home or business. Each has advantages (resistive heaters are cheap to buy, heat pumps use the least power of all heater types) and disadvantages (resistive heaters are inefficient, heat pumps are expensive to buy).

Consumer Classes

Different types of consumers purchase electricity. About half of all electric power is used in residences, which vary in the brands and types of appliances they own, and their daily activity patterns. Another fourth is consumed by commercial businesses both large and small, that buy electricity, having some similar end-uses to residential consumers (heating and cooling, and illumination), but that

have many needs unique to commercial functions (cash register/inventory systems, escalators, office machinery, neon store display lighting, parking lot lighting). Finally, industrial facilities and plants buy electricity to power processes such as pulp heating, compressor and conveyor motor power, and a variety of manufacturing applications. As a result, the load on an electric system is a composite of many consumer types and appliance applications. Figure 3.1 shows this breakdown of the peak electrical load for a typical power system by consumer class and end-use category within one class.

Appliances Use Electric Power to Satisfy End-Use Demand

The term *load* refers to the electrical demand of a device connected to and drawing power from the electric system in order to accomplish some task, e.g., opening a garage door, or converting that power to some other form of energy, such as a light. Such devices are called *appliances*, whether they are a commonly regarded household item, e.g., a refrigerator, lamp, garage door opener, paper shredder, electric fence to keep cattle confined, etc. To the consumer, these appliances convert electricity into the end product. But the electric service planner can turn this relation around and view an appliance (e.g., a heat pump) as a device for transforming a demand for a particular end-use - warm or cool air - into electric load.

The level of power they need usually rates electrical loads, measured in units of real volt-amperes, called *watts*. Large loads are measured in kilowatts (thousands of watts) or megawatts (millions of watts). Power ratings of loads and T&D equipment refer to the device at a specific *nominal voltage*. For example, an incandescent light bulb might be rated at 75 watts and 1,100 lumens at 120 volts, at which voltage it consumes 75 watts and produces 1,100 lumens of light. If provided with less voltage, its load (and probably its light output) will fall.

By this definition, an electric vehicle (EV) such as a Chevrolet Volt is an appliance. It converts electric power drawn from the grid into an end use (personal transportation – the ability to move from place to place). EVs are one example of how and why electric energy consumption and the value society as a whole sees from utility service continues to grow.

Load Curves and Load Curve Analysis

The electric load created by any one end-use usually varies as a function of time. For example, in most households, demand for lighting is highest in the early evening, after sunset but before most of the household members have gone to bed. Lighting needs may be greater on weekends, when activity often lasts later into the evening, and at times of the year when the sun sets earlier in the day. Some




Figure 3.2 End-use demand. Right, average demand for BTU of cooling among houses in one of the authors' neighborhoods on a typical weekday in June. Left, lighting lumens used by a retail store on a typical day in June.




Figure 3.3 Electric demand for each class varies hourly and seasonally, as shown here, with a plot of average coincident load for residential users in central Florida.




Figure 3.4 Different consumer classes have different electrical demand characteristics, particularly with regard to how demand varies with time. Here are summer peak day load curves for the three major classes of consumer from the same utility system.




Figure 3.5 Demand on an hourly basis (blocks) over 24 hours, and the actual load curve (solid black line) for a feeder segment serving 53 homes. Demand measurement averages load over each demand interval (in this case each hour) missing some of the detail of the actual load behavior. In this case the actual peak load (263 kW at 6 PM) was not seen by the demand measuring, which “split the peak,” averaging load on an hourly basis and seeing a peak demand of only 246 kW, 7% lower. As will be discussed later in this chapter, an hourly demand period is too lengthy for this application.

end-uses are quite seasonal. Air-conditioning demand generally occurs only in summer, being greatest during particularly hot periods and when family activity is at its peak, usually late afternoon or very early evening. Figure 3.2 shows how the demand for two products of electric power varies as a function of time.

The result of this varying demand for the products of electricity application, is a variation in the demand for power as a function of time. This is plotted as a load *curve*, illustrated in Figure 3.3. Typically, the values of most interest to the planners are the peak load (maximum amount of power that must be delivered). This defines directly or indirectly, the capacity requirements for equipment; the minimum load and the time it occurs; the total energy, area under the curve that must be delivered during the period being analyzed; and the load value at the time of system peak load.

Consumer Class Load Curves

While all consumers differ in their electrical usage patterns, consumers within a particular class, such as residential, tend to have broadly similar load curve patterns. Those of different classes tend to be dissimilar in their demand for both quality and quantity and the time of day and year when their demand is highest. Therefore, most electric utilities distinguish load behavior on a class-by-class

basis, characterizing each class with a “typical daily load curve,” showing the average or expected pattern of load usage for a consumer in that class on the peak day, as shown in Figure 3.4. These “consumer class load curves” describe how the demand varies as a function of time. While often an entire 8,760-hour record for the year is available, usually only key days – perhaps one representative day per season – are used for studies.

The most important points concerning the consumers’ loads from the distribution planner’s standpoint are:

1. *Peak demand* and its time and duration
2. *Demand at time of system peak*
3. *Energy usage* (total area under the annual load curve)
4. *Minimum demand*, its time and duration

Details of Load Curve Measurement

Demand and demand periods

“Demand,” as normally used in electric load analysis and engineering, is the average value of electric load over a period of time known as the *demand interval*. Very often, demand is measured on an hourly basis as shown in Figure 3.5, but it can be measured on any interval basis – seven seconds, one minute, 30 minutes, daily, and monthly. The average value of power during the demand interval is given by dividing the kilowatt-hours accumulated during the demand interval by the length of the interval. Demand intervals vary among power companies, but those commonly used in collecting data and billing consumers for “peak demand” are 15, 30, and 60 minutes. Load curves may be recorded, measured, or applied over some specific time. For example, a load curve might cover one day. If recorded on an hourly demand basis, the curve consists of 24 values, each the average demand during one of the 24 hours in the day, and the peak demand is the maximum hourly demand seen in that day. Load data is gathered and used on a monthly basis and on an annual basis.

Load factor

Load factor is the ratio of the average to the peak demand. The average load is the total energy used during the entire period (e.g., a day, a year) divided by the number of demand intervals in that period (e.g., 24 hours, 8,760 hours). The average is then divided by the maximum demand to obtain the load factor:

$$\text{Load factor} = \text{kWh/Hrs/Peak kW} = \text{Average kW/Peak kW} \quad (3.1)$$

$$= \text{KWhr}/(\text{Peak kW} \times \text{Hr}) \quad (3.2)$$

Load factor gives the extent to which the peak load is maintained during the period under study. A high load factor means the load is at or near peak a good portion of the time.

Load Duration Curves

A convenient way to study load behavior is to order the demand samples from greatest to smallest, rather than as a function of time, as in Figure 3.6. The two diagrams consist of the same 24 numbers, in a different order. Peak load, minimum load, and energy (area under the curve) are the same for both.

Load duration curve behavior will vary as a function of the level of the system. Load duration curves for small groups of consumers will have a greater ratio of peak to minimum than similar curves for larger groups. Those for very small groups (e.g., one or two consumers) will have a pronounced “blockiness,” consisting of plateaus – many hours of similar demand level (at least if the load data were sampled at a fast enough rate). The plateaus correspond to combinations of major appliance loads. The ultimate “plateau” would be a load duration curve of a single appliance, for example a water heater that operated a total of 1,180 hours during the year. This appliance’s load duration curve would show 1,180 hours at its full load, and 7,580 hours at no load, without any values in between.

Hourly Load Curve




Figure 3.6 The hourly demand samples in a 24-hour load curve are reordered from greatest magnitude to least to form a daily load duration curve.




Figure 3.7 Annual load duration curve for a proposed commercial site with a 90 kW peak demand, from a DG reliability study in which three 40 kW distributed generators will be used to provide the power. As shown here, the demand exceeds 40 kW, the capacity of one unit, 2,800 hours a year. During those 2,800 hours, service can be obtained only if two or more units are operating. Information like this helps planners determine the equipment needs of power systems and project future operating costs.

Annual load duration curves

Most often, load duration curves are produced on an annual basis, reordering all 8,760 hourly demands (or all 35,040 quarter hour demands if using 15-minute demand intervals) in the year from highest to lowest, to form a diagram like that in Figure 3.7. The load shown was above 26 kW (demand minimum) 8,760 hours in the year, never above 92 kW, but above 40 kW for 2,800 hours.

basis, an often important aspect of siting for facilities, including DG (Willis, 1996).

Most often, load duration curves are produced on an annual basis, reordering all 8,760 hourly demands (or all 35,040 quarter hour demands if using 15-minute demand intervals) in the year from highest to lowest, to form a diagram like that in Figure 3.7. The load shown was above 26 kW (demand minimum) 8,760 hours in the year, never above 92 kW, but above 40 kW for 2,800 hours.

Load Duration Curves Do Not Provide Detailed Load Behavior Data

The load duration curve is a traditional tool used to represent the larger aspects of behavior of electric load. It developed prior to WWI and used in conjunction with other pre-computer era analytical tools such as break-point analysis for generation planning and K-factor analysis in distribution planning. In the 21st century, a load duration curve provides a useful way to summarize the major aspects of a utility's demand behavior, but does not provide the information needed for modern planning methods, nor put it in a form compatible with modern analytical system engineering tools. It is not a tool recommended by the authors for anything more than screening studies.

The major limitation of the load duration curve is that it does not provide information on the *order in time* of the load values it plots. The 2,800 hours when load is above 40 kW, shown in Figure 3.7, could be clustered entirely in the May – September timeframe, or spread evenly through the year as 7.7 hour periods every day. Such details are important both for determining the type and operating cost of the DG units, and for applying modern IRP methods that would seek to use load control and other means to maximize energy efficiency at the site.

Spatial Patterns of Electric Demand

An electric utility must not only produce or obtain the power required by its consumers, but also must deliver it to their locations. Electric consumers are scattered throughout the utility service territory, and thus the electric load can be thought of as distributed on a *spatial* basis as depicted in Figure 3.9. Just as load curves show how electric load varies as a function of time, (and can help identify when certain amounts of power must be provided), so has Spatial Load Analysis been used since its development in the 1970s (Scott, 1972) to identify where load is located and how much capacity is needed in each locality.

The electric demand in an electric utility service territory varies as a function of location depending on the number and types of consumers in each locality, as shown by the load map in Figure 3.9. Load densities in the heart of a large city can exceed 1 MW/acre, but usually average about 5 MW per square mile over the entire metropolitan area. In sparsely populated rural areas, farmsteads can be as far as 30 miles apart, and load density as low as 75 watts per square mile.




Figure 3.8 Left, map showing types of consumer by location for a small city. Right, map of electric demand for this same city.

Regardless of whether an area is urban, suburban, or rural, electric load is a function of the types of consumers, their number, their uses for electricity, and the appliances they employ. Other aspects of power system performance, including capability, cost, and reliability, can also be analyzed on a location- basis, an often important aspect of siting for facilities, including DG (Willis, 1996).

3.3 ELECTRIC CONSUMER NEED FOR QUALITY OF POWER

As mentioned in this chapter's introduction, a central issue in customer value of service analysis is matching availability and power quality against cost. T&D systems with near perfect availability and power quality can be built, but their high cost will mean electric prices the utility customers may not want to pay, given the savings an even slightly less reliable system would bring. All types of utilities have an interest in achieving the correct balance of quality and price. The traditional, franchised monopoly utility, in its role as the "electric resource manager" for the customers it serves, has a responsibility to build a system whose quality and cost balances its customers' needs. A competitive retail distributor of power wants to find the best quality-price combination - only in that way will it gain a large market share.

While it is possible to characterize various power quality problems in an engineering sense, characterizing them as interruptions, voltage sags, dips, surges, or harmonics, the customer perspective is somewhat different. Customers are concerned with only two aspects of service quality:

1. They want power when they need it.
2. They want the power to do the job.

If power is not available, neither aspect is provided. However, if power is available, but quality is low, only the second is not provided.

Assessing Value of Quality by Studying the Cost of a Lack of It

In general, customer value of reliability and service quality are studied by assessing the "cost" that something less than perfect reliability and service quality creates for customers. Electricity provides a value, and interruptions or poor power quality decrease that value. This value reduction - cost - occurs for a variety of reasons. Some costs are difficult, if not impossible to estimate: re-scheduling of household activities or lack of desired entertainment when power fails¹, or flickering lights that make reading more difficult.

But often, very exact dollar figures can be put on interruptions and poor power

¹ No doubt, the cost of an hour-long interruption that began fifteen minutes from the end of a crucial televised sporting event, or the end of a "cliffhanger" movie, would be claimed to be great.




Figure 3.9 Left, cost of a week-day interruption of service to a pipe rolling plant in the southeastern United States, as a function of interruption duration. An interruption of any length costs about \$5,000 - lost wages and operating costs to unload material in process, bring machinery back to “starting” position and restart - and a nearly linear cost thereafter. At right, present worth of the loss of life caused by harmonics in a 500 horsepower three-phase electric motor installed at that same industrial site, as a function of harmonic voltage distortion.

quality including food spoiled due to lack of refrigeration; wages and other operating costs at an industrial plant during time without power; damage to product caused by the sudden cessation of power; lost data and “boot up” time for computers; equipment destroyed by harmonics; and so forth. Figure 3.9 shows two examples of such cost data.

Value-Based Planning

To be of any real value in utility planning, information of the value customers put on quality must be usable in some analytical method that can determine the best way to balance quality against cost. Value-based planning (VBP) is such a method: it combines customer-value data of the type shown in Figure 3.9 with data on the cost to design the T&D system to various levels of reliability and power quality, in order to identify the optimum balance. Figure 3.10 illustrates the central tenet of value-based planning. The cost incurred by the customer due to various levels of reliability or quality, and the cost to build the system to various levels of reliability, are added to get the total cost of power delivered to the customer as a function of quality.² The minimum value is the optimum balance between customer desire for reliability and aversion to cost. This

² Figure 3.10 illustrates the concept of VBP. In practice, the supply-side reliability curves often have discontinuities and significant non-linearities that make application difficult. These and other details will be discussed in Chapter 5.




Figure 3.10 Concept of value-based planning. The customer's cost due to poorer quality (left) and the cost of various power delivery designs with varying levels of quality (center) are computed over a wide range. When added together (right) they form the total cost of quality curve, which identifies the minimum cost reliability level (point A).

approach can be applied for only reliability aspects, i.e., value-based reliability planning, or harmonics, or power quality overall. Generally, what makes sense is to apply it on the basis of whatever qualities (or lack of them) impact the customer - interruptions, voltage surges, harmonics, etc., in which case it is comprehensive value-based quality of service planning.

Cost of Interruptions

The power quality issue that affects the most customers, and which receives the most attention, is cessation of service often termed "service reliability." Over a period of several years, almost all customers served by any utility will experience at least one interruption of service. By contrast, a majority will never experience serious harmonics, voltage surge, or electrical noise problems. Therefore, among all types of power quality issues, interruption of service receives the most attention from both the customers and the utility. A great deal more information is available about cost of interruptions than about cost of harmonics or voltage surges.

Voltage Sags Cause Momentary Interruptions

The continuity of power flow does not have to be completely interrupted to disrupt service: If voltage drops below the minimum necessary for satisfactory operation of an appliance, power has effectively been "interrupted" as illustrated in Figure 3.11. For this reason many customers regard voltage dips and sags as momentary interruptions - from their perspective these are interruptions of the end-use *service they desire*, if not of voltage.

Much of the electronic equipment manufactured in the United States, as well as in many other countries, have been designed to meet or exceed the CBEMA




Figure 3.11 Output of a 5.2 volt DC power supply used in a desktop computer (top) and the incoming AC line voltage (nominal 113 volts). A voltage sag to 66% of nominal causes power supply output to cease within three cycles.




Figure 3.12 CBEMA curve of voltage deviation versus period of deviation, with the sag shown in Figure 3.11 plotted (black dot).

(Computer and Business Equipment Manufacturer's Association) recommended curves for power continuity, shown in Figure 3.12. If a disturbance's voltage deviation and duration characteristics are within the CBEMA envelope, then normal appliances should operate normally and satisfactorily.

However, many appliances and devices in use will not meet this criterion at all. Others will fail to meet it under the prevailing ambient electrical conditions (i.e., line voltage, phase unbalance power factor, and harmonics may be less than perfect).

The manner of usage of an appliance also affects its voltage sag sensitivity. The voltage sag illustrated in Figure 3.11 falls just within the CBEMA curve, as shown in Figure 3.12. The manufacturer probably intended for the power supply to be able to withstand nearly twice as long a drop to 66% of nominal voltage before ceasing output. However, the computer in question had been upgraded with three times the standard factory memory, a second and larger hard drive, and optional graphics and sound cards, doubling its power usage and the load on the power supply. Such situations are common and means that power systems that deliver voltage control within recommended CBEMA standards may still provide the occasional momentary interruption.

For all these reasons, there are often many more "momentary interruptions" at a customer site than purely technical evaluation based on equipment specifications and T&D engineering data would suggest. Momentary interruptions usually cause the majority of industrial and commercial interruption problems. In addition, they can lead to one of the most serious customer dissatisfaction issues. Often utility monitoring and disturbance recording equipment does not "see" voltage disturbances unless they are complete cessation of voltage, or close to it. Many events that lie well outside the CBEMA curves and definitely lead to unsatisfactory equipment operation are not recorded or acknowledged. As a result, a customer can complain that his power has been interrupted five or six times in the last month, and the utility will insist that its records show power flow was flawless. *The utility's refusal to acknowledge the problem irks some customers more than the power quality problem itself.*

Frequency and Duration of Interruptions Both Impact Cost

Traditional Power System Reliability Analysis recognizes that service interruptions have both frequency and duration (See Chapter 4). Frequency is the number of times during some period (usually a year) that power is interrupted. Duration is the amount of time power is out of service. Typical values for urban/suburban power system performance in North America are 2.2 interruptions per year with 100 minutes total duration.

Both frequency and duration of interruption impact the value of electric service to the customer and must be appraised in any worthwhile study of customer value of service. A number of reliability studies and value-based

planning methods have tried to combine frequency and duration in one manner or another into “one dimension.” A popular approach is to assume all interruptions are of some average length (e.g., 2.2 interruptions and 100 minutes is assumed to be 2.2 interruptions per year of 46 minutes each). Others have assumed a certain portion of interruptions are momentary and the rest of the duration is lumped into one “long” interruption (i.e., 1.4 interruptions of less than a minute, and one 99-minute interruption per year). Many other approaches have been tried (see References and Bibliography). But all such methods are at best an approximation, because frequency and duration impact different customers in different ways. No single combination of the two aspects of reliability can fit the value structure of all customers.

Figure 3.13 shows four examples of the author’s preferred method of assessing interruption cost, which is to view it as composed of two components, a fixed cost (Y intercept) caused when the interruption occurred, and a variable cost that increases as the interruption continues. As can be seen in Figure 3.13, customer sensitivity to these two factors varies greatly. The four examples are:

1. A pipe-rolling factory (upper left). After an interruption of any length, material in the process of manufacturing must be cleared from the welding and polishing machinery, all of which must be reset and the raw material feed set up to begin the process again. This takes about 1/2 hour and sets a minimum cost for an interruption. Duration longer than that is simply a linear function of the plant operating time (wages and rent, etc., allocated to that time). Prior to changes made by a reliability study, the “re-setting” of the machinery could not be done until power was restored (i.e., time during the interruption could not be put to use preparing to re-start once it was over). The dotted line shows the new cost function after modifications to machinery and procedure were made so that preparations could begin *during* the interruption.
2. An insurance claims office (upper right) suffers loss of data equivalent to roughly one hour’s processing when power fails. According to the site supervisor, an unexpected power interruption causes loss of about one hour’s work as well as another estimated half hour lost due to the impact of any interruption on the staff. Thus, the fixed cost of each interruption is equivalent to about ninety minutes of work. After one-half hour of interruption, the supervisor’s policy is to put the staff to work “on other stuff for a while,” making cost impact lower (some productivity); thus, variable interruption cost goes down. The dotted line shows the cost impact of interruptions after installation of UPS on the computer system, which permits orderly shutdown in the event of an interruption.




Figure 3.13 The author's recommended manner of assessing cost of interruptions includes evaluation of service interruptions on an event basis. Each interruption has a fixed cost (Y-intercept) and a variable cost, which increases as the interruption continues. Examples given here show the wide range of customer cost characteristics that exist, and that residential costs are generally far less than commercial costs: the scale of the residential curve is $1/1000^{\text{th}}$ of those for the other classes shown. The text gives details on the meaning of solid versus dotted lines and the reasons behind the curve shape for each customer.

3. An acetate manufacturing and processing plant (lower left) has a very non-linear cost curve. Any interruption of service causes \$38,000 in lost productivity and after-restoration set-up time. Cost rises slowly for about half an hour. At that point, molten feedstock and interim ingredients inside pipes and pumps begins to cool, requiring a day-long process of cleaning sludge and hardened stock out of the system. The dotted line shows the plant's interruption cost function after installation of a diesel generator, started whenever interruption time exceeds five minutes.
4. Residential interruption cost function (lower right), estimated by the authors from a number of sources including a survey of customers made for a utility in the northeastern United States in 1992, shows roughly linear cost as a function of interruption duration, except for two interesting features. The first is the fixed cost equal to about eight minutes of interruption at the initial variable cost slope which reflects "the cost to go around and re-set our digital clocks," along with similar inconvenience costs. The second is a jump in cost between 45 and 60 minutes, which reflect inconsistencies in human reaction to outage, time on questionnaires. The dotted line shows the relation the authors' uses in their analysis, which makes adjustments thought reasonable to account for these inconsistencies.

This recommended analytical approach, in which cost is represented as a function of duration on a per event basis, requires more information and more analytical effort than simpler "one-dimensional" methods of reliability evaluation. However, advantages are that the method is more sensitive to actual customer needs, and that the results are more credible when used objectively, in justifying utility expenditures and efforts to improve reliability.

Interruption Cost is Lower if Prior Notification is Given

Given sufficient time to prepare for an interruption of service, most of the momentary interruption cost (fixed) and a great deal of the variable cost can be eliminated by many customers. Figure 3.14 shows the interruption cost figures from Figure 3.13 adjusted for "24 hour notification given." These are much lower than for situations where no notice is given.

Cost of Interruption Varies by Customer Class

Cost of power interruption varies among all customers, but there are marked distinctions among classes, even when cost is adjusted for "size" of load by computing all cost functions on a per kW basis. Generally, the residential class has the lowest interruption cost per kW and commercial has the highest. Table 3.1

Table 3.1 Typical Interruption Costs by Class for Three Utilities – Daytime, Weekday (dollars per kilowatt hour)

Class	1	2	3
Agricultural	3.80	4.30	7.50
Residential	4.50	5.30	9.10
Retail Commercial	27.20	32.90	44.80
Other Commercial	34.00	27.40	52.10
Industrial	7.50	11.20	13.90
Municipal	16.60	22.00	44.00




Figure 3.14 If an interruption of service is expected customers can take measures to reduce its impact and cost. Solid lines are the interruption costs (the solid lines from Figure 3.13). Dotted lines show how 24-hour notice reduces the cost impact in each case.

gives the cost/kW of a one-hour interruption of service by customer class, obtained using similar survey techniques for three utilities in the United States:

1. A small municipal system in the central plains,
2. an urban/suburban/rural system on the Pacific Coast and,
3. an urban system on the Atlantic coast.

Cost Varies from One Region to Another

Interruption costs for apparently similar customer classes can vary greatly depending on the particular region of the country or state in which they are located. There are many reasons for such differences. The substantial difference (47%) between industrial costs in utilities 1 and 3 shown in Table 3.14 is due to differences in the type of industries that predominate in each region. The differences between residential costs of the regions shown reflect different demographics and varying degrees of economic health in their respective regions.

Cost Varies among Customers within a Class

The figures given for each customer class in Table 3.1 represent an average of values within those classes as surveyed and studied in each utility service territory. Value of availability can vary a great deal among customers within any class, both within a utility service territory and even among neighboring sites. Large variations are most common in the industrial class, where different needs can lead to wide variations in the cost of interruption, as shown in Table 3.2. Although documentation is sketchy, indications are *the major differing factor is the cost of a momentary interruption* - some customers are very sensitive to any cessation of power flow, while others are impacted mainly by something longer than a few cycles or seconds.

Cost of Interruption Varies as a Function of Time of Use

Cost of interruption will have a different impact depending on the time of use, usually being much higher during times of peak usage, as shown in Figure 3.15. However, when adjusted to a per-kilowatt basis, the cost of interruption can sometimes be higher during off-peak than during peak demand periods, as shown. There are two reasons. First, the data may not reflect actual value. A survey of 300 residential customers for a utility in New England revealed that customers put the highest value on an interruption during early evening (Figure 3.15). There could be inconsistencies in the values people put on interruptions (data plotted were obtained by survey).

Table 3.2 Interruption Costs by Industrial Sub-Class for One hour, Daytime, Weekday (dollars per kilowatt)

Class	\$/kW
Bulk plastics refining	38
Cyanide plant	87
Weaving (robotic loom)	72
Weaving (mechanical loom)	17
Automobile recycling	3
Packaging	44
Catalog distribution center	12
Cement factory	8




Figure 3.15 Cost of a one-hour interruption as a function of when it occurs, as determined by surveys and interviews with 300 customers of a utility in the northeastern United States, was determined on a three-hour period basis. A high proportion of households in this survey have school age children at home and thus perhaps weighed interruption costs outside of school hours more than during school hours. However, in general, most households rate interruption cost as higher in early morning and early evening.

Recommended Method of Application of Customer Interruption Cost Data

As mentioned earlier, the recommended analytical approach to value-based planning of power delivery includes assessment of customer costs using functions that acknowledge both a fixed cost for any interruption, no matter how brief, and a variable cost as a function of duration. It is also important to acknowledge the differences in value of service among the different customer classes, and the time differences within those customer classes.

Often, the cost *per kilowatt* is highest off-peak. Then, only essential equipment such as burglar alarms, security lighting and refrigeration is operating. These end-uses typically have a high cost of interruption. While it is generally safe to assume that *total* cost of interruption is highest during peak, the same cannot be assumed about the cost per kilowatt of interruption.

Ideally, reliability and service quality issues should be dealt with using a value-based reliability or power-quality planning method with customer interruption cost data obtained through statistically valid and unbiased sampling of a utility's own customers (see Sullivan et al.), not with data taken from a reference book, report, or technical paper describing data on another utility system. However, in many cases for initial planning purposes the cost and time of data collection are not affordable.

Figure 3.16 provides a set of costs of typical interruption curves that the author has found often match overall customer values in a system. It is worth stressing that major differences can exist in seemingly similar utility systems, due to cultural and economic differences in the local customer base. These are not represented as average, or "best" for use in value-based planning studies, but they are illustrative of the type of cost functions usually found. They provide a guide for the preparation of approximate data from screening studies and initial survey data.

Cost in Figure 3.16 is given in terms of "one hundred times nominal price." *It is worth noting that in many surveys and studies of interruption cost, the cost per kW of interruption is on the order of magnitude one hundred times the normal price (rate) for a kWh.* Generally, if a utility has low rates its customers report a lower cost of interruption than if it has relatively higher rates. No reliable data about why this correlation exists has been forthcoming.³

³ It could be that value of continuity is worth more in those areas where rates are high (generally, more crowded urban areas). However, the authors believe that a good part of this correlation is simply because in putting a value on interruptions, respondents to surveys and in focus groups base their thinking on the price they pay for electricity. Given that a typical residential customer uses roughly 1,000 kWh/month, they may simply be valuing an interruption as about "one tenth of my monthly bill."




Figure 3.16 Typical interruption cost characteristics for customer classes.

Cost of Surges and Harmonics

Far less information is available on customer costs of harmonics and voltage surges as compared to that available on the cost of interruptions. What data is made available in publications, and most of the reported results and interpretations in the technical literature, were obtained in very customer-specific case studies, most done on a single customer basis. As a result there is very little information available on average of typical costs of voltage surge sensitivity and harmonics sensitivity.

Few customers suffer from voltage dip (voltage problems lasting minutes or hours), surge, harmonics, electrical noise, and similar problems. For this reason most studies of cause, effect, and curve are done on a single-customer or specific area basis. Here, results reported in technical literature are the best guide.

End-Use Modeling of Consumer Availability Needs

The customer-class, end-use basis for analysis of electric usage, discussed in Chapters 2 and 15, provides a reasonably good mechanism for study of the service reliability and power quality requirements of customers, just as it provides a firm foundation for analysis of requirements for the amount of power. Reliability and power quality requirements vary among customers for a number of reasons, but two predominate:

1. *End-usage patterns differ.* The timing and dependence of customers' need for lighting, cooling, compressor usage, hot water usage, machinery operation, etc., varies from one to another.
2. *Appliance usage differs.* The appliances used to provide end-uses will vary in their sensitivity to power quality. For example, many fabric and hosiery manufacturing plants have very high interruption costs purely because the machinery used (robotic looms) is quite sensitive to interruption of power. Others (with older mechanical looms) put a much lower cost on interruptions.

End-use analysis can provide a very good basis for a detailed study of power quality needs. For example, consider two of the more ubiquitous appliances in use in most customer classes: the electric water heater and the personal computer. They represent opposite ends of the spectrum from the standpoint of both amount of power required and cost of interruption. A typical 50-gallon storage electric water heater has a connected load of between 3,000 and 6,000 watts, a standard PC has a demand of between 50 and 150 watts. Although it is among the largest loads in most households, an electric water heater's ability to provide hot water is not impacted in the least by a one-minute interruption of power. In most cases a one-hour interruption does not reduce its ability to satisfy the end-use demands put on it.⁴ On the other hand, interruption of power to a computer, for even half a second, results in serious damage to the "product." Often there is little difference between the cost of a one-minute outage and a one-hour outage.

It is possible to characterize the sensitivity of most end-uses in most customer classes by using an end-use basis. This is in fact how detailed studies of industrial plants are done in order to establish the cost-of-interruption statistics, which they use in VBP of plant facilities and in negotiations with the utility to provide

⁴ Utility load control programs offer customers a rebate in order to allow the utility to interrupt power flow to water heaters at its discretion. This rebate is clearly an acceptable value for the interruption, as the customers voluntarily take it in exchange for the interruptions. In this and many other cases, economic data obtained from market research for DSM programs can be used as a starting point for value analysis of customer reliability needs on a value-based planning basis.

upgrades in reliability to the plant. Following the recommended approach, this requires distinguishing between the fixed cost (cost of momentary interruption) and variable cost (usually linearized as discussed above) on an end-use basis.

A standard end-use model used to study and forecast electric demand can be modified to provide interruption cost sensitivity analysis, which can result in “two-dimensional” appliance end-use models as illustrated in Figure 3.17. Generally, this approach works best if interruption costs are assigned to appliances rather than end-use categories. In commercial and industrial classes different types of appliances within one end-use can have wildly varying power reliability and service needs. This requires an “appliance sub-category” type of an end-use model. Modifications to an end-use simulation program to accommodate this approach are straightforward (see Willis, 1996, Chapter 11), and not only provide accurate representation of interruption cost sensitivity, but produce analysis of costs by time and location, as shown in Figures 3.18 and 3.19.




Figure 3.17 The simulation’s end-use model is modified to handle “two-dimensional” appliance curves, as shown here for a residential electric water heater. The electric demand curve is the same data used in a standard end-use model of electric demand. Interruption cost varies during the day, generally low prior to and during periods of low usage and highest prior to high periods of use (a sustained outage prior to the evening peak usage period would result in an inability to satisfy end-use demand).




Figure 3.18 Result of a power quality evaluation, using an end-use model. Top: the daily load curve for single family homes segmented into four interruption-cost categories. High-cost end-uses in the home are predominantly digital appliances (alarm clocks, computers) and home entertainment and cooking. The bottom figure shows total interruption cost by hour of the day for a one-hour outage - compare to Figure 3.19.




Figure 3.19 Map of average reliability needs computed on a 10-acre small area grid basis for a port city of population 130,000, using a combination of an end-use model and a spatial customer simulation forecast method, of the type discussed in Chapter 15. Shading indicates general level of reliability need (based on a willingness-to-pay model of customer value).

3.4 Two Q Analysis: Quantity and Quality of Power Are Both Important Aspects of Consumer Value

The amount of power used, and the dependability of its availability for use are both key attributes in determining the *value* of the electric power to each consumer. The values attached to both quantity and quality by most consumers are linked, but somewhat independent of the value attached to the other. In order to demonstrate this, it is worth considering two nearly ubiquitous appliances, which happen to represent opposite extremes of valuation in each of these dimensions: the electric water heater and the personal computer.

A typical residential storage water heater stores about 50 gallons of hot water and has a relatively large connected load compared to most household appliances: about 5,000 watts of heating element controlled by a thermostat. Its contribution to coincident peak hourly demand in most electric systems is about 1,500 watts. The average heater's thermostat has it operating about 30% of the time (1,500/5,000) during the peak hour.

Despite its relatively high demand for quantity, this appliance has a very low requirement for supply quality. A momentary interruption of electric flow to this device – on the order of a minute or less – is literally undetectable to anyone using it at that moment. Interruptions of up to an hour usually create little if any disruption in the end-use the device provides, which is why direct load control of water heaters is an “economy discount” option offered by many utilities. Similarly, the water heater is not critically dependent on other aspects of power quality. It will tolerate, and in fact turn into productive heat, harmonics, spikes, and other “unwanted” contamination of its supply.

By contrast, a typical home computer requires much less quantity, about 180 watts. But unlike the water heater, it is rendered useless by even a momentary interruption of its power supply. Just one second without power will make it cease operation. Battery backup power supplies, called *uninterruptible power supplies*, can be used to avoid this problem, but at an increase in the computer’s initial and




Figure 3.20 Electrical appliances vary in the amount of electricity they demand, and the level of continuity of service they require to perform their function adequately. Scales shown here are qualitative. Typically quantitative scales based on the most applicable customer values are used in both dimensions.




Figure 3.21 A Two-Q diagram of a consumer’s needs can be interpreted in either of two dimensions to identify the total (area under the curve) and range of needs in either Q direction, producing Q profiles of the importance of demand, or reliability to his particular needs.

operating cost. In addition to its sensitivity to supply availability, a computer is also more sensitive to harmonics, voltage dips, and other power quality issues than a water heater. Figure 3.20 shows these two appliances, and several other household energy uses plotted on a two-dimensional basis according to the relative importance of both quantity and quality. This Two-Q plot is a type of two-dimensional demand and capability analysis that treats both Qs as equally important. Traditionally, “demand” has been viewed only as a requirement for quantity of power, but as can be seen the true value of electric service depends on both dimensions of delivery. Recognition of this fact has been the great driver in moving the electric industry toward “tiered rates,” “premium rates,” and other methods of pricing power not only by amount and time of use, but by reliability of delivery. Either or both dimensions in the plot can be interpreted as single-dimensional profiles for quantity or quality, as the case may be, providing useful information in the analysis of consumer needs and how to best meet them.

(Figure 3.21). Two-Q analysis using these same two dimensions also applies to the design of utility power systems, as will be discussed in Chapter 4 (Section 4.4).

3.5 CONCLUSION AND SUMMARY

The whole purpose of electric utility systems is to serve the electric demands of individual households, businesses, and industrial facilities. Two useful mechanisms for evaluation of consumer requirements and planning for their provision are end-use analysis and Two-Q planning, both of which provide a useful basis for the study and analysis of both quantity and quality.

In addition to having a demand for electric energy, electric consumers value availability and quality of power too, seeking satisfactory levels of both the amount of power they can obtain, and the availability of its supply. Each customer determines what is satisfactory based on their specific needs. An interruption that occurs when a customer does not need power does not concern them. Harmonics or surges that make no impact on appliances are not considered of consequence. Among the most important points in customer value analyses are:

1. *Interruption of service* is the most widely affecting power quality problem. A majority of electric customers in most utility systems experience unexpected interruptions at least once a year.
2. *Voltage dips and sags create momentary interruptions* for many customers that are indistinguishable in impact on their needs with short-term interruptions of power flow.
3. *Analysis of customer needs for power quality is generally done by looking at the cost which less than perfect quality creates, not the value that near-perfect quality provides.* Thus, availability of power is usually studied by analyzing the cost incurred due to interruptions, and determining the costs created by those interruptions assesses harmonics.
4. *Power quality costs and values vary greatly* depending on customer class, time of day and week, region of the country, and individual characteristics.
5. *Price concerns often outweigh quality concerns.* Throughout the late 1980s and early 1990s, several independent studies indicated that roughly “30%-40% of commercial and industrial customers are willing to pay more for higher power quality,” or something similar, a statistic that seems quite believable and realistic. This led to considerable focus on improving power system reliability and quality. However, it is worth bearing in mind that these same

studies indicate that 20-30% is quite content with the status quo, while the remaining 30-40% would be interested in trading some existing quality for lower cost.

REFERENCES AND BIBLIOGRAPHY

- P. F. Albrecht and H. E. Campbell, "Reliability Analysis of Distribution Equipment Failure Data," EEI T&D Committee Meeting, January 20, 1972.
- R. N. Allan, et al., "A Reliability Test System for Educational Purposes – Basic Distribution System Data and Results," *IEEE Transactions on Power Systems*, Vol. 6, No. 2, May 1991, pp. 813-821.
- "Guide for Reliability Measurement and Data Collection," EEI Transmission and Distribution Committee, October 1971, Edison Electric Institute, New York.
- W. F. Horton, et al., "A Cost-Benefit Analysis in Feeder Reliability Studies," *IEEE Transactions on Power Delivery*, Vol. 4, No. 1, January 1989, pp. 446 - 451.
- Institute of Electrical and Electronics Engineers, *Recommended Practice for Design of Reliable Industrial and Commercial Power Systems*, The Institute of Electrical and Electronics Engineers, Inc., New York, 1990.
- A. D. Patton, "Determination and Analysis of Data for Reliability Studies," *IEEE Transactions on Power Apparatus and Systems*, PAS-87, January 1968.
- N. S. Rau, "Probabilistic Methods Applied to Value-Based Planning," *IEEE Transactions on Power Systems*, November 1994, pp. 4082 - 4088.

4

Power System Reliability and Reliability of Service

4.1 INTRODUCTION

“Reliability” as normally applied by power distribution-utilities means continuity of service to the electric consumers. A 100% reliable power system provides power without interruption, power that is available all of the time regardless of how much and if the consumer is using it at any given moment. Almost anything less than perfection in this regard will garner complaints that the power supply is “unreliable.” An availability of 99.9% might sound impressive, but it means eight and three quarter hours without electric service each year, nearly all electric consumers in North America would consider that a level of service unsatisfactory. That is why availability averages approximately 99.98%, or roughly 1.75 hours of unavailability per year. In the early 21st century, even this level of performance is becoming unacceptable to many energy consumers.

This chapter reviews power delivery reliability concepts in general and introduces the key components used in reliability analysis. Sections 4.2 and 4.3 cover equipment outages and service interruptions, the reliability indices used to evaluate reliability performance, and their application to electric utility systems. Section 4.4 covers various ways of setting reliability standards and design criteria. Section 4.5 discusses the cost of reliability, and the various types of tradeoffs made in design. Two-Q analysis of power systems is covered in Section 4.6. Section 4.7 concludes the chapter with a summary of the most important concepts in the chapter.

Table 4.1 Reliability Terms and Definitions

Availability:	The fraction of the time that service is available. The (steady-state) probability that power will be available.
Duration:	The total elapsed time of an interruption or outage.
Failure rate:	The average number of failures of a component or unit of the system in a given time (usually a year).
Failure probability:	The probability that a unit or system will be in a “failure mode” at a time selected at random.
Frequency:	How often interruptions or outages, as the case may be, occur.
Interruption:	A cessation of service to one or more customers, whether power was being used at that moment or not.
Mean time between failures (MTBF):	The average or expected time a unit of equipment is in service between failures.
Mean time to repair (MTTR):	The average or expected time to repair an outaged component or unit of the system once failed.
Outage:	The unavailability of part of the power supply system, whether due to unexpected or planned circumstances.
<i>Forced outage:</i>	The outage due to unexpected and unscheduled event (usually a failure or impending failure).
<i>Scheduled outage:</i>	An outage that was planned and scheduled in advance (usually for maintenance).
Reliability:	The fraction of time that a unit of equipment or portion of the system is capable of performing the required function. The (steady state) probability that it will be in service in a condition where it can function as required.
Reporting period:	The period of time over which reliability, interruption, and outage statistics have been gathered or computed.
Restoration:	Return of electric service after an interruption, because of repair of the outage causing the interruption, or because of re-switching of the supply source, or the starting of an alternate DG source.
Restoration time:	The time required restoring service to an interrupted customer or portion of the system.
Service:	Supply of electric power in <i>sufficient amount and quality</i> to satisfy customer demand.

4.2 OUTAGES CAUSE INTERRUPTIONS

Table 4.1 provides a list of key reliability definitions. Among the most important are “outage” and “interruption.” Although often used interchangeably, they have far different meanings. “Outage” means a failure of part of the power supply system. For example, a line down, transformer out of service, or a breaker that opens when it shouldn’t. “Interruption” means a cessation of service to one or more customers. Interruptions are almost always caused by outages, and this leads many people to confuse the two words and their meanings. However, to understand power supply reliability properly, and more importantly, to set criteria and design a system to achieve high reliability, the distinction between outage (the cause of service problems) and interruption (the result) is critical. Electric service to a customer is interrupted whenever equipment in the line of supply to that customer fails or is otherwise out of service. Interruptions are caused by outages, but not all outages cause interruptions.

Frequency and Duration

Two different aspects of reliability receive attention in any type of power supply reliability analysis, whether of interruptions or outages. These are the *frequency* (how often something occurs) and *duration* (how long it lasts). With respect to interruption, frequency refers to the number of times service is interrupted during a period of analysis – once a decade, two times a year, five times a month, or every afternoon. Duration refers to the length of interruptions - some last only a few cycles, others for hours, even days.

Extent

Frequency and duration are important factors in reliability analysis and design, but there is a third factor equally important to the electric service planner because it is the one over which the planner has considerable control, particularly in the design phase. The *extent* of an outage’s impact on service – how many customers or what portion of the plant’s loads are interrupted by the outage of a particular unit of equipment – is the key to the outage-interruption relationship. The design of the electric power supply system greatly influences just how much interruption occurs when an equipment outage occurs.

The Distribution Planner has control over extent to which outages lead to customer interruptions. Among other things it is a function of how the distribution system is laid out. In general, a system laid out to minimize extent of outages from any failure will be more expensive than one that is not, but for the planner this means that consideration of configuration and its influence on reliability is just one more alternative to be considered in planning. In many cases a higher-reliability layout may be less expensive than any other way of achieving the particular reliability goal.

Most distribution systems use some form of radial feeder and secondary layout, which means that a single failure anywhere in the electrical path between substation and customer will cause service interruptions to all customers downstream. One way to improve reliability in such systems is to lay out the feeder trunk and branches, and arrange the protection (fusing) so that outages have a low extent (i.e., so that the average feeder segment failure does not interrupt service to very large numbers of customers).

Types of Interruptions

Interruptions are classified by their duration. Table 4.2 lists typical interruption definitions (see IEEE Gold Book for more information), with the times given being averages obtained by the authors in a survey of definitions and usage throughout the power industry. Electric service planners should keep in mind that practice and interpretation of reliability definitions vary greatly from industry to industry, from utility to utility, and from engineer to engineer.

In the authors’ opinion, the definitions shown are all somewhat unsatisfactory. The names imply that they relate to interruption duration, yet they are defined by equipment and operating factors, such as whether an interruption is restored automatically or manually. This, along with differences in exactly how the definitions are interpreted, means that the application of a particular definition, for example “momentary interruption,” varies greatly both in definition and therefore in *how it is used*.

Table 4.2 Traditional Classification for Power Supply Interruptions

Type	Definition
Instantaneous:	An interruption restored immediately by completely automatic equipment, or a transient fault that causes no reaction by protective equipment. Typically less than 15 seconds.
Momentary:	An interruption restored by automatic, supervisory, or manual switching at a site where an operator is immediately available. Often defined as less than three minutes, but a proposal by an IEEE Task Force suggests standardizing on five minutes.
Temporary:	An interruption restored by manual switching by an operator who is not immediately available. Typically, 30 minutes.
Sustained:	Any interruption that is not instantaneous, momentary, or temporary. Normally more than an hour.

Table 4.3 Duration Limits Used to Define Instantaneous, Momentary, and Temporary Outages at Eight Utilities

Utility	Instantaneous	Momentary	Temporary
1	not used	< 5 minute	5 min. to 45 min.
2	< 15 sec.	15 sec. to 1 min.	1 min. to 45 min.
3	not used	up to 30 sec.	30 sec. to 30 min.
4	< 15 sec.	15 sec. to 2 min.	2 min. to 1 hour
5	< 60 sec.	60 sec. to 5 min.	5 min. to 2 hours
6	not used	< 1 min.	1 min. to 52 min.
7	< 1 min.	not used	1 min. to 2 hours
8	“cleared by inst. relay”	up to 2 min.	1 min. to 1 hour

Perhaps most importantly, since these distinctions mean nothing to electric consumers, they do little to help an electric service planner improve quality from the consumer’s perspective. A four-minute interruption is a four-minute interruption, whether classified as momentary or temporary or permanent. Beyond the distinctions shown in Table 4.2, interruptions are classified into scheduled interruptions (the utility scheduled the interruption for maintenance purposes) or forced interruptions (the interruption was not expected and not scheduled). Implicit in these definitions, and the most important factor in the authors’ minds, is whether customers are informed reasonably far in advance of the details of any “scheduled interruption.”

Every planner should pay close attention to the definitions used by his company and by utilities and other players in any planning situation of collaboration. What appear to be identical indices can vary because of differences in how interruptions are defined. Many utilities do not include scheduled outages, storms, and other “non-failure-related” interruptions in their reporting procedures.¹ Others leave momentary interruptions out of reported averages, or report them as a special category.

Every power system and DG Planner should bear in mind when working with reliability statistics that utilities that report a certain type of interruption statistic, such as number of momentary operations, vary widely in how they interpret that and other definitions. This is shown by Table 4.3, which lists interruption

¹ Most utilities remove “storms” from their statistics, arguing that natural disasters should be removed from statistics that are meant to measure their performance. However this practice has been carried to extremes and abused by some utilities. As Samuel Clemens once said, “There are three kinds of lies: Lies, damned lies and statistics.”

definitions from eight utilities in North America. They should also remember that utility customers and electric consumers do not really care about definitions and nomenclature. For them, an “outage” or “interruption” is anything that causes their equipment and facilities to stop working.

Such differences in definitions, interpretations, and reporting practices can frustrate planners who try to find some real meaning in comparisons between the reliability figures reported by different electric utility companies or by different industrial facilities. Worse, these differences often cause confusion and frustration among customers, regulators, and standards organizations.

Voltage Sags Often Look Like Interruptions

Many electric consumers report a momentary or “very short” interruption when all that really occurred was a brief “voltage sag.” For example, during a fault (short circuit) on a distribution line, voltage on nearby circuits might drop to less than 70% of normal. Typically the fault would be cleared within less than a second, but this “blink” may be sufficient to halt the operation of computers, robotic equipment, and other digital systems. In that case, this short “voltage sag” is indistinguishable from a short, total cessation of voltage (an interruption). Either way, some critical equipment stops functioning.

The Computer and Business Equipment Manufacturers Association (CBEMA) has a recommended guideline on what voltage sags and surges should be tolerable to office and business equipment, as shown in Figure 4.1. The “CBEMA envelope” establishes a relationship between degree of voltage deviation and time.

The curve’s application is simple: if an event’s voltage deviation and duration characteristics are within the CBEMA envelope, then normal appliances should operate normally and satisfactorily despite it. But many appliances and devices in use are more sensitive than specified by the CBEMA curves and will not meet this criterion at all. Others may fail to meet it under the prevailing electrical conditions (i.e., line voltage, phase imbalance, power factor, and harmonics may be less than perfect).

The manner of usage of an appliance also affects its voltage sag sensitivity. The small circle in Figure 4.1 indicates an event that “crashed” a local net server (computer) at a hosiery plant. The event falls just within the CBEMA curve, and the computer’s manufacturer probably intended for it to withstand that particular deviation-time combination. However, the computer in question had been upgraded with three times the standard memory, a larger hard drive, and optional network cards, doubling its power usage and the load on its power supply. Such situations are common and mean that power systems that deliver voltage control within recommended CBEMA standards may still create problems for computers, robotic assembly tools, and other sensitive equipment.




Figure 4.1 Data from 67 disturbances which occurred over a two-year period at a clothing factory, plotted against the CBEMA curve (solid line) and the actual requirement envelope of a digitally-controlled hosiery loom (dotted line). In many systems, about 40% of all voltage sags and 10% of all voltage surges lie outside of the CBEMA envelope. Small circle refers to event mentioned in text.

4.3 RELIABILITY INDICES

In order to deal meaningfully with reliability as a design criterion for distribution, it is necessary to be able to measure it and set goals. A bewildering range of reliability indices is in use within the power industry. Some measure only frequency of interruption, others only duration. A few try combining both frequency and duration into a single value, which proves to be nearly an impossible task. Some measures are system-oriented, looking at reliability averaged over the entire customer base. Others are customer- or equipment-oriented, meaning that they measure reliability only with regard to specific sets of customers or equipment. To complicate matters further, most utilities use more than one reliability index to evaluate their reliability. The large number of indices in use is compelling evidence that no one index is really superior to any other. In fact, the authors are convinced that no one index alone is particularly useful.

The basic problem in trying to measure reliability is how to relate the two quantities, frequency and duration. Are two one-hour interruptions equivalent to one two-hour interruption, or are two one-hour interruptions twice as bad as the one two-hour interruption? Most people conclude that the correct answer lies somewhere in between, but no two electric service planners, and no two electric consumers for that matter will agree on exactly how frequency and duration “add up” in importance. For one thing, the importance of frequency and duration vary

tremendously from one electric consumer to another. There are some industries where a one-minute interruption of power will cause over ninety minutes of lost productivity. Computer and robotic systems must be reset and restarted, or a long production process must be cleared and bulk material reloaded, before the plant can restart at the beginning of a production cycle. For these types of consumers, five one-minute outages are much more serious than a single outage, even if it is five hours' duration.

But there are other customers for whom short outages cause no significant problems, but who experience inconvenience during a sustained outage. This category includes factories or processes such as pulp paper mills, where production is only damaged if equipment has time to "cool down." Interruption of power to a heater or tank pressurized for a minute or two is not serious (although it might be inconvenient), but sustained interruption allows a cool-down or pressure drop that results in lost production.

Over the last two decades of the 20th century, frequency of interruption became increasingly important to a larger portion of the electric power market. Before the widespread use of digital clocks and computerized equipment, few residential or commercial customers cared if power was interrupted briefly (for a few minutes at most) in the early morning (e.g., 3:00 AM) while the electric utility performed switching operations. It was common practice for utilities to perform minor maintenance and switching operations during that period to minimize any inconvenience to their customers. Today, this practice leads to homeowners who wake up to a house full of blinking digital displays and wake up later than they expected, because their alarm clocks did not operate.

The amount of power interrupted is also an important factor. Some reliability calculations and indices are weighted proportionally to customer load at the time of interruption, or to the estimated energy (kWh) *not* delivered during the interruption. However, the four most popular indices used in reliability analysis make no distinction of demand size. They treat all customers in a utility system alike regardless of peak demand, energy sales, or class. These indices are SAIFI and CAIFI, which measure only frequency, and SAIDI and CTAIDI, which measure only duration. Usually these four are used in conjunction – four numbers that give a rough idea of what is really happening to reliability system-wide.

These four indices, and most others, are based on analysis of customer interruptions during some *reporting period*, usually the previous month or year. All count a customer interruption as the interruption of service to a single customer. If the same customer is interrupted three times in a year, that constitutes three customer interruptions. If one equipment outage caused simultaneous interruption of service to three customers, that too is three customer interruptions.

System Average Interruption Frequency Index (SAIFI) is the average number of interruptions per utility customer during the period of analysis.

$$\text{SAIFI} = \frac{\text{number of customer interruptions}}{\text{total customers in system}} \quad (4.1)$$

Customer Average Interruption Frequency Index (CAIFI) is the average number of interruptions experienced by customers who had at least one interruption during the period.

$$\text{CAIFI} = \frac{\text{number of customer interruptions}}{\text{number of customers who had at least one interruption}} \quad (4.2)$$

The “S” in SAIFI means it averages the interruption statistic over the entire customer base (the system), the “C” in CAIFI means it refers to only customers who experienced interruptions. Customers who had uninterrupted service during the period are precluded from CAIFI.

System Average Interruption Duration Index (SAIDI) is the average duration of all interruptions, obtained by averaging over all of the utility customers – those who had outages and those who did not.

$$\text{SAIDI} = \frac{\text{sum of the durations of all customer interruptions}}{\text{total customers in system}} \quad (4.3)$$

Customer Total Average Interruption Duration Index (CTAIDI) is the average total duration of all interruptions, but averaged only over the number of utility customers who had at least one outage.

$$\text{CTAIDI} = \frac{\text{sum of the durations of all customer interruptions}}{\text{number of customers who had at least one interruption}} \quad (4.4)$$

Momentary Interruption Index

Momentary Average Interruption Frequency Index (MAIFI) is the average number of momentary (and sometimes instantaneous) interruptions per utility customer during the period of analysis.

$$\text{MAIFI} = \frac{\text{number of customer momentary interruptions}}{\text{total customers in system}} \quad (4.5)$$

Often, momentary interruptions are not included in the SAIFI value reported by an electric utility. In such cases, total interruptions (from the consumer standpoint) are best estimated by the sum: SAIFI + MAIFI.

Load Curtailment Indices

In addition, a number of indices try to relate reliability to the size of customer loads, in order to weight the interruption of a large load more than a small load,

$$\text{Customer load curtailment} = \text{duration of outage} \times \text{kVA unserved.} \quad (4.6)$$

Customer Average Load Curtailment Index (CALCI) is the average kVA x duration interrupted per affected customer, per year.

$$\text{CALCI} = \frac{\text{sum of all customer load curtailments}}{\text{number of customers who had at least one interruption}} \quad (4.7)$$

Normal usage is often sloppy and informal

Even though the “S” in the terms SAIDI and SAIFI means, “averaged over the entire system,” some utilities internally report “SAIDI” and “SAIFI” values by feeder. While strictly speaking, SAIDI and SAIFI are the same for all feeders (they are computed by being averaged over the entire system), it is understood that what they mean is the average annual duration and frequency for customers on that feeder. Similarly, engineers at these and other utilities often use the two terms when they mean “duration” and “frequency” of interruptions for an area of the system or a small group of customers.

Analysis Using Reliability Indices

Reliability indices are used to evaluate historical and recent data to reveal trends and patterns, expose problems, and reveal how and where reliability can be improved. Any of the reliability indices discussed above can be tracked over time to identify trends that indicate developing problems. Figure 4.2 shows frequency of customer interruption (SAIFI) for one suburban operating district of a metropolitan utility on an annual basis over a 35-year period. The operating district was open farmland (no tree-caused trouble) up to 1962. At that time suburban growth of a nearby metropolitan region first spilled into the area, with much construction and with it the typically high interruption rates characteristic of construction areas (trenchers digging into underground lines, early equipment failures of new distribution facilities, etc.).

Over the following ten years, as the area plotted in Figure 4.2 gradually filled in with housing and shopping centers and construction and system additions stabilized, frequency of interruption dropped. Twenty years after the area first began developing, the interruption rate again began to rise because of two factors:

1. Gradually increasing cable failure rates due to age.
2. The effects of trees, planted when the area first developed, reaching a height where they brushed against bare conductors on overhead lines.




Figure 4.2 Three-year moving average of the annual number of interruptions for a region of a utility in the central United States. This area grew from open farmland into a developed metropolitan suburb in the period from 1962 to 1970, during which time interruption rate rose due to the problems common to areas where new construction is present (digging into lines, initial above-average failure rates for newly installed equipment). Twenty years later, interruption rate began to rise again, a combination of the effects of aging cable (direct buried laterals in some residential areas) and trees, planted when the area was new, finally reaching a height where they brushed against overhead conductors. The anomaly centered on 1982 was due to a single, bad storm.




Figure 4.3 Average hours out of service per interruption on a monthly basis for a rural utility in the northern United States varies due to random fluctuations and changes in weather. It appears that duration of interruptions (dotted line) is slowly increasing over time, but further study is necessary before this conclusion can be confirmed.

Another obvious comparison to look at is the average duration of an outage, which is simply the total duration of all interruptions divided by the total number of interruptions that occurred during a period. Average duration of interruption provides an indication of how the utility is managing the entire process from trouble call to service restoration, as shown in Figure 4.3. The increase in outage duration during winter for this utility in the northern United States is clearly visible (winter storms stress repair resources and make travel to repair sites lengthier). A trend upward is discernible. Without additional study, it is not possible to determine whether this is due to an increase in the number of outages, a change in type (perhaps “serious” outages are occurring more often as equipment ages), or a less effective (slower) repair process.

4.4 RELIABILITY AND CONTINGENCY CRITERIA FOR PLANNING

Indirect Reliability Engineering: Contingency and Margin Criteria

The vast majority of electric utilities engineer reliability of service into their system indirectly, through the use of *contingency design criteria*. The utility maintains reliability goals and still measures and reports SAIDI, SAIFI and similar indices, but these values are not used directly as goals in the design process. Rather, adequate performance is achieved by applying contingency margin criteria which have been interpreted from those SAIDI and SAIFI goals and which are hoped to be appropriate. For example, distribution system engineers will apply equipment loading standards, design and layout guidelines, and switching and fusing rules, which limit the extent of outages and afford extra capacity to pick up outaged load by alternative means.

Contingency margin standards are designed to allow sufficient leeway for any unit of equipment to pick up load interrupted due to the outage of a nearby unit, and to permit reasonably quick (e.g., 15 or 30 minutes to an hour) switching. For example, a utility may establish a policy saying that its substations will always have two transformers of equal size each loaded to no more than 80% of its rating. During the outage of either transformer, the other can pick up both loads with an overload of 60%, an amount it should be able to sustain for two to four hours, time enough, hopefully, for the utility to take steps to handle the problem over a longer period if necessary. Similarly, a utility may ensure that every feeder will have three switchable portions, or zones, each of which can be divided and transferred to other feeders if its primary feed is out of service.

Such a policy establishes a certain type of design (two transformers per substation), loading standard (a transformer cannot exceed 80% of its rating), and feeder layout (three switchable zones) in order to leave sufficient margin for equipment failure and switching. Similar and compatible criteria are established for conductor loading, switch placement, and other design practices, all rules and standards that *imply* certain reliability results.

Contingency-based Planning Methods Have Weaknesses for Aging Infrastructure Planning

This type of engineering, called *contingency-based engineering*, is heavily involved with the entire aging infrastructure issue. Contingency-based engineering has sensitivities to high utilization rates for system equipment, and other factors that make it not as dependable as it was traditionally with respect to use as a necessary and sufficient design criteria. Chapter 8 will discuss this in considerably more detail, with examples.

Direct Reliability Engineering

By contrast, reliability engineering achieves the targeted level for availability in the system by actually computing the expected reliability of the power system based on detailed analysis of its equipment and configuration. Areas of equipment that do not meet the criteria are re-engineered until they do. In this approach, frequency and duration of service interruption are computed for each potential alternative design and compared to planning reliability targets to assure compliance with all new designs, additions, and operating changes. Use of reliability values directly as design criteria is not a widely used practice at the time of this writing. Instead, most utilities still use indirect criteria based on contingency margin, switching zone counts, etc., as described above. However, it can be expected that in the future many will choose reliability directly as a criterion for three reasons:

- 1) It addresses reliability needs directly from the consumer standpoint, rather than indirectly from the utility (equipment criteria standpoint), leading to improved quality assurance.
- 2) It imposes no restrictions on design other than that the final result meet criteria. This can lead to lower costs.
- 3) Dependable, proven methods to accommodate reliability-based design are becoming widely available.

Duration Is the Most Popular Criterion

Regardless of whether an electric utility tries to ensure adequate customer service availability by using indirect or direct reliability engineering methods, almost all maintain specific reliability targets as their performance goal. In order of importance or popularity as a design rule, the most widely used are:

SAIDI = the expected total time in a year that the average customer in the area being studied will be without available power

MaxD = the maximum expected total time in a year that any customer in

the area being studied will be without available power

SAIFI = the expected number of times in a year an average customer in the system or area being studied will have power availability interrupted

MaxF = the expected maximum number of times in a year that any customer in the area being studied will have power availability interrupted

These reliability targets are applied within any study or design area as general targets that must be satisfied, much like standards of voltage drop, power factor, etc. In the design of any feeder, substation layout, or other addition to the distribution system, the planners use these criteria as design constraints. Any acceptable plan must have an expected average frequency and duration of interruption equal to or better than the corporate goals. Otherwise, the plan would dilute the system's overall reliability. In addition, the expected maximum frequency and duration for any customer covered in the planning case must be less than MaxF and MaxD.

In general, Distribution Planning studies focus on small portions of the system, at most a few substations at a time. If every such area of the system is planned so that its frequency and duration of interruption fall close to the corporate SAIDI and SAIFI targets, reliability will be as evenly distributed as is practical throughout the system. The utility might still compute CTAIDI and CAIFI in its assessment of actual operating experience, to make certain results match its goals, but planners do not need to apply more than these four criteria.

Duration Is Equivalent to “Probability of Unavailable Power”

Some utilities use a simple “probability of being without power” factor as a design target. This is equivalent, in practice, to planning with duration of interruption as the target. For example, a criterion that calls for the outage probability to be no more than .025%, means that the customer will be without power no more than

$$8,760 \text{ hours} \times .025\% = 2.16 \text{ hours/year} \quad (4.8)$$

Reliability Targets for DG Planning

Generally, distributed generation is not “distributed” over a large system like a distribution system, but is focused on one or a small number of customers. Thus, SAIDI and SAIFI are not really meaningful statistics to use in DG planning. Their equivalents in DG studies are the expected duration and frequency of unavailability. Again, often planners simply compute the expectation of total hours of outage, using this as the figure of merit in planning.

Tiered Reliability Areas

In a few cases, a utility will designate different frequency and duration targets for urban and rural areas. For example, its reliability criteria may specify part of the service territory as “urban” with SAIFI and SAIDI targets of 1.0 interruptions and 50 minutes/year for the system within that region. The rest of the system would be characterized as “rural” and designated with SAIFI and SAIDI rates of 2.2 and 150 minutes respectively. Such a tiered reliability design structure merely recognizes a common and long-standing characteristic of most utility systems – service reliability is lower in sparsely populated areas than in densely populated ones.

This disparity between urban and rural standards does not imply that rural customers are less important, nor is it due to inattention on the part of the utility or to poorer design practices applied to rural areas. The cost of reliability enhancement in rural areas is higher, because of the great distances involved, and the fact that the higher cost must be spread over fewer customers per mile of line. Most rate-design and regulatory processes do not consider it equitable to spread the cost of upgrading service in these rural areas to the same standards achievable in urban areas over the entire rate base (i.e., to those who live in the cities, too). Thus, a higher frequency of interruption and longer interruptions in electric service, are among the expected pleasures of rural life.

Predictive Reliability Analysis Methods

For planning purposes, reliability analysis normally focuses on comparing the expected reliability performance of different candidate designs. Plans are evaluated and compared to determine which is less likely to fall short of required capacity. Reliability is only one aspect of planning evaluation: cost and other factors enter into the final decision too, often with reliability vs. cost being the major design compromise the planner must consider.

Table 4.4 Reliability Evaluation Methods

Method	Reference
Simple State Enumeration	Used here
Explicit State Enumeration	Brown
Monte Carlo Simulation	Patton
Cut-Set Analysis	Billinton
Adequacy Equivalent	Zhang and Billinton
Markov Chain Analysis	Patton
Multi-Criteria Analysis	Atanackovic et al

Failure: when capacity is less than load

Generally failure states for a power system include all those conditions in which the demand for power (load) at a site exceeds the available supply of power to that site. At the distribution level, failure conditions are generally easy to identify. Since distribution circuits are usually radial, a simple “loss of connectivity” analysis is sufficient to identify outages that lead to failure. For distributed generation, the situation is far more difficult to analyze. Given several generators at a site, service might be maintained even if one or two are out of service. Various combinations of units in service, and their probability of occurrence, have to be evaluated against load curve shape. Methods to do so will be demonstrated in Chapters 13 and 16. Reliability evaluation of an entire power system is more challenging still, due to the number of combinations of sources and paths. Even a small power system with less than 70 generators and 100 lines, can become a considerable challenge (Zhang and Billinton, 1997).

Reliability Evaluation Methods

Literally hundreds of technical papers, and quite a few books, have been devoted to power system reliability analysis. Table 4.4 provides a summary of the methods along with references on each. But despite this diversity in the technical, numerical means of approach, reliability analysis for planning falls into two functional categories:

1. *Historical Assessment Methods* analyze historical operating results to determine the cause of past reliability performance problems and to infer the base equipment failure rates and reliability factors. For example, “Based on historical analysis overhead lines are likely to fail once every five years and require an average of five hours to repair.”
2. *Predictive Reliability Methods* estimate the reliability expected from a candidate design or system configuration. These are used by Planners to determine if a change in design, for example a “fix” to a reliability problem, will give sufficient availability of power.

Application of Probabilities

Almost all reliability prediction methods apply some form of analysis based on probabilities. A few rare methods based on adaptive pattern recognition do not. But the basis for what most methods apply as reliability, the way they apply it, and the detail and representations used in that application, vary greatly among methods. This makes it crucial to pick a method tailored to a situation’s particular needs. For instance, some methods analyze the probability or the expected amount of time that power is unavailable at a particular point in the power network. They estimate the total amount of time that power will be unavailable

(e.g., .03%, equal to 2.6 hours per year). Others predict the mean time between failures and expected time to repair or restore power in each case.

The best of these methods can estimate both the expected average frequency and duration of interruption for any particular candidate system /DG design (e.g., power is expected to fail on average twice a year and take an average of 1.3 hours to restore each time). Still more complicated methods predict the expected probability for both frequency and duration of interruptions. For example, power is expected to fail twice a year, but there is a 5% likelihood it could fail four times in a year - the average duration is expected to be 1.3 hours, but there is a 25% likelihood an interruption could last more than ten hours. But almost without exception, the more comprehensive methods require more data, involve more involved mathematical methods, and are more sensitive to details or implicit assumptions.

These various reliability prediction methods differ in the degree of approximation; the type of system they can analyze; and the type of reliability evaluation they provide; their complexity of analysis; and data requirements. In addition many methods are most appropriate for application to one type of system design (e.g., network, radial) or under certain conditions (only one source of supply but many available power transmission paths). The choice of reliability assessment method depends on all these factors, as well as the degree of detail sought by the planner.

In most places throughout this book, the authors will use a simple reliability assessment method that computes only the expected time power is unavailable, not the frequency and duration. This method is general enough for our examples and much simpler. In many cases, merely adding together the probabilities of rare occurrences, rather than considering that they could overlap further approximates reliability analysis.

4.5 COST IS USUALLY THE MAJOR CONSIDERATION

Quantity and quality *both* have value for the electric consumer. But so does the consumer's money and as a result, no factor in the decision about energy source and energy usage is more important to most consumers than cost. Chapters 4 and 5 will deal with planning, cost evaluation, engineering economics and decision making. The authors have included this short section to emphasize the point that quantity, quality, and above all, *cost*, form the basis for the market's evaluation of suppliers and technologies, and that a competitive, deregulated electric industry will no doubt respond by offering options with respect to reliability and power quality.

The three dimensions of power – quantity, quality, and cost – form a “value cube” that determines the overall benefit consumers see from the use of electric power (Figure 4.4). The importance of each in determining what is “best” will vary from one consumer to another.




Figure 4.4 Ultimately, the value electric power consumers see in their electric power supply is based on a “value volume” of three dimensions: quantity, quality, and cost.

Electricity Pricing Has Traditionally Been Based Only On Quantity

Traditionally, the price of power in the electric utility industry was based solely upon the *quantity* used. A consumer would pay more if he or she used more power and less if their use was less. Except in rare and specially treated cases, an electric utility would offer a “one size fits all” level of quality. Customers got what was available, whether or not they needed and would have paid for higher quality, or whether they would have preferred a discount even if it meant somewhat lower quality.

The Market Comb

Numerous surveys and customer focus meetings conducted by utilities, consultants, and research firms (including the authors’ firm) indicate that electric consumers differ widely in their need for, and their willingness to pay for, reliability of service. Consumers also differ in exactly what “reliability,” and in a broader sense, “quality,” means to them, although availability of service, quick and knowledgeable response on the part of their supplier, and power of a usable nature are always key factors in their evaluation.




Figure 4.5 The electric marketplace can be likened to a comb: composed of many small niches, each made up of customers who have a different need for and cost sensitivity to reliability of service. Even those who put a high value on reliability may differ greatly in how they define “good reliability,” one reason why there are no broad market segments, only dozens of somewhat similar but different niches.

As a result, the electric power demand marketplace can be likened to a comb as shown in Figure 4.5 with a series of small niches, that vary from the few customers who need very high levels of reliability, to those who do not need reliability, only power, and are motivated purely by lowest cost.

Offering Variable Levels of Reliability

Traditionally, electric utilities have not offered reliability as a price-able commodity. They selected equipment, and designed and engineered their systems based on engineering standards that were aimed at maintaining high levels of *power system equipment* reliability. These standards and methods, and the logic behind them, were actually aimed at minimizing utility equipment outages. The prevailing dogma maintained that this led to sufficiently high levels of reliable service for all customers and that reliable service was good. This cost of the reliability level mandated by the engineering standards was carried into the utility rate base. A utility’s customers basically had no option but to purchase this level of reliability (Figure 4.6). Only large industrial customers, who could negotiate special arrangements, had any ability to make changes bypassing this “one-size-fits-all” approach to reliability that utilities provided.

The traditional way of engineering the system and pricing power has two incompatibilities with a competitive electric power marketplace. First, its price is cost-based; a central tenet of regulated operation, but contrary to the market-driven paradigm of deregulated competition.




Figure 4.6 Generally, the cost of providing reliability in a power system is a non-linear function of reliability level: providing higher levels carries a premium cost, as shown by the line above. While reliability varies as a function of location, due to inevitable differences in configuration and equipment throughout the system, the concept utilities used was to design to a single level of reliability for all their customers.




Figure 4.7 A competitive market in power will recognize the demand for various levels of reliability and providers will have to broaden the offerings they provide (left) until the range of options available to the customer covers the entire spectrum of capabilities, as shown at the right. Actual options offered will be even more complicated because they will vary in the type of reliability (e.g., perhaps frequency or duration is more important to a specific customer; seven days a week or only normal business hours).

Secondly, any “one-size-fits-all” reliability price structure will quickly evaporate in the face of competitive suppliers. Someone among the various suppliers will realize there is a profit to be made and a competitive advantage to be gained by providing high-reliability service to consumers willing to pay for it. While others will realize that there are a substantial number of customers who will buy low reliability power, as long as it has a suitably low price. The concept that will undoubtedly emerge (in fact, is already evolving at the time this book is being written) in a deregulated market will be a range of reliability and cost options. The electric utility market will broaden its offerings, as shown in Figure 4.7.

For example, distributed generation (DG) interacts with the opportunities that the need for differentiated reliability creates in a deregulated marketplace in two ways. First, this new marketplace will create an opportunity for DG, which can be tailored to variable reliability needs by virtue of various design tricks (such as installing more/redundant units). Secondly, DG is a “threat” that may force utilities to compete on the basis of reliability, because it offers an alternative, which can easily be engineered (tailored), to different reliability-cost combinations. Utilities may have no choice but to compete in pricing reliability.

4.6 TWO-Q ANALYSIS OF POWER SYSTEMS

Two-Q plots similar to those shown in Chapter 3, but oriented toward measurement of system capability rather than consumer needs, can be used to analyze the ability of a power system to meet projected customer demands for both quantity and quality of power. They provide a useful tool in determining both the appropriateness of the system design, and how to tailor it to best fit consumer needs.

Figure 4.8 shows a two-Q plot of the capability of a set of feeders in a utility system in the mid U.S. The area involved is part of a small town of 40,000 population and its industrial area. Over the period 1989 – 1998, the peak load in this area grew by more than 10% (from 19.5 to 21.7 MW). Planners for the utility, constrained by pressures to minimize capital cost, accommodated the load growth largely through a process of transferring switching and making small changes in the system. From a “load flow” (capacity, or quantity only) perspective the planning was successful - the utility figuring out a way to connect consumers with sufficient capacity to satisfy the peak load, at least when everything was working as expected.

However, as Figure 4.8 shows, the various means taken to accommodate more quantity in power gradually eroded the system’s reliability capability. This was both due to use of contingency margin for service of peak load and use of switching and alternate paths for peak load service. Both means left no “backup” for contingencies, with the result that SAIDI in the region went from acceptable, to completely unacceptable by local consumers’ standards.




Figure 4.8 Two-Q plot of capability of a set of feeders in a large utility system in the central U.S, showing how increases made to accommodate higher load gradually eroded capability to deliver high reliability of supply. See text for details.




Figure 4.9 Two-Q plot of the design flexibility and options available in the region from Figure 4.8 showing results of Two-Q analysis. See text for details.

Figure 4.9 shows a two-Q diagram of design options determined from detailed evaluation of the system's capability and consumer demands in the area. This study was done using a computerized method that analyzes both capacity (load flow type factors) and reliability (frequency, duration, voltage sag probabilities) simultaneously (Two Q analysis). The dashed line in Figure 4.9 shows the actual migration path for the system as modified and operated by the utility over the period 1989 – 1998. Total capital cost during the period was \$41,500 (PW 1989).

The lowest of the solid lines, labeled "optimum zero cost" shows the locus of superior performance points obtainable through re-switching and no more than \$15,000 in cost. This indicates that the entire migration path for the period was inferior by roughly 40 minutes SAIDI or .66 MW quantity capability. This was due to the utility's lacking reliability-based techniques to hone the design well. Traditional single Q (quantity) design methods simply could not optimize the system to meet consumer needs at lowest cost.

The other solid lines in the drawing show the optimized capacity curves for various budget-levels (PW 1989 capital). The lines represent the locus of points for different designs or options, not the same plan. But all points on each curve are options purchasable at that budget amount. The utility's options and the trade-off between reliability of service and configuration for peak capacity is only clear.

4.6 CONCLUSION AND SUMMARY

Reliability of service is one of the major factors customers weigh in perceiving how well their electric supply system is doing its job. Whether that service is provided by an electric utility system, or a distributed generator, the planners of the supply system must set appropriate reliability goals and plan, engineer, and operate it to achieve targeted levels of customer service reliability. Various reliability indices can be used to track frequency and duration of customer interruptions over time and by location within the distribution system, in order to identify problems and ensure that goals are being met, but ultimately the only test that matters is consumer satisfaction.

Most utilities make sure that reliability targets are met, not through detailed reliability analysis in the planning stage, but implicitly through contingency and design criteria that are based on assessment of their reliability impacts. A more effective but complicated approach is to explicitly analyze reliability of planned designs, in an effort to determine if they are suitable and to engineer them to fully meet requirements at the lowest possible cost. Examples later in this book will illustrate this approach.

Electric consumers vary in their need for reliability of service as much as they vary in the amount of power they need. The electric power marketplace can be likened to a comb, a series of niches, each populated by consumers with different needs for both quality and quantity of power, and different preferences in how they would trade off cost against reliability of service, if forced to make sacrifices.

Value-based planning, although difficult, is also useful, particularly in cases where reliability is an issue to large or special industrial and commercial customers. An important element is to understand and apply how much the consumer values reliable service by analyzing their cost of interruption.

REFERENCES AND BIBLIOGRAPHY

- P. F. Albrecht and H. E. Campbell, "Reliability Analysis of Distribution Equipment Failure Data," EEI T&D Committee Meeting, January 20, 1972.
- R. N. Allan, et al., "A Reliability Test System for Educational Purposes – Basic Distribution System Data and Results," *IEEE Transactions on Power Systems*, Vol. 6, No. 2, May 1991, pp. 813-821.
- D. Atanackovic, D. T. McGillis, and F. D. Galiana, "The Application of Multi-Criteria Analysis to a Substation Design," paper presented at the 1997 IEEE Power Engineering Society Summer Meeting, Berlin.
- R. Billinton, and J. E. Billinton, "Distribution System Reliability Indices," *IEEE Transactions on Power Delivery*, Vol. 4, No. 1, January 1989, pp. 561-568.
- R. Billinton R., and R. Goel, "An Analytical Approach to Evaluate Probability Distributions Associated with the Reliability Indices of Electric Distribution Systems," *IEEE Transactions on Power Delivery*, PWRD-1, No. 3, March 1986, pp. 245-251.
- R.E. Brown, *Electric Power Distribution Reliability – Second Edition*, CRC Press, Boca Raton, 2009.
- R.E. Brown, *Business Essentials for Utility Engineers*, CRC Press, Boca Raton, 2010.
- J. B. Bunch, H.I. Stalder, and J.T. Tengdin, "Reliability Considerations for Distribution Automation Equipment," *IEEE Transactions on Power Apparatus and Systems*, PAS-102, November 1983, pp. 2656 - 2664.
- "Guide for Reliability Measurement and Data Collection," EEI Transmission and Distribution Committee, October 1971, Edison Electric Institute, New York.
- W. F. Horton, et al., "A Cost-Benefit Analysis in Feeder Reliability Studies," *IEEE Transactions on Power Delivery*, Vol. 4, No. 1, January 1989, pp. 446 - 451.
- Institute of Electrical and Electronics Engineers, *Recommended Practice for Design of Reliable Industrial and Commercial Power Systems*, The Institute of Electrical and Electronics Engineers, Inc., New York, 1990.
- A. D. Patton, "Determination and Analysis of Data for Reliability Studies," *IEEE Transactions on Power Apparatus and Systems*, PAS-87, January 1968.

- N. S. Rau, "Probabilistic Methods Applied to Value-Based Planning," *IEEE Transactions on Power Systems*, November 1994, pp. 4082 - 4088.
- A. J. Walker, "The Degradation of the Reliability of Transmission and Distribution Systems During Construction Outages," Int. Conf. on Power Supply Systems. IEEE Conf. Pub. 225, January 1983, pp. 112 - 118.
- H. B. White, "A Practical Approach to Reliability Design," *IEEE Transactions on Power Apparatus and Systems*, PAS-104, November 1985, pp. 2739 - 2747.
- W. Zhang and R. Billinton, "Application of Adequacy Equivalent Method in Bulk Power System Reliability Evaluation," paper presented at the 1997 IEEE Power Engineering Society Summer Meeting, Berlin.

5

Cost and Economic Evaluation

5.1 INTRODUCTION

As was stated previously, reliability and service quality are very important, but “cost is king” in almost all-electric distribution decisions, both those made by consumers and by utilities. Every decision contains or implies certain costs: equipment, installation labor, operating, maintenance, losses, and many others. Regardless what type of paradigm is being used, and whether traditional or new technologies are being considered, planners and decisionmakers must deal with cost as a major element of their evaluation.

Many planning decisions revolve around selecting from alternatives on the basis of *when* they call for money to be spent. Two plans might be equal with respect to capacity and reliability, but different in both *when* and *how much* they require to be spent: one might eventually spend a lot, but only in the long run; the other might spend less, but call for more initial investment. At what point is the planner justified in picking the larger total expense because most of it lies far into the future?

This chapter provides a brief tutorial on engineering economics for planning to provide a basis for the methods applied later in the book. Section 5.2 reviews the major elements of cost. Section 5.3 looks at the time value of money and future-cost discount methods. Benefit cost analysis is covered in section 5.4, and marginal benefit cost analysis, a cornerstone of aging infrastructure planning, in section 5.5. Only a quick review is given here. More comprehensive treatment is available in books listed at the end of this chapter.

5.2 COSTS

“Cost” is the total sacrifice that must be expended or traded in order to gain some desired product or end result. It can include money, labor, materials, resources, real estate, effort, lost opportunity, and anything else that is given up to gain the desired end. Usually, the costs of these many different resources and commodities are measured on a common basis – money – by converting materials, equipment, land, labor, taxes and permits, maintenance, insurance, pollution, and lost opportunity costs to dollars, pounds, marks, yen, or whatever currency is most appropriate.

In those cases where all these elements can be put on a common basis, the subsequent planning can be done in a *single-attribute* manner. The goal being to achieve the planner’s goals while minimizing the total cost by trading the cost of one item against the others to find the best overall mix. However in some cases, one or more types of cost cannot be converted to money, for example, aesthetic impact or other “intangibles.” In such cases, *multi-attribute* planning and cost minimization must be done, involving more complicated evaluation methods (D. Atanackovic et al, 1997).

Types of Cost

Initial cost is what must first be spent to obtain the facility or resource. It includes everything needed to put the facility in place, and is usually a single cost, or it can be treated as a single cost in planning studies if a series of costs over a period of time lead up to completion of the facility.

Continuing costs are those required to keep the facility operational and functioning the way its owner wants it to do. They include inspection and maintenance, fuel, supplies, replacement parts, taxes, insurance, electrical losses, and perhaps other expenditures. They persist as long as the facility is in use. Usually, continuing costs are studied on a periodic basis – daily, monthly, or annually.

Fixed costs are those that do not vary as a function of any possible variable element of the planning analysis or engineering study being carried out. For example, the annual costs of taxes, insurance, inspection, scheduled maintenance, testing, re-certification, and so forth, required to keep a 500 kW DG unit in service, do *not* vary depending on how often or heavily it is run. They are fixed costs.

By contrast, the DG unit’s fuel costs *do* vary with usage – the more it is run the higher the fuel costs. These are *variable costs*, a function in some manner of the amount of load served. Some types of maintenance and service costs may be variable, too. Highly used DG units need to be inspected and stress-caused wear and tear repaired more often than in units run for fewer hours.

		Timing of Expense	
		Initial	Continuing
Variability of Expense	Fixed	One-time “first” costs that do not depend on usage pattern	Periodic, on-going costs that do not depend on usage pattern
	Variable	One-time “first” costs that do depend on usage pattern	Periodic, on-going costs that do depend on usage pattern

Figure 5.1 Two characterizations, one based on when costs occur and the other on how they vary with usage, result in four categories of costs.

Figure 5.1 shows how these two sets of categorizations – initial, continuing, fixed, and variable - can be combined. This two by two matrix allows planners to identify fixed initial costs, variable initial costs, fixed continuing costs, and variable continuing costs. In this way, costs in studies can be characterized simultaneously by both their relationships to time and how they respond to assumptions about usage.

Sunk Cost

Once a cost has been incurred, even if not entirely paid, it is a “sunk cost.” For example, once a facility has been constructed and is in operation, it is a sunk cost, even if ten years later the company still has its capital depreciation on its books as a cost.

Embedded, Marginal, and Incremental Cost

Embedded cost is that portion of the cost that exists in the current system, configuration, or level of use. Depending on the application, this can include all or portions of the initial fixed cost, and all or parts of the variable costs.




Figure 5.2 Hourly cost of losses for power shipped through a substation transformer as a function of the load. Currently loaded to 45 MVA, it has an “embedded” loss cost of \$1.73/MVA, and a marginal cost of losses at that same point of \$3.70/MVA. The incremental cost of losses for an increase in load to 55 MVA load is \$5.90/MVA.

Often, the “embedded” cost is treated as a fixed cost in subsequent analysis about how cost varies from the current operating point. Marginal cost is the slope (cost per unit) of the cost function at the current operating point (Figure 5.2). This point is usually (but not always) the point at which current embedded cost is defined.

Incremental cost is the cost per unit of a specific jump or increment – for example, the incremental cost of serving an additional 17 MVA from a certain substation or the incremental cost of losses when load on a feeder decreases from 5.3 to 5.0 MVA. Marginal and incremental costs both express *the rate of change of cost* with respect to the base variable, but they can differ because of the inconsistencies and variances in the cost relationships.

In the example shown in Figure 5.2, the marginal cost has a slope (cost per unit change) and an operating point (e.g., 45 MW in Figure 5.2). Incremental cost has a slope and both “from” and “to” operating points (e.g., 45 MVA to 55 MVA in Figure 5.2) or an operating point and an increment (e.g., 45 MW plus 10 MVA increase).

5.3 TIME VALUE OF MONEY

Expenses Now and Then

Almost any electrical project has both initial and continuing costs. Figure 5.3 shows the costs for a new distribution substation, planned to serve 65 MW of residential load on the outskirts of a large city. It has a high initial cost that includes the equipment, land, labor, licenses and permits, and everything else required in building the substation. Thereafter, it has taxes, inspection, maintenance, electrical losses, and other continuing costs on an annual basis. Figure 5.4 shows the initial and continuing costs identified with respect to fixed and variable categories, too. Typical of most electrical planning situations, the initial cost is almost entirely fixed, and the continuing costs are mostly variable.

Decisions Based on When and How Much is Spent

Any utility planner must deal with two types of time versus money decisions. The first involves deciding whether a present expense is justified because it cancels the need for some future expense(s). For example, suppose it has been determined that a new substation must be built in an area of a utility's service territory where there is much new customer growth. Present needs can be met by completing this new substation with only one 25 MVA transformer, at a total initial cost of \$1,000,000. Alternatively, it could be built with two 25 MVA transformers installed, at a cost of \$1,400,000. Although not needed immediately, this second transformer will be required within four more years because of continuing growth. If added at that time, it will cost \$642,000 extra - \$242,000 more than if it is included at the time the substation is first built. The higher cost to add it later is a common fact of life in T&D systems' planning and operation. It's also a reflection of the additional start-up cost for a new project and of the higher labor and safety costs required for installing a transformer at an already-energized substation, rather than when the site is cold.

Which plan is best? Should planners recommend that the utility spend \$400,000 now to avoid a \$642,000 expense four years from now?

A second and related type of time-vs.-money decision is determining if a present expense is justified because it will reduce future operating expenses by some amount. Suppose two types of distributed generators are candidates for installation at a high-reliability customer site. The first, Type A, costs \$437,000 installed. It fully satisfies all loading, voltage regulation, and other criteria. On the other hand, the higher efficiency type B unit will also meet these same requirements and criteria and lower annual fuel costs by an estimated \$27,000 per year. However, it has an initial cost of \$597,000. Are the planners justified in recommending that \$160,000 be spent, based on the long-term fuel savings?




Figure 5.3 A new substation's costs are broken into two categories here, the *initial cost*, a one-time cost of creating the new substation, and its *continuing annual costs* which are required to maintain it in operation.




Figure 5.4 The new substation's costs can also be viewed as composed of both fixed costs – those that are constant regardless of loading and conditions of its use, and variable costs. Those costs change depending on load, conditions, or other factors in its application as part of the power system. In this example, the variable costs increase nearly every year, reflecting higher losses' costs as the load served gradually increases due to load growth.

These decisions, and many others in utility planning, involve comparing present costs to future costs, or comparing plans in which a major difference is *when* money is scheduled to be spent. To make the correct decision, the planner must compare these different costs in time on a sound, balanced basis, consistent with the electric utility's or DG owner's goals and financial structure.

Common Sense: Future Money Is Not Worth As Much As Present Money

Few people would happily trade \$100 today for a reliable promise of only \$100 returned to them a year from now. The fact is that \$100 a year from now is not worth as much to them as \$100 today. This "deal" would require them to give up the use of their money for a year. Even if they do not believe they will need to spend their \$100 in the next year, conditions might change and they might discover later that they do need that money. Then too, there is a risk that they will not be re-paid, no matter how reliable the debtor.

Present Worth Analysis

Present worth analysis is a method of measuring and comparing costs and savings that occur at different times on a consistent and equitable basis for decision-making. It is based on the present worth factor, P , which represents the value of money a year from now in today's terms. The value of money at any time in the future can be converted to its equivalent present worth as:

$$\text{Value today of } X \text{ dollars } t \text{ years ahead} = X \times P^t \quad (5.1)$$

Where P is the *present worth factor*

And t indexes the year ahead

For example, suppose that $P = .90$; then \$100 a year from now is considered equal in value to today's money of

$$100 \times (.90) = \$90$$

and \$100 five years from now is worth only

$$\$100 \times (.90)^5 = \$59.05$$

Present worth dollars are often indicated with the letters PW. Today's \$100 has a value of \$100 PW, \$100 a year from now is \$90 PW, \$100 five years from now is \$59.05 PW, and so forth.




Figure 5.5 Present Worth Analysis discounts the value, or worth, of a dollar the further it lies in the future. Shown here is the value of one dollar as a function of future year, evaluated at a .90 present worth factor (11.1% discount rate).

Alternatively, the present worth factor can be used to determine how much future money equals any amount of current funds. For example, to equal \$100 present worth, one year from now a person would need

$$\$100/.90 = \$111.11$$

Two years from now, one would need $\$100/ (.90)^2 = \$123.46$ to equal \$100 PW.

A continuing annual future cost (or savings) can be converted to a present worth by adding together the PW values for all future years. For example, the present worth of the \$27,000 in annual fuel savings discussed for the type B DG unit earlier in this chapter can be found by adding together the present worth of \$27,000 next year, \$27,000 the year after, and so forth

$$\begin{aligned}
 \text{PW of } \$27,000/\text{year} &= \sum_{t=1}^{\infty} (\$27,000 \times P^t) & (5.2) \\
 &= \$27,000 \times \left(\sum_{t=1}^{\infty} P^t \right) \\
 &\cong \$27,000 \times \left(\sum_{t=1}^{30} P^t \right) \\
 &= \$258,554
 \end{aligned}$$

Discount rate

Present Worth Analysis essentially *discounts* the value of future costs and savings *just* because they lie in the future, as shown in Figure 5.5. The discount rate used in an analysis, d , is the perceived rate of reduction in value of money from year to year. The present worth factor is related to this *discount rate*

$$P(t) = 1/(1+d)^t \quad (5.3)$$

Where d = discount rate

And t = future year

If d is 11.11%, it means that a year ahead dollar is discounted 11.11% with respect to today's dollar, equivalent to a present worth factor of $P = (1/1.111) = .90$. Therefore, \$111.11 a year from now is worth $\$111.11/1.1111 = \$100$ today. A decision-making process based on the values of $d = 11.11\%$ and $PWF = .90$ would conclude that spending \$90 to save \$100 a year from now was a break-even proposition (i.e., there is no compelling reason to do it). While if the same \$100 savings can be had for only \$88, the "deal" has a positive value.

A higher discount rate, corresponding to a lower present worth factor, renders the decision-making process less willing to trade today's costs for future costs. If the discount rate is doubled, to 22.2%, P drops to .82. Now, \$88 spent today to obtain a \$100 cost reduction a year from now is no longer viewed as a good investment: that \$88 must yield at least $\$88/1.222 = \$108$ a year from now to be viewed as merely break even.

Present worth analysis does not say "no" to truly essential elements of a plan

Power system planners with a commitment to see that service is always maintained often fall into a trap of wanting to build everything as early as practical in order to have plenty of capacity margin and to allow for unexpectedly high load growth. As a result, they come to view present worth analysis, and any other aspect of planning that might say "delay," as negative.

But this is not the case. If applied correctly, present worth analysis never says "no" to essential expenses, for additions or changes that must be done *now*. For example, if a hurricane has knocked down ten miles of critically needed lines; there is no question about the *timing* of the expense in putting it back up. That must be done as soon as conditions permit, and present worth analysis is not even appropriate to apply to the decision.

Present Worth Analysis should be used to evaluate and rank alternatives only when there is a difference in the *timing* of expenses (the substation could be built this year or next year), or when costs and savings that occur at different times

must be balanced (more spent today will lower cost tomorrow). Present worth analysis is an essential element of keeping costs as low as possible, and it should be applied in all situations where alternatives differ with respect to *when* expenses will be incurred.

How Are Present Worth Factors and Discount Rates Determined?

There is no one rule or set of guidelines that rigorously defines what contributes to the present worth factor (PWF), or its companion, the discount rate. Likewise, there is no inviolate formula like Ohm's law or Schroedinger's equation that lays out completely and with no exceptions exactly how to compute PWF from those factors that are deemed relevant. Quite simply, PWF is merely a number, a decision-making tool that allows planners to compare future expenses to present ones. The present worth factor is often arrived at through careful analysis and computation, but sometimes it is an empirical estimate "that just works well."

The present worth factor is simply a value that sums up all the reasons why an organization would prefer to spend money tomorrow rather than today.

There can be many reasons why a person may wish to limit spending today in favor of spending tomorrow. These reasons all influence the selection of a value for the present worth factor. Priorities vary greatly from one company to another and also change with time. Among them is one-reason planners often finding difficult to accept: "We just don't have the money, period." In most electric utilities, present worth factors and discount rates are determined by a utility's financial planners and based on the company's regulatory requirements for financing and use of money. The methods applied to fix these requirements quantitatively and their relation to PWF are beyond the scope of this discussion, and available in other texts (e.g., *IEEE Tutorial on Engineering Economics*). However, it is worth looking at some of the primary influences on the present worth factor and how and why they influence it.

Interest rate

The fact that it takes more future dollars to equal a dollar today is often attributed to interest rate: a person who has \$100 today can invest it at the prevailing interest rate, i.e., so that a year from now it will be $(1+i)$ times as much. Thus, \$100 invested today at an annual rate of 5% interest will be worth \$105 a year from now and \$128 five years from now. If the prevailing interest rate is 5%, then it is cost-effective to spend \$100 only if it will save or return a value a year from now that exceeds \$105. Otherwise, it would be better simply to leave the money invested and drawing interest. A PWF of .952 can be applied in a present worth analysis as described earlier, and will lead to decisions that reflect this concern.

But, in practice, present worth factor clearly represents something more than *just* interest rate, because the present worth factor as applied by most electric utilities is nearly always greater than what corresponds to the prevailing interest rate. For example, at the time of this writing, the inflation-adjusted interest rate on safe, long-term investments is about 5%, yet most utilities are using a present worth factor of about .11, equivalent to a 12.4% interest rate.

The difference, 12.4% versus 5.5% in this case, is attributable to other factors, although it may prove difficult or impossible for a distribution planner to determine and quantify all of them.

Inflation

Inflation is *not* one of the factors normally taken into account by the present worth factor, although this is often misunderstood, and new planners often assume it is a part of the analysis. Inflation means that what costs a dollar today will cost more tomorrow – 3% annual inflation would mean an item costing a dollar today will be priced at \$1.03 a year from now. While inflation needs to be considered by a utility's financial planners, distribution planners can usually plan with *constant dollars*, by assuming that there is no inflation.

The reason is that inflation raises the cost of everything involved in distribution planning cost analysis. If inflation is 3%, and then on average a year from now equipment will cost 3% more, labor will cost 3% more, as will paperwork and filing fees and legal fees and taxes and replacement parts and transportation and everything else. More than that, over time inflation will call for similar adjustments in the utility's rates, in the value of its stock, its dividends, and everything associated with its expenses and financing (including hopefully the planners' salaries). It affects nearly every cost equally, at least to the extent that costs can be predicted in advance.¹ From the standpoint of planning, since inflation makes no impact on the *relative costs* of various components, it can be ignored, making the planning process just a bit easier.

Some utilities do include inflation in their planning and in their present worth analysis. In such cases, an increment to account for inflation is added to the discount rate. Given a 5% interest rate and a 3% inflation rate discussed above, this would mean that the planners' present worth factor might be

$$P = 1/(1+5\%+3\%) = .926$$

Present worth analysis would now compare today's investment measured in today's dollars against tomorrow's investment measured in tomorrow's inflated

¹ In actuality, utility and business planners and executives all realize that inflation will not impact every cost equally, but small variations in individual economic sectors have proven impossible to predict. In cases where there is a clear indication that certain costs will rise faster or slower than the general inflation rate, that difference should be taken into account.

dollars. This type of accounting of future costs must sometimes be done for budgeting and finance estimation purposes.²

Inflation must be included in the analysis, but all planning can be done with “constant dollars”

However, while an “inflated dollars” present worth analysis sometimes has to be done to estimate budgets, it is rarely an effective planning tool. Planning and decision making are facilitated when expenses and benefits are computed on a common basis – by measuring costs from all years in *constant dollars*. Is plan A, which calls for \$1.35 million in expenses today, worth more than plan B, which costs \$1.43 million three years from now? If both plans have been calculated using their respective present worth measured in constant dollars, then it is clear that plan A is the less costly of the two. But if inflation were a factor in the present worth analysis, then one has to do further adjustment of those numbers to determine which is best. (At 3% inflation, \$1.35 million in three years will equal \$1.48 million in then inflated dollars, making plan B, which costs \$1.43 million *three-year ahead inflated dollars*, the better plan). Beyond this, errors or unreasonable cost estimates for future projects are more readily caught when expressed in non-constant dollars. For these reasons, present worth analysis in constant dollars is often used for planning, but an initial analysis, which accounts for inflation, may occasionally be needed for other purposes.

Earnings targets

The present worth factor for a dollar under the utility’s control should be higher than $1/(1 + \text{interest rate})$, because *the utility must be able to do better with its earnings than the prevailing interest rate*. If a utility cannot better the prevailing interest rate with its own investment, then it should liquidate its assets and invest the results in earning that interest rate. Frankly, it won’t get the chance to do so: long before that, its shareholders will sell their stock and invest their money elsewhere, in companies that *can* beat the prevailing interest rate through their investment of the stockholder’s money.³ Therefore, a goal must be to use the company’s money to earn more than it could by other safe means.

² For example, when estimating a future budget, the costs often do have to be put in dollars for the years being estimated, and those are inflated dollars.

³ Planners from municipal utilities may argue that this does not apply to them, but that is not necessarily true. If a municipal electric department cannot “earn” more from its investments than other businesses could, it is costing the city money to design and operate a system that needs to be subsidized by taxpayer money. Such subsidized operation may be the policy of the city government (some city governments deliberately subsidize commercial electric rates, using the low rates as a tool to attract employers to their region). Regardless, its planners should try to get as sound a return as possible.

As a result, while the prevailing interest rate may be 5%, the utility's financial planners may have determined that a 12% earnings potential on all new expenditures is desirable. Rather than a PWF of .952 (5% interest) the utility would use a PWF of .892 (12%) to interject into the decision-making process its unwillingness to part with a dollar today unless it returns at least \$1.12 a year from now. This is one tenet of many corporate plans, that all investments must return a suitable earning.

Risk

One hundred dollars invested with the promise of \$108 payback a year from now may look good, but only if there is a very small chance that the investment will go wrong, with the loss of the interest *and* the \$100 itself. Such cases are very rare, particularly if investments are made wisely, but there are other types of risks that a utility faces. For example, suppose that shortly after spending \$100 to save \$108 a year from now, a company gets hit by an unexpected catastrophe (e.g., an earthquake) that damages much of its equipment. It may desperately wish it had that \$100 to pay for repairs, work that it simply has to have done to re-establish production. It now has no choice but to borrow the money it needs at short-term interest rates, which might be 12%. In retrospect, its expenditure of \$100 to save \$108 a year later would look unwise.

In practice, a present worth factor often includes a bias or margin to account for this type of "bird in the hand" value. By raising the PWF from 8% to 10%, the utility would be stating that yes: perhaps \$108 is the year-ahead earnings goal. Simply breaking even with that goal is not enough to justify committing the company's resources: today's money will be committed only when there are very sound reasons to expect a better than minimum return.

Planning errors

In addition, the present worth factor often implicitly reflects a sobering reality of planning – mistakes cannot be avoided entirely. Under the very best circumstances, even the finest planning methods average about 1% "planning error" for every year the plan is extended into the future. Distribution expansion projects that must be committed a full five years ahead will turn out to spend about 5% more than if the planner could somehow go back and do things over again, knowing with hindsight exactly what the minimum expenditure is needed just to get by.

Adding 1% – or whatever is appropriate based on analysis of the uncertainty and the planning method being used – to the PWF biases all planning decisions so that they reflect this imperfection of planning. This makes the resulting decision-making process slower to spend money today on what *appears* to be a good investment for tomorrow, unless the predicted savings include enough

margin over the element of risk to account for the fact that the planning method could be wrong.

Present Worth Factor Is a Decision-Making Tool, Not a Financial Factor

Thus, the present worth factor and discount rates are merely tools. Present worth analysis is a decision-making method to determine *when* money should be spent. It can embrace some or all of the factors discussed above, as well as others. Two equally prudent planners might pick very different present worth factors, depending on their circumstances, as summarized in Table 5.1 with the inputs that determined the PWFs used in the mid-1990s by a large investor-owned utility (IOU) in the northwestern United States and a municipality in the southwestern United States.

But while the present worth factor can be affected by many factors, the simple fact is that such determinations are irrelevant to the electric service planner. Present worth factor is used to evaluate and rank alternatives based on when they call for expenditures. It is just a tool, part of the analysis, period.

A relatively low PW factor means that the planning process will be more likely to select projects and plans that spend today in order to reduce costs tomorrow. As the PWF is decreased, the planning process becomes increasingly unwilling to spend any earlier than absolutely necessary unless the potential savings are very great. A very low PW factor will select plans that wait “until the last moment” regardless of future costs.

Table 5.1 Discount Rate “Computation” for Two Utilities

Factor	IOU	Muni	Comment
Prevailing interest rate	5.7%	5.1%	Municipal bond rate is lower than prime.
Inflation factor	-	-	Both do constant dollar planning.
Earnings target	5.5%	-	Municipal has no need for “earnings.”
Risk	1.5%	3.0%	Political cost more serious than financial.
Planning error	1.0%	1.5%	This IOU has the better planning method.
“We just don’t have funds”	-	5.0%	This municipality is nearly “broke.”
Total discount	<u>13.7%</u>	<u>14.6%</u>	
Equivalent PWF	87.9%	87.3%	

Table 5.2 Comparison of Yearly Expenses by Category for an Eight-Year Period

Year	Load-MW	Alternative A Build initially with two			Alternative B Build initially with one		
		Capital	O&M&T	Losses	Capital	O&M&T	Losses
0							
1		20			20		
2		370			290		
3	12	1010	52	54	690	44	39
4	15.5		110	119		92	99
5	18.5		110	130		92	122
6	21.3		110	142	80	92	147
7	23.8	—	<u>110</u>	<u>155</u>	<u>562</u>	<u>101*</u>	<u>164*</u>
	1400 492		600	1642	317	666	
			Total	= 2492		Total	= 2633

*Taxes and losses for the transformer added in year eight (alternative B) are prorated for that year, since it is assumed here that it is installed just prior to summer peak (six months into the year).

Comprehensive Present Worth Example

Table 5.2 and Figure 5.6 show more cost details on the two-versus-one substation transformer decision mentioned at the beginning of this section. In the case shown, the decision is being made now (in year zero), even though the substation will actually be built three years from now – the decision must be made now due to lead time considerations. Alternative A calls for building the new substation with two 25 MVA transformers at a total cost of \$1,400,000 initially. Alternative B defers the second transformer four years (to year 7), reducing the initial cost to \$1,000,000. And again, the addition of the second transformer four years later is estimated to cost \$642,000 because of project start-up costs, separate filing and regulatory fees, and the fact that the work will have to be done at what will then be an energized site.

However, this decision involves other costs, which while secondary in importance, should be considered in any complete analysis. For example, no-load losses, taxes, and O&M will be greater during the first four years in Alternative A than in Alternative B. Table 5.2 and Figure 5.6 compare the expenditure streams for the eight-year period for both alternatives, including capital; operations, maintenance, and taxes (O&M&T); and losses.⁴

⁴ The table and figure cover the only period during which there are differences that must be analyzed in the decision making. At the end of the eight-year period, the two alternatives are the same – either way the substation exists at that point with two 25 MVA




Figure 5.6 Yearly expenses for the two alternative plans for the new substation (constant dollars).

Note also, that the capital for substation construction in either alternative is not spent entirely in the year the substation is completed. A small amount is spent two years earlier (for filing and preparing for the site), and the site, itself, and some equipment purchased a year earlier so that work can begin on schedule.

In comparing differences in the two alternatives, observe also that the second transformer in Alternative A increases O&M&T costs during the first four years of operation as compared to the other alternative, because taxes must be paid on its additional value, and it must be maintained and serviced.

The difference in losses' costs is more complicated. Initially, losses are higher if two transformers are installed, because this means twice the no-load losses, and the twin transformers are initially very lightly loaded, so that there is no significant savings in load-related losses over having just one transformer. But by year six the two-transformer configuration's lower impedance (the load is split between both transformers) produces noticeable savings due to reduced losses as compared to the other alternative.

Overall, Alternative A spends fewer dollars in the period examined \$142,000 less. It has the lower amount of long-term spending.

Table 5.3 and Figure 5.7 show the various cost streams for both alternatives converted to PW dollars using a present worth factor of .90 (equivalent to $p =$

transformers installed. Present worth thereafter is essentially the same.

11.11%). Note that dollar amounts have dropped—discounted future dollars just aren't as valuable as the undercounted dollar figures. More important, is the present worth evaluation reverses the order of cost preference for the two options. Even though its total spending is less, Alternative A is rated a *higher* present worth by \$51,000 PW, because many of its expenses are closer to the present.

The total difference in PW is slightly more than 3% of the average PW value of the alternative. Considering the straightforward nature of this analysis, and the high proportion of the various cost elements common to both alternatives, it is very unlikely that the analysis is in error by 3%. Therefore, this is very likely a dependable evaluation: Alternative B is the recommended “low-cost” option. Alternative A would cost more.

Looking at the sensitivity of this decision to present worth factor, the two alternatives would be considered equal if the evaluation had used a present worth factor of .937, equal to a discount rate of 6.77%. Present worth factor would have to be higher (i.e., the discount rate lower) to have reversed the determination that Alternative B is, in fact, the low-cost option.

Again, considering that nominal PWF values less than .92 are extremely uncommon, the recommendation that Alternative B is the preferred option seems quite sound. Therefore, the planner's recommendation should be to build the substation with one transformer initially, and add the other at a later time, as needed.

Levelized Value

Often, it is useful to compare projects or plans on the basis of their *levelized* annual cost, even if their actual costs change greatly from one year to the next. This involves finding a constant annual cost whose present worth equals the total present worth of the plan, as illustrated by Figure 5.7. Leveled cost analysis is a particularly useful way of comparing plans when actual costs within each plan vary greatly from year to year over a lengthy period of time, or when the lifetimes of the options being considered are far different. It also has applications in regulated rate determination and financial analysis. In general, the present worth, Q , leveled over the next “ n ” years is

$$\text{Leveled value of } Q \text{ over } n \text{ years} = Q (d \times (1+d)^n) / ((1+d)^n - 1) \quad (5.4)$$

Thus, the steps to find the leveled cost of a project are: 1) Find the present worth of the project and, 2) Apply equation 5.2 to the present worth. As an example here, using Alternative A's calculated total present worth of \$1,684,000 from Table 5.3, that value can be leveled over the seven-year period in the table as

$$\begin{aligned} &= \$1,684,000 (.1111(1.1111)^7 / (1.1111^7 - 1)) \\ &= \$358,641 \end{aligned}$$

Table 5.3 Comparison of Yearly Present Worth by Category for an Eight-Year Period

Year	Present Worth Factor	Alternative A			Alternative B		
		Build initially with two Capital	O&M&T	Losses	Build initially with one Capital	O&M&T	Losses
0	1.0						
1	.90	18			18		
2	.81	300			235		
3	.73	736	38	39	503	32	28
4	.66		72	78		60	65
5	.59		65	77		54	72
6	.53		58	75	43	49	78
7	.48		53	74	269	48*	78*
		1054	286	343	1067	244	322
		Total = 1684			Total = 1633		

*Taxes and losses for the transformer added in year eight (alternative B) are prorated for that year, since it is assumed here that it is installed just prior to summer peak (six months into the year).




Figure 5.7 Present worth of expenses associated with the two alternative substation plans given in Table 5.2 and Figure 5.6.

By contrast, Alternative B has a leveled cost over the same period of \$347,780. If the utility planned to pay for alternative A or B over the course of that seven-year period, the “payments” it would have to make would differ by a little more than \$10,800 per year.

Lifetime leveled analysis

Rather than analyze planning alternatives over a short-term construction period planners or budget administrators often wish to look at a new addition to the system over the period in which it will be in service or be financed.

For example, in the two-or-one transformer substation planning problem presented above, the substation would be planned on the basis of providing a minimum of 30 years of service, and it would be depreciated on the company books for 30 years or more. Therefore, planners should analyze cost over 30 years and compare alternative plans, A and B, on that same basis.

Figure 5.8 shows projected costs for Alternative A over a 32-year period, which includes the first through thirtieth years of service for the substation. The total present worth of the substation’s costs over the period is just over \$3 million (construction plus 30 years of O&M&T plus losses), which represents a leveled cost of \$348,500/year. By contrast, Alternative B (the less expensive of the two construction schedules, as determined earlier) has a leveled cost of \$342,200/year. It saves an average of \$5,900/year or about \$500/month in terms of a 30-year budget perspective.




Figure 5.8 Cost of Alternative A from Figure 5.7 over a 32-year period that includes its first 30 years of service. Present worth of all cost comes to slightly more than \$3,000,000. Leveled over years three through 30, those in which the substation is planned to be in service, this represents an annual cost of nearly \$350,000 per year, which can be viewed as the annual cost of providing the substation service.

Table 5.4 Lifetime Analysis of Illumination

Incandescent					
Year	PWF	Capital	Energy	Total	PW
1	.90	\$2.00	\$23.00	\$25.00	\$22.50
2	.81		\$23.00	\$23.00	\$18.63
3	.73		\$23.00	\$23.00	<u>\$16.77</u>
Total Present Worth					\$57.90
Levelized over 3 years =					\$23.74

Halogen					
Year	PWF	Capital	Energy	Total	PW
1	.90	\$8.00	\$13.00	\$21.00	\$18.90
2	.81		\$13.00	\$13.00	\$10.53
3	.73		\$13.00	\$13.00	\$9.48
4	.66		\$13.00	\$13.00	\$8.53
5	.59		\$13.00	\$13.00	<u>\$7.69</u>
Total Present Worth					\$55.11
Levelized over 6 years =					\$14.95

Fluorescent					
Year	PWF	Capital	Energy	Total	PW
1	.90	\$17.00	\$8.00	\$25.00	\$22.50
2	.81		\$8.00	\$8.00	\$6.48
3	.73		\$8.00	\$8.00	\$5.83
4	.66		\$8.00	\$8.00	\$5.25
5	.59		\$8.00	\$8.00	\$4.72
6	.53		\$8.00	\$8.00	<u>\$4.25</u>
Total Present Worth					\$49.04
Leveled over 6 years =					\$11.62

Comparing cost over different equipment lifetimes

Very often the lifetimes or periods of cost being considered for two or more options are not equivalent, and yet it is desired to compare their costs on a valid basis. Leveled lifetime cost is a good way to do so. Table 5.4 compares the annual costs of providing 1,000 lumens of illumination (about the equivalent of a standard 75 watt incandescent light bulb) for 2,500 hours per year, using incandescent, halogen, or compact fluorescent lamps, which have expected lives of three, five, and six years respectively. The initial costs, the annual operating costs, and the lifetimes of these three options all differ substantially. In each case, initial and continuing annual energy costs are evaluated on a present worth basis over the lifetime of the particular lamp; then the leveled cost over that period is

determined. This permits a comparison that identifies the compact fluorescent as having less than half the annual cost of the incandescent.

In this example, each alternative was evaluated over a different period of time – its lifetime – yet the results are directly comparable. The results are identical to those that would have been obtained from analysis of all over any much longer period, taking replacement costs and other factors into account.

For example, suppose that the analysis in Table 5.4 had compared all three options over a 30-year period. That would have required a series of ten incandescent bulbs, one every three years, six-halogen, and five-compact fluorescent lamps. Each of the sections of Table 5.4 would have extended for 30 years with appropriate PWF, computations, and sums, but the analysis would have reached the same conclusion about which of the three lamps was best.

5.4 DECISION BASES AND COST-EFFECTIVENESS EVALUATION

How does a planner know that a particular project is “best”? In many cases this is determined once present worth analysis has adjusted all costs and savings with respect to equipment lifetime, schedule, and other time-value-of-money considerations. Such decisions are straightforward – in many cases trivial. But determining what to do is not always merely a matter of *just* identifying the option with the lowest cost. Selection may be based on the relative merits of costs versus savings, the benefits versus cost or the relative gain per dollar invested. In integrated resource planning (IRP) situations, the decision may also depend on the applicable philosophy about what constitutes savings and who should benefit most from a public utility’s investment in the system.

This section briefly reviews cost-effectiveness concepts and methods. More comprehensive discussions of the concepts discussed here and the methods developed to apply them can be found in the references at the end of the chapter.

Lowest Cost Among Alternatives

In many cases and from many perspectives, the goal of planning is to identify an alternative with the lowest cost. If all pertinent costs (capital, losses, O&M&T, etc.) are considered, and the appropriate type of present worth adjustment is made for timing, lifetimes, and schedules, then the alternative with the lowest PW cost *is* the preferred option.

Payback Period

The period of time required for “long-term” savings to re-pay the initial investment is often used as a measure of the desirability of making an investment in a resource or other addition to the electric supply. Returning to the “type A

versus type B DG unit” problem discussed earlier in this chapter, a fuel savings of \$27,000 per year will require six years ($6 \times \$27,000 = \$162,000$) to repay the \$160,000 additional cost of the type B unit over type A. (If the break-even is computed based on PW dollars at a .90 PW factor, the payback of the original \$160,000 takes ten years.)

“Payback period” analysis is most often used informally, to provide an additional perspective on plans and financing. Rarely is it part of the formal planning criteria or a specific cost-effectiveness test called for by regulatory authority or management. However, it does address one question that can be quite important to any investor – “When will I get my money back?”

Benefit/Cost Ratio Analysis

In many planning situations, one or all of the options being considered have both costs *and* benefits, and there are constraints or interrelated considerations with respect to both that impact the selection of the plan. In such cases, planning consists of more than simply ranking the alternatives by cost and selecting the one at the top of the list. Alternatives vary in benefit, too. This situation is particularly common in IRP situations, where options such as DSM have a cost, but also benefits accruing from their use.

Very often in such cases, the planning decisions are based on an evaluation of the expected gain versus the cost. Benefit-cost ratio is a simple concept: a ratio less than one indicates the project will lose more than it gains; the higher the benefit/cost ratio above 1.0, the greater the potential gain relative to its cost. For example, suppose that it has been determined that a particular project, with a PW cost of \$160,000, would provide PW savings equal to \$258,600. A planner could evaluate the benefit to cost ratio of this proposal as

$$\text{Benefit/Cost} = \$258,600 / \$160,000 = 1.62$$

This alternative has more benefit than cost – 62% more.

Benefit-cost ratio is generally applied as a *ranking method*, to prioritize projects that are viewed as discretionary or optional, or in cases where there simply isn’t enough funding for all the projects that present worth analysis has indicated are worthwhile. It has drawbacks, however, as explained below.

5.5 BUDGET-CONSTRAINED PLANNING: MARGINAL BENEFIT/COST ANALYSIS

Example Problem from an Aging Infrastructure System

Table 5.5 shows the different planning options available for two feeders in a utility, which has aging infrastructure problems. Feeder 1 does not currently exist,

it is being planned for a growing part of a utility service territory where new customers are developing, customers who must be served in order to meet the utility's obligation to serve (or a wireco's obligation to connection). This new distribution feeder can be built in any of several sizes and designs, as shown, using conductors and layouts of various types, all with advantages and disadvantages (including differences in cost). The details of exactly what each alternative is, and the differences in costs are unimportant to this discussion. What matters is that there *are* alternatives for the construction of this new feeder, and that they differ in both ability to handle future growth without further construction and service quality. Note that "do nothing" is listed as an option. This is not a feasible option for the utility, but for completeness sake it is given a PW cost of \$20,000,000 – the consequences, at least in part, of not serving these new customers.

Feeder number 2 is an existing feeder with an average equipment age of 45 years, part of an aging infrastructure area of the system. Slow load growth (.3% per capita increase annually) has exacerbated the usual problems that aging causes, leading to below-average service quality. The table shows the capital costs for various rebuilding programs, along with their expected PW costs based on project O&M savings and restoration costs, calculated by analysis of age, equipment condition, load curves, and failure-likelihood. Again, "do nothing" is an option in this case.

Least-Cost Selection of Alternatives (Traditional Approach)

Traditionally, electric utilities have applied a planning paradigm that includes several fundamental assumptions:

1. All new customers will be connected to the system.
2. All new designs will meet standards.
- 3) Long-term (e.g., present worth) costs will be minimized.
- 4) Capital funding for any project that meets the above requirements will be forthcoming.

This paradigm can be labeled an engineering-driven paradigm, because engineering requirements (1 and 2 above) drive the decision-making process (3) and define the financial requirement (4). This is illustrated in Figure 5.9.

In this traditional approach, the two alternatives in Table 5.5 would be evaluated for total present worth cost over a long period (usually 30 years) or alternatively on a lifetime basis, adding together present worth costs for both initial and continuing operating costs over the period. If load projections and planning criteria showed that the feeder had to be upgraded during the period, the present worth cost of the future capital cost for that upgrade would typically be included.

Table 5.5 Two Feeder Project Alternatives Evaluated on the Basis of Least-Cost Present Worth Value System at 11% Discount Rate - \$ × 1000

Feeder Number One – New Growing Area of System					
Feeder Option	Description of Approach	Meets Criteria?	Year 1 Capital	PW Cost	Comments
0	Do nothing	No	\$0	\$20,000	Do nothing is not really an option.
1	4/0 large trunk	No	\$2,250	\$7,550	Must be rebuilt for growth in year 2.
2	4/0-4/0 branched	Not	\$2,350	\$7,150	High losses. Part rebuild yrs. 7 & 12.
3	336/4/0 branched	Most	\$2,625	\$6,360	Must be rebuilt for growth in year 8.
4	336 large trunk	Barely	\$2,850	\$6,100	Reinforcement needed in year 14.
5	636 large trunk	Yes	\$3,075	\$5,500	Reinforcement needed in year 17.
6	795 large trunk	Yes	\$3,450	\$4,900	Part rebuild, yr. 24. Lowest PW cost.
7	636/336 branched	Yes	\$3,550	\$4,800	Needs rebuild yr. 38.
8	795/336 branched	Yes	\$3,675	\$4,750	Never needs rebuild. Low losses.
9	795-636 branched	Yes	\$3,800	\$4,800	Low losses don't justify cost.
10	1113 large trunk	Yes	\$3,900	\$5,100	Too much equip. Higher O&M.
11	1113/633 branched	Yes	\$4,250	\$4,950	Gold plated, not justifiable.
Feeder Number Two – Aging Infrastructure Area of System					
Feeder Option	Description of Approach	Meets Criteria?	Year 1 Capital	PW Cost	Comments
0	Do nothing	Yes	\$0	\$6,850	Tolerate failures and poor service.
1	Fix most serious	Yes	\$550	\$5,500	Repair only most critical items.
2	Rebuild Plan A	Yes	\$900	\$4,700	Rebuild trunk.
3	Rebuild plan B	Yes	\$1,400	\$3,500	Rebuild trunk and branches.
4	Rebuild plan C	Yes	\$1,850	\$3,600	Rebuild all but <20 year old.
5	Rebuild All	Yes	\$2,750	\$3,835	Rebuild all.




Figure 5.9 The traditional engineering-driven approach to prioritizing projects and determining utility budgets.

In each case, the alternative with the lowest PW cost would be selected. In the example shown in Table 5.5, this process would select option 8 for feeder 1 (the 795/336 branched feeder design) at a capital outlay of \$3,695,000 and option 3 for feeder 2 (rebuild plan B) at a cost of \$1,400,000. The utility would need a total of \$5,075,000 in capital for the two projects, probably borrowing the money for this work, as is the typical financial method at most investor owned utilities.

Facing Budget Constraints

Suppose the utility that needed to do both projects listed in Table 5.5 was forced to reduce its capital budget by 30% in order to achieve financial targets for both rate and stock return reasons. This reduction would drop the \$5,075,000 capital total for the two projects to an allowable \$3,550,000.

One possible adjustment within the traditional paradigm to accommodate budget constraints is to change the time-value of money used in the PW analysis. The discount rate can be raised to 20%, 30%, or even higher, with the result that it “cuts” items from the approval list of projects. Each increase in discount rate reduces the perceived value of future savings from a project, and biases the decision-making process in favor of alternatives with lower initial costs, such as “doing nothing.”

To achieve the desired reduction to below, \$3.55 million in this example, a discount rate 38% is necessary. This will force the selection process for the two example projects to \$3,550,000 by selecting only option 7 for feeder 1, and selecting nothing for feeder 2. This provides a total cost of \$3.55 million (just meeting the budget) and selects “do nothing” for the refurbishment project.

Using this approach, the capital budget required by the planning process is

reduced. However, there are four practical problems with this approach:

1. It does not provide a mechanism to force the budget below any specific limit. By trial and error, the utility can find a discount rate that “works out” to the right budget, but it is not a simple process. More important, it is not a dependable process: next year the required discount could be different. This approach does not lead to a simple rule that can be “aimed” at a specific budget target.
2. No ability to make really difficult cuts. As discussed in Chapter 6, PW analysis never says “no” to essential projects. In the face of severe budget cuts, the traditional paradigm still decrees that all new construction must be done, and work must be done to standards.
3. Using a high discount rate is inappropriate. High discount rates means future savings are unimportant, which is not true: The utility’s current financial situation precludes it from *affording* long-term savings, but this is not the same as saying it does not want them. (The importance of this distinction will become clear in several more pages.)
4. Customers in older areas of the system might legitimately complain that they are not getting the service and attention they deserve. New customers in outlying areas get shiny new service built to high standards. They get old equipment and lousy service.

Raising the discount rate to accommodate this budget reduction needs to change only the relative importance of the future savings against which present costs are evaluated: The planning method still assumes that the goal is to balance long- and short-term costs. Truly deep budget cuts need to be addressed by looking at how necessary the various projects are *today*, not at how today’s costs weigh against tomorrow’s savings, in any measure. (This will be discussed at length in a few paragraphs.)

Using Benefit/Cost Ratio to Make the Decisions

Ranking projects based on the benefit/cost ratio is often used to decide among project alternatives. This method can compare the present value of savings over the life of the project to today’s cost to implement. One can also evaluate projects based on their payback period. The two methods usually produce comparable results, selecting or rejecting the same sets of projects. With all proposed projects ranked in this manner, the utility can then determine how much capital it would need to finance all worthwhile projects. Alternatively, it can choose to find only those above any arbitrary benefit/cost limit.

Figure 5.10 plots the results of such analysis for the two feeder projects from Table 5.5, and others that are in the present-year budget. This particular utility

needs \$88 million to finance all worthwhile projects, as determined by traditional (PW least cost) planning paradigms.

The utility could achieve a budget reduction of 30%, limiting capital spending to only \$61.6 million, by approving only projects with a benefit-cost ratio of 4.4:1 or better.

Table 5.6 shows the impact of using benefit cost ratio as the arbiter in prioritizing the two example feeder projects. The benefit cost ratio selected for each option for both projects is shown in the rightmost column. For feeder 1, option 6 is selected (795 large trunk) at \$3,450,000. For feeder 2, nothing is selected; all the options fall below the cutoff. The utility would have to have a budget of at least \$77 million (only a 12.5% reduction) before any of these projects would be approved.




Figure 5.10 Cumulative capital requirements versus benefit/cost ratio for all “worthwhile” projects. The utility needs \$88 million. Using B/C as the prioritization test, a reduction of 30% (to \$61.6 million) means it can afford projects only up to 4.4:1, and only those with a B/C ratio above that limit should be approved. Benefit/cost ratio provides a way to limit spending to within a severe budget constraint, but it does not spend that budget most effectively (see text).

Table 5.6 The Two Feeder Projects From Table 5.5 Assessed By B/C Ratio**Feeder Number One – New Growing Area of System**

Feeder Option	Description of Approach	Year 1 Capital	PW Cost	Benefit (PW savings)	B/C Ratio
0	Do nothing	\$0	\$20,000		
1	4/0 large trunk	\$2,250	\$7,550	\$12,450	5.5
2	4/0-4/0 branched	\$2,350	\$7,150	\$12,850	5.5
3	336/4/0 branched	\$2,625	\$6,360	\$13,640	5.2
4	336 large trunk	\$2,750	\$6,100	\$13,900	5.1
5	636 large trunk	\$3,075	\$5,500	\$14,500	4.7
6	795 large trunk	\$3,450	\$4,900	\$15,100	4.4
7	636/336 branched	\$3,550	\$4,800	\$15,200	4.3
8	795/336 branched	\$3,675	\$4,750	\$15,250	4.1
9	795-636 branched	\$3,800	\$4,800	\$15,200	4.0
10	1113 large trunk	\$3,900	\$5,100	\$14,900	3.8
11	1113/633 branched	\$4,250	\$4,950	\$15,050	3.5

Feeder Number Two – Aging Infrastructure Area of System

Feeder Option	Description of Approach	Year 1 Capital	PW Cost	Benefit (PW savings)	B/C Ratio
1	Do nothing	0	\$6,850		
2	Fix most serious	550	\$5,500	\$1,350	2.5
3	Rebuild Plan A	900	\$4,700	\$2,150	2.4
4	Rebuild plan B	1400	\$3,500	\$3,350	2.4
5	Rebuild plan C	1850	\$3,600	\$3,250	1.8
6	Rebuild All	2750	\$4,435	\$2,415	0.9

Results Still Not Completely Satisfactory

The use of B/C ratio as the prioritization for projects resulted in only slight real improvement over the use of a higher discount rate. Of the four objections listed earlier for that approach, there still exist numbers 1 and 4:

1. B/C ratio analysis does not provide a mechanism to force the budget below any specific limit. By trial and error, the utility can find a discount rate that “works out” to the right budget, but it is not a simple process. More importantly, it is not a dependable process: next year the required discount could be different. This approach does not lead to a simple rule that can be “aimed” at a specific budget target.

(Objections 2 and 3 were solved by the use of B/C analysis)

2. Customers in older areas of the system might legitimately complain that they are not getting the service and attention they deserve. New customers in outlying areas get shiny new service built to high standards. They get old equipment and lousy service.

Number 4 is a particularly damning objection in the long run, sure to get the utility into trouble with its customer base and regulators alike.

Budget Constrained Planning: Optimization Using Marginal Benefit Cost Ratio

A business rather than engineering-driven planning process

Even though many utilities may still operate within a regulatory framework, they may face limits on the amount of money they can borrow or other financial pressures that preclude as much capital spending as they would like. This constraint is completely incompatible with the fourth precept listed for the traditional paradigm earlier, which assumed that sufficient capital funding would be available for all projects that justify themselves on the basis of present worth economics.

Since the constraints facing the utility are business-related (goals related to borrowing limitations, required rate reductions or caps, and targets for satisfactory stock dividend performance), it makes sense to adopt a “business driven” decision making basis for capital spending. Many utilities are forced, or chose to, shift to the business-driven structure shown in Figure 5.11, in order to plan to meet the pressures of a competitive marketplace and industry. In this paradigm, a survival and expansion plan, based on business necessities and goals, defines the available T&D budget for the utility. Good or bad, that’s all the budget there is, and unfortunately experience throughout the 1990s proved that in general

the required reduction was on the order of 25%–35% less than traditional levels.

What went wrong with benefit/cost ratio prioritization?

There is, however, a problem with using the benefit/cost ratio as previously described. Consider the preceding example. The first option for feeder 1 had such a high benefit/cost ratio that all options that built on it also had high benefit/cost ratio. Each of the options for feeder 1 builds on the high B/C ratio of its option. For example, option 2 costs an additional \$150,000 and lowers PW cost by \$400,000. Those additional dollars spent over option 1 carry a benefit cost ratio of only 4:1, yet the option is ranked as having an overall B/C ratio of 5.5 to one because its costs and benefits include those of option 1 also. By contrast, all the options for feeder 2 are much lower in B/C ratio: feeder 2's highest option is less than feeder 1's lowest option.

As a more dramatic example of the “wrong” this method applies, consider adding an option to Table 5.6. Option 1b would be a variation of option 1 that costs \$500,000 more in capital, but for which that additional cost contributes nothing in terms of improved PW benefit. At a cost of \$2,850 and benefit of \$12,450, this combination would still have a benefit/cost ratio of 4.5:1, so it would be *still be approved* under the budget-restricted guideline and over any spending on feeder 2 - a total waste of \$500,000!

Marginal Benefit Cost Analysis

Table 5.7 shows the options for the two feeders evaluated on a *marginal* benefit/cost ratio basis. In each project, the additional present worth savings delivered by moving up from one option to the next is evaluated against that additional expense.

Given a budget constraint of \$3,550, planners “buy” the most effective use of their dollars by taking the options whose marginal benefit cost ratios are shown in bold (final column) in Table 5.7. This is option 3 for feeder 1 and option 3 for feeder 2. A total of \$3,525,000 in expenses for the two projects gains a total savings (over do nothing) of \$15,790. No other combination of choices delivers so much “bang” for the buck.

Figure 5.12 shows the budget requirements when all projects the utility is considering are re-ranked on the basis of their marginal present worth/marginal capital ratio. There, the total for every marginal B/C ratio is determined, a plot analogous to that of Figure 5.10 in how it is used, but using marginal rather than just benefit cost ratio. As before, slightly more than \$88 million is required to meet all projects with a ratio above zero. As indicated, when limited to \$61.6




Figure 5.11 Many utilities have shifted to a business-driven structure, in which business needs define the available budget. The traditional planning paradigm has difficulty adjusting to this external change in cases where the reduction from traditional budget levels is significant. Although this paradigm stretches some traditionally rigid standards, it obtains both better reliability and better economy from the limited budget.

million, the utility must approve only projects with a MPW/C ratio of 1.5:1 or better. But, unlike the case where it was using benefit/cost ratio as the guiding rule, it will now always “buy” an equivalent (and biggest overall) PW savings - the most it can afford.

Application of Marginal Benefit-Cost Evaluation

Generally, the way one proceeds with the decision making on two or more projects using this approach is shown in Table 5.8. There, the options for the two feeder projects (from Table 5.7) have been merged and sorted into order based on marginal benefit cost ratio. Planners make their selections against a budget total by moving down the table, selecting options and adding their marginal cost (delta cost) into their total, stopping when they have the appropriate budget total (equal to the limit).

Given a budget limit to work within, (\$3,550,000 in this example), planning begins by selecting feeder 1 option 1 (\$2,250,000), the highest item on the table. It then proceeds to add feeder 1, option 2, for an additional \$100,000, as well as option 3 for that same feeder, at \$275,000, for a total so far of \$2,625,000 in cost. The next item on the list is feeder 2 option 1 with a marginal cost of \$550,000, bringing the total spent so far to \$3,175. Finally, feeder 2 option 2 can be added, at an additional \$350,000, bringing the total to \$3,525, the best, with the options listed, totaling under \$3,550.


Table 5.7 The Example Two Feeder Projects From Table 5.5 Assessed By Marginal B/C Ratio

Feeder Number One – New Growing Area of System						
Feeder Option	Description of Approach	Year 1 Capital	Delta Capital	PW Cost	Delta PW Savings	Marg. Ratio
0	Do nothing	\$0		\$20,000		
1	4/0 large trunk	\$2,250	\$2,250	\$7,550	\$12,450	5.5
2	4/0-4/0 branched	\$2,350	\$100	\$7,150	\$400	4.0
3	336/4/0 branched	\$2,625	\$275	\$6,360	\$790	2.9
4	336 large trunk	\$2,750	\$125	\$6,100	\$260	2.1
5	636 large trunk	\$3,075	\$325	\$5,500	\$600	1.8
6	795 large trunk	\$3,450	\$375	\$4,900	\$600	1.6
7	636/336 branched	\$3,550	\$100	\$4,800	\$100	1.0
8	795/336 branched	\$3,675	\$125	\$4,750	\$50	0.4
9	795-636 branched	\$3,800	\$125	\$4,800	(\$50)	-0.4
10	1113 large trunk	\$3,900	\$100	\$5,100	(\$300)	-3.0
11	1113/633 branched	\$4,250	\$350	\$4,950	\$150	0.4

Feeder Number Two – Aging Infrastructure Area of System						
Feeder Option	Description of Approach	Year 1 Capital	Delta Capital	PW Cost	Delta PW Savings	Marg. Ratio
1	Do nothing	0		\$6,850		
2	Fix most serious	550	\$550	\$5,500	\$1,350	2.4
3	Rebuild Plan A	900	\$350	\$4,700	\$800	2.2
4	Rebuild plan B	1400	\$850	\$3,500	\$1,200	1.4
5	Rebuild plan C	1850	\$450	\$3,000	\$500	1.1
6	Rebuild All	2750	\$900	\$3,835	(\$835)	-0.9

Table 5.8 Options for Both Feeder Projects From Table 5.7 Merged Into One List, Ranked By Marginal B/C Ratio

Feeder #	Feeder Option	Description of Approach	Year 1 Capital	Delta Capital	PW Cost	Delta PW Savings	Marg. Ratio
1	0	Do nothing	\$0		\$20,000		
2	1	Do nothing	0		\$6,850		
1	1	4/0 large trunk	\$2,250	\$2,250	\$7,550	\$12,450	5.5
1	2	4/0-4/0 branched	\$2,350	\$100	\$7,150	\$400	4.0
1	3	336/4/0 branched	\$2,625	\$275	\$6,360	\$790	2.9
2	2	Fix most serious	550	\$550	\$5,500	\$1,350	2.4
2	3	Rebuild Plan A	900	\$350	\$4,700	\$800	2.2
1	4	336 large trunk	\$2,750	\$125	\$6,100	\$260	2.1
1	5	636 large trunk	\$3,075	\$325	\$5,500	\$600	1.8
1	6	795 large trunk	\$3,450	\$375	\$4,900	\$600	1.6
2	4	Rebuild plan B	1400	\$850	\$3,500	\$1,200	1.4
1	7	636/336 branched	\$3,550	\$100	\$4,800	\$100	1.0
1	8	795/336 branched	\$3,675	\$125	\$4,750	\$50	0.4
1	9	795-636 branched	\$3,800	\$125	\$4,800	(\$50)	-0.4
1	10	1113 large trunk	\$3,900	\$100	\$5,100	(\$300)	-3.0
1	11	1113/633 branched	\$4,250	\$350	\$4,950	\$150	0
2	6	Rebuild All	2750	\$900	\$3,835	(\$835)	-0.9

**Figure 5.12** Cumulative capital requirements for the utility, versus the marginal PW/capital ratio required for that level of requirement. The utility's \$61.6 million budget allows it to spend up to a 2.1:1 marginal benefit cost ratio. This approach provides a much better guiding rule for project approval and assures maximum overall "bang for the buck" using the limited capital budget.

New basic paradigm for T&D project decision making

Generally, this approach is applied in decision making with a portion of the future costs being the performance-based rate penalties, customer paybacks, or other monetary risk measures of poor customer service quality that might be incurred by the utility. More on this concept will be given throughout this book, as prioritization based on marginal benefit-cost ratio is one of the cornerstones of the results-driven management approach and a key factor in the procedures that will be outlined to balance and optimize spending over the entire utility. This approach represents two major shifts in the thinking pattern, or paradigm applied to evaluation of projects.

1. *Zero-base planning.* No project, even a new customer extension, is absolutely necessary. All planning is done on a “zero base” basis, by beginning with “do nothing” as the base assumption.
2. *Service-quality rather than standards-driven planning.* If the cost penalty due to poor reliability is high enough to justify “doing something,” then money will be spent. Otherwise nothing is approved no matter how far below standards the resulting design will be (excepting safety standards, which must be met).

Note that the method satisfies all of the four problems listed earlier with application of the higher discount rate for planning. In particular, customers in older areas of town now have the same cost per improvement of service quality directed at them as consumers in newer areas of the system.

Refurbishment alternatives receive a “marginal” advantage

The net effect of the changes made in prioritization due to use of this method is complicated. First, although this paradigm does not assume that new service extension will be done, it’s impossible to argue with the service contribution that *doing something* delivers in all new construction cases (building *something* avoids nearly 8,760 hours of interruption time, more “bang” than one can get with most other projects). As a result, from any practical perspective, this planning method will always decide to build *something* to serve new customers.

But there are substantial differences in *what* this process will decide to build. Taking into account the value of service, this approach will no longer build the new service extension project to a very high level meant to meet thirty-year needs, just because that is what meets standards. Instead, it adjusts both new customer construction and aging infrastructure refurbishment decisions toward equivalent levels of payback on money spent for customer service. Money is spent on making *some* refurbishment additions to raise poor service levels, and the capital is obtained by cutting what is built in new areas so that service levels there may be less than if built to traditional “lowest PW” levels.

The results indicate the types of sacrifices that are often spread around the system. Restricted by budget constraints to projects with a 2.5:1 marginal ratio, the feeder 1 project is built with only a limited design not meeting traditional standards, for a capital reduction of \$1,000,000 over the unconstrained (lowest PW cost) situation. For the refurbishment project, its option 2 (plan A) is selected, a reduction from the recommendation that flowed out of unconstrained PW cost minimization, of \$500,000.

The reductions are spread equitably over both projects: using higher discount rates or B/C ratio made all the cusps in the aging infrastructure area. Note that with this planning method both projects get some of the limited funds, though neither gets as much as in the unconstrained budget (minimize present worth).

From the perspective of the traditional approach, the decision made for the new feeder project in this situation makes no sense. It is built with a capacity and style so that it is stressed at peak conditions from the very day it is built. However, from this new paradigm's standpoint, the service quality provided by the new feeder is no worse (in fact, it is slightly better) than the service quality delivered to customers in older areas of the system.

Interestingly, the ultimate result of this planning method is not a drop in overall reliability compared to continuing with the traditional paradigm. The new paradigm results in old and new areas being more balanced in capital spending and service quality results. Due to the non-linearities in reliability-versus-cost relationships, very often the utility's service quality indices, such as SAIFI, actually improve slightly with this value system.

This planning method is recommended in situations where a distribution utility has capital budget limitations that must be addressed and there is a concern that inequities in service quality may become a real or perceived problem, particularly with allocation of funds for refurbishment of aging infrastructure areas. Its application in a comprehensive manner is called *asset management*.

5.6 ASSET MANAGEMENT AND PARETO ANALYSIS

Asset management is a way of making capital budget decisions throughout an organization that aligns all capital spending with strategic business objectives - of maximizing business performance while managing risk, budgets, resources, and key performance indicators. A key element of this approach is a recognition that:

- a) The company has multiple objectives (things it wants to accomplish) such as reliability, working down average equipment age, profit margin, safety, etc.
- b) A candidate project might contribute to many or several of these objectives, improving reliability and safety and efficiency.
- c) There are often multiple ways a particular detail could be handled. A set

of troublesome MV breakers could be replaced with new, or refurbished, or merely serviced, or ignored. Each option has benefits (or not) and costs (or not), and might be most preferable depending on if and what other projects are funded, or not.

Pareto Analysis

Pareto Analysis is the basic analytical and decision-support engine of asset management. It is a strengthened version of the budget-constrained approaches discussed in the previous sections. Suppose planners at Big State Electric (BSE) identify 100 major capital spending projects the company *could do* in the next year, each with a defined scope (what it involves and what it does), an estimated cost, and estimated impacts on corporate KPIs (key performance indicators) like SAIDI, safety, system losses, power quality, working down the age of critical yet old equipment, future profit margins, etc.

Taken together, all 100 projects might cost \$500 million. However, BSE would never fund all of them, for two reasons. First, it doesn't have anywhere near \$500 million – it is thinking of spending somewhere maybe up to \$150 million and might be able to stretch to \$200 million if pressed to the wall, but no further.

Further, some combinations of these projects don't make sense if done together. BSE would never refurbish the vaults in its downtown network *and* replace them, but planners have scoped out potential projects to do both, in order to comprehensively look at the company's options: they'd like to see them replaced, but things are complicated, there are a lot of other important goals, too, and money is tight. Maybe refurbishment is all BSE can afford. Maybe it can't even afford that. Asset management will help the planners decide what to recommend.

Portfolios

A *portfolio* is a combination of projects and programs the utility decides to approve and execute. Traditionally, the utility called the set of projects it approved each year “the capital budget,” a term that stressed what it was spending, or “the approved construction list,” a term that stressed that it was going to build. Calling it a portfolio instead stresses that these projects are investments, i.e., that the company expects something in return and in fact wants to maximize that return.

Each circle in Figure 5.13 represents a different *potential portfolio*, some combination of projects and programs that the utility could decide to fund. The vertical axis is a measure of its contribution to success (how much the utility gets toward its KPI targets) while the horizontal axis is a measure of its cost (how much it spends) for that particular set of projects and programs. Each circle is plotted based on its total “bang” (what it provides in total towards the KPI goals the utility wants) versus “bucks” (what all of its projects cost when added together).




Figure 5.13 Potential portfolios (asset use and spending plans) can be plotted by their cost and performance results as shown here and explained in the text. Only the best portfolios lie on the efficient frontier, also called the Pareto curve, which denotes the maximum performance possible at every budget level. Optimization used in asset management finds only the portfolios on the efficient frontier and draws the curve for the utility planners.

Some of the potential portfolios in Figure 5.13 are *inefficient portfolios* (unshaded circles). BSE could buy any of these but never will, because it knows it can get more for its money. For example, why fund the combination of projects called portfolio B when one can fund A instead and get roughly 30% more “bang” for the same money? Or the company could fund portfolio C and get the same bang but for considerably less money. Portfolio B is inefficient, whereas A and C are efficient portfolios: if you want as much bang as C gives, you have to spend at least what C costs – no other portfolio can give you that much or more for less money. Similarly, if you want as much as B provides, it is the least costly way to get it.

The *efficient frontier* is the set of all efficient portfolios, and can be thought of as the upper boundary of this set of portfolios. The efficient frontier represents the *set* of efficient portfolio choices BSE has with regard to how much performance it wants to buy. This efficient frontier is also known as a Pareto curve, since it identifies a set of optimal solutions based on the same efficiency concepts used by the Italian economist Vilfredo Pareto when he examined issues related to social welfare (Pareto, 1906).

Efficient Choices

A key concept of asset management is *efficient choice*. The Pareto optimization determines the information shown in Figure 5.13 and identifies the frontier. This gives BSE a set of choices which are all optimum in the sense they maximizing what it gets for the money it spends. But the optimization does not chose or even

identify a best plan. The utility executives must do that, deciding how much of one thing they want (money) versus the other (success).

Finding the Pareto curve is straightforward in concept even if technically intricate in practice. The major steps planners have to take to do so are:

- 1) Define all the projects, their costs, and KPI impacts.
- 2) Convert the KPI impacts to a single valued measure (benefit).
- 3) Rate each as to benefit/cost ratio (bang for the buck).
- 4) Sort them by benefit/cost ratio, highest to lowest.
- 5) Plot them cumulatively in a figure like that shown in Figure 5.13.

Figure 5.14 shows an actual set of 300 projects, taken from an asset management study for a large investor-owned electric utility in the central United States in 2004. Each project is represented by a rectangle, of height equal to the benefit it provides and width equal to the project's cost. The project at the lower left is that project, among all, that has the highest benefit cost ratio.

The result is the Pareto curve for this list of potential projects. At each point on the curve, no combination of projects taken from that list can provide more bang for the buck than the combination already plotted: the curve describes the efficient frontier for the utility. The curve has a decreasing slope: initially spending is very efficient, providing a lot of benefit for the money spent, but as more is spent each successive project is, at best, only as efficient as the one before it- the result of sorting them and putting them in order of bang for the buck.




Figure 5.14 Pareto curve representing major capital projects for an electric utility. See text for details.

Realities of Asset Management

Figure 5.15 shows an actual Pareto curve from this same set of projects that includes consideration of several practical considerations that have to be included for the resulting plan to have any practical usefulness to the utility. These are:

Must do projects: a utility has to do some projects it must do. For example, it must do road widening projects to allow the state to build a new, wider highway, etc. It may have agreed to put certain existing lines underground when negotiating a new franchise agreement and rate schedule with a municipality in its service territory, etc. These projects are a given: they will be part of the portfolio.

Dependencies and other constraints. Some projects require another project be done first (Smart reclosers can't be installed in the Eastwood division until a central control and communication hub is installed in an expanded Distribution Operations Center). Other projects are mutually exclusive (The utility can rebuild a set of breakers or replace them, but would never do both). Planners want both in the list of potential projects to be considered because they provide different amounts of long- and short-term benefit and have different costs. They want the optimization to decide which is best for their company.)




Figure 5.15 The same Pareto curve after practical issues like must-do projects and various practical constraints are considered. Achieving the minimum acceptable level of performance is about 20% more expensive than in Figure 5.14: the real world is just more expensive than theory sometimes estimates. On the other hand, the marginal benefit (slope of the curve = how much one gets for each additional dollar spent above the minimum) at the minimum point is about 25% more than in Figure 5.14. In the real world, the utility has a bit more incentive to spend more than the minimum necessary.

Goals are digital in many practical senses. As a optimization begins buying the best project first, every minute of SAIDI improvement is equally valuable up to the point where enough has been purchased to meet regulatory and business targets for reliability. After that it may not be valuable at all, particularly if other targets like working down equipment age, or reducing O&M costs, or meeting safety goals remain to be achieved.

As a result, the most effective Pareto optimizations never try to combine all the KPI impacts into a single-valued measure. Instead, a multi-target constrained optimization engine is used that is told to find the lowest cost portfolio that:

- a) Meets all the minimum targets
- b) Includes all must do projects
- c) Includes a combination of projects that satisfies all constraints

Asset Management using Pareto Analysis driven by a MIP (mixed integer type) optimization is recommended for utility planning that includes both durability (lifetime and aging issues) and reliability (customer service) as well as other performance and business targets. For more information on utility asset management and optimization of methods see Willis and Brown (2010), which provides an accessible overview or Brown (2012), which gives a very detailed presentation on asset management and methods to implement it, in the References and Bibliography.

5.7 CONCLUSION

Planning involves the selection of the best alternative from among those at hand. In the power industry, that selection depends to a certain extent on the values of the group making the decision. *All* cost evaluations for power delivery planning are based in part upon economic evaluation that is consistent over the range of all alternatives. Usually, this means that present and future costs must be compared. Present worth analysis, levelized costing, or some similar method of putting present and future value on a comparable basis will be used.

Intra-Project, Marginal Benefit-Cost Alternatives Comparison

The foregoing method of maximizing benefit from a limited budget by comparing alternatives within each project (i.e., for each, several alternative ways of doing the project were evaluated) against alternatives from other projects, is called *intra-project* alternatives analysis. Methods that compare only a single alternative for each project (i.e., the method used in Table 5.5) are called *inter-project comparisons*. They compare each alternative based only on its ultimate “lowest cost” alternative. By contrast, *intra-project comparison* ranks the various

alternatives for each project against one another.

It is the combination of comparing multiple alternatives in a project, against those of others, and the use of marginal benefit-cost analysis, that yields the additional savings. This method – intra-project comparison based on marginal benefit cost analysis – is called *budget-constrained planning*, of BCP for short.

BCP is more complicated to implement than traditional prioritization methods. However, the authors have helped several utilities implement it, where it has led to bottom line savings and reliability improvements as dramatic as the improvement shown in the foregoing examples. BCP delivers results. It will be used both in capital planning, and in maintenance and operations prioritization, later in this book. When applied uniformly over an entire utility (all regions, all functions), it is the basis for Results-Driven Management.

REFERENCES AND BIBLIOGRAPHY

- D. Atanackovic, D. T. McGillis, and F. D. Galiana, “The Application of Multi-Criteria Analysis to a Substation Design,” paper presented at the 1997 IEEE Power Engineering Society Summer Meeting, Berlin.
- T.W. Berrie, *Electricity Economics and Planning*, Peter Peregrinus Ltd., London, 1992.
- R.E. Brown, *Business Essentials for Utility Engineers*, CRC Press, Boca Raton, 2010.
- J.J. Burke, *Power Distribution Engineering: Fundamentals and Applications*, Marcel Dekker, New York, 1994.
- V.D.F.Pareto, *Manuale di economia politica*, Centre d'Etudes Interdisciplinaires, Université de Lausanne, 1906.
- H.L. Willis, *Power Distribution Planning Reference Book- Second Edition*, Marcel Dekker, New York, 2004.
- H.L. Willis and R. W. Powell, “Load Forecasting for Transmission Planning,” *IEEE Transactions on Power Apparatus and Systems*, August 1985, p. 2550.
- H.L. Willis and G. B. Rackliffe, *Introduction to Integrated Resource T&D Planning*, ABB Guidebooks, Raleigh, NC, 1994.
- H.L. Willis and R. E. Brown, “Substation Asset Management,” Chapter 19 in *Electric Power Substation Engineering Handbook – 3rd Edition*, CRC Press Boca Raton, 2012.
- World Bank, *Guidelines for Marginal Cost Analysis of Power Systems*, Energy Department paper number 18, 1984.

6

Equipment Inspection, Testing, and Diagnostics

6.1 INTRODUCTION

This chapter provides a summary of the basic methods with which an electric distribution utility can better judge the condition of its equipment: inspection, testing, and diagnostics. It is not meant to provide comprehensive coverage of electrical equipment testing methods and how to apply them: that would fill more than one book. This chapter is intended to review the scope of the field and to make points and provide a perspective on inspection, testing and diagnostics as part of aging power delivery infrastructure management programs.

This chapters begins with an overview of inspection and testing and their goals, in Section 6.2. The various types of tests and diagnostics applied to equipment are covered in Section 6.3. Equipment tests include a variety of methods to evaluate the electrical and mechanical capabilities of the units and to diagnose what might be wrong internally with a unit that gives Orr test results. Section 6.4 covers tests of insulating oil. Oil is used in a wide variety of power system equipment. Oil tests focus both on evaluating the condition of the oil itself, and on using oil and other test results to diagnose the internal condition of the equipment. Section 6.5 concludes with a summary of key points.

6.2 INSPECTION, TESTING, AND DIAGNOSTIC EVALUATION

To be effective, inspection, diagnostics, and testing must be coordinated and applied in a well-designed program that also uses the results from system operation (trouble reports) to identify overall equipment condition and

performance, and that prioritizes preventive maintenance, repair, rebuilding, and replacement. Optimum results are derived by then loading equipment to levels determined based on its assessed condition, the loss of lifetime caused by the usage, and the available economic alternatives. Procedures for organizing and managing such programs will be discussed in Chapter 10. This chapter focuses on condition assessment inspection, diagnostics, and testing methods themselves.

Inspection, testing and diagnostics are the means through which the condition of equipment is determined. These three types of condition verification all have different meanings.

Inspection involves evaluation of conditions based on visual, aural or equipment-supported inspection of a device or structure. Basically, one looks, and at times listens to or scans, the device to determine if it looks or sounds as if it is in good condition. This inspection may be augmented by equipment to enhance the detail or better quantify the results (infrared camera, microphone, and harmonic analyzer system). Good examples of this are the use of infrared scanners to identify equipment that is running hotter than normal, and devices that check the thickness of paint or galvanization of metal through scanning of the material.

Olfactory inspection can also be an aid in a few cases (“Do you smell something burning? Does this oil smell sooty to you?”) and is a big area for future development of inspection and test equipment. Sniffer machines, which can detect trace amounts of certain chemicals in the air or in a fluid, are being developed to augment inspection and testing of equipment. (At the present time the most frequently seen examples of these devices by the general population are the explosives sniffers used in security checks at airports.)

Another example of equipment-enhanced sensing for inspection is a practice at numerous utilities of leaving the radio of company cars tuned to the AM band while driving line routes for visual inspection of equipment. Cracked insulators, corroded brackets, and broken conductor strands create small amounts of radio interference, detectable as a sudden increase in “static” background noise by the radio when near them.

Testing involves checking the device and its associated support equipment (its system) to make certain it operates as it should, by checking to see that it can indeed function to specification, or for the presence of evidence that something is wrong. Thus, a breaker test involves making certain that it trips in the time and at the speed it should, and that it re-charges or re-sets itself as expected, and that its relay performs as “programmed.” A test of a line switch involves cycling it once to assure that it will open and close as expected. A transformer may be tested at full load, or merely monitored in detail during peak season and its data results (maximum temperature, rate of temperature rise, losses) compared to expected values. Any deviation

would trigger diagnostics to determine what was wrong.

Diagnostics involve “tests” that measure factors other than those directly related to a device’s function: tests whose results provide clues to the device’s internal condition, the performance of subassemblies inside it, or the possibility of impending failure. Usually, diagnostics probe the condition of materials (e.g., insulation) and are tracked over time to determine the rate of deterioration or contamination. The amount and type of gas, and the rate of its increase over time, identify whether there are internal faults, arcing, or overheating inside the unit – problems with something other than with the oil itself. This makes them a diagnostic test, not strictly speaking, an equipment test.

“Test” can have two meanings

The word “test” has a somewhat confusing dual meaning with respect to electrical equipment, a fact compounded by the inexact use of the word on a day-to-day basis by many power systems professionals. There is a fine line between testing and diagnostics and two experienced test engineers may not agree entirely on whether a procedure is purely a test or a diagnostic test. This is not important as long as the data collected are used fully, correctly, and well.

Strictly speaking “testing” of a device always refers to actions or steps performed to measure or verify its performance or characteristics against its functional specifications. Actions that check the unit to see that it has the correct turns ratio, that its losses, leakage, and thermal buildup are within spec, that it meets trip or operation timing, and that it meets any other specifications for its performance (noise) are tests of the unit itself.

However, many diagnostic procedures involve a “test” of a particular component or factor in the machine or equipment. For example, dissolved gas tests of transformer oil provide an indication of the condition of the oil itself and whether it can do its job, making them a “test” of the oil. However, dissolved gas tests of insulating oil done are done primarily because they are *diagnostic* – they provide powerful clues to the condition of a transformer’s winding insulation and the causes of problems therein: a clue to the condition of something other than what is being directly tested.

Thus a test of a component or part (transformer oil in this case) can take two forms:

1. *Testing of the oil* to determine if it performs its function. For example, dielectric breakdown test of transformer oil directly measures one aspect of the performance the oil is expected to provide (insulating ability).
2. *Diagnostic tests*. By contrast, the oil could be tested for acidity or by dissolved gas analysis (DGA). These tests respectively determine the

content of acids in the oil (formed by oxidation) and the amount of gas dissolved in the oil. The former gives an indication of both the condition of the oil and the likely amount of sludge buildup in the transformer. The second test is widely regarded as the best possible early indicator of impending problems with the transformer itself.

Three Purposes for Inspection, Testing, and Diagnosis

Inspection, diagnostics, and testing fall into three different categories, with three different purposes:

1. *Acceptance, or capability testing* – new equipment is often tested to verify that it both is undamaged and ready for service, and that it meets specifications as to capability, losses, impedance, and a host of other factors important in its performance. Some of these tests may be the same as performed later during routine inspections or condition evaluation, but many are unique, one-time evaluations, such as a test for no load and load-related losses performed on new transformers, to verify what was built to specification.
2. *Routine inspection* – typically equipment in service is inspected, and perhaps some diagnostics or tests performed, on a periodic basis – annually, every three or five years, etc. This is done for what might be called “general principles” and because warranties may require it. The overall goal is to identify deterioration or damage before it becomes too serious to perform preventive maintenance, and to identify equipment that is about to fail in time to remove it from service in an orderly fashion.
3. *Condition evaluation* – on occasion, more comprehensive tests than called for on a routine basis are performed on some or all equipment to establish more precisely the condition of these units. The purpose of these tests is to identify as much as possible the condition of the equipment and any problems that might preclude its equipment being left in service with a high degree of confidence that it will perform its function dependably.

Often the line between routine inspection and more thorough condition evaluation is blurred. For example, routine inspection policies at some utilities call for mandatory follow-up testing of any equipment found to have possible problems. Their condition evaluation is considered a second-stage in routine testing; done only if it appears needed. Other utilities include as part of routine testing, through condition evaluation of randomly selected units, the goal being to test enough units to identify any trends of deterioration or deficiency endemic to all units in a certain class or category. Policy on inspection and condition evaluation is appraised and set, using methods that will be discussed later.

Field or Laboratory Evaluation

Inspection, diagnostic tests, and testing can be carried out in the field (where the equipment is installed) or in a laboratory or test facility. The advantage of field tests is that they can be done on-site, usually taking less time and involving less skill and expense than laboratory tests.¹ However, generally the laboratory tests are more accurate and thorough.

For example, dielectric testing of transformer oil can be done either as a field test using the ASTM (American Society of Testing and Materials) D-877 disc electrode method, or as a laboratory test using the ASTM D1816 “VDE cell” method. The former involves testing samples of about one pint of fluid for dielectric breakdown strength, by placing the fluid in a small container in which two small disc electrodes are spaced .1 inch apart. Higher and higher voltages are applied until the oil breaks down. The entire test takes but half an hour and uses a portable device the size of a small suitcase that is taken to each transformer to be tested. By contrast, the D1816 “VDE cell” dielectric breakdown test of oil is done in a laboratory. It involves testing a larger sample of fluid drawn from each transformer and taken to the laboratory, in a “VDE cell” within which the oil is circulated, with test electrodes that more closely mimic the field strength and uniformity of voltage seen inside a transformer. This test is regarded as more accurate in depicting electric strength and scatter than the disc electrode test, but it takes more time for the test itself, and for collection of materials and running back and forth from the field to the laboratory.

Non-Destructive, Destructive or Forensic Tests

Testing methods can also be categorized by the condition of the device or equipment being tested before and after the evaluation, as shown in Figure 6.1.

Destructive tests ruin or damage the device or equipment. Afterwards, the entire device, or a portion of what was tested, is damaged and not suitable for service and may in fact have been completely destroyed. Examples are (for cables in this case) voltage breakdown testing of a section of cable carried out to the point of failure, to determine its dielectric strength and mode of failure, and inspection of samples of cable for treeing. The former destroys the cable section being tested. The latter removes part of the cable (which must be repaired).

Destructive tests are generally carried out by manufacturers on new equipment to verify strength and identify failure modes. Utilities use the tests on select samples of equipment that have been in service, in order to

¹ At least if the cost of equipment is not included. Often field tests require specialized diagnostic equipment whose cost can be quite high, whereas laboratory tests for the same or similar factors would involve only standard laboratory equipment.

		Condition Before Testing	
		Good	Ruined
Condition After Testing	Good	Non-Destructive Tests	<i>Don't exist (repair)</i>
	Ruined	Destructive Tests	Forensic Tests

Figure 6.1 Diagnostic and testing methods can be categorized by the condition of the equipment before and after the tests.

Table 6.1 Standards Bodies with Guidelines or Recommended Testing Methods for Electrical Power Equipment

Acronym	Meaning	Country
AFNOR	Association Francaise de Normalisation	France
ANSI	American National Standard Institute	Mainly USA
ASTM	American Society of Testing and Materials	Mainly USA
BSI	British Standard Institute	United Kingdom
DIN	Deutsche Industrial Normal	Germany
IEC	International Electrotechnical Commission	Mainly Europe
IEEE	Institute of Electrical and Electronics Engineers	Mainly USA
JIN	Japanese Industrial Standard	Japan
VDE	Verband Deutscher Electrotechniker	Germany

determine the rate of deterioration of various components. These tests are often applied to cables, switches, and service transformers. Usually, equipment that is destructively tested is damaged to the point of having no further value. Regardless, this equipment should never be returned to service after the tests, unless it is part of an experiment on durability, etc.

Non-destructive tests are intended to cause no damage.. However, while such tests rarely fail a unit at the time of testing, some experts believe the added stress of some tests increases the subsequent failure rate. One example is high-energy pulse (thumper) testing for underground cable fault location. In this cable fault test, a failed cable section is pulsed from one end with a series of high frequency, high-voltage pulses. Field personnel use portable aboveground sensors to detect the field generated by these pulses along the cable route. This field will be generated only by the cable up to the location of the fault, not beyond it, making it (relatively) easy to find the fault location. The pulses from these tests are suspected of causing high stress on cable insulation, contributing to its premature deterioration.²

Regardless, a non-destructive test or diagnostic is intended to make no impact on the equipment's subsequent performance. Most non-destructive tests certainly fall completely into this category. Only a few are known or suspected of causing large amounts of equipment stress.

Forensic tests may be carried out on failed equipment after an event in which "something went wrong," in order to determine the cause of the failure. Often, the equipment was destroyed in the event, and the testing involves a type of detective work to determine how and why the failure occurred. Much of forensic evaluation involves trying to identify the exact sequence of events of the failure by examination of the failed device (what is left of it) – where did the problem begin and how did it spread?

Comprehensive forensic tests are often done after a failure that caused damage to public property or life, often in preparation for a lawsuit. However, forensic tests can be important to a utility trying to manage aging infrastructure because they can establish if there is a pattern in the failures of equipment. For example, if forensic evaluation of all failed transformers of a certain type shows that they fail for what appears to be a voltage breakdown in a certain part of the unit, then diagnostic testing can focus on that in the future.

² Deterioration of insulation in cables is very roughly inversely proportional to frequency. Thus the same cable should last about 16% longer in service in a European system (50 Hz) than in an American system (60 Hz), all other things being equal. Pulses used in some cable fault testing have a high voltage and a very high frequency content – the pulse edges are of course very high frequency but the pulse repetition rate is often quite high as well. In combination with a pulse voltage raised in order to make the field more detectable, this often results in calculated aging rates of up to 4,000 times normal service.

Standards and Recommendations for Inspection, Diagnostics, and Testing

There are a number of standards bodies that have defined recommendations (standards) for inspections, diagnostic tests, and testing of electrical equipment. Not all of these standards groups have an applicable standard in all categories of tests or equipment. Not surprisingly, where two or more organizations have overlapping standards, they may differ in the details of their recommended testing methods. Table 6.1 lists some of these organizations.

General Categories of Inspection, Diagnostics, and Testing

Voltage withstand and dielectric strength

Voltage capability tests including dielectric strength and breakdown tests are performed to determine if the voltage withstand capability of the equipment is up to requirements. Slight degradations in insulation are detectable as leakage or by other indications given by some types of diagnostics (e.g., partial discharge classification based on field uniformity, power factor tests).

Thermal testing

Thermal tests determine if the device overheats during normal service conditions. Two processes can lead to overheating. First, the unit may be generating more heat inside it than is normal – as from internal faults. Secondly, its ability to cool itself might be degraded, as from sludge covering internal components or blocking cooling ducts. Heat being the major cause of early deterioration of just about everything involved in electrical equipment, such problems can lead to short equipment lifetimes if left uncorrected.

Thermal testing of a unit such as a transformer can be done by subjecting the unit to special full-load tests in a laboratory – subjecting a “cold unit” to a step increase to full load and measuring its temperature over a 24-hour period thereafter. However, often such tests are unfeasible for equipment in the field. The same information can be determined by analysis of load-cycle and temperature monitoring for a period of several weeks, and a subsequent analysis of the results, as will be discussed in Section 9.3.

Mechanical tests

Tests are often done on devices with moving parts, to determine that those parts do move with the speed, smoothness, and other properties required. Thus, line switches are checked to make certain they open and close (corrosion might prevent operation), breakers are checked for speed of operation and re-set, tap changers for proper movement, etc.

In addition, poles and structures may be tested to determine their mechanical strength. Both poles and cables can be tested using a type of penetrating probe that is pushed into the wood or insulating surface, the resistance to the penetration of

the probe giving an indication of the strength of the wood or insulating material.

Physical condition determination

This is often entirely from visual inspection. Physical condition determination aims to identify signs of past, present, or forthcoming problems. It seeks to identify any noticeable corrosion, deterioration or damage (evidence that a car hit a pole, or an insulator has been shot), or outward signs of internal malfunction (leaking oil around seals). Physical inspection of the device usually includes visual or infrared examination of the exterior of the unit while in service, but can include detailed examination of its internal parts (as for example when disassembling a circuit breaker). It may also include evaluation of noise and even odor (“Smell something burning?”) using the senses, or special equipment to augment them (infrared scanners, acoustic signal analyzers, so called “sniffers” made to detect trace gases in the atmosphere).

Materials evaluation

Most diagnostic tests evaluate the chemical or mechanical properties of sub-components in an electrical device. These tests determine if the component material (e.g., XLPE cable insulation) itself has deteriorated or is defective in some way (treeing, internal voids). Performing diagnostic tests can also provide data that indicate there may be problems with other parts of the device (dissolved gas analysis of transformer oil).

Timing, Synchronization, and Automation Sequences Tests

Many types of power equipment have moving parts that are designed to operate in very specific ways, one of which is to operate within a specific time or at a specific speed. Timing tests verify that they operate as designed with respect to both the actual timing as well as other timing-related aspects of that operation.

Traditionally, timing tests were considered “just another test,” and in earlier editions of this book were grouped in the section above as one of a category of tests performed on equipment. However, the widespread use of automation, remote control and autonomous equipment in “Smart Grid” system has created a new meaning and requirements for tests of timing. In the 21st century, “timing” can mean not only that the equipment itself operates at the speed required, but that that operation is synchronized with or coordinated to the timing of other equipments action. This might mean that two devices operate not only each within specification, but in the intended sequences, device A before device B, with a certain minimum or maximum time in between. It might mean that a sequence of events occurs in the right order: A opens and then B opens and then C closes. Finally, it might mean that the timing, order, and sequence of events is accurate, as is required to verify the timing of automated systems: if a fault occurs and [this condition is met] then A opens and then B opens and if [this condition is seen] C closes, etc., all within certain specified times, perhaps.

Such “testing” for timing is quite similar to relay testing of protection schemes, and quite far from the types of tests most useful to managers and operators interested in aging equipment and systems. The authors include a summary here, for completeness.

Electromechanical action tests

A particular circuit breaker might have been designed so that it operates within a certain amount of time after receiving a signal to open. In order to perform the function for which it was designed and has been installed in the power system, it must do so. Definitions of this operation will be very precisely defined: it starts within so many milliseconds or cycles, it completes its operation and is fully open within a certain time. Testing methods, often called “Doble testing” after the company that pioneered the methods, will verify its performance in detail.

Breaker testing may include electrical tests of the timing as well as mechanical tests. Since the purpose of the breaker and the required timing is to successfully interrupt an electric current, the test may involve doing so – perhaps not during a high-current fault but in a manner that leads to diagnostic information on the electrical performance relevant to the test (the contacts and all electric-flow related aspects are also within specifications).

Synchronization tests

Often the action of a particular device is required to coordinate precisely with the timing of other devices, or with electric flow. An example is the testing of switches intended to open only at the zero-crossing of AC current flow. Testing of a three-phase zero-crossing switch will verify that it:

- a) operates fast enough overall, opening its contacts within a certain period after being told to open,
- b) opens each of its phase switches at the zero crossing of the current for that phase,
- c) opens them in immediate succession, i.e., it does not open phase A at the first zero crossing but leave phase B closed for two or three cycles before opening it at a zero crossing.

Automation Programming Testing

Smart equipment such as the switching on a set of distribution circuits may have been set up in an intricate scheme verification of which requires that:

- a) Breakers and switches operate correctly (within specifications as to timing, synchronization with zero crossings, etc.), equipment with staged operation (voltage regulators) does as it should.
- b) Sensors and monitoring equipment are set correctly to detect conditions relevant to their operator.

c) The programming of operation is correct in all cases.

Testing to verify that all this is correct is generally done in two stages. Individual equipment is tested or verified to operate within specifications for it. The programming and joint operation of the devices is tested via real-time simulation, using a Real Time Digital Simulation Test facility to replicate the conditions (signals and readings from monitors) as the equipment would see it, and then creating the conditions under which the equipment is supposed to operate (or not) and verifying it does.

6.3 EQUIPMENT TESTING AND DIAGNOSIS METHODS

This section reviews the various types of electrical test procedures used on cables, conductor, and electrical equipment. Most tests and diagnostics for electrical equipment apply a voltage of some type (DC or AC) to the equipment in order to measure current, resistance, or the rate of change in one or both over time. The methods of detection or measurement range from simple and direct (apply the voltage, measure the current, compute the resistance) to clever and complicated (use an AC voltage to excite a partial discharge at places where there are weaknesses, detect the location by timing of the return pulses).

Electrical testing methods fall into three categories: tests that use AC voltage; those that use DC voltage; and partial discharge (PD) tests, which use an AC voltage of some type but in a different manner than in “AC” tests. Generally, these testing methods require de-energization of the device, although most can be done while the equipment is in place (i.e., they are not solely laboratory tests but can be done in the field).

DC or AC?

Tests of an electrical device using an applied voltage are done for a variety of reasons:

- To determine resistance of insulation.
- To determine if internal faults exist, to identify overall condition.
- To infer amount and type of contaminants, and in some cases, to simply stress the device in the hopes that if there is a flaw, failure will occur now, not months from now while it is service.

Some tests use DC voltages, others use AC voltages of standard or non-standard frequency. Usually test current is kept low through various means of arranging and setting the impedances. Some tests (high potential voltage) can be done using either DC or AC, with slightly different diagnostics results from each.

AC and DC voltages have different effects on a device, a fact that determines which approach is suitable for some tests, and which alters the exact effect of a high potential test, which can use either approach. AC voltage tests apply a

voltage similar to that seen by the device when in service. If done at normal system frequency (e.g., 60 Hz) these tests see impedances and produce stresses similar to those that the equipment is expected to endure. AC voltage tests are therefore a good proxy for the stresses and responses expected during service and can be used to verify that a machine or cable’s condition matches that needed for service.

DC tests create no alternating field and “see” no effects of reactance. Thus they measure different aspects of a device’s behavior than do AC tests. DC tests can be used to measure capacitance of a device, or to infer AC impedance or the “complex” condition of insulation, through measurements of time constants or rates of change in field-related factors during certain tests. One advantage is that they put less stress on insulation in some situations: voltage stress is frequency related.

It is useful to consider electrical testing in the frequency spectrum: DC testing is merely testing with voltage whose frequency is zero; various AC and PD tests use a higher frequency, one that may be lower, the same, or higher than system frequency, depending on the specific test procedure. Electrical equipment varies in its impedance as a function of frequency. So do the characteristics of flaws created in manufacturing or by the processes of deterioration and damage that the diagnostic testing may wish to identify.

Therefore, AC tests using a frequency other than standard system frequency (e.g., 7,200 Hz) excite different modes of response from the system and can be used to “tune” a test to examine only certain aspects of a device’s insulation strength or operating characteristics. For example, Table 6.2 shows the difference in stress that AC or DC hi-potential testing of the windings of a transformer puts on a 15 kV transformer. By selecting either DC or AC frequency, a diagnostic test can be tailored to focus on a particular part of a device of greatest interest, or to avoid putting stress on a particular portion suspected of weakness.

Table 6.2 Division of Stress between Paper-Cellulose and Oil Insulation in a Transformer Depending on Type of Test

Insulation Type	AC Test Results	DC Test Results
Paper-cellulose	25%	75%
Oil	75%	25%

Source: Gill, Electrical Power Equipment Maintenance and Testing, 2008

Comparing DC and AC testing, qualitative differences are:

1. DC testing is generally considered slightly less stressful on transformer and cable insulation than AC testing. Therefore, DC

testing is less likely to do harm to a device, but AC testing is more likely to reveal flaws that might otherwise be revealed while in service.

2. Generally DC test equipment is lighter, more portable, easier and faster to use.
3. DC testing avoids problems with high charging capacitance, as for example, when testing cables.
4. Generally, DC test results are more sensitive to the temperature of the device than AC tests. They often require adjustment of their raw test results with extensive correction formulas or tables of result-versus-temperature. These correction factors require accurate knowledge of the equipment temperature, which might not always be available (as when testing cables where one cannot measure the duct temperature along the entire length of the cable).
5. AC testing provides a more realistic test of the dispersion of voltage in the device than DC testing and tends to stress insulation more in proportion to stresses expected in service. For example, in transformer tests DC testing tends to stress paper insulation much more, and oil much less than in actual service.

Resistance Tests

Several types of resistance tests are performed on transformers. These are all variations on testing the resistance of the winding insulation. Resistance tests are essentially DC voltage tests applied with a megohmmeter.

1. *Insulation Resistance Tests* measure the resistance through the transformer insulation. They are carried out with a hand-cranked megohmmeter or modern digital test unit that applies a DC voltage of higher-than-normal level for a brief period, and has a relatively low potential to cause damage or high stress. Lower than expected values of measured resistance reveal flaws in the insulation level.
2. *Dielectric Absorption Test* is a version of the insulation resistance test in which the test is carried out over a period of up to 10 minutes, with the resistance measurement plotted over that time. Good insulation will reveal a slightly rising insulation level over time in a resistance test, poor insulation will not. Plotted as resistance versus time on log-log paper, a good insulation system will give a straight line, a system contaminated, or with deteriorated insulation, will have a curved profile.

3. *Dielectric Absorption Ratio (DAR)* test is an older, field version of the dielectric absorption test, which does not plot resistance with time but only evaluates the ratio of resistance measured at 60 seconds to that obtained at 30 seconds. A value below 1.25 is considered unacceptable.

Power Factor Tests

Power factor tests evaluate the insulation strength of a material by measuring both the capacitive and dielectric components of the test current from an application of AC voltage. It then evaluates the condition of the component based on the magnitude of both, and the measured power factor of the test. As in power system analysis, power factor in the test refers to the ratio of the complex product of voltage and current to their scalar product. The voltage and current in this case are the applied voltage and the resulting current flow produced by an AC test applied to a cable, across the terminals of a device, or to a sample of oil. One advantage of measuring PF is that it is independent of sample size (e.g., the cable length). Therefore, it is accurate when one is uncertain of the size of the sample as compared to the size of the device being tested. For example, when testing cable, power factor is invariant as a function of cable length. This results from the fact that both the resistive and capacitive components increase with the length of the cable.

When applied to a cable or oil or other insulating material the power factor test assesses dielectric condition by determining if there is a lag between applied voltage in a test and the current it produces. Simplifying the concept somewhat by neglecting a few details of secondary importance, if dielectric strength is very high, the only current flow produced by the applied voltage will be a capacitive flow, completely out of phase with the voltage, producing a power factor of 0.0. By contrast, if the dielectric strength is weak, current flow through the oil will be resistive – a current will flow in phase with the applied voltage, and power factor will be far greater than 0.

In actuality, a power factor test can provide a bit more information about condition. By varying test voltage, the type of contamination affecting dielectric strength can be identified. For example, dielectric losses due to water or carbon tracking tend to vary as the square of test voltage, whereas those due to corona effects increase exponentially. Therefore, varying voltage can develop a profile that indicates whether moisture or carbon tracking damage has occurred.

PF tests are prescribed as both routine tests, and as among the first of a series of comprehensive tests done on units suspected of having problems, because in most cases the PF test causes no stress on the insulation. Thus, it is suitable for use on equipment suspected of having weaknesses that partial discharge or high-potential tests might exacerbate. Yet the PF test is reasonably dependable for indicating the level of contamination or deterioration in insulation. However, a power factor test does not identify what the *cause* of the degraded capability is,

and further tests may be needed to determine what it is and to interpret that, to find out if the oil can be reconditioned or if it should be replaced.

An applicable standard for testing transformers is ASTM D-924, which specifies a type of “cell” used for PF tests of the insulating oil, and gives the results expected. Good oil (e.g., new) has a test power factor of .05. Contamination by the products of deterioration (moisture, dissolved gases, carbon, etc.) will increase the power factor observed. Values above .5% and up to 2% indicate further investigation is in order. Values above 2% indicate a need to do further tests of make repairs, depending on the type of device. Table 6.3 lists typical PF values considered to be the limits for “good” insulation for various types of power equipment insulating materials.

Table 6.3 Limits for “Good” Classification of Various Insulation and Equipment Types Based on PF Test

Insulation or Equipment	PF - %
Oil	< .5
Paper	< .5
Cross-link polyethylene	.05 – 1.0
Ethylene/propylene rubber	.5 – 1.0
Rubber	3 – 5
Varnished cambric	4 – 8
Air circuit breaker	< 1

Polarization Recovery Tests

Deterioration of cellulose insulation and insulating oil inside a transformer can reduce the time constant they exhibit for establishing a space-charge polarization across their insulation or through a sample of oil. If a DC voltage is briefly applied across insulation to build up a charge across it, and then removed, the rate at which the impressed voltage grows and/or dissipates can be tracked. The time constant of the rate of dissipation gives an indication of the quality of the insulation. It will be much shorter if the insulation has more leakage.

Thus, a test of how quickly an insulator or insulated device reacts to an applied DC (capacitive) charge applied to it (i.e., its polarization time constant), is a measure of its deteriorated or contaminated condition. Note that this test is broadly similar to the power factor test in one way: it keys on capacitive/resistive properties of the insulation to identify its strength.




Figure 6.2 Polarization test determines a time constant of the insulation, which correlates with condition. This plot shows time constants derived from laboratory tests of oil-impregnated paper with known quantities of moisture.

Polarization recovery tests of a transformer’s winding insulation are performed on a de-energized unit. The low side windings are joined and test cables are connected to the high side terminals. A known DC voltage (e.g., 2,000 volts) is applied for a period of time that is varied (minimum and maximums are typically 1/50 to 5 seconds). By varying the time of applied voltage, the time constant associated with the polarization as it develops across the insulation can be determined (Gill, 2008). Figure 6.2 shows the results of tests on transformer-type oil-impregnated paper insulation in tests of insulation with known quantities of moisture. Testing for cables of other power equipment is similar in nature.

Hi-Potential (Hi-Pot) Tests

Transformer winding insulation strength, both turn-to-turn and phase-to-phase-to-ground, can be evaluated for leakage and possible internal faults using an hi-potential test, in which a higher than normal voltage is applied briefly to the device in order to “stress test” its insulation strength. The test can be carried out either with DC or AC voltage.

Generally, DC rather than AC hi-potential tests are used when testing cable, as DC avoids any complications due to charging capacitance. Either, or both, AC and DC may be used on transformers, the major difference in results being that the DC test tends to stress paper-cellulose insulation more and the oil less, while AC tests do just the opposite (see Table 6.2). As a rough rule of thumb, DC voltages

are preferred to test lower voltage transformers, while AC voltages are used on units rated above 25 – 35 kV because DC testing might over-stress their paper insulation. Thus, the AC hi-potential test is used on most power transformers, DC on most service transformers.

In an AC test, an AC voltage at standard operating frequency but somewhat higher than the standard operating voltage is applied briefly to the transformer terminals, putting the insulation under considerable stress to test it for flaws. Application of this test varies depending on manufacturer's recommendations and practices, the standard being followed and individual preferences that differ from utility to utility. If done as a check after maintenance, an AC voltage of about 33% (but at least 5 kV) above nominal operating voltage is applied for a period of three minutes. Slightly higher voltages might be used on new units or after a complete rebuild (acceptance tests). Higher than expected currents during this test indicate a weakness in the insulation. DC test procedures are similar in nature.

Hi-potential tests, particularly those involving AC, do put a good deal of stress on a device and fall into that category of "non-destructive" which can push a nearly failed device over the edge, discussed in Section 6.2. The very high test voltages put considerable strain on the unit and will accelerate deterioration of any incipient flaws.

Thus, high-potential voltage tests can be viewed either as an "if it can handle this it should do well in service" or as a "fail it now, not later" test. Either perspective is legitimate. For this reason, some maintenance experts (Gill, 2008) recommend against hi-potential tests as routine procedure. Instead, recommendations are usually to restrict their use to new or rebuilt equipment undergoing acceptance testing. Regardless, common sense dictates that they be applied after less stressful forms of testing a cable, device, or transformer core have been successfully passed, such as resistance and PF tests and the TTR and winding resistance tests in transformers. Those might reveal problems that hi-potential testing would worsen.

Induced Potential Test

The induced potential test can be viewed as a variation on the hi-potential test, in that it too, is a quality evaluation stress test of winding insulation. Carried out much like the hi-potential test, it is done with higher than normal operating frequencies. As with the hi-potential tests, there are slightly different forms of this test, carried out with different voltage levels and frequency. One common form involves applying twice the nominal operating voltage at 7,200 Hz for up to one minute. Others vary the voltage, frequency, and time, with lower voltages and frequencies being applied (e.g. 65% and 120 Hz).

This test puts units under considerable stress due to high voltage and high frequency (see footnote 2, earlier) and is recommended only on 5 year intervals or greater. It also must be carried out with great care to avoid creating resonance or induced voltages in other parts of the transformer and its auxiliaries.

Partial Discharge (PD) Tests

If an AC voltage is applied to a cable, it may cause a “partial discharge” inside any voids, cracks, or gas bubbles in the insulation. Generally, the volume inside the void has a lower breakdown limit than the insulation around it. Therefore, the induced voltage, if high enough, will cause this “weaker” zone to discharge without causing any discharges in “good” parts of the cable.

This partial discharge creates a small, very brief (measured in nanoseconds) electromagnetic pulse, which can be sensed by a suitable type of detector if placed on the cable. The pulse is detected by sensing it at one or both ends of the cable, and the computed time lag between pulse initiation and detection is used to determine distance from the point of measurement to the point of the defect. PD tests excite a cable using an appropriate AC voltage, and use a sensitive detector to track the pulses given off by the partial discharges. The PD test equipment analyzes the timing of both the test voltages and pulse arrival (i.e., both travel the length of the cable to the point of the flaw) to determine the location of the partial discharge. In this way, this test is similar to *time domain reflectometry*.

There are a host of complications to be overcome, however. The pulse, never very strong (these are only partial discharges, and relatively weak), suffers attenuation as it radiates away from its point of origin, so that the detector must be very sensitive to detect it. And, being of very short duration, the pulse is composed of all many frequencies.³ The attenuation of its radiating signal is frequency dependent. Essentially, it radiates through a filter with different amounts of attenuation in different frequency bands. Attenuation in the different bands depends on the test equipment being used, the cable type, duct types, soil, and other conditions that might influence the impedance as a function of frequency in the cable. Splices and branches along the cable also affect pulse propagation and attenuation, and may create “echo” pulses. If the neutral is badly corroded or damaged, the partial discharge will not pass down the cable. If the test is measuring a branched circuit, it becomes difficult to determine which branch a particular pulse is measuring.

These problems are solved for the most part in a wide variety of commercially available sources of PD tests. While based on this overall concept, each approach has specific characteristics designed to give advantages in some types of situations and/or to work around some of the most common problems encountered. Voltages used are generally about at the normal operating voltage or slightly higher (voltage must be high enough to excite partial discharges). Some test methods use standard frequency, others very low (.1 Hz), others high frequencies. Some use DC and then switch to AC of standard or varying frequency. Other variations include a host of novel and innovative ways to detect and measure pulse characteristics and infer location from them.

³ From a Fourier analysis standpoint: a pulse of “zero” width contains an equal amount of all frequencies. The actual pulse is so short that it has a very wide spectrum.

The wide range of characteristics involved makes it impossible to generalize about either the results or the limitations of PD methods. The amount of stress that is put on the equipment also varies greatly depending on the exact type of test procedure used. However, some PD tests, which apply an alternating voltage up to 1.5 times nominal, can put considerable stress on cables and associated equipment.

6.4 TESTS AND DIAGNOSIS OF INSULATING OIL

A major component in many electrical devices is insulating oil. It provides two functions. First, it provides dielectric strength, acting as an insulator, limiting current flow to where it is intended to flow. Secondly, it provides cooling capability, absorbing heat and transferring it away from where it is produced (e.g., a transformer core). The oil in major equipment such as oil-filled transformers, circuit breakers, and voltage regulators should be checked periodically. This can include both tests to determine that it is capable of doing its jobs (insulation and cooling), and diagnostic tests for clues to the condition of components inside the device (sludge build up, indications of arcing or overheating). Analysis of the insulating oil is one of the best condition assessment tools available for devices such as transformers.

Simple, Fast, and Reasonably Dependable Diagnostics

Tests of insulating oil are among the chief means of diagnosis for power system equipment for several reasons. First, a lot of electrical equipment contains insulating oil – it cannot be ignored and in fact should be checked, meaning that most testing has to deal with it. Secondly, such tests are performed on samples that can be drawn from the device easily without de-energizing it or making any special arrangements. Most power transformers are fitted with an oil sample valve or outlet to facilitate such testing. The equipment itself is not fitted with sensors or otherwise affected and can be left in service, a big advantage. Third, the tests are then carried out away from the device, in a non-energized (except for any test equipment) situation, which is safer and faster.

Finally, and perhaps most importantly, some types of insulating oil diagnostics (e.g., dissolved gas analysis) provide dependable diagnostic information about the severity and type of internal problems in a device. For this reason alone, tests of insulating oil are the cornerstone of condition assessment, particularly for older equipment.

Of course, oil tests and diagnostics are not enough by themselves. Oil tests do nothing to assure that other aspects of a device (bushings) are in good shape. Condition of other equipment can often be determined only by detailed exterior inspection and monitoring (for leakage in bushings, etc). But any program aimed at condition assessment of transformers and similar equipment will have to rely heavily on analysis of the insulating oil, about the only means of non-intrusively diagnosing the conditions *inside* the device.

Oil Samples Required

ASTM standard D923 provided a recommended standard for collection and handling samples of transformer, breaker and other power system equipment oils and is widely followed within the United States. It recommends sampling fluid from the bottom of enclosures when the fluid (oil) has a specific gravity less than 1.0 (as for most insulating oils), and from the top of the tank for those with a specific gravity above 1.0. This standard also specifies the types of materials to be used for containers and equipment (generally glass, with no rubber in the seals) and how to clean them. It gives recommendations for how much (usually 2 quarts) and under what conditions samples should be drawn (avoid humid days), and how to handle and ship samples, if needed. A good summary of this standard is included in Gill, 2008, Chapter 5 (see references).

Freeze test

A quick, inexpensive, but inexact, sample test for insulating oil is to place the oil sample in a transparent glass bottle and put it into a freezer set to slightly below 0°C (32°F) for several hours or overnight. Water in the oil will freeze, while the oil will not. Oil with any noticeable moisture content will appear milky or foggy.

Color test

This test involves passing a light of known strength and frequency content through the oil and measuring the spectrum of transmission (amount and color of light passed). The resulting spectrum is compared to a table of known colors for new oil and known changes due to various types of contamination and/or deterioration. The color test can identify quite well the degree of carbonization from switching (tap changers, breakers, oil-filled switches), which correlates with switching operations/condition. It can simultaneously evaluate oil for a number of other conditions too, but is considered inexact, so that it is often not used if more comprehensive tests for acidity, gas, and moisture are to be done.

Water Content (Karl Fischer) Test

Water can exist in the oil inside a transformer or other device in one of three ways. First, it can be dissolved into the oil. Second, it can be emulsified; a state in which it is distributed throughout the oil as minute droplets. Third, it can pool at the bottom of the tank or device. Water has several effects, all detrimental. First, it can reduce dielectric strength. Second, it can be absorbed from the oil into paper insulation and other materials in the equipment, reducing its strengths (both mechanical and dielectric). Third, it can promote corrosion (rust) of the case and other materials.

A small amount of moisture can actually be dissolved in the oil without having a highly deleterious effect on dielectric strength. However, this amount is temperature dependent – typically oil can contain up to 20 ppm at 0°C, rising to

100 ppm at 45°C and 200 ppm at 50°C. If dissolved moisture builds up when the unit is very hot and the unit is then cooled (as could happen during high loading periods). It will then “emulsify” out as micro-droplets throughout the oil – emulsified water. Even a very small amount of emulsified water in oil will reduce its dielectric strength greatly – the tiny droplets form pathways for discharge.

Thus problems caused due to moisture in the oil often occur after the load has passed peak load period and as the oil is cooling from the temperatures reached during peak conditions, but some internal elements of the windings are still warm. Emulsified oil stays that way (unless temperature rises again and it dissolves back into the oil). Typically, it will not “demulsify” (accumulate into larger drops and fall to the bottom of the tank) but will stay mixed into the oil as micro-droplets until removed.⁴

Tests for water include specific gravity tests, which measure the oil’s specific gravity against that of new oil of the same type. These are non-exclusive tests for water - deviations from a typical new value of .84 may indicate something other than water is present. Worse, the presence of other contaminants (lighter than oil) can counteract the presence of water (heavier than oil) in this test.

A water content test, often called the Karl Fischer Test, can be applied. There are three ways to apply the test, but the concept in all forms is to use iodine to absorb moisture chemically, then to detect and measure its presence in the sample. These tests are inexpensive, fast, quantitative, and dependable.

Dielectric Breakdown Tests

The oil can be directly tested for insulating strength using any of several types of test procedure. Two common types defined by the American Society for Testing and Materials (ASTM) are the D877 test, often called the “cup test” or disc electrode test, and the D1816 test, and often called the VDE test. Both identify at what voltage gradient strength the oil will break down.

The D877 test is performed in the field with portable machinery and involves putting about a pint of oil in an insulated container holding two round-faced disc electrodes separated by a .1 inch gap. Voltage is then raised at a rate of 3 kV/second until breakdown occurs. The test is usually conducted on five samples. A value of 23 kV is considered satisfactory.

The D1816 test is performed in a laboratory using about a quart of fluid in a VDE cell, a type of container within which are spherically capped electrodes with a 36 mm sector diameter and 25-mm radius with a very narrow (.08 inch) gap. The oil can be circulated in the cell, simulating oil movement inside a transformer.

⁴ Emulsified water is best removed by vacuum dehydration. Cooling the oil will precipitate some but not all emulsified water has droplets large enough that they will drop to the bottom of the unit in all conditions (the amount of water that can emulsify in water is a function of temperature). Typically, centrifuges or filtration will not remove emulsified water. Chemical means could also be used, but most of those methods leave traces of other chemicals which are just as detrimental, if not worse.

This test is considered more sensitive to moisture in the oil, and a better indicator of electric strength as the transformer sees it, than the D877 test.

Acidity and Sludge

Oxidation of the oil (chemical breakdown/combination of the oil with oxygen) will take place anytime the oil comes in contact with air and water, both of which contain oxygen that will chemically combine with the oil to form acid; eventually this acid precipitates out as a sludge. Some materials commonly used in power equipment promote oxidation, acting as a catalyst. One is copper.

Sludge buildup not only can clog radiators but also effectively “removes” oil from the reservoir, reducing the net amount available for cooling. Thus, most transformers are sealed to limit the amount of oxygen (air and moisture) that can get inside the unit. Sludge will stick to internal surfaces of a device, not necessarily accumulate on the bottom of the oil container. Inside a transformer it will coat the core, cooling ducts, and radiator lines, thereby reducing thermal transfer capability and resulting in more rapid and higher heat rise during peak operation. A transformer or breaker with indications of sludge buildup should be given priority for removal of sludge. The sludge can be removed by treating/filtering the oil (slow, uncertain) or by circulating hot, new oil to remove it. Regardless, heavy sludging leads to failure.

Oxidation is accelerated by temperature. A rough rule of thumb is that the *rate of oxidation doubles for every 18°F (10°C) rise in temperature*. Therefore, overloading a transformer can greatly affect the rate of acid formation. Tests have indicated that acidity is proportional to the oxygen absorbed in the oil.

Generally, insulating oils used in power equipment have “antioxidation” compounds such as DBPC (2,6 di-tert-butyl persesol) added to them to prevent acid formation. These act as “sacrificial” oxidation targets, chemically combining with oxygen to remove it before it can attack the oil itself. Oxidation of the oil begins when these materials are completely consumed.

Acidity tests

There are two ASTM laboratory tests (D974 and D664) and one field test (D 1534) for acidity of oil, all of which determine a “neutralization number” which is the number of milligrams of potassium hydroxide required to neutralize the acid (return it to neutral pH) in 1 gram of fluid. A test value of .4 or less is considered acceptable. Values above 1.0 indicate severe problems. The D1534 field test involves drawing out small samples of the insulating oil and adding it to a “wash solution” already contained in a small glass cylinder. Pre-packaged, sealed small “doses” of KOH (potassium hydroxide) are then added and their number counted. The wash solution has a color-changing property (like litmus paper) – it turns pink when one too many vials of KOH have been added.

Interfacial tension tests

An interfacial tension test can determine if the acid accumulation has or is leading to accumulation of sludge. Acid buildup in an oil leads to sludging, but as shown in Figure 6.3, the latter may follow the other by a good deal of time, and acid alone is not an indication of the presence of sludge.

Interfacial tension tests (IFT), such as ASTM D971 and ASTM D2285 measure the sludging characteristics of the oil and give some indication of the amount of build up. These test for sludge implicitly, but measure the tension between oil and water (i.e., akin to the surface tension between liquid and air). Interfacial tension between the two is reduced by the presence of “polar molecules” (sludge) in the oil. A reading of less than 15 dynes/cm is indicative of sludge buildup (a good reading is 22 dynes/cm).




Figure 6.3 Relative rates of antioxidant consumption, acid accumulation, and sludge accumulation, and the resulting change in interfacial tension in one transformer test. A drop in interfacial tension signals that acidity is already increasing in the oil, but is usually a precursor to, rather than an indicator of, sludge formation.

Analyzing Gases Produced by Insulation Deterioration

The different processes of deterioration or damage occurring inside a transformer or other oil-filled device create various gases, such as CO_2 , acetylene, hydrogen, and methane, which dissolve into the oil or pool in the space above the oil. While dissolved gases can affect the properties of the oil, the primary reason to test for the amount and type of dissolved gases in oil is that they give an excellent indication of the type and state of deterioration taking place inside the device.

The chief cause of gas being generated inside a transformer is from deterioration of the paper-cellulose and oil insulation – deterioration that is greatly exacerbated by heating. Deterioration rates are an exponential function of temperature. Deterioration-to-failure times for “cold units” are often measured in centuries, but lifetime at full load might be only a few years.

Cellulose (paper) insulation overheated to 150°C or above begins a breakdown process that produces carbon monoxide (CO), carbon dioxide (CO_2), traces of hydrogen (H_2) and methane (CH_4), and water (H_2O). The presence of ethylene (C_2H_4) or ethane (C_2H_6) indicates the oil has been heated to temperatures close to its absolute physical limit of around $550\text{--}600^\circ\text{C}$. Acetylene (C_2H_2) indicates the oil has been heated to even higher temperatures, usually by electrical arcs passing through it (very bad in a transformer, not so bad in a circuit breaker). Thus, the type and amounts of gas present give an indication of both the amount and cause of deterioration or damage that have occurred inside the device.

Total combustible gas test

For the most part, the gases produced by deterioration and damage (except for water vapor and CO_2) are combustible: they will freely combine with oxygen in the presence of a flame or sufficiently high temperatures. Total combustible gas analysis (TCGA) measures the amount of these gases by drawing off some of the gas from the space above the oil in a transformer and measuring what portion of the gas sample burns. Such tests can only be performed on sealed transformers which have a nitrogen blanket (nitrogen filled space) or conservator system, but this includes most types of power transformers. Many power transformers are fitted with a gas-sampling line or an external valve to facilitate such tests.

Total combustible gas analysis tests can be performed in the field with a portable unit. These units use a catalytic burner and a Wheatstone Bridge circuit to measure resistivity. The burning of the gases causes changes in the resistance of one element of the circuit, which is detected and correlated with the amount of gas. These tests are quick, inexpensive and precise as to total *amount* of gas, but non-specific as to the *type* of gas, and hence the exact nature of the problem. However, large amounts of combustible gases clearly indicate a problem.

Dissolved gas tests

Dissolved Gas Analysis (DGA – Dissolved Gas Analysis) is done in a laboratory

and will identify exactly what and how much of each type of gas is present. They are done not from a sample of the top-oil gas, but from a sample of the oil itself. A vacuum pump is used to extract the dissolved gas from an oil sample and pass it to a gas chromatograph. The quantity of each type of gas is then determined by gas chromatography. This gives a better indication than TCGA of both the amounts of deterioration and the type and cause of it.

Use of insulating oil tests

The presence of CO, CO₂, and water alone, determined by a DGA test, indicates operation at high but not necessarily undue temperatures, and what might be called standard deterioration – aging with time in service. Sound transformers running at peak loading typically generate small but measurable quantities of these gases. However, traces of the other gases, ethylene, ethane, and methane, indicate problems (non-standard deterioration) and give an indication that the equipment is not a sound unit.

However, the most commonly used gas test is the total combustible gas analysis (TCGA), because it is quick and can be performed in the field. It does not provide accurate data on *which* gases are present and in *what amounts*, and so cannot provide the information needed to determine the cause of any problems detected directly. However, by doing TCGA analysis frequently (it is a very inexpensive test) the *rate* of increase of these gases can be tracked, and that rate is useful in distinguishing among the different gassing causes. Creation of gases from deterioration of paper insulation, even if at very high “normal temperatures” is a very gradual process. A high measurement might mean only that the unit is old and has not been tested/purged in the past. Subsequent tests several months later might reveal a low gassing rate and thus indicate this is the cause.

By contrast, impending-failure conditions that are caused by the very high temperatures generated by short circuits or by arcing through the oil create gas at a much higher rate. A high rate of increase over a short period of time (a day in the case of a unit thought to be in jeopardy, a quarter year otherwise) can provide a good clue as to the cause of the gassing: slow – normal deterioration; fast – an internal problem possibly indicating imminent failure.

How the Test is Carried Out is Important

It is important to bear in mind that the cost, time and resources required, as well as the sensitivity and quality of results obtained will depend greatly on both the exact nature of the equipment being used and the procedure for its application. This means keeping in mind what is grounded and what is not, measurements of equipment temperature before and after, etc. Similarly, the level of stress put on the equipment can depend on instructions being followed precisely and how the test was applied.

Table 6.4 Power System Electrical Test Procedures

Test	Type	Where?	In serv?	No-Go	Condition	Stress level
Thermal load tests	AC	Field	Yes	No	Yes	-
Resistance test	DC	Field	No	Yes	Maybe	-
Dielectric absorp.	DC	Field	No	No	Yes	-
Power factor tests	AC	Field	No	No	Yes	-
Polarization recov.	DC	Field	No	No	Yes	x
Hi-potential test	both	Field	No	Yes	No	X
Induced potential	VF	Field	No	No	Yes	X
Partial discharge	AC	Field	No	No	Yes	x
Transf. turns ratio	AC	Field	Yes	Maybe	Maybe	-
Oil acidity tests	-	Lab/Field	Yes	No	Yes	-
Interfacial tests	-	Lab	Yes	No	Yes	-
TCGA tests	-	Field	Yes	No	Yes	-
DGA tests	-	Lab	Yes	No	Yes	-
Disassembly/inspection		Field	No	Yes	Yes	-

Comparison of Tests and Diagnostic Methods

Table 6.4 compares the various test procedures for equipment and oil on the basis of various characteristics. From left to right, column one describes the type of test, column two the type of voltage used (AC, DC, or VF for non-standard or variable frequency), and column three identifies whether it is performed in the field or laboratory. Column four tells whether the unit can be left in service or must be de-energized while the test is being run. Column Five identifies whether the test is typically used as a “Go/No-Go test” – one that provides a definite indication that the unit is unfit for continued service. Column six (conditions) explains whether the test provides relative diagnostics on condition and deterioration that can be tracked over time. Finally, column seven gives the level of stress caused by the test, the size of the X indicating relative level of equipment stress caused by the testing.

6.5 ON LINE AND REAL-TIME MONITORING AND TESTING

Any test that can be performed on an as-needed basis in the field can be performed on a continuous and ongoing basis by equipment installed to test and remotely report the test data. Whether such online or real-time testing is justifiable from a benefit, cost, and operational basis is another matter entirely. It is technically feasible, and in the case of most types of tests, equipment is available or can be

assembled to do so when needed. Online means the data is sent back remotely but perhaps not instantly: the maximum temperature seen during yesterday's 24 hours, etc. Real-time means the data is sent nearly instantly, with no delay except that very slight time required to assimilate and transmit the data.

Online Monitoring versus Testing

A distinction should be made between online and real-time monitoring of equipment and testing.

Monitoring is the remote measurement and reporting of data relevant to its function and operation. Current, voltage, and power factor measured at a power transformer or through a metering point provide data needed to observe how the equipment is operating and make good decisions as to the continued operation of that equipment. For example, temperature is often monitored in a transformer and used for operational decisions on loading. Similarly, breakers and switches might be monitored to verify their position (open, closed, breaker recharged and ready to operate, etc.), and voltage regulators, load tap changers, or switched-bank capacitors monitored as to position, etc.

Testing is remote measurement and reporting of data specific to the diagnosis of the condition or status of the equipment. For example, online reporting of the TDG (total dissolved gas) amount in a power transformer provides nothing useful to operators of the system, other than an indication that the unit might be at jeopardy of failing. However, a watchful eye on the total over time will provide good information about the general condition of the unit. Similarly, remote collection of information on the timing of each breaker operation or close can give an indication of the condition of the breaker itself. That information will be of very limited use to operators of the system but can be used to diagnose condition and if and when units need maintenance.

Self-diagnosis and incipient failure diagnosis involves local collation and analysis of test results and operation monitoring data to determine testing and diagnosis results which are then transmitted. A system might analyze DGA readings on a transformer as well as acoustic and radio-frequency monitoring of noise, along with temperature, and deduce that the unit was suddenly changing condition rapidly and/or approaching a boundary or pattern for one of more of these readings that indicates concern. An alarm, perhaps in conjunction with all the readings, would be sent. This local analysis might be done by a unit attached to and at the equipment, or a local hub, for example one system that analyzes all four power transformers at a major substation.

A Slow and Long Term Trend Toward More Online Testing

Remote online or real-time reporting of TDG readings and switch operating times is technically feasible and equipment is available (much of it based on methods and equipment originally developed for military applications). But it is seldom used, because it is often not cost justifiable. At many utilities, benefit/cost analysis of online TDG readings shows it will provide little additional benefit

beyond that derived from routine manual testing, and at a higher cost. Breaker and other equipment testing via remote collection is similarly not justifiable. That said, three trends will all lead to a gradual increase in the amount of online/real-time testing done remotely.

- 1) *The benefit derived from online testing will gradually improve.* As summarized in Chapter 1, the problems caused by aging equipment do not increase dramatically and quickly, but instead slowly but steadily over the long run. Sound management will mean a good deal of old equipment will always remain in systems. Therefore, in the long run, the value seen from online test result reporting will gradually increase – perhaps to two or three times that at present.
- 2) *Cost will drop* for the sensors, data systems and hub receiving systems, and analysis and tracking software. There is no reason to think that this will not also be a gradual but steady trend in the reduction of the price of equipment and support software for such applications, along the lines typical in any field of electronics.
- 3) *High rates of automation and smart systems will reduce the marginal cost of adding testing.* Finally, Smart Grid systems will result in a nearly omnipresent data communications grid, meaning that the cost of adding testing to the current system will be small compared to the cost of creating the system for that purpose today.

That said, the authors do not anticipate anything other than a *very* slow move toward online testing and self-diagnosis for equipment. This will occur in a top down manner, with major units such as large auto-transformers and key choke points for system security receiving comprehensive online testing and real-time incipient failure coverage first, with the system gradually deploying to equipment of lesser criticality.

Smart Grid Automated Testing

One key performance goal for many Smart Grid distribution systems is a “self-healing” feature. When breakdowns or failures occur the system will re-configure its switching to isolate failed equipment and connect sections of the system isolated by the failure to a still-working source of power. This is done in a manner that causes the least possible interruption to customer service.

Currently, most self-healing system plans and concepts envision the self-healing capability as only reactionary – the smart equipment waits until a failure has occurred to take action. A potential improvement would be a Smart Grid *proactive self-healing capability*. This would use automated and remote testing of key equipment to determine where and when it is likely that an

equipment breakdown or failure is likely to occur. Switching would then re-configure the network to isolate and deactivate that equipment, avoiding the failure in a way that would eliminate or mitigate service interruptions to customers.

Such pro-active systems will very likely be deployed eventually, at least in places where the value they provide in increased reliability of reduced equipment damage justifies their addition. However, considerable progress needs to be made in the durability and dependability of automatic remote testing sensors, and in the cost of such units, and in evolving central or peer-to-peer control systems to properly handle the outages.

Therefore, it will be the reactionary self-healing systems that drive deployment of Smart Grid self-healing capabilities. A consequence is that their deployment will create a lower marginal cost for the *addition* of automated remote testing and pro-active self-healing capabilities. The authors expect that while self-healing capabilities will make their way quickly into distribution systems in most first-world countries, pro-active capabilities will deploy much later, and much more slowly.

6.6 SUMMARY AND FINAL COMMENTS

This chapter summarized inspection, testing, and diagnostic procedures for power system equipment. The major goal of these activities, as it relates to aging infrastructure, is to determine the condition of power system equipment already in the field. Properly carried out, such inspections and tests can identify if that equipment is “good to go” for service and give some estimate of remaining service lifetime. They also catch some failures before they occur, permitting the utility to make repairs, saving equipment that might otherwise be ruined.

While this “ounce of prevention” function is worthwhile in reducing equipment replacement costs by saving some equipment from failure, the real value of inspection, testing, and maintenance is its ability to reduce the number of outages that occur in service. It is worth considering that exact knowledge of when a unit would fail would be worth a great deal to a utility, allowing it to schedule replacement, but at exactly the latest moment possible. While even the best of condition-assessment methods can not exactly determine expected lifetime, the information provided is a step in the right direction and worth a good deal toward improving service quality and allowing optimal management of equipment lifetimes.

The methods covered in this chapter are worthwhile only if applied in well-managed program aimed at improving and maintaining service quality at high levels while keeping costs as low as possible. Chapter 9 will discuss application of these methods to power system equipment and identify how the tests are best organized and used in an inspection and test program. Chapter 10 will discuss prioritization methods for O&M, and how utilization of units can be optimized so the utility gets the most overall “bang for the buck.”

REFERENCES AND BIBLIOGRAPHY

- J. J. Burke, *Power Distribution Engineering – Fundamentals and Applications*, Marcel Dekker, New York, 1994.
- P. Gill, *Electrical Power Equipment Maintenance and Testing – Second Edition*, CRC Press, Boca Raton, New York, 2008.
- H. H. Malik, A. A. Al-Arainy, and M. I. Qureshi, *Electrical Insulation in Power Systems*, Marcel Dekker, 1998.
- O. C. Seevers, *Management of Transmission and Distribution Systems*, Fairmont Press, Lilburn, 1995.
- W. A. Thue, *Electrical Power Cable Engineering –Third Edition*, CRC Press, Boca Raton, 2012.

7

Aging Equipment and Its Impacts

7.1 INTRODUCTION

Aging equipment is undoubtedly the primary cause and effect associated in most people's minds with the term "Aging Infrastructure." And while there are other significant factors involved as were summarized in Chapter 1, there is no doubt that aging infrastructures present an incredible challenge largely due to the effects of aging equipment. Aging equipment has high failure rates and requires, proportionally, more inspection and maintenance cost. Often parts and service equipment are hard to come by (and expensive).

This chapter looks at equipment aging – what it is, how it impacts the various types of equipment, and how those changes affect the reliability and economy of a power system's operation. Invariably failure rates increase with age, usually in an exponential manner. But characteristics vary depending on equipment type (aging impacts poles very much differently than say, cable) and conditions of use. Most important to the modern distribution utility is how an aging base of equipment changes the performance (reliability and economy of operation) of the system as a whole, and impacts the quality of customer service.

This chapter begins in Section 7.2 with a qualitative look at aging and its impacts on equipment, discussing the different categories of aging impact, and the general characteristics of aging on equipment such as transformers, poles, cables, etc. Section 7.3 then looks at quantitative trends in failure rates as equipment gets older – how and why do failure rates increase? This section raises an interesting concept – what would exact information on when a particular unit of equipment

will fail be worth, and how could a distribution utility utilize that information? Section 7.4 delves into the impact that aging and failure rates have on the installed base of equipment in a utility system and the repercussions on cost and service quality that result. Section 7.5 concludes with a summary of key points.

7.2 EQUIPMENT AGING

All electrical equipment kept in service ages. During every year of service the equipment becomes a year older. One consequence of this aging is gradual deterioration of the physical and electrical strengths of the equipment, until at some point failure occurs one way or another – the end of the equipment’s useful lifetime. Table 7.1 shows several types of deterioration that affect older equipment. The “caused by” columns show which of the three primary types of aging, covered in Table 7.2, lead to the respective type of deterioration.

Table 7.1 Types of Deterioration Causing By Aging

Type of Deterioration	Caused By			Comments
	CA	CSS	AES	
Corrosion	X	x	x	Chemical decomposition, or combination with oxygen or other ambient elements, until the material loses its required mechanical or electrical strengths or qualities.
Dielectric loss	x	X	x	A number of different mechanisms (treeing, contamination) can lead to the loss of electrical withstand strength.
Shrinkage/Hardening	X	x		Paper, rubber, and synthetic gaskets and seals harden and/or shrink with age, losing their ability to keep out moisture or contain pressure.
Wear		X	X	Mechanical components lose tolerance and bind, or do not hold with the same bond they once did.
Moisture retention	X			Water is gradually absorbed into a material, degrading its mechanical or electric strength

Table 7.2 Categories of Equipment Aging Impact and Their Meaning

Category	Meaning and Impact
1. Chronological age (CA)	Age since construction. Certain materials deteriorate over time due to natural causes most directly associated with chronological age (paint chalking and corrosion, hardening of rubber and gasket materials, etc.).
2. Cumulative service stress (CSS)	The cumulative effect of the time the unit has been energized or in use, and the load (mechanical, electrical) it has served in that time.
3. Abnormal event stress (AES)	The cumulative impact of severe events generally not considered as “normal service.” This includes through-faults for transformers, storm and auto-accident stress for poles, etc.
4. Technical obsolescence (TO)	Digital and data communications equipment, as well as other “hi-tech” equipment, can become “too old” by virtue of not being compatible with new systems and equipment.

Types of Aging Impact

Aging can have any of several particular meanings with respect to power system equipment, as shown in Table 7.2, which lists four specific categories of aging and their impacts. Generally, the terms “aged or old” and “aging” are used to refer to some non-specific combination of all four effects, with the understanding that the equipment is not as useful or dependable as similar equipment of lesser age.

Chronological age

Certain materials such as paint, paper and fabrics, rubber and synthetic gaskets and seals, and insulation, deteriorate over time. Time since manufacture alone is the direct cause of this type of deterioration. Thus, seals and gaskets begin to harden, and paint to chalk, as soon as a transformer or breaker is built, progressing at much the same rate regardless of whether the unit is installed and in service or kept in inventory in an equipment yard. A creosote pole will dry out and begin the very slow process of chemical breakdown of its preservative whether or not it is stored in an equipment yard or put in the ground. It will however, begin to rot at

the ground line only after being put in the ground, particularly if the soil around it is moist and of the right pH.

As the pole example in the last sentence illustrates, very often the rate of deterioration over time is exacerbated or mitigated by ambient conditions. Transformers, poles and cables stored outside and subjected to summer heat, winter cold, direct sunlight and ambient moisture, will deteriorate faster than if kept in sealed, temperature-controlled and dehumidified warehouses. Deterioration due to chronological aging can be cut by two thirds in some cases if the equipment is kept in a controlled environment. Generally, “chronological aging” rates and average lifetimes assume that the equipment is installed in typical ambient service conditions (i.e., poles and transformers are kept outside whether in service or not).

In many cases, the deterioration that would occur due to time alone is accelerated by the stresses caused by putting the equipment in service. Hardening of gaskets and seals, and paper insulation, will occur at a faster rate when a transformer is operating (heat generated by electrical losses will accelerate the deterioration processes).

In some cases, however (i.e., paint chalking and some types of corrosion), the rate of deterioration does not depend on whether the unit is in service. And in a few situations, putting a unit in service will reduce the rate of chronological deterioration. For example, energizing a power transformer being kept for standby service creates a low level of heating (due to no-load losses) that “cooks” any moisture build-up out of the unit, slowing that rate noticeably.

Cumulative service stress

Certain types of deterioration occur only, or primarily, due to the use or operation of the unit for its designed purpose, and are proportional to the time or the cumulative level of use in service, not chronological time. There are three primary types of processes involved:

1. *Electromagnetic field stress* established when a cable or unit of equipment is placed in service can lead to degradation of dielectric strength and in some cases, promote corrosion. Essentially, applying voltage to the device produces a stress that will eventually break down insulation and accelerate some chemical processes (particularly if the unit is not grounded well or cathodically protected). This type of stress depends only on the voltage. It does not depend on the amount of current (power) the unit is handling.

Every unit of electrical equipment – transformer, cable, bell insulator – is always designed to withstand a nominal voltage, that which it can withstand safely and dependably over a long period of time (decades). But, given enough time, the electromagnetic stress from that

voltage will lead to various types of insulation breakdown, which can include treeing, attraction of impurities that build flashover paths, etc. Eventually, over a century (post insulators), or only a period of several decades (UG cable), the insulation will fail for this reason alone.

In most cases the deterioration rate due to this cause is very voltage sensitive. Raising voltage even slightly will increase stress a good deal. Thus, higher voltages, due to operating the unit at the high end of its nominal range, or to misoperation, switching surges, or lightning, can accelerate this deterioration, sometimes quite dramatically. Similarly, lowering the voltage on equipment that is suspected of being weak (old cables) to the low end of the permissible scale (e.g., .95 PU) can often reduce stress, deterioration and failure rates significantly. And of course, voltage that reaches too high a level – above the tolerable range – leads to immediate failure.

2. *Wear.* Parts that move against one another, in devices such as tap changers, capacitor switches, circuit breakers, load break switches, etc., wear. The movement gradually erodes material in the moving junction between parts, loosening tolerances and often scouring smooth bearing surfaces. Such wear can lead to binding of moving parts and failure to operate, or to operate slower than necessary (breakers).

The term “wear” can also be applied to a type of mechanical deterioration that occurs in parts that are not technically moving against one another. An example is in overhead conductors and their splices, terminations, and support brackets. The conductor “moves” in the sense that it sways in the wind, and expands and shrinks, and sags more or less, with temperature (and loading). Over decades, this movement can lead to loosening of the bond between clamps and conductor, or within splices or termination fittings, and even to cracking of the conductor material itself. Eventually something breaks and the line falls.

This gradual deterioration is accelerated greatly by *aeolian vibration*. Wind blowing across the conductor sets in motion a resonant vibration, much like blowing on a harp string might cause it to hum. This vibration stresses the bonds in splices, brackets and clamps, and can fatigue the metal strands of the conductor, over a period of only months in extreme cases. Various types of vibration dampers are available to fit to conductors to stop such vibration, but these are usually fitted only to known or suspected cases of heavy vibration. Most overhead spans experience a slight amount of this vibration, and over decades it leads to a degradation of mechanical strength in the conductor and its fittings.

3. *Heat stress.* Higher temperatures accelerate many of the physical and chemical mechanisms involved in material’s deterioration, to the point

that heat can be considered electrical equipment's worst enemy. In electrical equipment like transformers, regulators, and motors, heat is generated by electrical losses in the equipment. The higher temperatures created as a result cause deterioration of materials in insulation, gaskets, seals, the transformer oil itself, and in some cases, the metal in conductors and/or mechanical components. Similarly heat affects all electric equipment, "aging" or ruining the functionality of various materials or their components. Parts that should be soft and expanded, so they seal well, shrink. Parts that should insulate lose dielectric strength. Metal expands, binding and wearing, or, at very high temperatures anneals so that it becomes brittle and of low strength – essentially a different type of alloy.

In almost all cases, the rate of deterioration of a component or material is a function of the temperature, generally increasing exponentially with a temperature up to some maximum-tolerable temperature beyond which the material fails immediately. The relationship is exponential – a device might be able to operate for decades at 80°C, for several years at 90°C, for a week at 100°C, and for only a few hours at 110°C.

In addition, temperatures that reach certain levels can cause internal physical changes in some types of electrical equipment, changes that promote misoperation. For example, in a power transformer, extreme heating of the windings (due to losses from very high loadings) can cause hot spots on the core that are hot enough to create gassing or boiling of the oil. The gas introduced into the oil by the bubbling contaminates the oil's ability to act as an insulating medium, and a "bubble path" can provide a route for flashover, leading to immediate and catastrophic failure.

In mechanical devices such as breakers, load tap changers, and motors, high temperatures can cause swelling of mechanical parts, leading to binding (bearing tolerances are reduced to the point that the device will not operate) and misoperation, and/or high rates of wear. In overhead conductor, high enough temperatures (typically caused when losses generate heat sufficient to create a 100°C rise above ambient) will anneal the conductor material. As a result it hardens, becomes brittle and loses its elasticity and mechanical strength. Vibration from wind and the natural expansion of heating and cooling from diurnal temperature variations then quickly leads to minute cracking and mechanical failure and the conductor breaks and falls to the ground.

The deterioration caused by high temperature is cumulative – i.e., a period of high temperature will cause deterioration that will not be "undone" when temperatures fall. A device will suffer gradual

deterioration of its components, equal to the cumulative stress of high temperature over its service life to date. Since heat (the cause of high temperatures) is due to load (the load current causes losses that create the heat), power engineers often associate deterioration rate directly with load: higher levels of load lead directly to higher rates of deterioration. The term “loss of life” refers to the rate of equipment lifetime loss associated with a particular loading (heat and temperature) level for the equipment.

Many of the chronological aging processes discussed earlier are accelerated if the device sustains high temperatures for long periods of time (as is common for many electrical devices). Thus, chronological aging and cumulative service stress is not, in practice, completely independent.

Abnormal event stress

Over its service life, any particular unit of power system equipment is likely to see a number of “events” which lie outside the normal conditions expected in service. These include electrical events such as through-faults, switching surges, and lightning strikes, and/or harsh mechanical events such as automobiles striking the device (i.e., a pole or pad mounted device), high ice loadings from freak winter storms that mechanically stress overhead conductors, and similar situations.

Auto accidents and storm damage are prime causes of wood pole and metal structure failure, and somewhat unavoidable given that these structures are often placed on public easements and parallel roadways. However, most of these events are reported (many result in immediate failure of the pole, or equipment at the top of it) so the utility is at least aware of the event and can inspect and repair any damage done.

Ice loadings are another stressful mechanical factor, but one that is more difficult to track. During severe winter weather, ice will accumulate on overhead conductors. Given extreme conditions, several inches of ice can accumulate, weighing many hundreds of pounds per span. The cumulative weight of conductor and ice can exceed the strength of the line, and it will part, falling to the ground. However, in the vast majority of winter storms, icing does not lead to failure – the weight is tolerated without any damage or deterioration.¹ When the ice melts after the storm, the conductor is, if not

¹ Conductors, particularly those with steel reinforcement (ACSR – Aluminum Clad Steel Reinforced Conductor) are designed to have high mechanical strength to deal with periodic severe ice loadings. In some areas of the world (e.g., parts of Saskatchewan) the distribution utility uses an entirely steel wire in its OH distribution sacrificing the lower resistance of aluminum for the superior ice loading capability of pure steel wire.

“good as new,” as good as it was before the storm.

But in between these two situations – outright failure on one hand and no damage done on the other – is a narrow range of “damage-done” cases that are particularly vexing to utilities. The ice loading on a particular span of an overhead conductor might reach a level that does not cause immediate failure, but is sufficient to stretch the conductor to its elastic limits, or over-stress clamps, splices, etc. This leads to accelerated deterioration and failure soon afterward. The utility has no way of knowing if any portion of its overhead lines were so affected and if so, which portions.

Through-faults are the most stressful event routinely seen by many transformers, from the standpoint of gradual deterioration leading to failure. A device downstream of the transformer experiences a fault, and for a brief time until a circuit breaker clears the fault, the transformer sees current flow through it that is from three to fifty times normal maximum. Heating caused by this high current is usually not the major impacting process – in most cases the fault current lasts but a few cycles and little heat energy is generated.

Instead, the most common damaging effect of through-faults is the magnetic field created by the very high current levels, and its shock impact on internal components. Transformers (and motors and voltage regulators as well) are designed so that the current flow through them causes an intense magnetic field – required for their operation. But a through-fault multiplies that field strength in direct proportion to its greater current level, creating a tremendous magnetic field and compressing or pulling apart nearby components, etc.

The magnetic force from a severe fault can create tremendous force. If high enough, this force can literally rip connections loose inside the transformer, leading to immediate and perhaps catastrophic failure. But typically, in a well-planned power system, fault currents are limited by design and operation to levels not so severe that they lead to such problems. Still, the mechanical impulse from a “tolerable” through-fault will make a large power transformer, weighing many tons, shake and “ring” as if it had been dropped several feet, and has an overall effect on it similar to such abuse. The cumulative impact of a number of such severe shocks over a transformer’s lifetime can be a significant loosening of the core stack, and stretching and twisting of windings and connections – a general weakening of the mechanical integrity of the unit.

The cumulative impact of through-faults is thought by many power equipment experts to be the leading cause of power transformer failure. There are recorded cases where transformers (and motors, and voltage regulators) fail when a through-fault occurs, generally when such faults occur during peak loading conditions (when the transformer is already quite

hot from high loadings and thus under a good deal of stress). Previous events have weakened the overall integrity of the unit until this one last through-fault is sufficient to cause failure.

However, more often than not, failure does not occur during or immediately after a through-fault. Since peak loads occur less than 5% of the time most through-faults happen during off peak times when conditions in the transformer are at low stress levels. Often the unit tolerates the fault at that time, but the damage has been done. The unit subsequently fails at a later time – hours, days, or even weeks afterward – when it is first exposed to prolonged peak load stresses.

Few utilities have dependable, easy to access records on the through-faults experienced by major transformers during their lifetimes. As a result, one of the most useful records of cumulative stress needed to determine present condition is unavailable to engineers and planners trying to deal with aging equipment and the estimation of its remaining lifetime.

Lightning strikes are severe-stress events that often lead to immediate failure of a device, almost a certainty in cases where the device is struck directly. Lightning is a pure current source, with the current flow being from ten to one hundred times the normal fault levels seen in power systems. A lightning strike can cause an immense “through-fault” like failure, or lead to a voltage flashover that causes other problems.

Failure modes due to lightning strikes are complex, and not completely understood, but failure occurs mainly due to heat. Although very brief, the incredible magnitude of the lightning current, often over 500,000 amps, creates heat – trees hit by lightning explode because the water in them is instantly vaporized. Similar impacts occur in electrical devices. In other cases the magnetic shock of high current does the damage.

There is considerable evidence that a good deal of damage can be done to electric equipment by indirect lightning strikes. A lightning strike to the ground near a substation can cause high voltages and or erosion of grounds. A strike in the ground may also travel to a nearby UG cable or grounding rod, burning it badly.

Lightning strikes to overhead lines often cause immediate failure, but can instead cause a flashover, which, due to breaker operation and an interruption, leaves no outaged equipment. However, the equipment may be damaged to the extent that it will fail in the next few weeks or months. In some utility systems, there is a statistically significant increase in service transformer and line failures after the annual “lightning season.”

Abnormal stress events can occur from time to time in the life of any electrical equipment. In some cases, the severity of one event leads to immediate failure, but often it is the cumulative impact of numerous events, over years or decades,

that gradually weakens the device. Sometimes these events cause deterioration similar to that caused by time or service, but most often the degradation in capability is different – through-faults loosen transformer cores and twist internal fittings in a way that no amount of normal service does.

Technical obsolescence

Most types of equipment used in power systems are from “mature” technologies. Although progress continues to be made in the design of transformers, breakers, and so forth, these devices have existed as commercial equipment for over a century, and the rate of improvement in their design and performance is incremental, not revolutionary. As a result, such equipment can be installed in a utility system and be expected to “do its job” over its physical lifetime – obsolescence by newer equipment is not a significant issue.

A forty-year old transformer in good condition may not be quite as efficient, or have as low maintenance costs as the best new unit, but the incremental difference is small and not nearly enough to justify replacement with a newer unit. However, in most cases, if a unit is left in service for a very long time, the issue of spare parts will become significant. Some utilities have some circuit breakers that have been in service 50 years or more. Replacement parts and fittings are no longer manufactured for the units, and have to be customized in machine shops or scavenged from salvaged equipment.

But a noticeable segment of the equipment used in utility power systems does suffer from technical obsolescence. The most often cited cases are digital equipment such as computerized control systems, data communications and similar “hi-tech” equipment associated with automation, and remote monitoring and control. Typically anything of this nature is eclipsed in performance within three to five years, often providing a utility with a new level of performance that it may wish to purchase. The older equipment still does its job, but is vastly outperformed by newer equipment. Economic evaluation of the benefits of the new systems can be done to determine if the improvement in performance justifies replacement. Often it does.

But beyond performance issues alone, there are other factors that mitigate against this equipment remaining in service for decades. First, since in some cases spare parts are made for this type of equipment for only a decade or so after manufacture, there is no viable prospect of leaving the equipment in place for many decades. Secondly, even if parts are available, newer equipment requires different test and maintenance procedures. Qualified personnel to make inspections and repairs become more difficult (and expensive) to find.

While digital equipment used in control and automation is the most obvious case of “technical obsolescence,” there is a much more dramatic, and significant case that is causing perhaps the most contentious issue facing the power industry as it moves toward full deregulation. The issue is a stranded assets and the cause is

technical obsolescence of generating plants. Many utilities own generating units that are fifteen to twenty years old. These units are not fully depreciated – when the units were built, the utilities borrowed money to pay the expenses, financed over thirty-year periods, and they still owe significant amounts on those loans.

But the technology of generator design has improved considerably since those units were built. Newer generators can significantly outperform those older units, producing power for as much as a penny/kWh less and requiring less frequent and costly O&M and fewer operators. In a deregulated market, merchant generators can buy these units and can compete “unfairly” with the utility’s older units. Thus, these older units are “stranded” by the shift to deregulation – the utility was promised cost recovery when they were purchased, but now faces competition due to technical obsolescence, and a drop in the sales that were expected to provide revenues to pay off the loans.

In the T&D arena, technical obsolescence will continue to be an important factor in control, automation, and metering areas. In addition, a high rate of technical progress will continue to impact power system equipment areas such as distributed generation, power electronics/power quality equipment, and home and building automation and control systems.

Lack Of Proper Care Can Accelerate Aging Effects

Lack of proper maintenance or inattention in setup and/or operation can, in some cases, lead to excessive deterioration due to secondary impacts (e.g., neglected seals let in moisture which leads to dielectric failure). The degree to which neglect leads to premature failure depends on the type of device and the extent to which its design depends on periodic maintenance.

Equipment with moving mechanical components tends to be most sensitive, that sensitivity proportional to the frequency of mechanical operation. Thus, tap changers and voltage regulators are most prone to early failure due to neglect. However, upon failure they can often be repaired by performing the maintenance or replacement of parts that has been deferred.

On the other extreme, most service transformers and residential power meters are essentially no-maintenance devices – they can be installed and expected to last for a normal lifetime without any maintenance. The first attention they get subsequent to installation is when they fail to do their job for some reason. At this point they are usually replaced. Equipment requiring little or no maintenance generally has only one “service mode” replacement.

Neglect or stupidity in setup or operation is another matter entirely. A neglected aspect of equipment use, often leading to rapid deterioration is poor grounding. Equipment may not be well grounded when installed, or the quality of the ground might degrade over time due to corrosion of clamps, etc. In either case, a weak ground can cause higher than normal and/or abnormal currents during “normal operation” accelerating the deterioration caused by normal service

stresses. More critical yet, poor grounds often exacerbate the impact of lightning, switching surges, and through-faults to the point that the stress caused is two to five times greater. Failure is nearly assured.

Other types of poor setup include mismatch of transformer impedance in multi-phase banks composed of single-phase units, incorrect settings for regulators, and similar inattention to details in equipment utilization.

Loss of Life and Load-Related Lifetime

Overall, of all the deterioration modes mentioned above, the major contributor to aging of electrical components in a power system is high temperature, a service-stress-related cause. High enough temperatures will cut short a device's life instantly. More typically, the cumulative effect of high temperatures over years of service leads to deterioration of key elements inside the device until they fail, leading to cessation of function, and often ruining the device.

The heat generated by electrical loss causes high temperatures. Those losses are a function of the electric load on the device. Therefore, the major component of aging is related directly to the loading on the device. While other types of deterioration proceed in parallel, usually this is the one that leads to failure (unless a high incidence of abnormal events fails the unit first).

For this reason, power system equipment engineers often speak of load-related lifetimes and "loss of life" of a transformer, motor, regulator or other component, as a function of its loading. Figure 7.1 shows the expected lifetime of one particular transformer if loaded to various levels. The unit can be expected to last 40 years (350,000 hours) if loaded continuously to 80% of its rated load, or only 7.5 years if loaded to 100% on a continuous 24 hour a day basis (solid line). If serving a typical utility load (time-varying load with a 60% annual load factor and peak periods of four hours duration) the unit has an expected lifetime of 40 years when serving a peak load equal to 100% of its rating (dotted line).

Thus, operation at 150% of rated load (estimated lifetime, 350 hours) can be likened to accelerating the unit's normal 40-year (in utility service) lifetime by a factor of 1,000. The unit loses 1,000 hours of expected operational service lifetime for every hour of operation at 150%, i.e.; every hour of operation at this loading level uses about 40 days normal service lifetime. Most other types of power system equipment have qualitatively similar lifetime versus loading relationships.

While this is a dramatic acceleration in loss of life, most utilities make it policy to operate a transformer at much higher loads during contingencies. Under emergency conditions, transformers are often operated at load levels up to 166%




Figure 7.1 The expected lifetime (time to 50% likelihood of failure) for one particular type of power transformer, as a function of the loading in percent of nominal rating. Solid line shows lifetime for constant loadings. Dotted line shows expected lifetime if serving an annual load curve of 63% load factor with a peak load measured at 100% percent of rated capacity.

of rating. At this load level, expected lifetime is only 120 hours and every hour of operation uses the equivalent of four months of normal operation (see Chapter 8 for a discussion of contingency loading policies).

The reason is one of economics, and makes sense in almost all cases. Given that peak loads occur only occasionally, a transformer that has been installed so that it is expected to take *up to* 166% of its normal peak load during a contingency will probably see only 20 hours of such intense service during a 30 year plus service lifetime. Most contingencies will occur off peak when loading and thus the stress is far lower. Those 20 hours of loading will cumulatively lead to a loss of 7 years of the expected 40-year lifetime.

But the benefit gained (contingency coverage of other transformers) is considerable, because other options for that coverage are costly. And the savings lie entirely in the present – it has a present worth factor of 1.0. By contrast, the

loss in life is realized at the end of the transformer's life, 33 years into the future (40 years minus 7). Present worth factor, in that timeframe, is minute (about 3%) so the lost service life has minimal value by comparison. Simply put, benefit exceeds cost.

Of course, a utility can get into trouble, as many have, by putting old transformers under such high levels of stress during contingencies. Many utilities have transformers that have successfully made it through forty or more years of service, during which time the units were never exposed to such high contingency loadings, due to the more relaxed loading standards that existed in the traditional utility industry. The units are partly deteriorated and weaker than new – they may have only a decade of remaining expected life. But due to financial pressures, the utilities have to switch to a policy of intense contingency loading. The problem created is that the stress of this use may absorb most of what life remains.

“Old” Depends on the Expected or Design Lifetime

Finally, it is worth noting that whether equipment is old or not depends on the expected equipment lifetime, as designed. A 20-year old wood pole is not really old (expected lifetime is at least twice that), but a 20-year old substation RTU is. A 40-year old span of ACSR is not yet really old, but a forty-year-old set of capacitor switches is.

Electrical equipment can be designed to withstand any degree of deterioration, by building it from a higher grade of materials that assures longer life (slower deterioration rates, or in some cases immunity from deterioration). Similarly, heavier material (e.g., thicker paint) can be used to allow more deterioration to occur before failure occurs, or because harsher conditions are expected. This results in a longer expected period of time, service, or abnormal events before degradation leads to failure or misoperation.

But such measures cost money. Thus, ultimately equipment lifetime is somewhat a function of cost. In general, most electrical equipment like transformers, breakers, and regulators is designed to withstand the chronological, service, and abnormal events associated with about 40 years of service under typical ambient conditions and loading. Overhead conductors and its ancillary equipment generally have a lifetime of 60 years or more (the required mechanical strength pretty much assures immunity from electrical, thermal, and/or rapid deterioration unless abused by very harsh service). Wood poles have about a 40 – 75 year expected lifetime, depending on quality and soil conditions and steel structures typically last up to 100 years. Underground cable often has the shortest

Table 7.3 Effects of Power System Equipment Aging, in Order of Importance

1.	Higher likelihood of failure
2.	Increased likelihood of malfunction
3.	Reduced capacity
4.	Increased losses
5.	Greater noise




Figure 7.2 Example of the impact of “de-rating” a unit in order to increase expected lifetime. See text for details. (Note: This curve is not general; it does not accurately apply to all transformers.)

expected lifetime. Although paper insulated, lead-covered (PILC) cable has an expected lifetime of up to 60 years if well-cared for, most other types of cable have lifetimes between 20 and 35 years.

Worn out equipment usually is really worn out

Usually, all of the types of deterioration discussed above are occurring simultaneously in any electrical device or line. Corrosion, dielectric breakdown, wear, and loss of mechanical strength all proceed at their own paces. Failure of a unit like a breaker or transformer is a type of horse race between the various failure modes that could fail it. The unit will fail when the first of these various types of deterioration reaches the point where it causes misoperation or catastrophic failure (the device destroys itself).

One element of good design for power system equipment is a balance among the expected lifetimes of all these various components. This just makes sense: if a transformer's case is known to fail within 50 years, it is a waste of money to install a core that can easily last 100 years. Good design tries to balance the various strengths and amounts of materials in a unit so that all have roughly the same lifetime in typical service. This means that when failure occurs or is imminent due to one particular type of deterioration, the unit is most likely worn out in all its other deterioration modes, too. It may not be worth repairing because such repair essentially requires rebuilding and or replacing *everything*.

Effects of Equipment Aging

Table 7.3 lists the effects of equipment aging, in rough order of importance to most electric utilities. As units age they have a higher likelihood of failure and misoperation (all equipment). In some cases aging results in a reduced capacity or rating (nearly all equipment, whether recognized or not – this will be discussed in the next section). Some types of devices also see increased electrical or mechanical losses with age (some motors and transformers). In other cases, the device may produce more sound pollution as it may age (motors and transformers, due to core or stator stack loosening, etc.).

By far the most significant factor, and the major concern to a distribution utility, is the top effect listed in Table 7.3: increasing likelihood of failure. The next section discusses this in great detail.

De-rating to Preserve Remaining Lifetime

Table 7.3 lists de-rating as one impact of aging. Since the major modes of deterioration and failure are often heat-related or other service-stress created causes, one option available to an equipment operator is to lower the loading (i.e., cause of stress) in order to increase the expected lifetime of the equipment.

Figure 7.2 shows a de-rating curve for a power transformer that has an expected 40-year lifetime and has already seen 30 years of service at designed load levels. At this point, it has an expected 14 years of service left if kept at its present service level where peak load is 100% of its rating.² If the owner wishes to extend the expected time-to-failure to 40 more years of service, he must reduce loading by 29%, to only a peak of 71% of the unit's rating. Similarly, if this unit is loaded to 110%, it will very likely fail within a few years. Thus, the owner can balance the annual value to be obtained from the unit (how much power it handles) against the number of years that value is expected to be provided by the unit.

7.3 EQUIPMENT FAILURE RATE INCREASES WITH AGE

All equipment put into service will eventually fail. All of it. Some equipment will last longer than others will and often these differences in lifetimes will appear to be of a random nature. Two apparently identical units, manufactured to the same design, from the same materials, on the same day at the same plant, installed by the same utility crew on the same day in what appear to be identical situations, can nonetheless provide very different lifetimes. One might last only 11 years, while the other will still be providing service after 40 years. Minute differences in materials and construction, along with small differences in ambient conditions, in service stress seen by the unit, and in the abnormal events experienced in service, cumulatively lead to big differences in lifetime.

In fact, generally the only time one finds a very tight grouping of lifetimes—i.e., all of the units in a batch fail at about the same time—is when there is some common manufacturing defect or materials problem—a flaw—leading to premature failure.

Predicting Time to Failure

High failure rate and uncertainty make for a costly combination

As a conceptual learning exercise, it is worth considering how valuable exact knowledge of when any particular device would fail would be to a distribution utility. Suppose that it were known with certainty that a particular device would fail at 3:13 PM on July 23. Replacement could be scheduled in a low cost, low impact on customer manner prior to that date. There would be no unscheduled outage and no unanticipated or difficult to manage costs involved: impact on both

² It does not have only ten years of expected service left. This is a “survivor” of 30 years service. As such, statistically, it has a likelihood of lasting more than the average 40 years expected when it was new. Essentially it is making up for the few other units of its type that failed prior to 30 years of service, keeping the average at 40 years service.

customer service quality and utility costs could be minimized.

It is the *uncertainty* in the failure times of power system equipment that creates the high costs, contributes to service quality problems, and makes management of equipment failure so challenging. The magnitude of this problem increases as the equipment ages, because the failure rates increase: there are more “unpredictable” failures occurring. The utility has a larger problem to try to manage.

Failure time prediction: an inexact science

With present technologies, it does not seem possible to predict time-to-failure exactly. In fact, capability in failure prediction for equipment is about the same as it is for human beings.

1. *Time-to-failure can be predicted accurately only over a large population* (set of units). Children born in 2000 in the United States have an expected lifetime of 76 years, with a standard deviation of 11 years. Service transformers put into service in 2000 have an average expected lifetime of 43 years with a standard deviation of 7 years.
2. *Assessment based on time-in-service can be done*, but still leads to information that is accurate only when applied to a large population. Thus, analysts can determine that people who have reached age 50 in year 2000 have an expected 31 years of life remaining. Service transformers that have survived 30 years in service have an average 16 years of service remaining.
3. *Condition assessment can identify different expectations* based on past or existing service conditions, but again this is only accurate for a large population. Smokers who have reached age 50 have only a remaining 22 years of expected lifetime, not 31. Service transformers that have seen 30 years service in high-lightning areas have an average of only 11 years service life remaining, not 16.
4. *Tests can narrow but not eliminate the uncertainty* in failure prediction of individual units. All the medical testing in the world cannot predict with certainty the time of death of an apparently healthy human being, although it can identify flaws that might indicate likelihood for failure. Similarly, testing of a power transformer will identify if it has a “fatal” flaw in it. But if a human being or a power system unit gets a “good bill of health,” it really means that there is no clue as to when the unit will fail, except that that failure does not appear to be imminent.
5. *Time to failure of an individual unit is only easy to predict when failure is imminent*. In cases where failure is due to “natural causes” (i.e., not

due to abnormal events such as being in an auto accident or being hit by lightning), failure can be predicted only a short time prior to failure. At this point, failure is almost certain to advanced stages of detectable deterioration in some key component.

Thus, when rich Uncle Jacob was in his 60s and apparently healthy, neither his relatives nor his doctors knew whether it would be another two years or two decades before he died and his will was probated. Now that he lies on his deathbed with a detectable bad heart, failure within a matter of days is nearly certain. The relatives gather.

Similarly, in the week or two leading up to failure, a power transformer generally will give detectable signs of impending failure: an identifiable acoustic signature will develop, there will be internal gassing, and perhaps detectable changes in leakage current, etc.

6. *Failure prediction and mitigation thus depend on periodic testing as units get older.* Given the above facts, the only way to manage failure is to test older units more periodically than younger units. Men over 50 years of age are urged to have annual physical exams in order that possible failures are detected early enough to treat. Old power transformers have to be inspected periodically in order to detect signs of impending failure in time to repair them. Table 7.4 summarizes the key points about failure time prediction.

Table 7.4 Realities of Power System Equipment Lifetime Prediction

1. Time to failure can be predicted accurately only over large populations.
 2. Past and present service conditions can be used to narrow the expected uncertainty, but only slightly.
 3. Even under the best of conditions, a considerable standard deviation of remaining life and failure rates will exist.
 4. That standard deviation will indicate heteroscedastic behavior – it will increase as age increases.
 5. Testing provides accurate “time to failure” information only when it reveals flaws that mean, “failure in the near future is nearly certain.”
 6. The only way to assure relatively accurate prediction of time to failure is to periodically test a unit as it gets older.
-

Quantitative Analysis of Equipment Failure Probabilities

Figure 7.3 shows the classic “bathtub” lifetime failure rate curve, qualitatively representative of equipment failure likelihood as a function of age for just about all types of devices, both within the power systems industry and without. The curve shows the relative likelihood of failure for a device, in each year over a period of time. Failure likelihood is not constant from year to year.

This curve and its application are probabilistic. As stated above, in a practical sense it is simply impossible to predict accurately when any particular unit will fail. However, over a large population one can predict with near certainty the overall characteristics of failure for a group as a whole. Curves such as Figure 7.4 represent the expected failure of devices as a function of age, over a large population. They can be applied deterministically to equipment only when it is in bulk, or to individual units only as expectations. Thus, based on Figure 7.3, 4 out of 1000 units that have lasted 18 years will fail in that 18th year of service. Similarly, for any unit that has survived 35 years of service, it has a probability of .1 of failing in its next year of service.




Figure 7.3 The traditional bathtub failure rate curve, as discussed in the text. Here, the device has a useful lifetime of about 30 – 35 years.




Figure 7.4 A bathtub curve showing impact of periodic rebuilding and refitting.

Figure 7.3 identifies three different failure periods of the device's lifetime. These are:

1. *Break-in period.* Failure rate for equipment is often higher for new equipment, as shown in Figure 7.3. Flaws in manufacturing, or installation, lead to quick failure of the device once placed in service.
2. *Useful lifetime.* Once the break-in period is past, there is a lengthy period when the device is performing as designed and failure rate is low, during which the design balance of all of its deterioration rates is achieving its goal. During this period, for many types of equipment, the failure rates due to "natural causes" (deterioration due to chronological aging and service stress) may be so low that the major causes of failure are abnormal events.
3. *Wear-out period.* At some point, the failure rate begins to increase, due to the cumulative effect of time, service stress, and abnormal events. From this point forward it increases with time, reaching the high double-digit percentages at some point.

Often periodic maintenance, rebuilding or refitting, can "extend" lifetime or lower expected failure rate, as shown in Figure 7.4. This concept applies only to units designed for or conducive to service (i.e., this applies to breakers, perhaps, but not overhead conductors). Here, the unit is serviced every ten years or so.

Note however, that failure rate still increases over time, just at a lower rate of increase, and that these periodic rebuilds create their own temporary “infant mortality” increases in failure rate.

Failure Rate Always Increases With Age

All available data indicates that inevitably, failure rate increases with age in all types of power system equipment. Figure 7.5 shows data on the failure rate of underground equipment for a utility in the northeast United States, all of which is qualitatively similar to the failure performance seen on all equipment, in any power system.

Failure rate escalation characteristics

Figure 7.6 illustrates how the *rate of increase* of failure rate over time can exhibit different characteristics. In some cases the failure rate increases steadily over time (the plot is basically a straight line). In other cases the rate of increase increases itself – failure rate grows exponentially with a steadily increasing slope. In yet other cases, the rate climbs steeply for a while and then escalation in failure rate decreases – a so-called “S” shaped curve. Regardless, the key factor is that over time, failure rate always increases.

Eventually the failure rates become quite high

Figure 7.7 is data representing a large population in an actual and very representative electric utility aging equipment investigation done in the late 1990s. Note that the failure rates for all three types of equipment shown eventually reach values that indicate failure within 5 years is likely (rates in the 15 – 20% range). In some cases failure rates reach very high levels (80% - failure in the next year or so is almost a certainty. To power system engineers and managers who have not studied failure data, these values seem startlingly high.³ However, these are typical for power system equipment – failure rates *do* reach values of 15%, 25% and eventually, 80%. But these dramatically high values are not significant in practice, as will be discussed below, because few units “live” long enough to become that old. What really impacts a utility is the long period at the end of useful lifetime and beginning of the wear-out period when failure rate begins to rise, slowly at first, but inevitably, to two to five times normal (useful lifetime) rates.

³ In many cases the older age/high failure rate portions of a curve fitted to a utility’s data are not statistically sound. Perhaps there was only one unit or one failure of a unit older than 50 years in all the data. Fitting a curve to that data and then extending it out to 100 years provides questionable estimates of failure rate in the 65+year range. Here, there were sufficient old units, and failures, to provide statistically significant sample of old units and trends which their failure rate follows in this time frame.




Figure 7.5 Data on failure rates as a function of age, for various types of underground equipment in a utility system in the northeastern United States. Equipment age in this case means, “time in service.” Equipment of different voltage classes can have radically different failure characteristics, but in all cases failure rate increases with age.

7.4 IMPACT OF ESCALATING FAILURE RATES

What Do Aging Utility Equipment Bases Really Look Like?

The paragraph above stated that the very high failure rates reached after 50 years or more (see Figure 7.7) make relatively little impact on a utility's cost or service quality. The reason is that very few units survive to be so old, so there are few units that have such very high failure rates.

Example 1: A Typical Failure Rate Escalation and Its Impact on the Installed Equipment Base

Figure 7.6 illustrates a very simple example that will begin this section's quantitative examination of failure, installed base characteristics, and overall impact on the utility. In this example, the group of 100,000 service transformers is installed at one instant, a rather unrealistic assumption but one that has no impact on the conclusions that this example will draw.

As a group, this set of 100,000 units has the statistical failure rate characteristic shown in the top part of Figure 7.6. That plot gives the probability that an operating unit of any particular age will fail in the next 12 months of service. In this case there is no high break-in-period (infant mortality) failure rate. The base rate during normal lifetime begins at 1.5% per year, rising to 2.5% by year 24, 6.6% by year 30, and to 9% annually by age 40. This curve is representative of actual service transformer failure rate curves (see bottom of Figure 7.6), if noticeably more extreme than is typical for electric utility equipment.

The bottom diagram in Figure 7.6 shows, as a function of age, the percent of the 100,000 units installed in year zero that can be expected to remain in service each year as units fail according to the expectation defined by the top curve. In year one, 1.5% of the units' fail, meaning 99% are in service at the beginning of year two. At the end of a decade, 85% are still in service.

The failure rate is initially 1.5%, increasing slightly above that value each year. Despite this rise, only 15% of the units (ten times 1.5%) fail in the first decade. The reason is that the number of initial units left to fail decreases each year – there are only 98,500 units left at the end of the first year, etc., so 1.5% is not exactly 1,500. The number of actual failures decreases slightly each year for the first decade, to a low of only 1,440 failures in year ten, because the number of units remaining to fail drops faster than the failure rate increases.

At the end of 20 years, 71% of the units remain, and at the end of thirty, only 53% remain. The 50% mark is reached at 32 years. By the end of year 50, failure rate has escalated to more than 15%, but only a paltry 10.3% of the original units remain. Only 500 (.7%) make it to year 60. Fewer than two are expected to be in service by year 70, when failure rate has escalated to 50%. The average unit ends up providing 43.8 years of service before failure.




Figure 7.6 Top, failure rates as a function of age for a group of 100,000 service transformers installed in one year. Bottom, the percent remaining of the original group as a function of time.




Figure 7.7 Number of failures occurring each year in the population that originally numbered 100,000 units.

Failures in intermediate years are the real culprit

The very high failure rates that develop after five or six decades of service make little real impact on the utility's quality of service. Units that are 70 years old have a 50 percent likelihood of failing in the next year, but as shown above, there are only two units out of every 100,000 that make it to that age – essentially an anomaly in the system. They are not going to make a big impact on utility operations.

Instead, the high impact failure levels that would plague a utility operating this equipment would be caused by failures at intermediate age, even though the failure rate in that age range is far lower. Figure 7.7 shows the number of units that fail each year in this example. As mentioned earlier, the number of failures is initially 1,500 units per year and the count actually *drops* slightly during the first decade as the product (failure rate x number of units remaining) decreases slightly. The number of failures peaks in year 44, with 2,576 expected failures. Thereafter, even though the failure rate keeps increasing every year, the number of remaining units this rate applies to is so low, that the net number of failures decreases rapidly.

Failure-Count Diagrams

Figure 7.7 is a very interesting and significant diagram, because it has another interpretation or use: this is a plot of contribution to failure by age of transformer for any utility system built with this type of unit, regardless of when they were installed. The curve's shape shows, in relative terms, how much transformers of a particular age contribute to the overall installed base's failure rate.

In the example used above, all the transformers were installed in the same year, but Figure 7.7 can be applied to cases where a utility installed them in various years. It still has the same meaning, showing the relative contribution to failures of units of various ages. The majority of the failures, nearly two thirds, will be units that are between 15 and 45 years old.

Thus, in this system, with these units, in this year, transformers installed 44 years ago contribute the most to the utility's annual failure count for service transformers. *Given only that the utility added the same number of units each year in the past, this will always be the case.* Ten years from now 44-year-old units will be those giving it the most problems, then, only assuming that, 34 years earlier, the same number were installed as 44 years ago. In a system that adds a constant number of units annually in the past, it will *always* be 44-year-old units that are causing the most problems.

Scenarios involving replacement

In the example given above, units that fail are not replaced. This is not representative of an actual utility situation. In the operation of any utility system, failed units are replaced, the replacements are newer, and thus a more complicated situation will develop. In a "failures are replaced" scenario, in year one, the 1,500 units that fail are replaced with new units so the system of 100,000 goes into the following year (number two year) with most two-year-old units but 1,500 one-year-old units. The following year, a portion of the two-year-old and one-year-old units fail and are replaced with new, etc. Over time, a range of ages from a few new to many old will develop, even though initially all were installed at the same time. Add to this the fact that in real utility systems the units are not installed all at once but over time, and the distribution of ages and the failure-rate over time trends become quite complicated to track and predict. This will be the subject of later chapters and the model discussed in Appendix A. Regardless, the point here is that no matter when they are installed, really old equipment is usually not the equipment causing the most concern to the utility. It will be equipment in "middle age," which might be usefully defined as in that period when

$$(\text{number that have survived until now}) \times (\text{failure rate for this age})$$

is at a maximum, that matter most. This is significant, for it means that aging infrastructure programs should focus on reducing failure rate in this period rather than for older equipment: *life extension programs should begin earlier, not later.*

Example 2: A More “Real World” Case

As noted above, Example 1 did not provide for replacement of units when they fail, but tracking only the survivors of that initial batch of 100,000 units. Example 2 builds on Example 1, using the same type of service transformers. It:

1. Assumes, as before, 100,000 units installed in year 0.
2. Assumes a period of 70 years.
3. Assumes the failed units are replaced immediately with new ones, keeping the overall count at 100,000.
4. Assumes these new units follow the same failure rate curve as the original units.

This is a more realistic example, as it represents what a utility has to do – keep the same number of units in service, replacing failed units when they fail.⁴ The major change in the overall equipment base and its interaction with failures in this example is that *replacement units can fail, too*.

What does this installed equipment base look like with respect to age distribution of units in service, average failure rate, and failure count? Figure 7.8 shows the distribution of ages of the 100,000 units in the system after 70 years. In year 1 (the first year) 1,500 units failed, and were replaced with new units which are now 69 years old. Those 1,500 units failed with the same characteristic trend as the original set, meaning that 1.5-%, or 23, failed and were replaced in their first year of service (year 68). In that year, 1,489 of the original 100,000 units failed (1.5% of the 998,500 remaining original units). Thus, a total of 1,512 replacement units were installed in year 2. They began to fail along the same trend as the original units, so that in year 3 there were failures of units that had been installed in years 0, 1, and 2, etc.

The net result, when all of these replacements, and replacements for replacements that failed, etc., are added up, is that the utility had to install about 60,000 additional replacement units during the 70-year-period. And, following the failure count contributions of Figure 7.7, most of those replaced were of “middle aged transformers” – those in the 15 –45 year old range. Thus, Figure 7.8 gives the relative age of units replaced by the utility over the course of its annual O&M in each year.

⁴ The fact that this example assumes the system is created from scratch in year zero makes no impact on the results given here. As shown earlier, by year 70 only 2 of the original units are left. The population consists of units that have been replaced, and in some cases, failed and replaced again. As a result after 70 years there is only an insignificant “start effect” involved in the data – the model’s results in year 70 represent a fairly stable look at what year-to-year operation for the utility would be with respect to service transformers.




Figure 7.8 Distribution of ages of units in the Example 2 system, 100,000 units that have been replaced as they failed over the last 70 years.

Figure 7.8 shows the resulting equipment bases distribution of transformer ages, after 70 years. It has nearly an even distribution of transformers from age 0 (new) to 30 years. Thus, despite the escalating failure rate as transformers near age 30, 30 years can be considered a reasonable definition of the “useful life” of these units. At about 35 years the count takes a rapid plunge – this is the period (about 35 years in service) during which the bulk of failure counts occur (see Figure 7.4), and thus the age when a good deal of replacements had to be made.

In this system, the average unit in service is 22 years old. However, due to the escalation of failure rates as units age, those older than the average contribute a good deal more to the average failure rate. The entire population has an average failure rate of 3.15%, or more than twice that of new units. This corresponds to the failure rate of a unit that is 29 years old (see Figure 7.8, top).

Replacement policy analysis

Suppose this utility decided to replace all units when they reached fifty years in service. It would have to replace only about 75 units annually – not a tremendous cost, particularly considering the units will have to be replaced pretty soon, anyway (they will most likely fail in a few years, at that age). However, the impact on the overall failure rate is insignificant, making no real difference in the average age or failure rates for the total population. The high impact on annual

failure count comes from units that are aged 25 to 45 years, because there are so many of them. Replacement at age 50 gets to the units after too many have failed.

Replacement of units at age 40 has a far different effect. First, nearly 1000 units a year have to be replaced, so the annual cost is roughly 12 times that of a 50-year replacement policy. But, noticeable portions of unexpected failures are avoided. The average failure rate drops to 2.6% (from 3.1%), a reduction in unexpected failures of nearly 20%, wrought by replacement on only 1% of the units in the system annually.

Average age of a unit under this policy falls to less than 18 years (if all units were replaced at age 40 and none failed until that age, the average age of units in service would be exactly 20 years). The population's average failure rate of 2.6% is equivalent to the failure rate of a 25-year-old unit. Whether this policy makes economic sense or not is something for management to weigh in its evaluation of how best to spend the (always-limited) monies it has for system improvements. That will be discussed in Chapter 12.

7.5 SUMMARY OF KEY POINTS

All equipment installed in an electric system will eventually fail and need to be replaced. Key points brought up in this chapter are:

1. Failure is caused by deterioration, due to the effects of age, service, and infrequent but stressful abnormal events.
2. This deterioration eventually leads to failure.
3. Due to efficient design of most equipment, a unit that is near failure from one area of deterioration is probably fairly "shot" in other areas, too.
4. Regardless of equipment type, the failure increases, as a unit grows older.
5. For an installed base of equipment, the bulk of failures will come not from very old equipment but from "middle aged" equipment.
6. In all cases, due to failure rate escalation with age, the equivalent failure rate of an installed equipment base will correspond to a much older age than the average age of the equipment.
7. "Early replacement" policies can be worked out to determine if and how equipment should be replaced at some specific age (e.g., 40 years) and what that impact that would have on the net failure rate of the equipment base.

REFERENCES AND BIBLIOGRAPHY

- P. F. Albrecht and H. E. Campbell, "Reliability Analysis of Distribution Equipment Failure Data," EEI T&D Committee Meeting, January 20, 1972.
- R. E. Brown, and J. R. Ochoa, "Distribution System Reliability: Default Data and Model Validation," paper presented at the 1997 IEEE Power Engineering Society Summer Meeting, Berlin.
- J. B. Bunch, H.I. Stalder, and J.T. Tengdin, "Reliability Considerations for Distribution Automation Equipment," *IEEE Transactions on Power Apparatus and Systems*, PAS-102, November 1983, pp. 2656 - 2664.
- "Guide for Reliability Measurement and Data Collection," EEI Transmission and Distribution Committee, October 1971, Edison Electric Institute, New York.
- Institute of Electrical and Electronics Engineers, *Recommended Practice for Design of Reliable Industrial and Commercial Power Systems*, The Institute of Electrical and Electronics Engineers, Inc., New York, 1990.
- A. D. Patton, "Determination and Analysis of Data for Reliability Studies," *IEEE Transactions on Power Apparatus and Systems*, PAS-87, January 1968.
- N. S. Rau, "Probabilistic Methods Applied to Value-Based Planning," *IEEE Transactions on Power Systems*, November 1994, pp. 4082 - 4088.
- A. J. Walker, "The Degradation of the Reliability of Transmission and Distribution Systems During Construction Outages," Int. Conf. on Power Supply Systems. IEEE Conf. Publ. 225, January 1983, pp. 112 - 118.
- H. B. White, "A Practical Approach to Reliability Design," *IEEE Transactions on Power Apparatus and Systems*, PAS-104, November 1985, pp. 2739 - 2747.

8

Obsolete System Structures

8.1 INTRODUCTION

This chapter looks at the impact that “making do” with limited numbers and sizes of substation sites and sub-transmission rights of way has on the reliability and operation of a power delivery system. As discussed in Chapter 1’s introduction to aging T&D infrastructures, the second of the four interacting aspects of the aging infrastructure issue is the limited facilities which power distribution utilities invariably have in the core of thriving metropolitan areas. Difficulties in getting new sites, in expanding old sites, and in getting rights of way often mean that the utility’s system can not expand in company with the demand of an ever-growing downtown area.

This chapter looks at this problem and the effects that it has on power system performance. Section 8.2 presents the basic problem with some examples. Section 8.3 then examines the impacts that “outdated” or limited system structures have on the engineering decisions that planners must make with regard to the sub-transmission and substation levels of the system, and on the operating characteristics and daily constraints faced by the people managing that part of the system. All of these effects compromise the system’s reliability. Section 8.4 looks at similar impacts on the primary feeder system, which degrade its potential to contribute economy of operation and reliability to the system. Section 8.5 summarizes the various solutions and improvements recommended for the distribution utility faced with this problem. Key points are summarized in Section 8.6.

8.2 OBSOLETE SYSTEM LAYOUTS

Example System

Figures 8.1 and 8.2 depict the plight of a fictional electric distribution utility, Metropolitan Power and Light, which represents the average experience of eight major metropolitan utilities on the east and west coasts and the Midwest United States.

In the period 1965 to 2000, Metropolitan Power and Light saw the peak demand in its central operating district go from 1380 MVA to 3058 MVA, an average annual growth rate of 2.3%. But while the peak load and demand for quality of service have increased greatly, the basic topology of its system did not keep pace. The original (1965) system, shown in Figure 8.1, consisted of 16 substations in the central operating district, with an average capacity of 128 MVA each, for a total installed substation capacity of 2048 MVA. This resulted in a utilization ratio of 67% (1380 MVA/2048 MVA). The majority of these substations were originally fed by 69 kV transmission loops (two, or in some cases, two sets of two), lines along the same right of way, providing a parallel and redundant (contingency backup) feed to each substations.

During the intervening 35 years, Metropolitan added one new substation site and expanded the size of four others. It also upgraded a number of 69 kV sub-transmission lines to 138 kV and completed one “loop” transmission route by adding a right of way to connect two substations (Figure 8.2). In addition, it increased the capacity of the 138 kV circuit lines in the backbone right of way through downtown by re-conductoring and bundling of conductor, and by upgrading part of the route to 345 kV.

Table 8.1 Comparison of Metropolitan Power and Light
Downtown Systems in 1965 and 2000

System Measure	System in 1965	System in 2000	2000 as a % of 1965
Peak load	1380	3058	222%
Number of substations	16	17	107%
Substation capacity - MVA	2048	3400	166%
Peak load/substation - MVA	86	180	209%
Peak load/sub-trans. - MVA	52	101	194%
Number of feeders	261	386	147%
Load/feeder	5.2	7.9	150%
Substation utilization at peak	67%	90%	134%
Feeder utilization at peak	57%	74%	128%
MW-miles (x 1000)	2150	4967	230%




Figure 8.1 In 1965, the central operating district of Metropolitan Power and Light used 16 substations averaging 128 MVA capacity each (total 2,048 MVA capacity) to serve a peak downtown load of 1,380 MVA. This diagram shows only the overall structure of the system. Many of the apparently radial lines feeding substations are in fact loops (two circuits) or double loops (four circuits).




Figure 8.2 (2000) Thirty-five years of slow but steady load growth result in a year-2000 peak demand of 3,058 MVA, more than double 1965's. In that time, Metropolitan Power and Light has added only one substation site and expanded the area of four others, as well as upgraded various 69 kV lines to 138 kV lines. Total substation capacity now stands at 3,400 MVA, (200 MVA each, a 56% increase), for a utilization ratio of 90%.

However, this system growth did not keep pace with peak load growth. Year 2000 capacity is 3,400 MVA, an average of exactly 200 MVA per substation, giving a utilization ratio against the 3,058 MVA peak load of 90%. Utilization ratio of the transmission lines has similarly increased. Equally indicative are the facilities ratios, load/substation and load/transmission route, as shown in Table 8.1. Metropolitan Power and Light has a lot of eggs in very few baskets.

At first glance it might appear that the system in Figure 8.2 is fully as capable of providing high reliability of service as the system in Figure 8.1. In actuality, the system is more vulnerable to certain types of outages at the sub-transmission and sub-station levels. This increased vulnerability is due to the following factors: 1) higher utilization ratio; 2) “a lot of eggs in one basket and; 3) design choices dictated by limited options for facilities.

Why Did This Happen?

Before turning to a discussion of the impacts these changes had on the cost and reliability of electric service, it is worth looking at why Metropolitan got into this situation. Over a period of three and one half decades, as its peak demand grew, its system did not keep pace in some ways, out of both choice and constraint. This did not occur because the utility planners did not forecast the growth: for the most part they saw it coming. It happened because, incrementally, the utility allowed itself to be backed into the situation. The major reasons were:

No sites available

In 1965, the downtown area served by Metropolitan Power and Light was already fully developed, in the sense that there was no vacant, never-built-upon land. However, much of it was not what would, 35 years later, be considered “downtown development.” There were areas of vacant warehouses, some run down old hotels, a number of gravel-covered parking lots scattered about, an old rail yard about to be redeveloped, and other areas of downtown where land was available, if rather pricey by standards of the time.

Over the next 35 years, this price and availability situation gradually escalated and seemed to stay one step ahead of Metropolitan. As the city and the economic region around it grew the vacant warehouses, old parking lots, available rail yard, and other developable sites in the downtown area gave way to new development, and the various sites that could have become new substations disappeared. Simultaneously, the price of real estate climbed. By year 2000, very few sites suitable for construction of a large capacity (nothing else will help, given the load levels) substation built upon traditional lines exist period. Those that do are very expensive, not just in terms of money, but in time and political capital that the utility might have to expend in order to obtain the site and all permissions and licenses required to built on it.

Political pressure

Any utility has to be cognizant of local community values and try to accommodate the wishes and preferences of local municipal leaders. Metropolitan Power and Light often found itself facing tremendous political opposition, both in city hall and at a local grass roots level, when it attempted to obtain new substation sites or rights of way for sub-transmission lines. Most of this opposition was directed at the overhead line rights of way it preferred to build to any new site it obtained, or at the inconvenience and traffic tie ups that would be caused by the construction required to build UG transmission lines in the central city. In combination with the often-perceived high costs of new sites, this political pressure blocked most moves to obtain new sites.

What really failed, in the authors' opinion, was communication between utility and municipal authorities on a key point: the utility can easily do without any one of the sites or rights of way it requests, but in the long run it cannot make do if it is denied them all. In this respect, the utility had the responsibility to communicate this issue well to municipal leaders and community alike. It failed to do so. As a result, the incremental, repetitive process of overwhelming opposition on a site-by-site basis led, after three and one half decades, to a poor-performing system.

Financial pressure

Cost was always a major element of every decision made by Metropolitan Power and Light. As a regulated utility, it had an obligation to manage its affairs in a least-cost manner. Management throughout the 35-year period was never incompetent or inattentive; it wanted to make the best decisions regarding cost and system efficiency.

But a new substation site and the new transmission line right of way and the construction required to put those facilities into service always represented a very large budget item. As such, every go/no-go decision regarding a new substation received a great deal of scrutiny. Generally, Metropolitan Power and Light's management was always pre-disposed to find some other way to bolster the system, some other change or upgrade that would avoid, or defer for a few years, the need for a very expensive new site.

In addition, in the late 1980s and throughout the 1990s, Metropolitan Power and Light, like most large investor-owned utilities, began to respond to the uncertainty created by deregulation, and increasing pressures to raise financial performance, by reducing operating and capital budgets by up to 30%. An executive goal was to see utilization ratio on the system rise – "to get more out of existing equipment."

This led on several occasions to changes in loading policy, re-rating of transformers, cables, and overhead lines, and adoption of slightly less

conservative rules with regard to available contingency margin and operating policy. Not building substations fit right into this push for budget reduction, because it avoided both large outlays of capital, and construction of new facilities whose maintenance would have swelled the O&M budget. In addition, the re-rating of equipment mentioned in the paragraph above “created” a good deal of new capacity at many existing sites, permitting the utility to keep pace with the load growth without adding new facilities, and still meet its engineering criteria.

No additional room available at existing sites

When it was not allowed to gain more sites, Metropolitan Power and Light had no choice but to serve the growing load from the sites it had. Load/site naturally had to go up, and it did, from 86 MVA in 1965 to 200 MVA/substations in year 2000. *But this does not imply that utilization ratio had to go up.* Metropolitan could have kept utilization ratio constant at 67% by increasing installed capacity at its 17 downtown substations to an average 270 MVA each. It did not do so for a variety of reasons, increasing it only to an average of 200 MVA, and letting peak utilization rate of its downtown substations rise to 90%.

Metropolitan Power and Light made its decision to let utilization ratios rise as a result of several mutually reinforcing trends. First, land adjacent to existing sites, for expansion, was simply not available in many cases, and its price when it was for sale was astronomical compared to what the utility was willing to pay. Secondly, the lengthy period for and uncertainty involved with negotiation and approvals for such special sites made trying to get them risky.

Thirdly, in the late 1980s and early 1990s, Metropolitan’s executive management decided to cut back on capital spending. It wanted to reduce budgets and try to “get more from existing equipment.” It had an inclination to let utilization ratios rise, and the situation in the central operating district played right to that goal. Perhaps Metropolitan would not have let the situation become so extreme had more expansion space been more available, but the fact is that because of its budget-reducing focus it did not fight for expansion room as hard as it otherwise might have.

These effects limited the number of situations where the utility added space to existing sites (it expanded only four of its sixteen original sites during the 35-year period). Thus, while the load grew by 122% during the thirty-five year period, Metropolitan increased capacity of its substations only to an average of 200 MVA/substation, up from 128 MVA, a 56% increase. Section 8.3 will show how this was even more damaging to reliability potential than the lack of new sites.

Other Options

Every time that Metropolitan Power and Light needed a new site, there were always other options available to the utility. Among them were reinforcement re-design and re-configuration of existing facilities, and other means that could

provide capacity and contingency capability. Many of these were only “stop-gap” measures to defer the need for new sites for a few years, but in every case they had a lower first cost and met the existing criteria. While a very long-term view would have shown that a few new sites were necessary, the paradigm at work led Metropolitan to take other routes in nearly every case.

A contributing factor to these decisions was the characteristics of the traditional N-1 planning criteria and tools Metropolitan’s planners used to apply it, as will be discussed in Chapter 8. The utility planners’ tools did not alert them to the weaknesses that were developing in their system. The limitations of the “outdated system structure” they were evolving were not completely visible to them. As a result they were expecting better performance from the system than it could actually deliver.

Incremental Mistakes and Boiled Frogs

Traditional anecdotal wisdom holds that if one throws a frog into a bowl of very hot water it will immediately hop out. But, if that same frog is placed in tepid water, which is then slowly heated over a fire, it will float in the water until it boils to death. The authors have not verified that this is true (it seems unusually cruel to the frog and the authors have no desire to have a boiled frog on their hands). However, it is a good analogy to the way many utilities, including Metropolitan, gradually backed into a situation that they never would have accepted had they foreseen it. Certainly the situation with regard to substation sites was very gradual, with no one incremental event ever so dramatic as to trigger concern.

Another example of this is the gradual erosion of feeder system strength seen by many utilities in the 1980s and 1990s. Section 8.4 will discuss in detail the impact that the period 1965 to 2000 saw with respect to Metropolitan Power and Light’s primary feeder system. As the load grew and the number of substation sites did not keep pace, the feeder system had a gradually increasing burden placed on it. From 1965 to 2000, the load it had to distribute out of every substation grew by 109% (86 MVA in 1965, 200 per substation in 2000). At times this necessitated adding new feeders – Table 8.1 shows that the number of feeders grew by 47%. But to a large extent the load growth was accommodated by allowing the loading (utilization rate on feeders) to escalate.

Table 8.1 showed that average feeder loading increased by 50% during the period; an annual compounded growth rate of 1%. Such a small annual increase is easy to lose among the annual fluctuations in loading due to weather and load transfers, and never enough to raise alarm. This gradual increase in utilization of feeders eventually degraded much of the feeder system’s potential to provide high quality customer service, as will be discussed in section 8.4.

Metropolitan Power and Light’s inability to see fully the impact of increased feeder loading on its system reliability was, as in the cases cited earlier in this

section, partly a failure of its planning tools and engineering paradigm. Traditional distribution planning tools focused on “feeder at a time” short range analysis, techniques which cannot fully assess the feeder *system* impacts that increasing loadings and reduced contingency reach have on the emergency-backup capabilities among feeders.¹ These tools fit the engineering and planning paradigm for distribution planning, which was to examine feeders in detail, but usually only on a one-at-a-time or small-area basis. Feeder system studies involving all the feeders in a region, modeled in company with the substations and sub-transmission, were never done. Therefore, Metropolitan's planners never realized that the degradation of feeder-level reliability, amounting to about a 25% increase in contribution to SAIDI, was occurring. Intuition and judgment alone did not alert the planners to the problem largely because of the very slow rate of change.

Poor Planning was a Major Contributor

The discussions above twice cited the limitations of traditional planning tools as a contributing factor in Metropolitan Power and Light's getting into the outdated system structure situation without fully appreciating its implications with respect to reliability of service. Also cited were Metropolitan Power and Light's lack of effectively applying long-range planning – of using too short-term a focus and not looking at the long-term implications and costs of a de facto policy of constant deferrals and opposition to any major projects.

A longer-range view of its system, even without modern planning tools, would have revealed that problems were going to occur. Even without the best planning tools, such a perspective would have drawn attention to the issue, at least qualitatively. In addition, a longer planning horizon would have permitted Metropolitan Power and Light to identify potential substation sites farther ahead and perhaps obtain them when they were affordable (or at least, less pricey).

Thus, ineffective planning – done with tools of limited capability, and on too short a timeframe – was a contributing factor in Metropolitan Power and Light getting into the situation that it found itself in as the new century began.

8.3 IMPACTS ON THE SUB-TRANSMISSION – SUBSTATION LEVEL

Section 8.2's changes in system structure versus load level have a number of profound impacts on the performance of the electric system, not all of which can be appreciated by reference to utilization ratios or facilities counts alone.

¹ For a comprehensive discussion of these issues see the *Power Distribution Planning Reference Book*, Chapters 13 - 15 (Willis, 2004).

Site Congestion

Metropolitan Power and Light was unable to get either as many new substation sites as it needed, or to obtain expansion space it could have used at many of its existing substations. As such, many of the substations in the central operating district became congested, in the sense that space was completely used, and a lack of available space was the limiting factor on installed equipment at most sites. This lack of space caused two interrelated design compromises.

Forced to chose capacity over configuration

Figure 8.3 illustrates a common problem and its solution for many utilities. An existing substation site must be expanded in capacity. While some additional room is available, it is insufficient. Faced with the lack of room for transformers (capacity) Metropolitan's planners made the change shown at the right of Figure 8.3. Some breakers and buswork were eliminated, reducing switching flexibility, but freeing room for more transformer capacity. This permitted Metropolitan to double the installed capacity at the substation.²

Metropolitan traded switching flexibility during contingencies for more capacity. Lines in the system at the right are "hard wired" to each transformer (UG cable trunk connected). If *either* a transformer or a sub-transmission line fails, *both* have to be taken out of service. The failure rate of both is increased to the sum of the failure rates for both. Although this design may meet traditional N-1-planning criteria, reliability has suffered. The expectation of this outage occurring has increased, and the flexibility of switching when it occurs is more limited. (Chapter 8 will discuss this in much greater detail.)

Measures like those shown in Figure 8.3 were common in the evolution of many large metropolitan utility systems throughout the last quarter of the 20th century. Often, recognizing the need for capacity, the utility built the substation originally with a compromise in favor of capacity over configuration. A situation almost exactly like that depicted in Figure 8.3 was a contributory factor in the outages that plagued ComEd of Chicago during the summer of 1999.

Figure 8.3 shows only one possible way that configuration can be compromised in favor of capacity. It is not meant to represent in detail the exact way congestion problems were solved by any particular utility. But it represents fairly the type of compromises forced on utilities by crowded conditions in central-system substations and rights of way. Numerous other types of design compromises are possible, all sacrificing some or all configuration flexibility for capacity increases. The point is: lack of new sites and expansion room at existing sites forces utilities to deal with load growth by making sacrifices which always

² Breakers at the "sending end" of the transmission lines feeding this substation now provide protection for faults on the high side at this substation.




Figure 8.3 An existing site (left) needs to be doubled in capacity. Only a limited amount of expansion space (middle) can be obtained. The only viable solution is to delete high side buses and breakers and simplify the substation's configuration, so more capacity can be installed (right). The utility chose capacity over configuration.

have consequences similar to the changes shown in Figure 8.4. Capacity increases, but interconnected strength and contingency support capability decreases, and with it the reliability potential of the system.

Higher Equipment Utilization

In spite of the efforts to often choose capacity over configuration, during the last third of the 20th century capacity growth did not keep pace with peak load growth in the Metropolitan Power and Light example given here, or in most large metropolitan systems in the United States. A primary effect of the changes in this case was the increase in utilization ratio, from a ratio of peak load to capacity for substations of 67% in 1965, to 90% in 2000. Again, as discussed in detail in section 8.2, this occurred primarily because the utility could not get enough new sites, or enough room at existing sites, to provide all the capacity it might like. A contributing factor was management's commitment to reduced capital spending, which favored increasing utilization of existing equipment anyway.

Regardless, an increase in utilization ratio has a tremendous negative impact on expected reliability of the system, one, which can be completely mitigated by a combination of artful engineering changes to the structure of the system (interconnection and configuration of switching) and operating policies. However, if not accommodated with such changes the increased utilization ratio will typically result in degraded reliability of service to customers.

In 1965, each substation transformer or line was loaded during peak periods to about 66% of its rating. This meant that if a substation transformer, or bus, or breaker fails, that 67% loading can be distributed over two neighbors (of equal capacity) who each take half, raising their loading during the contingency from the normal design peak of 67% to 100%.

If brief overloads of equipment to 133% of rating are accepted during a contingency, the failed unit's load can be fully supported by only one neighboring unit. Peak load occurs only a few hours per year (typically about 350) so this high level of stress is expected to occur rarely (only when peak load *and* a contingency occur). Off peak, the outage of even two units could often be supported by temporarily overloading only one remaining unit of equipment.

But in a power system with 90% loading, during peak load conditions a failed transformer or line now needs *nine* neighboring units, not two, to support its outage if contingency loading is limited to 100%. (The 90% loading on the unit has to be distributed among nine neighbors, who each have only 10% capacity margin to aid the situation.)

If equipment overloads are accepted during contingencies, the situation is still worse than before. At a 135% overload, two neighboring units are required to support the outage of the failed unit, not one as before. Table 8.2 shows required "contingency support count" for various combinations of utility utilization ratio (designed peak load/capability ratio) and contingency overload limit. It plots the number of units required for contingency support as a function of utilization ratio, for systems composed of identical capacity units.

High utilization ratios mean more interconnection is needed

Assuming those neighbors required for contingency support are available, they will only provide the required support if they can be connected to the load of the failed unit during the contingency. Thus, in a system loaded to 67% at peak, and permitting 133% overloads during contingencies, every substation transformer, bus, and breaker needs contingency interconnection to only one neighboring substation-bus-breaker combination in order for its load to be served even if it is out of service. But in a system designed to a 90% utilization criterion, even if contingency overloading is raised slightly to 135%, interconnection to two neighbors is required.

A system with a high equipment utilization ratio relies on interconnection capability to provide the measure of the contingency withstand strength not being provided by a healthy capacity margin.

Although it is a generalization and there are rare exceptions, a utility desiring to maintain reliability of service at traditional levels must increase the "strength" of its system structure. It must provide both more static and switchable paths to

Table 8.2 Contingency Support Count – Number of Neighboring Units Required to “Pick Up” the Load of an Outaged Unit (All Units Same Capacity)

Utilization Ratio	If Neighboring Units Are Permitted to Overload To This Level (%)					
	100%	112%	125%	133%	150%	166%
67%	2	2	2	1	1	1
75%	3	3	2	2	1	1
83%	5	3	2	2	2	1
90%	9	5	3	2+	2	2
100%	-	9	4	4	2	2




Figure 8.4 Number of contingency support units required to pick up the load of an outaged unit, as a function of designed utilization ratio for the system equipment (i.e., all units).

neighboring equipment (as the case may warrant), and make certain there is sufficient capacity in those interconnections to handle the contingency loading. Ironically, the trends covered in Section 8.2 and the substation space problems depicted in Figure 8.3, forced utilities to do exactly the opposite. Configuration was often reduced as utilization ratio was increased, for two reasons. First, a lack of growth in overall system structure (Figures 8.1 and 8.2) occurred simultaneously with the forced rise in utilization rate. Configuration of the system layout overall did not keep pace with load growth.

Second, this capacity-over-configuration approach often compromises the interconnection strength that had been present. In the past, some utilities have failed to recognize the seriousness of the problem they are backing into, due to the limitations of traditional contingency-based design methods they are using to evaluate the situation.

More frequent and higher periods of equipment stress

Another point worth considering is the impact that the high utilization rate has on the frequency and degree of emergency overload for equipment in the system. In a system with a utilization ratio of 67% and contingency limits of 133%, loading one neighbor of a failed unit to 133% can cover every equipment outage.

In a system loaded to 90%, every failed unit requires two neighboring units to be loaded to 135% of rating. Assuming for the moment that the number of equipment failures is the same in both systems (it would probably be higher in the 90% loaded system), this means that twice as many units are exposed to overloads in the 90% system as in the 67% system. Viewed from the standpoint of the system as a whole, twice as many units are overloaded in any given period of time, and loss of life due to emergency use of units has at least doubled.

Despite this, many utilities responded to the pressures discussed earlier by raising the contingency loading limits for their equipment. For example, one utility in the central United States raised its limits from 115% to 166%. This move was necessitated by the lack of system structure to provide the increased “contingency interconnectiveness” of the system discussed above. The only remaining way to cover the contingencies of many units of its equipment was to accept loading of some equipment to 166% during contingencies.

Chapter 7 discussed equipment loss of life as it remains in service, and the fact that overloads greatly accelerate the loss of life process, so that a unit loaded for several hours at 133% might lose the equivalent of several days or even weeks of normal service lifetime. The change from 67% to 90% loading, a 34% increase, doubles the number of “high stress events” for units like transformers and tie lines. If a utility responds to the increased utilization ratios it desires by increasing contingency loading limits so existing equipment is used at higher loading during contingencies, equipment lifetimes and failure rates suffer.

The result can have a major impact on system performance

Evidence from several large utility systems indicates that the effects covered in this section can cumulatively double the SAIDI contribution made by the substation and sub-transmission levels of the system, producing an increase of 15% - 20% in total system reliability indices.

Utilization Ratio Is Not the Culprit

High utilization ratio, per se, is not necessarily an undesirable element in power systems. Reliable and economical power systems with peak utilization ratios up to 100% can be designed and operated. However, the rise in utilization ratio must be accompanied by changes in the system's ability to isolate outaged equipment and interconnect during emergencies to available contingency margin. The requirement for interconnection capability and flexibility rises along with system utilization ratio. Configuration must in some sense provide the strength that capacity margin does not. Increasing the utilization ratio without having sufficient configuration flexibility to accommodate it, is a mistake that will lead to deteriorated reliability of service.

8.4 FEEDER SYSTEM IMPACTS

This section examines the impacts the changes discussed in Section 8.2 make on the distribution utility's primary voltage feeder system. It first provides a background discussion of feeder systems and reliability, then shows how section 8.2's factors result in a general weakening of the feeder systems ability to provide reliability.

A well-designed and operated feeder system provides reliability of service to energy consumers in two ways, both of which are quite important in maintaining low SAIDI and SAIFI scores:

1. *Feeder level reliability.* It provides inherently good feeder level reliability by having immunity from outages of feeder system equipment, through switching capability combined with sufficient contingency margin to support reasonably expected outages of equipment in the feeder system itself.
2. *Substation – sub-transmission contingency support.* A strong feeder system can provide a measure of contingency support for failures of buses, breakers, and transformers at the substations, via load transfers away from the outaged equipment to other feed sources. This materially aids good reliability of service and is often the only viable approach to provide high reliability when utilization rates are high.




Figure 8.5 Allocation of SAIDI by level of the system for a large investor owned utility in Texas. Results here are based on the average of three years, 1996 – 1998.

Outdated system structure issues impact both areas. This section will examine in turn the impact that outdated system structures make on each of these two reliability aspects of feeder system performance.

Feeder Level Reliability

The primary voltage feeder system plays an essential role in the power system chain, taking the power from the substations to the vicinity of all customers. As such, its ability to perform its function in a reliable manner – to recover from outages of circuit elements and equipment – is essential to a good reliability record for the utility. Aging equipment, cables, and conductors in the feeder system all greatly impact reliability.

In almost all utility systems, the majority of customer service reliability problems is caused by problems on the distribution system, either the primary feeder system or the service-level equipment it feeds. Figure 8.5 shows typical operating results – about 45% of all customer outage minutes are due to problems that occurred on the primary feeder system, nearly twice as much as on the entire system-transmission-substation level.

The fact the feeder system is the major source of service interruptions in a power system is not the least surprising. In fact it seems inevitable. To begin, there is simply more of the primary distribution system than anything else in the power system. It consists of thousands of miles of line and tens of thousands of units of equipment.³ There is simply more equipment to fail. Additionally, in all but very rare cases, the primary feeder is radial in design, consisting of a single

³ The authors are counting service transformers as part of the primary feeder system in this example, because they are energized on the high side at the primary voltage level.

serial chain of lines and equipment leading to each customer. Any failure in this type of system will lead to an interruption of service to one or more customers.

And finally, the feeder system is distributed – in fact it is the ultimate distributed resource belonging to any electric utility – a system of many, many units of equipment scattered about the system in roughly the same pattern as its customers. As such, locations of failures and repairs are scattered throughout the system, too. Unlike the substation level, the failed equipment will not be at a few well-identified sites. It often takes time – as much as an hour of searching in some cases – for repair crews to find the location of a failure.

Thus, the feeder system is the source of nearly half of a utility's service reliability problems, because of its great extent, radial characteristics, and distributed nature. As such it is particularly sensitive to the increasing failure rates of aging equipment. The aging equipment trends (Chapter 6) mean that both "clear weather" and storm-related failure rates on the feeder system tend to increase rapidly with equipment age. Attention to the feeder level is critical just because so many failures occur there during the best of circumstances. So here is the potential for the increasing failure rates of aging equipment to cause very noticeable increases in SAIDI and SAIFI – even a small increase in such a large portion of the system will make a very perceivable difference.

Outdated System Structure Impacts on the Feeder System

Basically, many of the trends discussed in Section 8.2 lead directly or indirectly to changes on the feeder system that limit its ability to transfer loads, both to restore outages that occur in the feeder system, and to provide inter-substation transfer capability to mitigate substation-level outages.

Non-optimal feeder getaways and routes

Section 8.2 alluded to the fact that one consequence of outdated system structures is that existing substations are called upon to serve more loads than would ideally be the case. Figure 8.6 shows one consequence of this. The higher load requires more feeders (Recall that in the example in Section 8.2, the utility added 47% more feeders than originally intended to some of its substations, a typical situation.)

Getting so many more feeders out of a substation than originally designed may create a lot of headaches with regard to routing, configuration, and exposure. Often, feeder getaway constraints – limits in how much space is available above and below ground to run feeder circuits out of a substation – is a limiting factor in its total load. Less than ideal (from the standpoint of economics, esthetics, and reliability) design compromises may be required simply to accommodate the sheer number of separate circuits. Reliability suffers due to a host of secondary effects.




Figure 8.6 Left, as originally planned, one of Metropolitan Power and Light's substations was intended to serve 86 MVA of peak load with 13 feeders, each loaded to about 6.6 MVA. Right, as eventually operated by the utility due to outdated structure constraints and pressures, the substation served 180 MVA through 23 feeders (7.82 MVA). The additional feeder getaways often represented relatively costly designs compared to feeders planned in the initial layout of the substation, and suffered from compromised esthetics and more importantly, lower reliability than needed.

Non-optimal routing

Similar to the getaway problem, a utility may find that additional circuits it has to "work into" an area beyond the number originally planned forces it to use routes and layouts that normally would not have been acceptable. Otherwise, more circuits and higher level of loading could have been accommodated well if the entire area had been originally planned with those in mind.

However, typical if an area is planned to a lesser density and then upgraded, it is not possible to re-engineer it to the same level of economy and reliability after the fact. Circuit routes have been committed and constraints accepted in certain ways in the original design. Subsequent development in the area has reduced additional routes available to the planners.

The additional feeders often have to take routes that are less reliable than desired (heavy tree populations in areas difficult to trim or exposed to auto traffic along roads). Double-circuit pole routing of feeders (two feeders on one set of poles) often has to be used. Common mode failures (one event taking out more than one feeder) become much more likely. Switch locations and switching patterns often cannot be well arranged with so many feeders in such a small region.




Figure 8.7 Left, four feeders (lines) pass through a large cable vault (circle). Right, in the expanded feeder system, eight now populate the vault. Exposure has increased greatly: there is both more likelihood of a catastrophic failure (there is more equipment in the fault) and any problem potentially hazards twice as many feeders.

Duct crowding/exposure

In underground distribution systems, the attempt to fit as many feeders as possible into existing duct space will reduce reliability through two secondary effects (secondary effects make a substantial impact on performance and cannot be avoided using ordinary practices and means):

1. *Heating.* More cables in a highly populated duct will lead to greater amount of losses and more heat for all cables. Good engineering will assure heat buildup stays below limits, but nonetheless thermal stress on cables will be higher than it would be otherwise. Lifetime of cables may be noticeably shortened and failure rates increased.
2. *Common mode failure* exposure goes up greatly, as shown in Figure 8.7. Here, more feeders means more feeders passing through each cable vault, where a fire on one can now endanger more feeders.

In addition, there are two tertiary effects that exacerbate problems (tertiary effects make a noticeable impact but generally can be mitigated or avoided entirely through extra attention or effort to the application of standard practices).

3. *Complexity.* There can be so many cables in vaults that it inhibits the progress of work. Identifying which cable is which, accessing one while work on others is being done, can become a noticeable issue.
4. *Old collapsing ducts.* In the central parts of many cities, old ducts are collapsing or warping, making it difficult to pull cable, etc. Fitting as many feeders as possible in older ducting is an example of putting too many eggs in one basket.

The net effect: a noticeable increase in failure rate

Operating experience at one large investor owned utility showed that areas where additional feeders had to be added “after the fact” had an 11% higher frequency of outages than those areas where the system corresponded to the original planning. This was an 11% increase in the frequency of outages of equipment on the feeder system, leading to an 11% increase in the contribution that feeder system makes to the overall system SAIFI. The causes seemed to be an amalgamation of all the reasons listed above. Assuming that Figure 8.5’s statistics showing that 45% of all reliability problems are feeder-related applies here, this 11% increase in problems on the feeder system means that feeder-related problems caused *just* by outdated system structure problems contribute to $45\% \times 11\% =$ a 5% increase in the overall system SAIDI.⁴

Using the Feeder System to Mitigate Problems Caused by Substation-Level Outages

One of the most effective, if not the most effective, ways to support substation transformer outages in systems with high substation and sub-transmission utilization ratios is through the feeder system. When a transformer is loaded to 90% of its rating, as it is during peak in many utilities forced to high utilization rates by the issues covered in Section 8.2, avoiding lengthy customer outages when it fails means transferring a lot of load to other equipment, quickly. If the failed unit has only one neighbor that can interconnect to provide that support, that unit must go to 180% of rating – a level considered too high in all but extreme cases – or all of the load cannot be served. Concern over such high contingency loading levels is largely the reason why traditional utilization ratios for substation transformers were targeted at 66%. When one of two units at a substation failed while running at that level, the other had to go to no more than 133% of normal rating – a high but acceptable loading during a contingency.

One way to design a system so that any transformer has more “neighbors” that can support it during an outage, is to have more than two transformers at each substation. Some utilities (Houston) traditionally designed most of their distribution substations with four transformers. Loss of any one, even if loaded to 100% at peak, meant the load could be distributed among the remaining three with only a 133% overload of each. Transformer, bus, breaker, and similar contingencies could be contained within each substation, with no support needed for neighboring substations.

⁴ Strictly speaking, Figure 8.5 refers to contribution to SAIDI by level of the system – duration of outages – whereas this discussion refers to frequency of outage, a difference factor. However the authors assume the two are comparable here for the purpose of demonstrating the degree of importance of this concept.

But in many cases, three- or four-transformer substations are not an efficient investment. In most cases, the economy of scale for transformers is greater than that for multi-transformer substation layout. Thus, for a given required capacity, two large transformers prove to be much more cost-effective than four smaller ones.⁵ As a result, as much as there is any “standard” substation layout used in the United States, it is the two-transformer substation, composed of two like units (same capacity). In a power system based on these types of substations, any outaged transformer-bus-breaker combination unit would have one remaining neighbor to provide contingency support (in the substation). The utility has three choices for a viable contingency plan for these outages.

1. *Limit loading on the substation to what one transformer can take under contingency conditions.* A few utilities have shifted to permitting 166% loading during contingencies (for brief periods), permitting up to 83% loading of the substation. Most, however, limit loading to 75% under normal conditions with a 150% contingency loading, still a very aggressive contingency loading level.
2. *Provide support with a mobile transformer/substation.* The plan being that a mobile spare could be brought to the site “within a few hours.” The utility accepts a brief outage during that time as the nominal result of a transformer failure. This is often not a palatable option because the average mobile substation-response time is actually 12-18 hours.
3. *Provide support for transformer outages through the feeder system.* The concept here is that the load of the failed transformer is split – the remaining transformer at the substation picks up part of the load, and the rest of the substation’s load is transferred to neighboring substations via re-switching of the feeder system.

This can be implemented via “programming” of the protection and breakers scheme in one of two ways.

1. Upon a failure of one unit, the remaining transformer picks up the entire load it can within its contingency rating. For example, if loading on both transformers was at 90%, the remaining unit goes to 135%, picking up half the failed unit’s load. The other half is dropped temporarily. Feeders are then switched to transfer load from the substation to neighboring substations, so as to pick up the load that was dropped. This

⁵ Houston Lighting and Power (Reliant Energy) has a somewhat unusual situation in that its load density is much higher than average due to the intense amount of air-conditioning demand in the Houston area. Four transformer substations are therefore quite efficient in Houston’s case.

may be done with remote switching, but is most often accomplished by manual switching, taking up to two hours.

2. The remaining transformer immediately picks up the entire load. This means it would go to 180% of rating during peak, a level at which it is very quickly giving up lifetime to serve the load. As quickly as possible, feeders are now switched to transfer load to neighboring substations to bring loading down to a lower level for the duration of the contingency.

This scheme puts transformers (and buses and breakers associated with them) under high stress for brief periods, but avoids customer interruptions. If the switching is well managed, as for example with a computerized restoration system at the utility dispatch center, it can be done quickly enough to avoid large stress-induced loss of life on the one remaining operating transformer.

Figure 8.8 uses a simple example to illustrate the “transfer to neighbor through the feeder system” concept.⁶ At the left, a two-transformer substation with a hexagonal service area has boundaries with six (also hexagonal) neighboring substation areas. Transformers in this example are loaded to 90% of rating at peak, by design. Also shown are six feeder areas (three per transformer) alternating radially between each of the substation’s two transformers.⁷

In this example, when one transformer is out of service during peak conditions, the other picks up what load it can, going to 135% of its normal rating. The remaining load, corresponding to 45% of a transformer’s capacity, is transferred to neighboring substations through feeder re-switching. As shown, the concept is to transfer the peripheral load of the substation – the load closest to the neighboring substations. This assures the shortest electrical paths for the contingency-switching feeders (for voltage drop consideration). See Chapter 8, section 3, for a fuller discussion of how switching capability like this is designed into a distribution system and what its benefits can provide. Support of transformer outages through the feeder system results in no, or very short duration, interruptions of service compared to the multi-hour or day-long outages if mobile transformers or other options have to be implemented.

⁶ See Chapters 13-16 in the *Power Distribution Planning Reference Book – Second Edition* (Willis, 2004).

⁷ This radial alternation scheme promotes feeder contingency support. See the *Power Distribution Planning Reference Book – Second Edition*, Chapter 14, particularly the discussion with respect to Figure 14.16, (Willis, 2004).




Figure 8.8 Left, hexagonal substation area with borders to six neighboring substation areas. This substation has two transformers that loaded to 90% of their rated capacity at time of peak demand. Under normal circumstances the two transformers, A and B, each serve half the load. When transformer A fails, unit B picks up half of A's load, going to 135% of its rating, and 22.5% of each of the transformers' loads is transferred to feeders from neighboring substations. In this way the utility can "get by" with loading transformers at two-transformer substations to 90% of rating at peak demand, and still have a viable contingency capability for the failure of either one. See text for details.

A Useful Concept: Feeder System "Strength"

In order for the feeder system to work well for contingency support in the example given, the feeder system must be able to transfer 22.5% of the substation's load to neighboring substations (45% of one transformer's load equals 22.5% of the substation's). If the feeder system is not capable of transferring 22.5% of the load, the system dispatcher is forced to choose between dropping load (interrupting service to some customers) or loading the remaining transformer to an unacceptably high, potentially damaging, level.

On the other hand, if the feeder system can transfer more than 22.5%, then that capability can be used for support transformer contingencies even if design loading of the substation is raised beyond 90%. If it can transfer 35% of each substation's peak load, then utilization ratio can be raised to 100% of transformer rating while still preserving contingency support capability.

The contingency strength of a feeder system is the portion of the substation load that it can transfer to neighboring substations while keeping voltages and circuit loading on the feeders within the utility's standards for contingency operation.




Figure 8.9 Feeder system strength (inter-substation tie capacity) versus utilization ratio of transformers under normal conditions that can be tolerated while still preserving contingency capability for the outage of a transformer at a two-transformer substation. For example, the graph indicates that if a contingency loading of 166% is allowed by standards, then normal target loading must be limited to 83% if the feeder strength is zero, but can go to 100% if feeder strength is 17%.

Figure 8.9 shows the utilization ratio for a two-transformer substation that can be supported in contingencies versus feeder system strength – feeder system inter-substation transfer capability. The term “inter-substation transfers” refers to load transfers between substation service areas. The term “intra-substation transfers” refers to load transfers among feeders in the same substation area. Several different lines are plotted representing various levels of contingency loading limits. For example, the “133%” line represents the combination of utilization ratio and feeder system strengths needed to assure full contingency capability when contingency loading is limited to 133%. The Y-intercept values for each of the lines represent the maximum utilization ratio that a utility can use at a two-transformer substation given that no load will be transferred to neighboring substations in the event of a transformer outage. As can be seen, greater transfer capability through the feeder system means greater utilization ratios can be applied in the normal operation of the system.

Loss of Contingency Reach Due to Obsolete System Layout

“Reach” is the distance that a feeder system can move power while staying with the voltage standards applicable for its operating conditions. For example, a 12.47 kV, three-phase, overhead feeder of typical industry design, using 4/0 ACSR conductor, can move power corresponding to its thermal limit (the maximum amount of power possible within its current rating, 340 amps/phase, corresponding to 7.3 MVA) a distance of 1.8 miles before voltage drop on the line reaches 7.5% - the maximum permitted on the feeder system according to many electric utility voltage standards. This distance is 4/0’s *thermal reach*.

Substitution of another conductor will not materially change this distance: a 336 conductor will move its thermal limit (530 amps corresponding to 11.5 MVA) 1.8 miles; a 636 AA (770 amps, 16.6 MVA) will move its thermal limit 1.7 miles; a #2 ACSR (130 amps, 2.8 MVA), 1.8 miles. The larger conductor has lower impedance (less voltage drop per ampere) but a higher thermal limit. The two factors compensate for one another. Thus, the *thermal reach* of a 12.47 kV feeder system and the distance it can move power if operating at its thermal limits is 1.8 miles, regardless of conductor size.

Economic reach refers to the distance that power will travel on feeder segments before encountering the maximum allowed voltage drop, in situations conductor size and loading have been selected so that all conductor operates at or near its most *economical* loading level over the course of a year.⁸ The economic factors, load characteristics, and design standards that affect determination of economic loading vary from utility to utility. But again, the economic reach of each type of conductor turns out to be about the same. Under typical conditions economic reach works out to about 3.6 miles – about twice thermal reach (economical loads are usually about ½ the thermal limit).

Contingency reach refers to the distance that power can be moved on the under contingency conditions, when higher loading levels and more liberal voltage drop limits that are permitted. Again, standards vary from one utility to another. But as an example the ANSI C84.1-1989 standard specifies range A voltages as design criteria for normal conditions on the system, with a 7.5% maximum voltage drop on the primary feeder system permitted (9 volts on a 120 volt scale). It states that Range B voltage ranges apply in emergency situations, permitting up to 13 volts on a 120 volt scale (10.8% voltage drop).⁹

Thus, according to this standard, during a contingency a distribution utility is permitted to operate feeders in the affected area with up to $10.8/7.5$ equals 144%

⁸ For a detailed explanation of how economic loading is determined for conductors and for a distribution system, see the *Power Distribution Planning Reference Book – Second Edition*, Chapters 11 and 12 (Willis, 2004).

⁹ See the IEEE Red Book - *Recommended Practice for Electric Power Distribution for Industrial Plants* (IEEE, 1990).

more voltage drop than under normal conditions. This 44% additional voltage drop can be “spent” by either moving more power (more current equals more voltage drop) or by moving power over a longer distance (longer distance equals more voltage drop), or in some combination of more power moved a greater distance which is the normal case.

Generally, due to a bit of margin built into a feeder system for a variety of secondary and tertiary reasons, there will be a bit more voltage drop margin than minimum.¹⁰ Thus, under most circumstances it would be reasonable to expect about 50% – 60% to be the permissible additional margin, over and above normal peak load conditions, rather than just the 44%.

Consider again Figure 8.8 and its example of transfers to other substations. Note that 22.5% of each feeder’s load is transferred to neighboring feeders. Roughly speaking, the additional distance for power flow is about 20% farther (assume for the moment that the other substation areas/feeder areas are the same size and shape as those in the drawing). Thus, additional loading times additional distance ($122.5\% \times 120\%$) is very close to the 144% limit. If this feeder system is well designed and operated such that normal loadings are kept in the most economical range for all conductors (as it ideally would be) then no substantial reinforcement of its elements should be required to provide the required level of contingency support.

Little if any additional capacity needs to be installed to provide this contingency support. A well-designed feeder system with all conductors sized to be most economical, and with switching designed for inter-substation support, can be built at very little if any additional cost (beyond the extra planning and engineering work required) over a system that does not have these capabilities.

However, consider what happens when the loading on the feeder system is permitted to creep upwards, as occurred to Metropolitan Power and Light example case in Section 8.2. Over the course of three and one-half decades an increased loading on the feeder system was accepted as an inevitable result of changes the utility could not stop. Although reinforcement of the feeder system was done from time to time, substituting larger conductors where necessary, many portions of the feeder system ended up running at above their most economical loading ranges. As a result, they lose reach.

The net result is that a situation like that depicted in Figure 8.10 can occur during a contingency. There, two feeders for neighboring substations are shown (each a triangular area corresponding to the triangular feeder areas in Figure 8.8).

¹⁰ There are only a limited number of different conductors available. Often in “optimally sizing” a particular segment engineers determine that one size is a bit too small, the next size a bit too large, to be exactly optimal. Necessity dictates they pick the larger size, which provides a bit lower voltage drop than technically necessary – additional margin for contingencies, among other things. In typical systems this adds about 10% to the available voltage drop margin.

The rightmost feeder's substation transformer has failed. Its outer 22.5% load has been un-switched from the substation (opening switch B), and the feeder tie switch at point A closed, so that the feeder at the left can pick up the outer most load from feeder B. Due to higher feeder loadings, the feeder system does not have the contingency reach to provide this capability. The dotted line in Figure 8.10 shows where voltage drops below unacceptable levels.

The outdated system structure issues discussed in Section 8.2 created two effects that led to this situation:

1. *Loss of contingency capability.* A combination of higher substation utilization ratios and higher loadings on the feeder system lead to a large shift in the ratio of capability available/capability needed. Higher substation ratios raise the amount of load that needs to be transferred over the feeder system during contingencies. Also, higher loadings on the feeder system reduce the voltage drop margin available for contingency reach. As a result, the number of situations and hours per year (near peak) when contingencies cannot be supported increases. Reliability of service suffers.
2. *Contingency capability preempted by planning.* In many systems, local pockets of load growth are accommodated without reinforcement of the feeder system in the area, via load transfers. For example, suppose that in Figure 8.9, the "contingency" that occurred was a planning contingency: the load in the feeder area on the right grew to be 10% more than that feeder could handle while staying within standards, but the feeder on the left had 10% margin. Utility planners can avoid any reinforcement cost by performing the switching operation shown, as a *permanent solution* to the "overloading" problem. This "fixes" the normal operating problem but leaves insufficient contingency capability in the system.

Feeder Problems Often Don't Get the "Respect" They Are Due

While the feeder system is critical to any utility's establishing a good record of reliability, failures there never cause the widespread, "headline grabbing" problems that failures at the substation and sub-transmission level can cause. The failure of an entire feeder might drop service to between 500 and 1,500 customers. By contrast a bad problem at a substation can drop twenty feeders at once, causing a widespread outage that garners serious attention from newspapers and community leaders. For this reason, efforts to fix aging infrastructure problems often focus on levels of the system other than the feeder system. Still, the feeder system is an important part of the system, and as shown can be an important resource in limiting the impact of substation – sub-transmission outages. Therefore, attention given to it is more important in the overall reliability equation than is often realized.




Figure 8.10 Here, the feeder on the right has experienced a failure at the substation (shown by an X). Switch B has been opened, switch A has been closed, so the feeder on the left can support both all of its load and the outlying 22.5% of the other feeder's load (the rest is supported by switching the outaged feeder truck to another feeder near the substation). The feeder on the left is now serving 122.5% of its normal load, over a distance of transmission that is as much as 15% more. Dotted lines show the points where voltage drop reaches range B limit.

For this reason, efforts to fix aging infrastructure problems often focus on levels of the system other than the feeder system. Still, the feeder system is an important part of the system, and as shown can be an important resource in limiting the impact of substation – sub-transmission outages. Therefore, attention given to it is more important in the overall reliability equation than is often realized.

Outdated Structure Impact on the Feeder System: Reliability

Regardless of whether one or both of the situations discussed above develop the effect on reliability is the same. A growing portion of the equipment outages, which could *potentially* lead to widespread, lengthy interruptions (i.e. outages of major equipment at the substations) now have no viable contingency support.

In several large metropolitan systems outdated system structure issues (Section 8.2) eroded feeder contingency support strength due to these two reasons, to less than 40% of what was required to meet system needs, a loss of 60%. Figure 8.5 showed that substation – sub-transmission related problems cause about 25% of system SAIDI. Assuming again that Figure 8.5 is reasonably

close to the distribution of frequency-of-outage causes in the system (see Footnote 6) and that degradation in results is proportional to degradation of capability (it is actually worse), this 60% reduction in capability to support substation-level problems would lead to something along the lines of a 15% increase ($60\% \times 25$).

This impact on reliability is roughly three times the impact on system reliability that issues related to feeder system reliability cause (that came to 5%). Overall then, outdated system layout problems (section 8.2) lead to a 20% net increase in reliability problems, just due to their impacts on the feeder system alone.

8.5 FIXES FOR OUTDATED SYSTEM STRUCTURES

There are no simple, inexpensive solutions to problems of the magnitude discussed in this chapter. However, there are solutions, which are more effective and less expensive than trying to reverse the various impacts of aging infrastructure by adding more sites, more facilities, and reducing utilization ratios to traditional levels and practices. These require innovative approaches that often break with traditional concepts and practices.

Improved Communication with Community and Regulators

Electric utilities have to do a better job of identifying to those who approve their plans that new facilities and sites are needed if long-term reliability and service quality is to be kept at high levels. The quality of results that utilities obtain from their communication with regulators and community varies greatly. All utilities would be wise to emulate the methods employed by utilities that are particularly effective in this manner, such as PSE&G.

Compact Substation Designs

The compromises of capacity over configuration as illustrated by Figure 8.4 are simply unacceptable in a system where utilization rate is kept high and reliability is a priority. A better solution is to find a way to “cram” more capacity into existing sites. This can be done to two levels of improvement:

1. *Compact designs.* Various means exist to reduce the space required for air-insulated substation equipment. Usually this entails re-building the substation upward rather than outward. Among the items used are “low footprint” transformers, and high rise buswork designed to use little horizontal real estate (which as a consequence has a much higher profile). Compact designs can typically increase the capability of a substation (as measured by additional capacity added without any sacrifice in configuration) by up to 25 – 30%. Often this is not enough to provide all the improvement needed.

2. *Gas-insulated switchgear.* The solution of favor for crowded conditions in urban areas throughout Europe is the gas-insulated switchgear (GIS) substation, the so-called substation in a bottle. Slightly more than half the metropolitan substations throughout Western Europe utilize GIS technology. GIS substations have all of their bus-work and breakers contained inside a set of linked insulating capability. Voltage clearances between equipment can be much tighter than with air-insulated equipment, and the weather-protection rendered by the enclosures and gas means that much more compact designs can be used for all equipment. As a result GIS equipment requires far less space than its air-insulated equivalents.

GIS substations can increase the capability of a substation site over that obtainable from the tightest possible use of traditional air-insulated equipment, by a factor of at least three. If need be, an old 160 MVA site could be rebuilt with close to 500 MVA of capacity. Very often full use of this capability is not needed, but usually a doubling of capability, or potential for eventual growth, is implemented when GIS is retrofitted to a site. GIS also makes sense as the construction standard for new substations in aging infrastructure areas, because it enables the utility to have a much wider choice of available sites. Its compact nature and the fact that it can be put inside medium-sized buildings, means it will fit in many places that previously had to be rejected.

GIS substation equipment costs about 15% – 30% more than air-insulated alternatives. However it is slightly more reliable than, and requires much less maintenance cost than, air insulated designs of equivalent capability, basically more than canceling out this high initial cost if evaluated over ten or more years of operation. In addition, its minimization of real estate needs also contributes to lower cost for land, a major consideration and often more than enough to offset the higher initial equipment cost.

Reinforcement of Feeder Trunks and Switching

Given the typical effects of outdated system structures on the feeder system, providing the capability required for support of substation-level contingencies, will mean reinforcing feeder trunks with larger conductors, and installing additional switches for improved flexibility of operation. This is usually money well spent, and if planned well and coordinated with operation planning, feeder reinforcement for improved contingency capability is among the most effective ways to spend limited funds to obtain improved reliability of service.

Improved Planning

Improved planning can contribute significantly to a utility's both avoiding the situation depicted in this chapter, and mitigating its effects if it is already in such a situation. Good planning provides foresight, rational consideration of options, and under the right conditions, innovation. The following recommendations should be taken as a whole: they provide the proper type of planning for aging infrastructure areas as well as for utilities who need to maximize both customer service and stockholder (financial) performance. Chapters 12 and 13 will discuss this in more detail.

Section 8.2 cited ineffective long-range planning as a contributing cause to many of the causes that ultimately lead to outdated structure and poor reliability. These problems have to be fixed before lasting solutions can be put in place. Chapter 8 covers the nature of several problems with traditional planning tools. Chapters 11 and 12 discuss how improved planning methods and tools appropriate for modern high utilization-ratio systems can be effectively used.

But more is needed. Section 8.2 showed that a large measure of the problems utilities created for themselves were due to simply not looking far enough ahead or not using long-range planning capability that was there. Thus, in addition to changes in method and technique, utilities must simply pay attention to and use the longer-term perspective and evaluation of long-term consequences that flow out of long-range analysis of the consequences of growth and aging.

Two-Q approach

Chapter 4 introduced the concept of Two-Q planning, in which both quantity and quality of demand are viewed as burdens that set requirements for the system capability, and through which capability for both quality and quantity is arranged in an optimized fashion.

Two-Q planning explicitly identifies situations where new sites and new equipment are needed not for capacity (kVA) reasons alone, but for reliability reasons, and shows how much and why these needs arise. It makes such reliability-related requirements very clear and puts them in a context that is both communicable, where their justification is easier to make to executives and regulators alike. Two-Q is recommended both because it leads to a better understanding of reliability-related needs and the role elements of a long-range plan may have with regard to reliability, and because it provides a much better format and more documentation of need for their justification.

Reliability-based feeder system planning

The feeder system can be designed in a reliability-based manner to provide both better reliability itself, and more contingency support for substations, using methods discussed in Chapters 11 and 12.

8.6 SUMMARY OF KEY POINTS

“Outdated system structure” refers to a host of problems created for a utility when it cannot, or will not, expand its system’s layout and number of key substation sites and sub-transmission rights of way in company with the increasing demand for quantity and quality of power that a growing local economy and population provides. Key aspects of this impact are:

1. Operation of the system from fewer than an ideal number of substation feed points.
2. Compromises in design that sacrifice configuration for capacity.
3. Acceptance of very high utilization rates, even if not the utility’s economic standard, just because there is no room for other options.
4. Degradation of feeder system reliability and its ability to provide support during substation and sub-transmission level outages.
5. The situation developed gradually over many decades, the result of many small decisions and compromises none of which was recognized for what it was: failure of a long-term strategy to maintain a sound structural strategy for the system.

The net effect of outdated system structure issues, above and beyond the impacts of aging equipment itself, is an increase in reliability indices (SAIDI and SAIFI) due to impacts on the substation – sub-transmission and feeder levels of the system. Table 8.3 lists typical levels of impact. These structural impacts exacerbate the reliability of service and system efficiency problems generated by the deterioration of old equipment.

Table 8.3 Overall Impact of Outdated System Structure Issues on Customer Service Reliability

Level System	Typical Amount Of of Degradation
Substation/sub-transmission	15% - 20%
Feeder system reliability	5% - 15%
Feeder system contingency support	<u>15%</u> - <u>20%</u>
Total	35% - 55%

REFERENCES AND BIBLIOGRAPHY

- J. J. Burke, *Power Distribution Engineering – Fundamentals and Applications*, Marcel Dekker, New York, 1994
- R. E. Brown, S. S. Venkata, and R. D. Christie, "Hybrid Reliability Optimization Methods for Electric Power Distribution Systems," *International Conference on Intelligent Systems Applications to Power Systems*, Seoul, Korea, IEEE, July 1997.
- R. E. Brown, S. Gupta, S. S. Venkata, R.D. Christie, and R. Fletcher, 'Automated Primary Distribution System Design: Reliability and Cost Optimization,' *IEEE Transactions on Power Delivery*, Vol. 12, No. 2, April 1997, pp. 1017-1022.
- R. E. Brown, S. Gupta, S. S. Venkata, R.D. Christie, and R. Fletcher, 'Distribution System Reliability Assessment: Momentary Interruptions and Storms,' *IEEE Transactions on Power Delivery*, Vol. 12, No. 4, Oct 1997, pp. 1569-1575.
- Institute of Electrical and Electronics Engineers, *Recommended Practice for Design of Reliable Industrial and Commercial Power Systems*, The Institute of Electrical and Electronics Engineers, Inc., New York, 1990.
- H. L. Willis, *Power Distribution Planning Reference Book – Second Edition*, Taylor and Francis, Boca Raton, 2004.
- H. L. Willis and R. W. Powell, "Load Forecasting for Transmission Planning," *IEEE Transactions on Power Apparatus and Systems*, August 1985, p. 2550.

9

Traditional Reliability Engineering Tools and Their Limitations

9.1 INTRODUCTION

The third of the five aspects of aging T&D infrastructure discussed in Chapter 1 is Engineering Paradigms and Approaches. Traditional power system engineering methods, and the paradigms built around them, are as much a part of existing power systems and their strengths and weaknesses, as old equipment and obsolete system layouts. Many of the key concepts and engineering approaches used in “modern” power systems were created in the mid-20th century, at a time when electric consumer demands, societal expectations, and engineering criteria were far different from today’s. Like all engineering methods, the mid-twentieth century’s had their limitations and advantages. Those limitations were for the most part well understood by their developers, but were far outweighed by their advantages, among the most important being that the methods could be implemented with the very limited computing resources available at the time. The power industry has carried many of these engineering methods and the concepts built around them through several generations of power system engineers. Not only are these methods in widespread use, but several of these paradigms have become dogma.

This chapter will begin by examining the traditional power system reliability engineering tool: the N-1 criterion and the contingency-based planning approach. This method was first used in the mid-20th century and had been developed into

efficient computerized form by the late 1960s. It, and the conceptual approach to power system planning that goes along with it, are at the core of nearly every utility's overall power delivery planning and reliability engineering procedures. Similarly, many of the core concepts and criteria concerning the layout of distribution feeder systems were developed in the 1930s through the 1960s, and are largely unchanged in their application as the electric utility industry enters the 21st century. These will be covered in Chapter 10.

These traditional power system-planning methods proved more than adequate to meet the industry's needs through the 1950s to early 1990s. However, as experience has shown, and this chapter will explain, they are less than completely adequate for application to aging infrastructure areas in some of today's high-stress utility systems. This is due to an incompatibility between these traditional planning methods and the way that modern utilities need to plan, design, and operate their systems. As a result, systems planned and designed along proven but traditional lines, using tools that engineers could once confidently use, often provide poor customer service reliability and experience severe operating problems.

This chapter begins in Section 9.2 with an examination of the basic N-1 contingency-based reliability design criterion and the typical methods of its application. Section 9.3 then explores this methodology's limitations with respect to modern needs, and discusses how they interact with the characteristics of modern utility systems, and discusses how that often results in a system that does not provide the expected level of reliability. Section 9.4 looks at some other planning-related issues that have created real challenges for aging infrastructure utilities, most notably load forecasting errors and the way they interact with system reliability. Section 9.5 provides a reminder that high utilization rates and cost-reductions, per se, are not the reason that modern power systems and particularly aging systems, tend to give poor results – rather it is the inability of traditional planning tools to fully analyze the reliability implications of design in those areas. Section 9.6 rounds out the chapter by summarizing key points and gives five recommendations for effective planning procedures to be applied to aging infrastructure areas.

9.2 CONTINGENCY-BASED PLANNING METHODS

The N-1 Criterion

The traditional power system planning method used to assure reliability of design at the sub-transmission – substation level is the N-1 criterion. In its purest form, it states that a power system must be able to operate and fully meet expectations for amount (kW) and quality of power (voltage, power factor, etc.) even if any one of its major components is out of service (a single contingency). The system has N components, hence the name N-1.

This criterion makes a lot of sense. Unexpected equipment failures happen. Expected equipment outages (maintenance) are a fact of life. A prudent approach to design reliability should include making certain that the system can perform to its most stressful required level (i.e., serve the maximum demand, the peak load) even if a failure or maintenance outage has occurred. Just how much reliability this assures depends on a number of factors that will be addressed later in this section and the fact that it does not always assure reliability is the major topic of this chapter. However, from the outset it seems clear that this criterion sets a *necessary* requirement for any power system this is expected to provide reliable power supply.

Extension of the concept to multiple failures

The criterion can also be applied as an “N–2” criterion or “N–3” criterion in which case the system must be able to perform to peak requirements even if any two or three units of equipment are out of service, rather than one. Generalized, this becomes the N–X criterion, the power system will satisfy expectations even if any set of X and its components are out of service. Regardless, the method is generally referred to, and will be referred to here, as “the N–1” concept and criterion, even if X is greater than one.

Application of the Basic Contingency Planning Concept

Typically, this concept is applied as a criterion in the planning and engineering of the transmission/sub-transmission/substation portion of an electric distribution (T&D) system – the portion from 230 kV down to 69 or possibly 34.5 kV. At most electric distribution utilities, it is not applied to the distribution feeder system. Instead, techniques and analytical planning methods devised for application to radial power flow systems are used (see Section 9.3).

The base case

Application of the N–1 criterion begins with a *base case*, a model of the power system as designed or planned, with all equipment in place and operating as intended. An appropriate engineering description of all this equipment along with a set of expected peak loads that it will serve, forms the complete base case. In all modern utility planning procedures, this base case is a data set representing the system, to be used in a load-flow. A digital computer analysis will determine for the system represented by that data set, the power flows, voltages, power factors and equipment loadings that will result when that equipment set is asked to serve that demand. Various “minus one” or “contingency” cases are then done using this model as the base, literally deleting one element of the data set at a time and “resolving” the model to see what effect that loss had on voltages, currents, etc.

At the time the N–1 method was first developed, (which was prior to the

availability of digital computers), the base case was an analog computer model built using patch-cord connections and numerical settings of rheostats and switches on a network analyzer. Essentially, an analog computer built for simulation of power system behavior. Since the mid-1960s, digital computer programs that solve the load flow computation using a set of simultaneous equations have been used. By the end of the 20th century, these programs had become very specialized, with features and analytical tricks employed to make them fast, robust, and dependable when applied to contingency analysis.

But regardless of the type of engineering computer being used, studies that build upon, or more properly “delete upon” a base case are the foundation of the contingency-analysis method. As a first step, a base case representing the system in “normal” form, i.e., with all equipment in operation and fully functional, is set up and solved, making sure that it (a system with all equipment operating) fully satisfies all loading, power quality, and operating criteria.

Contingency cases

Variations from this base case are then conducted as a series of “contingency studies.” In each contingency study, one particular unit of equipment – a key transformer or line or bus, etc., is removed from the system database, and the remaining system’s performance studied using the engineering analysis model (load flow) applied to this “contingency case model.” The analysis determines if the system can still serve all the demand, while remaining within specified operational criteria (see Table 9.1), with this one unit out of service, or “outaged.” If not, additions or upgrades are made to the system model until the case does meet the criteria. Once the first contingency case is completed (the study for the first component in the system), the method proceeds to study the outage of the second. It begins with a “fresh copy” of the base case and removes the second unit in the equipment list, again performing its analysis. In this way, it proceeds through all N components, outaging each one and identifying whether performance in that situation is sub-standard, thus giving planners an indication of where problems in the system lie, and what the problems are.

Relaxation of design standards for contingencies

In most cases, electric distribution utility planners allow the contingency cases to meet less stringent requirements for loading, voltage, or other design goals, than for the “base” (no contingencies) case. For example, loading criteria may state that in the base case, no component can be loaded to beyond 100% of its normal rating. However, during any single contingency, a loading of 115% might be accepted, during a double contingency, a loading of 125% might be accepted. Table 9.1 lists the voltage and loading requirements for several utilities in the U.S. as a function of contingency situation.

Table 9.1 Transformer Loading Limits and Voltage Criteria for Various Contingency Situations Used By Four Electric Utilities in the U.S.

Utility Number	Base		N-1		N-2	
	Voltage	Loading	Voltage	Loading	Voltage	Loading
1	.97 – 1.03	83%	.96 – 1.04	125%	.95 – 1.05	133%
2	.96 – 1.05	75%	.95 – 1.07	115%	.95 – 1.05	125%
3	.97 – 1.05	90%	.96 – 1.05	135%	.94 – 1.05	166%
4	.95 – 1.04	90%	.94 – 1.05	135%	.94 – 1.06	166%

Application of N–1 using a Computer Program

As originally conceived, prior to the existence of really powerful digital computers for power system studies, this process was done with an analog computer. Each case was set up and studied on an individual basis by the utility’s power system planners, by adjusting settings and patch-cord connections on the analog computer. For this reason, initially (1950s) often only the 100 or so most important components (out of several thousand in a large power system) could be studied for contingency outage. However, beginning in the late 1960s, programs on digital computers were developed which would automatically check all single contingencies in the system [Daniels, 1967]. These programs became a staple of utility for system planning.

A modern contingency analysis program works along the lines shown in Figure 9.1. It is built around a load-flow program – a digital program that takes a data set describing the power system and the loads it is to serve, and solves a set of simultaneous equations to determine the voltages, flows, and power factors that can be expected in that system. In an automatic contingency analysis program, the basic load flow program is augmented with an outer loop, which automatically cycles through this power system data set, removing each unit of equipment and line, in turn, and solving the load flow analysis for that particularly contingency case.

For each such case, the program checks the results of that contingency case and reports any loadings or voltages that are out of acceptable range. It then “restores” that outaged component in the database, removes the next unit in turn, and runs that contingency case, cycling through all components in the system.




Figure 9.1 Basic approach behind the N-1 contingency planning approach. Engineering studies cycle through all components of the system and outage each one, studying what loadings and voltages would result. The system is considered to meet “N-1” criterion when all such contingency cases result in no out-of-range loadings or voltages.

A system plan was considered acceptable when this type of evaluation showed that no voltage and loading standards violations would occur in every one of these instances of single contingencies. Figure 9.1 illustrates this basic approach.

A Successful Method in the Mid and Late 20th Century

From the early 1960s through the early 1990s, the vast majority of different electric utilities applied this basic approach, with a number of slight variations here and there. Typically, a utility would design its system to completely meet N-1 criterion (the system can perform despite the loss of any one component) and to meet certain N-2 criteria (a set of specific two-failure conditions, which are thought to be critical enough to engineer to tolerate). Systems designed using this approach produced satisfactory performance at costs not considered unreasonable.

Supposedly Good Systems Begin Giving Bad Results

In the late 20th and early 21st century the operating record of large electric utilities in the United States revealed increasing problems in maintaining reliability of service to their customers. During the peak periods in the summers of 1998 and 1999 a number of power systems that fully met the traditional N-1 criterion experienced widespread outages of their power delivery systems. Table 9.2 lists only some of the more significant events in 1999, as identified by the U.S. Department of Energy. In particular, the ComEd system's (Chicago) inability to provide service was disturbing, because that system was designed to N-1 levels of

Table 9.2 Major System Outages Identified By US DOE in 1999

Area	When	Interruption Cause	Power Delivery?	Aging Infra?
New England	Jun 7-8	Wholesale generation shortage		
Chicago	Aug 12	Multiple power delivery failures	x	x
	Jul 30	Multiple power delivery failures	x	x
New York	Jul 6-7	Multiple power delivery failures	x	x
Long Island	Jul 3-8	Power delivery and grid problems	x	x
Mid Atlantic	July 6-19	Wholesale generation and grid problems	x	
N. New Jersey	Jul 5 - 8	Multiple power delivery failures	x	x
So.-Central US	Jul 23	Wholesale generation shortages		x
Delmarva	Jul 6	Wholesale grid and generation shortages		

Contingency standards and even met an N-2 and N-3 criterion in some places. There were a sufficient number of these events to make it clear that ComEd was not an isolated or atypical situation [U.S. DOE, 2000].

Perhaps more significantly, the level of reliability-related problems on U.S. systems had been growing for several years prior to 1999. Based on analysis of industry survey data, the authors first noticed the trend in 1994. While at that time there were few widespread customer outage events, throughout the 1990s there was a growing “background noise level” of operating emergencies, equipment overloads, and frequent if small customer interruptions on many utility systems. Even without the major events cataloged in the DOE report (Table 9.2) the industry’s record of customer service quality was falling in many regards prior to 1999.

9.3 LIMITATIONS OF N-1 METHODOLOGY

This section explains the limitations that N-1-power system planning and design techniques encounter when applied to modern (high utilization factor, lean margin) power systems. Contingency-based planning, like all engineering methods, is based on certain assumptions, uses approximations in key areas, and has limitations in its accuracy and range of application. And like other engineering methods, if applied to situations within which these assumptions, approximations, and limitations do not seriously hinder its accuracy and effectiveness, the method provides very good, dependable results. But if applied outside that range, its results may prove undependable, and the system plan may provide unsatisfactory performance.

Gradually, throughout the last quarter of the 20th century, the electric utility industry changed how it used and operated its power systems. Some of those changes meant that their systems operated in states and in situations for which the N-1 approach was not completely valid. The resulting incompatibility of method versus use of the resulting system contributed to many of the reliability problems experienced by utilities.

Of necessity, in order to fit within the space available in this book, this chapter’s discussion is somewhat shorter and contains certain simplifications with respect to a “real power system.” This abridging has been made in order to shorten the discussion, allowing the reader to quickly get to the heart of the matter, and to diminish distractions from the main theme due to secondary and tertiary factors. Therefore, this discussion makes the following “assumptions” or simplifications with respect to the system being discussed as an example here:

- All equipment of a specific type will be of the same capacity, e.g., all substation transformers are the same capacity.
- All equipment is loaded to the same peak level, that being the

average utilization ratio for the system at peak.

- All equipment units of any one type (e.g., all substation transformers) have the same failure rate.

The reader familiar with power systems will recognize all three as great simplifications of the many details that complicate power system planning and reliability engineering. However, real-world variations from these assumptions do not diminish the limitations and phenomena that will be discussed in the next few pages. In fact they slightly exacerbate them: these complexities generally worsen the problems explained here.

Overall Summary of the Problem

N-1 methods and the N-1 criteria assure power system planners and engineers that there is *some feasible way* to back up every unit in the system, should it fail. However, they make no assessment of the following:

- How likely is it that such backup will be needed?
- How reasonable is the feasible plan for each contingency situation or is the planner actually building a “house of cards” by expected “too many things to go right” once one thing has gone wrong?
- How much stress might the system be under during such contingency situations, and the long-term implications for both equipment life and operating policy?
- How often will conditions occur which cannot be backed up (e.g., multiple failures) and how bad could the situation become when that is the case?

As a result, systems that meet the N-1 criteria may be far less reliable than needed, even though the N-1 criteria “guarantees” there is a way to back up every unit in the system. This is much more likely to happen in modern power systems than it was in traditional, regulated power systems, due to changes in utilization and design made in the period 1990–2000, and the more volatile operating environment of deregulation.

Utilization Ratio Sensitivity

The major culprit that led to problems that “N-1 could not see” was an increase in the typical equipment utilization ratio used throughout the industry. When it is raised, as it was during the 1980s and 1990s, an N-1 compliant system, which

previously gave good service, may no longer give satisfactory reliability of service, even if it continues to meet the N-1 criterion. Similarly, power systems designed using the N-1 criterion and intended to operate at higher than traditional levels of equipment loading are very likely to have higher customer interruption rates than expected.

There is nothing inherently wrong with this trend to higher loading levels. In fact it is desirable because it seeks to make the utility financially efficient, which is potentially beneficial to both stockholders and customers. A power system that operates at 83% or 90% or even 100% utilization of equipment at peak can be designed to operate reliably, but something beyond N-1 methodology is required to assure that it will provide good customer service reliability.

N-1 is a necessary but not a sufficient criterion

Due to the success that N-1 methods had throughout from the late 1950s through the late 1980s, during which it led to power systems that provided good reliability of service, most power system planners and most electric utilities treated the N-1 criteria as a necessary and sufficient criteria. Design a system to meet this criterion and it was, by definition, reliable. But at higher utilization factors, while the N-1 criterion is still a *necessary* criterion, it alone is not *sufficient* to assure good quality of service. The reasons for this change in reliability as a function of utilization ratio are far subtler than is typically recognized. This section will review the limitations that N-1 has when applied to high-utilization-ratio systems and explain what happens, and why. The authors want to make clear that they definitely are not labeling high utilization ratios as the cause of all the industry's problems. Rather, it is the incompatibility between traditional ways of applied the N-1 criterion, and the way these systems operate, that created the problem.

Traditional utilization levels

In the 1960s through early 1980s, electric utilities typically loaded key equipment such as substation power transformers and downtown sub-transmission cables to only about two-thirds or a little more (typically about 66%) of their capacity, even during peak periods. The remaining capacity was kept as an "operating reserve" or "contingency margin." Engineers and planners at distribution utilities designed their power systems using the N-1 and other criteria, while counting on this margin.

In such systems, when a transformer or line failed, it required one neighboring transformer of equal capacity, perhaps already loaded to up to 66%, to be available to pick up its load. Depending on how close the system was to peak demand at the time of the outage, this unit might have to accept as much as 133% of its normal load (its normal 66% of rating and its neighbors 66%, too). Such overloading was tolerable for brief period. Power equipment can be run




Figure 9.2 Annual load duration curve for a utility system. Risk periods for high contingency loading of a traditional power system occurs only 10% of the time (shaded area). See text for details.

above rating for brief periods without significant damage, if this is not done too often.

And in fact it was unlikely that the loading would be as high as 133%, because that would only occur if the outage occurred during a peak load period. Only when loading was above 75% that of peak demand would overloads occur ($75\% \times 66\% = 50\%$, so at any load level below 75% of peak, one transformer can handle the load of two without going over 100% of its rating). In the system whose load duration curve is shown in Figure 9.2 (a typical U.S. utility system), such load levels occur only 10% of the year. As a result, it was (and still is) very likely that when equipment failures occur, they will occur at some time when loading is not near peak and hence stress on the equipment picking up the outaged units load was not unduly high.

Higher utilization rates

Beginning in the 1980s, increasingly in the 1990s, and through to today, utilities pushed equipment utilization upwards to where in some systems the average substation transformer was loaded to more than 100% of its nameplate during peak periods, by design. As was discussed in Chapter 8, in aging areas of this system, utilization rates can average close to 100% under “normal” peak conditions. Table 9.3 shows industry averages obtained in 2010 and past years

Table 9.3 Design loading guidelines for normal and contingency loading

Average maximum peak load as planned, under normal conditions, as a percent of nameplate thermal rating.	Percent of Nameplate				
	2010	1998	1988	1979	1970
Among utilities with > 1,000,000 customers	96	85	75	70	64
Among utilities with 500,000 and 1,000,000	90	85	75	70	65
Among utilities with < 500,000 customers	88	75	70	70	65
In rural areas	75	70	65	65	62
In suburban areas	90	85	80	75	67
In urban areas	95	90	82	75	72

Maximum planned peak load expected under contingency conditions (at least four hours duration)	Percent of Nameplate				
	2010	1998	1988	1979	1970
Among utilities with > 1,000,000 customers	155	145	140	135	133
Among utilities with 500,000 and 1,000,000	150	140	135	130	127
Among utilities with < 500,000 customers	145	135	130	130	127

in a comprehensive survey of loading practices across the industry.¹

The engineers in charge of these higher-utilization rate power systems knew, and in fact, planned for, their system to accommodate these higher utilization rates in several ways. To begin, these higher utilization rates required considerable augmentation of system configuration and switching. In a system loaded to an average of only 66% at peak, each transformer and line required one neighboring unit to “stand by” to pick up its outage. This meant, for example, that in each two-transformer substations, there had to be buswork and switchgear (breakers, switches) configured so that if one of the units failed, the other could automatically pick up its load. Alternately, the load at the substation had to be partially transferred to neighboring substations (onto their transformers) through the feeder system, as will be discussed in Section 9.4. Usually, some combination of stronger substation and sub-transmission-level buswork and switching flexibility, and increased reliance on feeder-level transfers was used.

¹ From *Electric Power Distribution Practices and Performance in North America – 1998* (“The Benchmark Report” by H. L. Willis and J. J. Burke, ABB Power T&D Company, Raleigh, NC). “Design loading” as used here, refers to the peak load on a transformer, above which it is considered so highly loaded that it should be upgraded in capacity, or load transferred elsewhere. “Emergency rating” refers to the maximum load permitted to on the substation during an equipment outage or excessive load contingency.

Those plans for these higher-utilization systems were created using N–1 methods, which “took into account” those higher utilization rates in their analysis. These N–1 analysis applications assured that *there was a way* to back up every unit in the system, should it fail. The system plans developed as a result fully met the full N–1 criteria everywhere, and N–2 criteria in critical places, even though they were operating at these higher utilization rates. Thus, these systems did have well-engineered contingency capability. The equipment was there and it would, and did, work as intended. Any problems lay elsewhere.

Looking at N–1’s Limitations

In order to understand where the problem with N–1 criterion application lies, it is important to first understand that a power system that meets the N–1 criterion *can and routinely* does operate with more than one unit of equipment out of service. Consider a power system that has 10,000 elements in it, each with an outage expectation of .16% – a value lower than one would ever expect on a real power system. One can expect that on average, about 16 elements will be out at any one time. Yet the system will usually continue to operate without problems. The reason is that the N–1 criteria has guaranteed that there is a backup for every one of these failed units, as shown in Figure 9.3.

The system will fail to serve its entire load only if two of these multiple outages occur among neighboring equipment. For example, if a transformer *and* the transformer designated to back it up both fail at the same time, as shown in Figure 9.4, then, and only then, will a customer service interruption occur.

Contingency support neighborhood

“Neighboring equipment” as used in the paragraph above means the equipment in the vicinity of a unit that is part of the contingency support for its outage. This can be more accurately described as its *contingency support neighborhood*: the portion of the system that includes all equipment that is part of the planned contingency support for the unit’s outage. For a substation power transformer, this might include at least one neighboring transformer (usually at the same substation) which would provide capacity margin during its outage, along with portions of the high-side and low-side buswork and switchgear, which would operate in a non-standard configuration during its outage.

Figure 9.5 illustrates this concept, showing several “contingency support neighborhoods” as in the example system used in Figures 9.3 and 9.4. As stated in the introduction to this section, this discussion simplifies the real world somewhat. Here, every unit is the same size and contingency support is always grouped exclusively in sets of neighbors. Actual design is more complicated, but the complications do not make any substantial net impact this discussion.




Figure 9.3 One-line diagram for a small part of a large power system. Four equipment outages are shown, indicated by an X, two transformers, one high-side bus, and one sub-transmission line. Each outaged unit has a neighboring unit (filled in) that has picked up its load: the N–1 criteria assured that this was the case. The system continues to operate smoothly because no two of the outages occur close enough to one another.




Figure 9.4 One set of dual failures in the same contingency support neighborhood, as illustrated here with the failure of two neighboring transformers (each was the designated backup for the other), will lead to interruption of service to consumers. Here, the shaded circle indicates the rough area of the system that would be without power.




Figure 9.5 Every unit in the system has a “contingency support neighborhood” that includes all the equipment that provides contingency support for the unit. Shown here are two transformers (filled in) along with their neighborhoods (circled). Equipment in a neighborhood provides contingency margin (capacity) as well as connectivity flexibility (switching, flow capability) during the outage of that unit.

Problems in this N–1 system that lead to customer outages occur only when *two or more* equipment units fail simultaneously within *one* contingency support neighborhood. Such a “double failure” does not have to be among just the unit and its like-type of support unit, i.e., the failure of both transformers at a two-transformer substation. Failures of one transformer and a line, a breaker, or a bus needed to support the contingency re-configuration of the system to support its outage can also lead to a failure to maintain service. Still, such occurrences are very rare. While there are perhaps 10 to 15 units out of service in a system of 10,000 elements, it is most likely that they are scattered singly throughout the system. The likelihood that two are concentrated in a one-minute neighborhood is remote.

Traditional power systems had “small” contingency support neighborhoods

In traditional power delivery systems, those whose utilization ratio for power transformers and sub-transmission lines was nominally targeted to be about 66% of equipment rating during normal (design) peak conditions, the contingency support neighborhood for any unit of equipment was small. As discussed earlier, every unit in the system needed one backup unit of like size. A 32 MVA

transformer would be loaded to 22.5 MVA (66%) at peak. If it failed, its partner at a substation, also already serving 21 MVA, would pick up its load too, briefly running at 133% (45 MVA) loading so that all demand was served. Therefore, the contingency support neighborhood for both units was a small locality that included the other transformer and the various switchgear and buswork needed to connect the two to the other's load during a contingency.

Systems with high utilization rates have larger contingency support neighborhoods

Suppose that the area of the power system being considered has 88.5% loadings on all transformers, instead of 66%. In that case, when any transformer fails, and if the utility is to keep within a 133% overload limit, a failed unit's load has to be spread over *two* neighboring transformers, not just one. The size of the "contingency support neighborhood" for each unit in the system has increased by a factor of fifty percent. Previously it included one neighboring transformer, now it includes two.

More importantly, the probability that an outage will occur among the designated support units for each transformer is double what it was in the system loaded to only 66%. Previously, whenever a transformer failed, there was only one unit whose failure stood in the way of good service. Now, if *either* of its two designated support units fails, an interruption of service to the utility's customers will occur. Two possible failures, each as likely to occur as the one failure that could have taken out the 66% loading system.

Thus, in a system where utilization rate has been pushed upward, every contingency support neighborhood is proportionally larger, and thus a greater target for trouble to occur: There is more exposure to "simultaneous outages." In a system loaded to 66%, there is only one major target. Failure to serve the load occurs only if a unit of equipment and *one specific neighbor* designated as its contingency support are both out of service. In a system or area of a system loaded to 88.5%, it occurs if a unit and *either one of two neighbors* is out. In an area of a system loaded to over 100%, (as some aging areas are) it occurs whenever the unit and *any one of three* designated neighbors is out (Figure 9.6).

Basically, the whole problem boils down to this: the contingency support neighborhoods are larger. But there are still "N-1" neighborhoods: each can tolerate only *one* equipment outage and still fully meet their required ability to serve demand. A second outage will very likely lead to interruption of service to some customers. In these larger neighborhoods, there are more targets for that second outage to hit. "Trouble" that leads to an inability to serve customer demand is more likely to occur. The analysis below estimates the relative likelihood that this occurs in example systems loaded to different levels.




Figure 9.6 At 100% loading, each transformer from Figure 9.5 needs three nearby units to cover its load, expanding the “contingency support neighborhood” involved. See text for details.

A system as discussed earlier with “10,000 major elements” might contain 1,200 substation transformers. Assuming that the outage rate for them is .25%, this means:

1. In a 66% utilization system, there are 600 two-transformer contingency support neighborhoods. Failure to serve the load occurs only if both transformers of this pair fail. That is:
 - Failure probability = $.0025^2 = .000625$
 - Hours per year = $.000625 \times 8760 \text{ hours/year} \times 600 \text{ pairs} = 32.9 \text{ hours/year}$.
2. In an 88.5% utilization system, with three-transformer contingency support neighborhoods, failure to serve the load occurs only if all three or any two transformers of this triplet fail. Over the whole system, annually, that is
 - Failure probability = $.0025^3 + 3 \times (.0025^2 \times (1 - .0025)) = .0001871872$
 - Hours per year = $.0001871872 \times 8760 \text{ hours} \times 400 \text{ triplets} = 65.6 \text{ hours/year}$.

3. In a 100% utilization system, with four-transformer contingency support neighborhoods, failure to serve the load occurs if all four, any three, or any two transformers of this triplet fail. Over the whole system, annually, that is
 - Failure probability = $.0025^4 + 4 \times (.0025^3 \times (1 - .0025)) + 6 \times (.0025^2 \times (1 - .0025)^2 = .000037375$
 - Hours per year = $.000037375 \times 8760 \text{ hours} \times 300 \text{ quadrals} = 98.22 \text{ hours/year}$

By comparison to a traditionally loaded system, a power system at a higher utilization rate is two to three times as likely to experience a situation where a pattern of equipment outages fall outside of the N-1 criterion. For example, one of the “N-2” situations that might lead to an interruption of load. Systems run at higher equipment utilization rates are more likely to experience events that could put them in jeopardy of being unable to serve all customer loads. N-1 analysis does not measure or evaluate this in any manner.

N-1 criterion assures planners and engineers that a feasible way to handle every equipment outage has been provided. It does nothing to address how often situations outside of that context – i.e., those that will lead to unacceptable service quality, might occur.

High Utilization Coupled with Aging System Equipment Leads to Greatly Increased Service Problems

Chapter 7 discussed the effect that aging has on equipment failure rates. In aging areas of a power system, the failure rate for equipment is three to five times that of normal areas of the system. Coupled with the high utilization rates common in these aging areas and the result is a ten to one or slightly worse incidence of customer service interruptions due to equipment outages.²

² Here the authors will bring in a real world factor. Chapter 8 discussed why utilization rates are usually above system average in aging infrastructure areas of the system. A system where the average has gone from 66% to 88.5% may have seen only a modest increase in the utilization rate for equipment in newer areas of the system, while increases in aging areas make up for those below-average statistics. Thus, the aging part of the system has much higher utilization than other parts. Its customer service problems stand out both because of the higher failure rate in this area, and the higher likelihood that outages lead to customer interruptions. As a result aging areas often have a customer interruption rate up to twelve times that of newer areas of the system.

Increased High-Stress Levels and Periods

The situation is slightly worse than the perspective developed above when one looks at the stress put on the system's equipment, and the portion of the year that the system is likely to see: high- and medium-stress events due to equipment outages.

In a 66% utilization system, every transformer is paired with one other: whenever the unit it is backing up fails, it must support the load of two transformers. Given the transformer failure rate of .0025%, this means each transformer can expect to have its partner out of service and thus be in this "contingency support mode" about 22 hours ($.0025 \times 8760$ hours) per year. Given that the system is at peak demand about 10% of the time, this means that a transformer can expect about two hours of severe loading time per year.

When utilization ratio is 88.5%, each transformer is partnered with two other units: failure of either one will put it in a contingency support mode. Thus, neglecting the slight amount of time when both its partners are out of service (and thus customer service is interrupted), the amount of time it can expect to be in this mode is twice what it was in the 66.5% system: or about 44 hours, 4.4 of that will be high stress. Similarly, in 100% utilization, each transformer will see about 66 and 6.6 contingency support and high-stress hours period year. *Stress put on system equipment is much higher in high utilization systems.*

"Low standards" Operating Hours are Increased

It is also worth considering that standards on loading, voltage regulation, and other operating factors are relaxed during contingency situations (see Section 9.2). Since the amount of time that the system spends in these "medium stress times" is greater in high utilization systems, which means that the distribution system spends more time in "sub-standard" situations – twice is too much if the utilization rate is 88.5%, three times as much if utilization is 100%.

The Result: Lack of Dependability as Sole Planning Tools

The limitations discussed above can be partly accommodated by modifications to the traditional approaches and changes in N-1 criteria application. But the overall result is that resource requirements (both human and computer) rise dramatically, and the methods become both unwieldy and more sensitive to assumptions and other limitations not covered here. The bottom line is that N-1 and N-2 contingency-enumeration methods were, and still are, sound engineering methods, but ones with a high sensitivity to planning and operating conditions that are more common today than in the mid-1960s when these methods came into prominence as design tools. These limitations reduce the dependability of N-1 analysis, and the use of N-1 as a reliability criterion, as a definition of design sufficiency in power system reliability engineering.




Figure 9.7 Contingency-based analysis (solid lines) determined that a particular power system could sustain the required peak load (10,350 MW) while meeting N-1 everywhere and N-2 criteria in selected places. This means it met roughly a 30 minute-SAIDI capability at the low side bus level of the power delivery system.

Figure 9.3 illustrates this with an example taken from a large utility system in the Midwestern U.S. Traditional N-1 analysis determines that a power system operating at an average 83% utilization factor can serve a certain peak load level while meeting N-1 criteria (defined as sufficient reliability). Basically, the “rectangular” profile on the diagram given by N-1 analysis indicated that the system passed N-1 criteria everywhere, and N-2 criteria at a set of selected critical points, while the demand was set to that of projected design levels for system peak load.³

Actual capability of the system shows a rounded corner to the load versus reliability of service capabilities of the system. It is capable of delivering only 9,850 MW with the required 30-minute SAIDI. If expected to serve a peak load of 10,350 MW, it has an expected SAIDI of four times the target, 120 minutes per year.

By contrast, an analysis of the system’s capability using a reliability computation method that does not start out with an assumed “everything in service” normalcy base, which accommodates analysis of partial failures of tap

³ N-1 analysis does not determine an actual estimated reliability value, but in this case subsequent analysis showed that a valid N-1 criterion was equivalent to about 30-minute SAIDI, and that value is used here as the target reliability figure.

changers, and which accommodated some (but not all) uncertainties in loads and operating conditions, determined the profile shown by the dotted line. At high loading levels (those near peak load), the system is incapable of providing the reliability required – the rectangular profile is actually rounded off. The system can provide the peak load but with much less reliability than expected.

9.4 OTHER PLANNING RELATED CONCERNS

Partial Failures

Traditional N–1 contingency planning methods use “zero-one” enumeration of failures. In the contingency case analysis method (Figure 9.1), every unit of equipment and every line in the system is modeled as completely in service. In each contingency case, a unit is modeled as completely out of service. But modern power systems often encounter partial failures:

- A transformer may be in service but its tap changer has been diagnosed as problematic and is locked in one position, limiting system operation.
- An oil-filled UG cable’s pumps are disabled and the cable has been de-rated, but is still in service.
- Concerns about a ground that failed tests have dictated opening a bus tiebreaker to balance fault duties.

At the traditional loading levels that existed when contingency analysis was developed, such partial equipment failures seldom led to serious operating limitations and were safely ignored, while the contingency analysis still remained valid. In systems operating at higher loading levels, partial failures cause problems under far less extreme situations, and often cannot be ignored in reliability planning. For example, a power transformer loaded to 85% at peak, whose tap changer is locked into one position, is subject to voltage regulation problems that can easily reduce its ability to handle load by close to 15% in some situations.⁴ The contingency margin (100% - 85%) that the typical N–1 method assumes is there may, in fact, be mostly nonexistent.

⁴ Loss of its load tap changer does not cause any lowering of the transformer’s capability to carry load. However, flow through it is now subject to variation in voltage drop – higher flows result in higher voltage drops. It may be unable to do its job, within the electrical confines of its interconnection to other parts of the network, due to this variable voltage drop, which may limit it to partial loading only. Ignoring this and accepting the higher voltage drop that goes with full capability during a contingency would lead to serious problems of another type (unacceptably low service voltages to customers, or higher demands on other transformers).

Connectivity Sensitivity

As mentioned earlier, success in handling a contingency depends on the power system in the “neighborhood of contingency support” around the failed unit being connected in such a way that the neighboring units can provide the support while remaining meeting all electrical standards and satisfying all operational requirements. At higher equipment utilization rates, this neighborhood is generally larger everywhere within the power system.

This greater size is not, *per se*, the cause of problems for traditional N–1 analysis. Standard contingency-based analysis and the engineering methods that accompany it *can* fully accommodate the detailed electrical and capacity analysis of any and all contingency support neighborhoods, regardless of their sizes.

But each of these wider contingency neighborhoods involves more equipment and interconnections. Thus, accurate modeling is sensitive to more assumptions about the exact amount and location of loads in the surrounding areas of the system, and the way the system operator has chosen to run the system at that moment, and myriad other details about operating status. There are *more assumptions* involved in accurately depicting the status of each of the N components’ contingency support neighborhood.

And the analysis of each of these N contingency support neighborhoods is more *sensitive* to these assumptions. The range of uncertainty in many of these factors about future area loads and operating conditions is $\pm 5\%$ to $\pm 10\%$. Such ranges of uncertainty are not critical in the N–1 contingency analysis of a system operating at 66% utilization. The available contingency margin (33%) is considerably larger than the range. But when operating at 90% utilization, the uncertainty ranges of various factors involved often equal the assumed available contingency support capacity, and there is a larger neighborhood, within which it is more likely something will be different than assumed in the N–1 analysis.

Aging High Utilization Systems Are Sensitive to Forecasting Errors

A projection of future need is the first step in power delivery planning. The forecast of future peak load defines requirements for the capability of the system, starts the evaluation of alternatives as to feasibility, value, and cost, and defines the constraints for selecting the alternative which best meet requirements. Poor load forecasting has been a contributor to a significant number of aging infrastructure system problems around the nature – in the authors’ experience roughly half. Two areas of forecasting deserve special attention.

Weather normalization

Peak demand levels depend very much on the peak seasonal weather. In summer, the hotter the weather, the higher the demand. In winter, colder weather increases demand levels. Therefore, a projection of electric demand can, and should,

include an assessment of the impact of temperature on demand. For example, Figure 9.8 shows the peak demand vs. peak temperature relationship for a small municipal electric system. The peak demand is:

- Summer Peak Load (MW) = $495 \text{ MW} + (T - 57) \times 12.5 \text{ MW}$ (9.1)

- Winter Peak Load (MW) = $495 \text{ MW} + (57 - T) \times 8.2 \text{ MW}$ (9.2)

Recommended practice for electric load forecasting is to adjust historical weather data to a *standard set of weather conditions* to which the system design is targeted, then to project future demands under this same constant weather criteria for all planning purposes. In this way all weather data and forecasts are based on comparable situations: increases or decreases due to “real” reasons versus those due to variations in weather are distinguishable. Similarly, all planning should target a specific “design weather standard.” The forecast load, which defines requirements for the system plan, should be adjusted to this weather criterion. Essentially the system is being designed to serve peak-demand for weather this extreme, but not worse. Weather conditions (and peak loads) that exceed those conditions are treated as contingencies, just like other contingencies.




Figure 9.8 Peak daily loads and temperatures are related with a “jackknife” function (solid line). Shown is peak daily load versus peak daily temperature for all Tuesdays in a year (several Tuesdays thought non-representative because they were holidays or similar special events were left out of the analysis). Only Tuesdays are used in order to reduce the effects that different weekday activity patterns may have on load variation. See *Spatial Electric Load Forecasting- Second Edition*, Chapters 5 and 6, (Willis, 2002).

What temperature should planners select for this standard weather condition? Temperatures vary from year to year. Setting the design conditions at the mean, or most expected temperature, means that the forecast loads will be exceeded, and the system's capability exceeded by the demand, roughly every other year. On the other hand, it isn't cost-effective to install equipment to handle the worst possible weather conditions: "the heat storm of the century," etc. Generally, the recommended practice is to define a set of "design weather conditions" "extreme enough" to be rare but not so extreme as to be totally unexpected. Situations and needs vary, but a reasonable criterion is "design weather conditions are defined so that they *will be exceeded no more than once every ten years.*" See Willis 2002, Chapters 5 and 6 for a detailed discussion of both the techniques used to determine such adjustments and for recommendations on what constitutes "extreme-enough weather."

Impact of Mistakes in Weather Normalization on Reliability

The weather normalization method used in planning and the design weather targets set for the system are among the easiest matters for "rationalization" when efforts are being made to cut costs. For example, if planners re-set their design weather conditions from a criterion of once in ten years to once in five, or lower the forecast target in some other manner, the budget requirements that flow out of their system planning process will fall. As an example, the utility whose load is diagrammed in Figure 9.8 has an annual growth rate of nearly 1.0%. Summer peak demand sensitivity is 1.25%/degree F. Reducing the design weather target by one degree Fahrenheit, about the equivalent of going from once-in-ten to once-in-five years, will reduce the load growth forecast for a four year period to that forecast over five years. Assuming for the sake of this analysis that budget directly corresponds to amount of growth, that means it results in an annual budget reduction of 25% over the next four years. For this reason, a number of utilities succumbed to the temptation to change weather normalization too much.

- *Weather normalization* that targets a "too average" weather condition puts the power system in high-stress situations too often.

A low load forecast results in several detrimental impacts. First, it generally leads to a situation where the system is serving more load than intended. Usually, this does not create severe problems when all equipment is functioning, although it does age equipment somewhat faster than expected (accelerated loss of life). One can view operation at loads above design conditions as a contingency. Poor load forecasts used in the planning or operation of a power delivery system effectively "use up" its contingency capability (see Willis et al, 1985). Poor normalization of weather data for forecasting or poor spatial forecasting (poor correlation of loads with areas and equipment) results in deterioration of a system's contingency

withstands capability. This greatly exacerbates the reliability-of-service problems discussed up to this point in this book.

Equally important, and far less frequently recognized as a key impact of poor forecasting, the system serving a load above that for which it was designed will operate for many more hours of the year in a state where service quality is at jeopardy if complete or partial failures occur, or if “things don’t go exactly right.” Figure 9.9 compares the annual load duration curves for a “average year” as used in design of the system and 1999 (a one-in-ten year), for a large investor-owned utility in the central United States. The difference in peak demand between an average year and an extreme year is 4.4%. However, as shown, not only peak load changes but, annual load factor. The period of time when the system is above 75% of peak (defined as “high stress” earlier in this chapter) increases by 28%. As a result SAIDI increases significantly.




Figure 9.9 When weather is above average, it is usually somewhat above average for an entire season (summer, winter) or a good portion of it, not just for a day or a week. Shown above are forecast versus actual annual load duration curves for a utility in the mid – United States, for 1998. As a result of the higher temperature summer weather, peak demand was 3.3% higher than the mean-weather peak projected for the system. But in addition and perhaps more serious, the number of hours that the system was in a “high stress” loading situation – where contingency loading and switching are a concern – was 30% higher than for an average year. This, rather than the higher one-time peak load, was the major impact weather that year would have on reliability and operation.




Figure 9.10 Maps of peak annual demand for electricity in a major American city, showing the expected growth in demand during a twenty-year period. Growth in some parts of the urban core increases considerably, but in addition, electric load spreads into currently vacant areas as new suburbs are built to accommodate an expanded population. Forecasts like this, done for two-, four- and similar periods out to twenty years, set the requirements for both short- and long-range power delivery system planning.

Table 9.4 Percent of Utilities in North America Using Some Type of Formally Recognized Spatial or Small Area Load Forecasting Method

	2010	1998	1992	1988	1979
Among utilities > 1,000,000 customers	70	45	75	70	35
Among utilities 500,000 - 1,000,000 customers	75	50	70	65	30
Among utilities < 500,000 customers	20	12	15	15	10

A number of methods are in use for spatial forecasting, from simple trending methods (extrapolation of weather-adjusted substation and feeder peak load histories) to quite comprehensive simulations involving analysis and projection of changes in zoning, economic development, land-use, and customer end usage of electricity. Results vary greatly depending on method and resources used, but engineering methods exist to both determine the most appropriate methods and forecast characteristics needed for any utility application, and to evaluate the efficacy of a forecast. The most important factor is that a utility employs some legitimate means of studying and projecting load on a detailed enough location-basis to support its planning needs.

Traditionally, good spatial forecasting required both considerable labor and above-average engineering skills and was considered a “high-expertise” function within state of the art distribution planning methods. The best traditional methods worked very well but had rather high labor and skill costs methods (see Willis and Northcote Green, 1983 and Engel et al, 1996). Many utilities cut back on both the quality of the technique used and the effort devoted to data collection and forecasting study, when they downsized professional staffs during the 1990s. Table 9.4 illustrates the reduction the number of utilities using the best class (simulation-based) spatial forecast methodology, but does not reflect reductions in the data or time put into the forecasting effort. As a result, at a time when spatial forecasting needs are at an all time high (see below), the quality of local-area forecasting done at many utilities deteriorated sharply.⁵

In the very late 1990s, new forecasting methods were developed that reduce labor and skill requirements considerably, but these have limited availability and are not widely used (Brown et al., 1999). However, this means methods that can provide the information needed within reasonable labor and skill limits are available to the industry.

Impact of spatial forecast errors on reliability

In some manner, every T&D plan includes a spatial forecast, for the total load growth allocated in some manner among the various parts of the system. Classically, the viewpoint on forecast sensitivity of T&D systems has been that if the spatial element of the forecast is done poorly, the result is a very poor use of capital. A projection of the wrong locations for future load growth identifies incorrectly those portions of the system that need to be reinforced. Capital additions are made less effectively than possible.

But in addition, a large effect of poor spatial forecasting is a loss of contingency capability. Normally a power system designed based upon a mildly incorrect spatial forecast (i.e., one pattern of “where load is”) will suffer from less

⁵ Load forecasting problems related to local area forecasting were identified as a major contributing problem in six events (Table 9.2) investigated by DOE’s P.O.S.T. report.

than the planned contingency withstanding capability (i.e., provide less reliability in service) than expected. It will operate well enough during times when “everything is going well” but suffer from problems that are both more serious and take longer to fix than expected during contingencies. Essentially the poor forecast “uses up” the contingency capability built into the system [Willis and Tram, 1984].

Systems with high utilization ratios are more sensitive to this degradation of contingency planning due to peak load and spatial forecasting errors. Although subtle, the effect is best described this way: The contingency neighborhoods described earlier increase in size as a result of the higher utilization ratios being used for major equipment (Figures 9.5 and 9.6). While it may seem that this makes the planning less in need of detailed spatial forecasts (there are fewer “units” – contingency support neighborhoods – and they are on average far larger), the opposite is true. Contingency capability is very sensitive to the allocation of load *within* each support neighborhood. For example, the analysis given earlier assumed the load in each neighborhood was evenly split among units in that group: if it is even slightly unbalanced, the system’s contingency capability is greatly degraded. The forecast of where load is located *within* each contingency support neighborhood is critical to planning of contingency schemes that will prove successful when needed. Again, as with poor weather normalization, a poor spatial forecast “uses up” the contingency capability of the system, something in very short supply in aging areas of the system. SAIDI increases.

Interconnection Complexity

In the slightly simplified power systems used as examples earlier in this chapter, the small contingency support neighborhoods needed at 66% loading required interconnection with only two neighboring units assured success without overload during N–1 conditions. But at higher utilization ratios, contingency support neighborhoods grew in size and number of mutually supporting components. Interconnection of more equipment, into a scheme where it could support one another during contingencies, was necessary for success of the contingency plans.

In aging areas or where for other reasons planners and engineers have accepted near 100% utilization of equipment, there is a requirement for an even stronger and more widespread interconnection. Everywhere in a high-utilization system, each of its N units must have a strong-enough electrical tie to a wider neighborhood of equipment around it, to support its outage.

At higher equipment utilization rates, the importance of configuration and operating flexibility in the design of the system becomes more critical to reliability.

Traditional contingency-based study methods *can* deal with the *analysis* of

these issues relating to the wider neighborhood of support around every one of the N units in the system and its great complexity. They can determine if the required number of neighbors are there, if they have enough margin of capacity to accept the load without overloads, and if the system's electrical configuration makes it possible for them to pick up the demand that was being served by the failed unit. Basically, modern $N-1$ methods can and will determine if the failure of each unit in the system is "covered" by some plausible means to handle its failure and still provide service.

Again, $N-1$ methods do not work with probabilities nor determine the system's sensitivity to multiple failures, so they cannot determine the failure sensitivity or failure likelihood of these complicated interconnected schemes. At higher utilization ratios, complexity of the contingency backup cases has increased. Determining if a particular contingency backup plan is really feasible, if it is really connected with sufficient strength to survive the likely failure states, or if it depends on too much equipment operating in exactly the wrong way, is something $N-1$ methods do *not* fully address. Configuration needs to be studied from a probabilistic basis – is this entire scheme of rollover and re-switching likely to really solve the problem?

The problem is not high utilization rates

A point the authors want to stress again is that high equipment utilization is *not* the cause of poor reliability in aging infrastructure areas of a power system. It is possible to design and operate power systems that have very high (e.g., 100%+) utilization factors and provide very high levels of service reliability. This is accomplished by designing a system with the configuration to spread contingency burden among multiple equipment units, with the flexibility to react to multiple contingencies, and that can apply capacity well in all the situations most likely to develop. Such designs require detailed analysis of capacity, configuration, configuration flexibility, failure probabilities, and the interaction of all these variables, and careful arrangement of circuits to build in sufficient reliability.

Equipment utilization is only one factor in the design of a power system. In some cases, the best way to "buy" reliability along with satisfactory electrical (power flow) performance is to use capacity – to build a system with low utilization ratios. But in other cases, particularly where the cost of capacity is very high (as it is in many aging infrastructure areas), good performance comes from using the equipment to its utmost. Achieving high reliability of service even in these situations where equipment is highly stressed and contingency margins are small or non-existent requires using configuration and interconnection flexibility in an artful manner.

9.5 SUMMARY AND CONCLUSION

Traditional Tools Have Shortcomings With Respect to Modern Needs

Many of the problems faced by a utility owner/operator of an aging power T&D infrastructure are compounded by the fact that the tools being used by its planners and engineers cannot directly address one vital aspect of the required performance: reliability. Traditionally, reliability was addressed in power system design by engineering contingency backup capability into the system: every major unit of equipment could be completely backed up should it fail. This was termed the “N–1” criterion and methods that engineer a system based on this criterion were often referred to as “N–1” or “N–X” methods.

Such methods addressed a key and *necessary* quality for reliability: there must be a feasible way to do without every unit in the system, some way of switching it around during its outage and/or picking up the burden it was serving during its outage. Realistically, no power system can be expected to provide reliable service to its energy consumers unless it possesses this necessary qualification of having complete N–1 contingency capability.

But through many years of use during periods when equipment loading levels of equipment were lower than is typical in the 1990s and 2000s, the power industry came to view the N–1 criterion as necessary and sufficient. For power systems with roughly a 33% redundancy (contingency margin), the criterion is effectively a necessary and sufficient criterion to assure reasonable levels of reliability. However, when a power system is pushed to higher levels of equipment utilization efficiency, N–1 is still a necessary criterion, but no longer *sufficient* to assure satisfactory levels of reliability.

Basically, the shortcoming of the N–1 criterion, as well as engineering methods based upon it, is that they do not “see” (respond to, identify problems with, or measure) anything with respect to reliability of service except the “yes/no” satisfaction of this one criterion. Therefore, they cannot alert engineers to a flaw in the design of a power system with respect to its reliability, due to either the likelihood that events may get worse than the system can stand. “Yes, you have a feasible backup plan, but for this system it is very likely that while this unit of equipment is outaged, something else will go wrong, too.” Or, because some parts of the system are key elements for reliability of the system around them, to the extent that they are “important enough” that they need more reliability built into their design. “This unit is more important than you realized: A second backup is really needed for the reliability level you want to achieve.”

Furthermore, they cannot provide quantitative guidance to planners and engineers on what to fix and how to fix it in a satisfactorily effective but economical manner.

A Deeper Problem, Too

Finally, there is a more subtle but perhaps more fundamental problem associated with the use of N-1, one that is difficult to fully demonstrate in a textbook, but nonetheless real. N-1 engineering procedures solve reliability problems by using *capacity*: It is a criterion and method that uses contingency margin as the means to achieve reliability. This is partly the nature of the criterion and the tool, but also the fault of the paradigm – the way of thinking – built for the engineers around the N-1 tools.

By contrast, experience with planning and engineering tools that directly address reliability will quickly show any engineer that very often *configuration* is the key to reliability. In fact, capacity and configuration are both key factors in achieving reliability and artful engineering of power system reliability requires combining both in a synergistic manner.

Traditional N-1 analysis tools (Figure 9.1) can analyze the functional and electrical characteristics of configuration. However, they cannot do so on a probabilistic basis, analyzing how the likelihood that the interconnection will be

Table 9.5 Seven Desired Traits of the Ideal Power System Reliability Planning Method (that N-1 Methods Do Not Have)

Trait	Comments
1. Analysis of probabilities	Computes actual likelihood of all equipment in any area of the system being in service at one time.
2. Partial Failures	Can model the de-rating of equipment in stages, rather than just zero-one outage
3. Connectivity Sensitivity	Can model impact of variations in local loads, possible changes in feeder switching, etc., on expected performance.
4. Quantitative reliability	Computes actual expected reliability of service at points throughout the network in a quantitative manner.
5. Contingency Prioritization	Quantitative assessment includes both likelihood and importance against requirements in a way that permits prioritization of contingency results.
6. Cost sensitive	Relates or can relate results directly to money, profit margin, and business risk.
7. Solutions focused	Can determine or recommend possible fixes based on both effectiveness and economy.

there and determining how that interacts with the probabilities that the capacity will have failed. Table 9.5 summarizes the limitations of N–1 analysis that need augmentation in order to meet modern power system reliability planning needs. Such methods will be discussed in Chapter 13.

Explicit Reliability-Based Engineering Methods

What is needed to assure the *sufficient* requirement in power system reliability engineering along with the necessary, is a method that addresses capacity, configuration with an explicit, quantitative evaluation of the reliability of service they provide. Ideally, this must be a procedure that computes the reliability of service delivered to every element of the system, in much the same manner that a load flow computes the current flow and voltage delivered to every point in the system. It should be a procedure that can then be used by planning engineers to explore the reliability performance of different candidate designs, and that, in the same manner that load flows identify equipment that is heavily loaded (i.e., key to electrical performance), would identify equipment that is heavily loaded from a reliability standpoint (i.e., key to reliable performance).

Modern planning needs are best met using power system design and planning techniques that directly address reliability performance: Explicit, rather than implicit, methods. Such techniques have been used for decades in the design of systems where reliability is of paramount importance – nuclear power plants for commercial and shipboard use, spacecraft design, and throughout the aircraft industry. Adaptation of these methods to power systems provided a much more dependable design method to achieve operating reliability.

Such analysis begins with the same “base case” model of the power system as traditional techniques did. But probabilistic analysis starts out with a true normalcy base – a recognition that the natural condition of the system is “some equipment out” and that that condition will always be in a state of flux, with equipment being repaired and put back in service, and other equipment out of service. Like contingency enumeration methods, probabilistic analysis determines the consequences of every failure or combination of failures – can the system continue to operate and how close to the edge will it be during that time? But unlike traditional methods, probabilistic analysis determines if and how every combination of two or more simultaneous failures could interrelate to create problems, whether they are nearby or not in the immediate neighborhood, and it determines if that combination of failures is likely enough to be of concern. By tracing reliability through the configuration of the system while analyzing expectation of failure-to-operate, it effectively analyzes configuration and its relationship to reliability, too.

Depending on the exact method used (Chapter 14 will cover three basic types), partial failures can be accommodated using conditional capacity levels, partial failure states, or by recognizing sub-units within each main unit. Assumptions

about loads, operating conditions, and other aspects of system operation can be modeled using probability distributions with respect to those variables.

Table 9.6 Overall Recommendations for Planning Methods

Method	Comments
1. N-1 contingency analysis	No reason to not use it, as it tests for a necessary criterion of reliability and it explicitly demonstrates for each contingency how the system will react and handle the demand during that contingency. However, some of the better explicit reliability analysis methods implicitly merge N-1 into their analysis, so an explicit N-1 analysis is no longer needed.
2. Explicit reliability analysis	Probabilistic analysis of both capacity and configuration, including contingency interconnection, with computation of the expected reliability of service at all customer and intermediate points in the system. Can be based on a very modified form of N-1 analysis (dynamic analytical simulation) which combines (1) above with this step into one comprehensive method.
3. Economic cost comparison	All reliability engineering is done on the basis of “bang for the buck.”
4. Good load forecast	All the reliability engineering possible is wasted if the system is designed to a different demand level than required. Good weather normalization and spatial allocation methods are necessary.
5. Coordinate feeder plans	Chapter 10 will explore engineering of reliability on the primary distribution system. Effective, economical design of reliability requires coordinating sub-transmission – substation level reliability with primary distribution reliability, as explained there.

Evaluations determine and prioritize problems in the system as they will occur in the actual operation. Areas of the system, or operating conditions, that are especially at risk are identified. A big plus is that the method can be fitted with optimization engines to solve the key identified cases – “Find the lowest cost way to make sure this potential problem won’t be a real problem.”

Such methodology for reliability-based design of a power system can be

created using a highly modified form of N-1 (Figure 9.1) analysis in which probability analysis is used at every stage, or by combining N-1 with reliability-analysis methods. What is best for a particular utility or situation depends on a number of factors specific to each case. But the important point is that there are proven, widely used methods to perform this type of work available to the power industry, and in use by selected utilities. Such methods are a key aspect of solving aging infrastructure problems: the reliability of such system areas must be well analyzed and solutions to it well engineered. These types of methods will be covered in Chapter 14. Table 9.6 summarizes the overall recommendations on planning method improvement needed to meet modern reliability engineering needs.

REFERENCES AND BIBLIOGRAPHY

H. A. Daniels, *Automatic Contingency-Based Analysis of Power Systems*, PhD Dissertation, University of Texas at Arlington, 1967.

M. V. Engel, editor, *Power Distribution Planning Tutorial*, IEEE Power Engineering Society, 1992.

H. L. Willis, *Spatial Electric Load Forecasting – Second Edition*, Marcel Dekker, New York, 2002.

H. L. Willis and J. E. D. Northcote-Green, "Spatial Electric Load Forecasting – A Tutorial Review," *Proceedings of the IEEE*, February 1983.

H. L. Willis, T. D. Vismor, and R. W. Powell, "Some Aspects of Sampling Load Curve Data on Distribution Systems," *IEEE Transactions on Power Apparatus and Systems*, November 1985, p. 3221.

United States Department of Energy, *Power Outages and System Trouble (POST) Report*, March 2000, Washington DC.

10

Primary Distribution Planning and Engineering Interactions

10.1 INTRODUCTION

This chapter is in essence a continuation of Chapter 9's discussion about engineering paradigms and methods. Here, the focus is on the distribution system. The power distribution system (primary feeder system) accomplishes more of the power delivery function in most electric utility systems than any other level of the system. Power is routed into a few locations in fairly large amounts, and the distribution system then routes and sub-divides that power into smaller allotments delivered to many different locations. No other layer of the power system accomplishes so much dispersion over area or so much sub-division of power from large to small.

The primary distribution level also has more impact on customer power quality than any other layer. It is both immediately downstream of the lowest-level voltage regulation equipment in the system (substation voltage regulators or load-tap changers), and by the nature of its topography, the source of the majority of voltage drop that the customer sees. That topography also means that the majority of outages that the customer sees are caused by the distribution system – there are simply more miles of exposed line at jeopardy to weather, accident, and random failure. Thus, in most systems, the distribution system is the source of the majority of both the voltage-related and availability-related power quality problems seen by energy consumers.

This alone would guarantee that distribution should receive a good deal of attention anytime a utility needs to improve reliability, as for example when hard pressed by aging infrastructure issues.

However, the distribution system is also a resource that can be used to bolster reliability of the entire system: if combined artfully with the sub-transmission – substation level designs, can result in improved performance in terms of “bang for the buck.” Such planning requires, as was the case in Chapter 9, engineering methods different from traditional approaches, but methods that are proven, and widely available.

This chapter will begin with a look at distribution, the approach to its planning and operation that is typically taken in the industry, and what can be accomplished by making a shift to a different way of looking at and engineering the distribution system, all in Section 10.2. This will introduce two key concepts. First, the use of the distribution system to bolster reliability at the substation area through a built-in ability to transfer large loads between substations. Second, the use of marginal benefit-cost ratio balancing as an optimization tool in power system planning. Together these lead to a much improved cost effectiveness of the entire system. Section 10.3 will then summarize the engineering of distribution layout – how the spatial or geographic configuration of trunks, branches, and switching can be used to build a distribution system that has the characteristics needed. Many distribution engineers are unaware of the extent to which this aspect of distribution systems can be engineered. Section 10.4 concludes with a summary of key points.

10.2 DISTRIBUTION PLANNING AND THE PERCEIVED ROLE OF DISTRIBUTION

At its most basic, a power distribution system must be designed to deliver power over wide areas while maintaining reasonable voltage drop, and with some means of providing redundancy (contingency switching) should any one flow pathway fail. Naturally, it is desirable to perform this function at the lowest cost compatible with providing satisfactory power quality within standards. This defines the traditional goal of power distribution planning, and reflects the traditional viewpoint on the role of power distribution system.

During the early to the mid-20th century, every utility developed its own distribution planning method to satisfy these needs. At each utility these evolved during subsequent decades into very tight institutionalized standards for “the way we do things.” In many cases, these were nearly inflexible “template” or “cookie-cutter” approaches to layout style, conductor size set, and switching zone and conductor sizing. In all cases, they resulted in a distribution system that adequately fulfilled the traditional needs of distribution well enough to get by. The system routes power from substations to customers without violation of voltage, loading or power factor standards. It does not cost too much. It is

serviceable and expandable as needed in the future.

Remarkably, these approaches to distribution differ substantially among utilities, although all those distribution systems used the same components and equipment and all fulfill much the same role. At many utilities, power distribution planning is very parochial, built upon principles developed in-house decades ago, lacking any external influence, and never comprehensively re-examined and revised with respect to modern needs. The reasons have a lot to do with the distribution aspect of the utility business. Compared to utility's transmission planners, distribution planners seldom have to interact with neighboring utilities through intertie meetings or power pool planning groups. They are seldom exposed to the different ideas and their particular design habits and guidelines can be challenged as that could be improved. "We do it differently."

As a result, nothing is quite as tightly institutionalized, nor varies quite as much, throughout the power industry, as the quality of distribution planning. As an example, one large IOU in Florida has standards for conductor policy and feeder layout that are the exact opposite in several regards of the practices at a large IOU in the western United States. What is required at one utility is not permitted at the other. In fact, neither utility is correct – both paradigms are appropriate in some cases, but not all. A more effective and lower cost system results when flexibility is introduced into the design: when layout, switching, and conductor sizing are *engineered*, rather than designed by guideline. Section 10.3, later in this chapter, will summarize several concepts related to such engineering.

Frankly, engineering of a distribution system so that it performs the basic electrical functions that were traditionally required of it (i.e., electrical flow within standards, feasible manual switching capability for contingencies) is rather straightforward. At many utilities, this was done for decades with a standardized layout rule, a set of tailored nomographs and tables and some rough calculations (slide rules at best), all guided by a series of rote-memorized rules of thumb. When digital computers became available, load flow computations and spreadsheet-level tools were adapted by each utility to its particular distribution paradigm. The basic concepts underlying each utility's method -- how things are done and the fact that the engineering is largely based on a set of inflexible rules – remained in place.

Both because adequate distribution planning is so straightforward to do, and because it is done at many utilities with the fairly straightforward template-guideline-type of approach, one often hears the terms "it isn't rocket science" applied to distribution planning. This belief that distribution is easy to plan, and that significant and high-expertise effort is wasted in its engineering, is exacerbated by the facts that distribution is of relatively low voltage compared to other parts of the power system, that the equipment involved is small and inexpensive (compared to other levels) and, that the system is normally laid out in

radial form (which makes power flow easy to calculate). All of this tends to make distribution seem simplistic, and in fact it is, if viewed only from the traditional perspective, that the task is just to “get it so it works.”

Distribution Can Be Asked To Do A Lot More

However, that traditional interpretation of the distribution system’s job is obsolete and incomplete with respect to modern societal needs and technical capabilities. When reliability of supply is a major performance metric, and when cost must be minimized, not just kept within traditional bounds, distribution planning becomes very challenging. The important point to realize, however, is that the effort required meeting this challenge more than pays for itself: the gap between the performance of really good distribution systems, and merely “adequate” ones, is immense. This gap manifests itself in two ways:

- *Economics:* Advanced engineering and flexibility of approach can yield considerable reduction in cost.
- *Reliability:* The potential of the primary system to maximally contribute to improved reliability is partially untapped in the vast majority of cases.

A key factor in both is that the distribution system is both engineered well (optimized) and its design is coordinated carefully with that of the sub-transmission and substation levels in ways not traditionally done.

Table 10.1 (same as Table 1.3) shows the results generated by this approach at nine utilities with which the authors have worked. As a group these utilities are reasonable representatives of the power industry throughout North America. Together they demonstrate both the margin of improvement that can be expected by updating and modernizing primary distribution planning and design paradigms, and the substantial variation in improvement and type of improvement that can be seen from one system to another.

Table 10.1 lists four statistics for each system. The first, “Electrical - \$,” is the reduction in overall (lifetime) cost of distribution that could be wrought with respect to the traditional primary distribution paradigm of moving power. This is the saving in *new* capital additions that results from improving how the distribution system is designed to do its job – moving power with good reliability of service.

The second statistic, “Existing - \$” shows how much the utilization of an *existing* distribution system can be economically improved with respect to that traditional paradigm. Existing systems encompass both the good and bad points of their past engineering and design, and those systems cannot be thrown away and re-designed based on new, improved guidelines. However, improvement can be

Table 10.1 Improvements in Distribution System Cost and Performance Derived from Revision of Traditional Design Guidelines and Engineering Methods

Case	Utility System Type	Electrical \$	Existing \$	Reliability \$	System Reli.
1	Urban UG, Atlantic coast	8%	2%	79%	3.4:1
2	Rural, small towns, Central US	12%	0%	25%	.8:1
3	Metropolitan, Midwest US	6%	3%	40%	1.8:1
4	Entire system, IOU, Central US	2%	3%	36%	1.5:1
5	Rural, small towns, Southern US	12%	11%	25%	1.6:1
6	Entire system, N.E. US	22%	16%	33%	.92:1
7	Metro suburban system, US	3%	5%	62%	1.8:1
8	Metropolitan system, central US	7%	15%	42%	1.5:1
9	Statewide system, western US	4%	9%	34%	1.4:1

made in the MW/\$ capability of these systems. But such improvement is only valuable where useful (if one improves a 5.0 MVA feeder so it can carry 5.4 MVA, the increase is useless unless it is in an area where one *needs* a .4 MVA increase in capability). The values shown in Table 10.1 reflect both what is possible and what was found useful in each system.

The third statistic, “Reliability - \$,” represents the reduction in the cost of improving reliability that was attained by revising design guidelines and applying the most effective reliability-based planning and engineering methods to these systems. These improvements while rather dramatic, are typical. As stated earlier, distribution systems are simply not designed from a reliability standpoint. When this is done and the engineering is optimized, the results are a very considerable improvement.

The final statistic “System Rel .(Reliability)” represents the margin in “bang for the buck” that the revised distribution reliability improvement made over the existing cost of improving reliability that was in effect before the revision. This statistic goes to the heart of the use of distribution as a resource in aging infrastructure areas, and of the authors’ overall theme of using all of the power system chain optimally to obtain the greatest “bang for the buck.” In four of the six cases listed, reliability improvement was between 20% and 40% less expensive to buy at the primary level than at other levels of the system. In these cases, spending on the distribution system could deliver reliability improvement for less cost than spending on other levels of the system.

In two other cases, this was not the case, and the value shown is negative. Of note here is that the utilities were unaware of this entire issue, because the traditional tools they were using for the design of all three levels were not capable

of directly engineering reliability or of even measuring such things as reliability improvement per dollar. Thus, in all cases, improvements between 9% (1/92 – see final column) and 70% were possible by adopting reliability-based engineering methods and using them well at the distribution system and for coordination of design among the distribution and sub-transmission-substation levels.

A Quantum Difference in “Bang for the Buck”

Table 10.1 made it clear that revision of distribution system guidelines and adoption of reliability-based engineering methods often (actually in all cases the authors know of) provides a good improvement. The increases in cost effectiveness of 10% and 6% versus the traditional distribution function are noticeable and quite valuable. However, it is the improvement from the “modern perspective” – that which views distribution’s role as one of both performing the electrical duties and fitting into the optimum reliability versus cost structure of the entire system – where the improvement represents a quantum leap. By looking at the distribution system as a reliability resource, and optimizing its use for that purpose, an average 40% increase in cost effectiveness of reliability improvement is obtained at the distribution level. By coordination with the planning and design of other levels of the system, a 70% increase in overall effectiveness of money spent on reliability improvement is achieved.

The authors are aware that many readers, particularly some with long experience as distribution planners at power delivery utilities, will dispute that the improvements shown are real (they were), or will argue that their system is different (not likely,) and that “it won’t work here” (it will). The six cases shown are all real utilities implemented long enough in the past that time has proven the expected savings and performance increases to be valid.

The 6% and 10% savings against normal needs come about because modern techniques can find better ways to reduce costs than the formalistic methods of the past. But the quantum gains in reliability and reliability improvement per dollar come from the fact that because of its dispersed nature, the feeder system can often be reinforced at reasonable cost to cover weaknesses in reliability (contingencies) at the substation level.

The concept is easy to understand, if not to engineer (at least with traditional tools). Suppose that for whatever reason, reliability at one substation is a particularly degraded level of expected reliability. Perhaps the utility has no choice but to accept very high equipment utilization at the substation and there is a less than ideal configuration of buswork and switchgear there to provide contingency flexibility. The feeder system in this substation area could be reinforced to provide “more than the normal amount of primary feeder level reliability strength” to make up for this deficiency. Specifically, since this

substation will be out of service more than usual, the feeder system can be designed to transfer its loads to other neighboring substations, maintaining service during those contingencies. This is really what is necessary to provide good service in high-utilization systems, particular in the aging areas where substations are congested and highly flexible switching and buswork is not an option.

Good Distribution Planning /s Rocket Science

Despite the fact that what might be called “commodity” design can create a power distribution system that will work “well enough,” the power distribution system is by far the most difficult of all layers in the power system to plan and engineer to a truly high standard. Meeting this challenge is difficult, because distribution systems are like piranha. A single piranha is small, and not of great concern, but in aggregate, a group of piranha becomes something entirely different.

Like the piranha, almost every unit in a distribution system is small, of low voltage, commodity design, and low price. But in aggregate a distribution system becomes large, expensive, and most importantly, *very* complicated. This is due to both the sheer number of components (perhaps millions) and their interconnectedness or interaction with one another (particularly if viewed from the reliability perspective). A distribution system presents a tremendous challenge in *combinatorial decision-making* – how to select equipment, route lines, site key equipment and switches, taper conductor, arrange normal and contingency flow patterns, and makes a myriad of other decisions needed. The *combinatorial* aspect is tremendous, resulting from the fact that not only are there dozens or even hundreds of individual decisions to make in even a small distribution layout problem, but also because these individual decisions all interact with one another. No decision can be made in isolation: they all influence one another.

It is possible to sidestep this combinatorial challenge by using a set of standardized layout rules, templates on sizing and design, and inflexible guidelines on equipment and engineering. This was how utilities approached distribution in the 1930s, by sacrificing the gains in cost and performance that might be obtainable, in exchange for a tremendous simplification in the process of finding a workable plan for every situation. This template approach does assure that minimum electric requirements are met, and that cost, while perhaps not optimum, is not unduly high (perhaps no more than 10% above the minimum possible).

However, maximizing benefit, and particularly, tailoring distribution in each area of the system to fit the specific needs and the characteristics of other levels of the system there, means facing up to that combinatorial challenge. The possibilities within all that flexibility have to be exploited. That means dealing with all those interrelated decisions and obtaining plans that are close to optimality. Distribution planning, done to the most effective manner possible, *is*

rocket science. Section 10.3 will summarize design and planning methods that accomplish these goals in a practical sense - they deliver results in the real world. However, they require new ideas and slightly more work than traditional methods, which brings up a final point about modern distribution planning.

Feeder System Strength

The “strength” of a feeder system will be defined here as the percent of the substation load that can be transferred to neighboring substations during contingencies. This is a characteristic of a distribution system that can be designed to have a value of anywhere from zero to one hundred percent. The authors have worked with systems at both ends of the scale and everywhere in between.

There is no single, recommended target for the feeder system strength. It depends on the voltages, capacities, demands, configurations and constraints on the sub-transmission – substation level, the capabilities of the type of distribution system being built, and the need for reliable service as interpreted by the consumers and the utility alike. Figure 10.1 shows the customer reliability plotted




Figure 10.1 Feeder strength (% of the average substation’s load that can be transferred to three large metropolitan central systems in the United States). The “strength” of the feeder system, in terms of its ability to transfer a portion of each substation’s load to neighboring substations, is a key factor in reliability of most power systems.

against feeder transfer strength for the central (metropolitan) division of the electric utilities that serve three North American cities. As can be seen, in all three systems the transfer capability of the system is a key variable in determining the overall capability of the system. Further, feeder system strength required to achieve a specific overall level of service reliability for the entire system (e.g., a SAIDI of 100 minutes) is far different in each of these systems.

In fact, usually there is a need for different levels of feeder strength in different areas of the same utility system. Very often in older (aging infrastructure) areas of the system, a stronger feeder system is the optimal (lowest cost, most “bang for the buck”) solution to problems of reliability due to combinations of aged equipment and limited or congested substation sites.

So, while there is no recommended value for feeder system strength, what is always recommended is that the characteristics of the feeder system be tailored to, and its planning both on an electrical and reliability basis be coordinated with, that of the rest of the system.

The Systems Approach

It is possible through the use of an appropriate type of reliability analysis (Brown, 2009) to determine the impact of various levels of feeder strength on the overall customer-reliability delivered by the power system. For example, determining curves like those shown in Figure 10.1, and engineering the system’s reliability to achieve the desired reliability target at the total lowest cost is the goal. This means balancing how much is spent on reliability and exactly how it is designed into the sub-transmission, substation, and feeder levels, realizing that through appropriate arrangement of configuration, each layer can to some extent provide contingency capability for reliability problems encountered in the other layers.

A key aspect of applying the systems approach to reliability engineering is the use of *marginal benefit cost analysis* – computing for each level of the system how much a small increment of reliability *increases* costs, based on the present configuration and the options available. Comparison of the values among levels and areas of the system leads to optimization.

Figure 10.2 illustrates a simple example. In this diagram, one sees the reliability versus cost curves for each level of a candidate system design – how much it costs to build the system that will provide various levels of reliability (so many minutes of outage annually due to events or failures on that level). In every case there is an exponentially increasing cost to buy increasing levels of reliability (decreasing levels of outage time).

The top diagram has large dots indicating the cost versus reliability point of each level of the system for one particular candidate design, a plan which is expected to provide a total of 119 minutes SAIDI, at an annualized cost of \$172/kW peak demand (119 and \$172 being the sums, respectively, of the SAIDI




Figure 10.2 Top, cost versus SAIDI curves for the different layers of a power system. Dots show the status of each layer in a candidate design, which in aggregate provide 119 minutes of SAIDI annually at an annualized cost of \$172. Bottom, by adjusting the feeder system design to “buy” ten minutes of SAIDI improvement for \$5.00, and saving \$5.00 by cutting at the substation level, a net improvement of five minutes SAIDI is obtained at no cost: SAIDI is now 114 minutes and cost is still \$172.

improvements and costs of all the dots representing a potential system design). The important items for improving reliability are not the actual values (i.e. positions of the dots) for any of these points. Rather, the optimum plan depends upon the *slope* of the cost versus reliability curve at those points. For example, the primary feeder level design is at 57 minutes and \$24/kW annualized cost. However, the slope at that point indicates that spending an additional five dollars will improve reliability by ten minutes (moving the dot along the curve to the left by ten minutes raises cost by five dollars). By contrast, buying ten minutes at the substation level (see curve and dot for that level) would cost \$32.00. Clearly, the way to improve reliability at a bargain price is to buy improvements in the feeder system.

But beyond this, in this candidate design, one can begin improving reliability without spending any additional budget by “trading” reliability and cost between the feeder and substation levels. This can be done by “buying” ten minutes of reliability at the feeder level for five dollars, and “selling off” five dollars worth of reliability at the substation level, which because of the slope of that curve, means giving up only five minutes of SAIDI. Overall performance is improved by five minutes at no change in overall cost.




Figure 10.3 The candidate design is adjusted until all the levels have the same marginal cost (slope) of reliability (dotted lines). SAIDI for this design is 91 minutes at the same overall cost of \$172/kW.

Thus modified (bottom plot), the slopes of the feeder and substation levels are still not equal, so this “trick” can be applied again, buying 13 minutes at the feeder level in exchange for eight minutes at the substation. This brings SAIDI down to 110 minutes at no increase in cost – annualized cost is still \$172. Far beyond this, the points in all the curves can be adjusted in this way, with the result shown in Figure 10.3. The final result is a SAIDI of 91 minutes at the same \$172/kW annualized cost.

The example shown is a slightly simplified example based on Case 5 in Table 10.1. In the real world, the curves for each level have discontinuities and small but frequent discrete jumps, and there is some interaction and interdependence among the curves. But the concept shown is a valid and workable approach, and it leads to viable, noticeable improvements in reliability, at no net increase in cost. In this case, to accomplish these improvements, more was spent on the distribution system and less on substation configuration and redundancy: distribution was both designed to a different style (see large trunk versus multi-branch in section 10.3) and its strength was augmented in selected areas.

Changing the Distribution Paradigm

Chapter 1 (pages 28–30) discussed a number of “cultural barriers” that often stand in the way of improving power delivery system performance using “non traditional ways.” In the authors’ experience, the technical methodology required to successfully implement the distribution system performance improvements discussed above are far less challenging to most utilities than the cultural changes required to fit it into their organization. Considerable commitment from executive management, and steady effort to accomplish the change, are required.

10.3 FLEXIBILITY AND EFFECTIVENESS IN FEEDER LEVEL PLANNING

Generally, N–1 criteria and planning methods were traditionally applied to only to the sub-transmission – substation portion of the power delivery system, everything from the low side (feeder) breakers upward toward the wholesale grid (Figure 10.4). It was and still is uncommon to see that method applied to the distribution feeder system, although concepts allied with N–1 criteria often are. The primary reasons for this difference in methods on each side of the substation feeder breaker are:

Tradition – in vertically integrated electric utilities of the 1930s through 1990s, the Transmission Engineering Department most typically engineered all of the system above the feeder breaker, the Distribution Engineering Department engineered the feeder system and everything downstream of it. At most utilities they evolved different methods.




Figure 10.4 N-1 criterion (Chapter 9) is applied to the power system at the transmission, sub-transmission, and substation level analysis. Generally, feeder switching studies are done to assure “reliability” at the feeder level.

Applicability – The N–X study method (Figure 9.1) can not be applied to distribution feeder systems without significant modification of the traditional “modified load flow” approach. Distribution level application would involve representation of manual and/or automated switching operations, in precedence, which requires a much more involved algorithm, both in terms of internal complexity and user effort.

Different paradigm – Generally, at the distribution feeder level, reliability is addressed through switching redundancy design, as will be discussed below.

Feeder Contingency Switching

Traditionally, reliability-assurance methods for feeder level engineering varied greatly throughout the power industry, there being no standard or norm to which most utilities adhered. However, reliability engineering of feeders recognized certain characteristics at the feeder level:

- Feeder circuits are radial, and as such, an equipment failure leads to an interruption of flow of power to customers.
- Switching can restore service in advance of repair. Switches to isolate failed sections and to switch feed of power to alternate routes are provided at points throughout the feeder system, so that line crews can restore service prior to making repairs. Figure 10.5 illustrates this basic concept.

Multiple Switches Divide Each Feeder into “Switch Zones”

Although practice varied greatly among electric utilities both within the United States, and internationally, the traditional concept for “building reliability” into a feeder system design was to place several switches at points on each feeder to segment it, with portions transferred to neighboring feeders. Figure 10.6 illustrates this point. The tie switch near the feeder permits the feeder to be fed from an adjacent feeder leaving the same substation. (If the design of the substation-feeder system is sound, the adjacent feeder will emanate from a low-side bus different than this feeder’s, permitting service to be restored with this type of switching, even if the outaged feeder’s low-side bus is out of service.)

Alternatively, a second feed option is usually available on most feeders. The tie switch at the end of the feeder permits it to be supported from a feeder from another substation, as was shown in Figure 10.5.




Figure 10.5 Example of feeder switching. Top, feeder from substation 1 and feeder from substation 2 are both in service. In the middle, a line failure at the location shown interrupts service to all customers on line now shown as dotted (the entire feeder). Bottom, a normally closed feeder switch is opened to isolate the failed section, and a normally open feeder tie switch between the two feeder end points is closed, restoring service to a majority of the customers on the outaged feeder. Only those portions shown as dotted lines are now without service (they must await repair of the damaged line). After repairs are made, the switching sequence is reserved to go back to normal configuration.




Figure 10.6 Top, a single feeder in a feeder system, in this case one of four identical (symmetrical) feeders from a substation. It has normally closed (n.c.) and normally open, (n.o.) switches which tie it to other feeders. Bottom, these four switches segment it into “switch zones” in which each is isolated (if a failed line section is in that zone) or transferred to adjacent feeders during a contingency, so service is maintained despite the failure in one zone.

Voltage, Loading Criteria, and Line Upgrades for Feeder Contingencies

During a contingency, when loads have been transferred, the feeder picking up additional load will experience higher than normal loading levels and voltage drops on the re-switched portion of its system. Referring to Figure 10.5, bottom, in that switched contingency configuration, the feeder from substation 2 is supporting roughly about 180% of its normal load (all of its customer demands and about 80% of feeder 1's). Current flow on its trunk is thus 80% higher than under normal conditions, and voltage drop correspondingly 80% higher. In addition, the power flow to the portion of feeder 1 that has been picked up during the contingency is over a far longer route than under normal conditions – voltage drop will be even more than 180% of normal: it might reach over two times the normal voltage drop. Figure 10.7 illustrates this.




Figure 10.7 Two large-trunk feeders tied together for contingency support. The feeder on the right has experienced a failure of equipment at the substation (shown by an X). Switch B has been opened, switch at A has been closed, so that the feeder on the left supports both sets of feeder loads. The feeder on the left is now serving twice its normal load. Even though its conductor has some margin over peak load, the 200% load pushes it far beyond acceptable loading limits. Dotted lines show the points where voltage drop reaches range B limit using normally sized conductor. Voltages are symmetrical about the trunk, but the profiles are shown only for the northern side. This means voltages beyond the dotted line are unacceptably low, even when compared to the already relaxed standards applied during contingencies. Thus, to avoid unacceptable loading levels and voltage drops during contingencies, the conductor size on both feeders must be increased to a larger wire size, far beyond that needed for normal service.

Again, although standards and practices vary greatly throughout the power industry, often utilities will reinforce the trunk of each feeder (use larger conductor or cable than needed to serve the peak demand), so that capacity and voltage drop capability are available for contingencies. This reinforced capacity is not needed during normal operations; it exists purely to augment contingency capability. Its cost is justified on the basis of reliability improvement

Generally, utilities also apply different standards for what is acceptable loading and voltage drop during contingencies, than they do for what is acceptable under normal circumstances. While practices vary, many utilities use ANSI C84.1-1989 range A-voltages as design criteria for normal conditions on their system, and range B for contingencies scenarios. Range A specifies a 7.5% maximum voltage drop on the primary feeder system (9 volts on a 120 volt scale, maximum) during normal conditions (no failures). Range B, which applies during contingencies, permits voltage drops of up to 13 volts on a 120 volt scale (10.8% voltage drop).¹

Similarly, loading criteria during emergencies typically permit excursions past the limitations set on expected loading during normal service. A particular conductor might be rated at a maximum of 500 amps, a substation transformer at 24 MVA for normal service – whenever everything is operating as planned. Under emergency or contingency operation, however, the conductor might be permitted a loading of 666 amps (133%) for six hours, and the transformer up to 30 MVA (125%) for four hours. While not all utilities follow these specific standards, most follow something similar. The establishment and application of emergency loading and voltage standards is a recommended practice, so that while the same rigid requirements of normal service need not be met during contingencies, some minimally-acceptable standards do have to be met, to provide a uniform target for contingency planning and operations.

Traditional Engineering and Planning Feeder Systems for Reliability

Reasonable levels of feeder system reliability can be assured engineering each feeder into an appropriate number and location of switchable zones, and arranging the line pathway of neighboring feeders so that they have the capacity to support these switchable zone loads [Willis, 1998].

During the latter fourth of the 20th century, practices at the feeder system level among traditional regulated, vertically-integrated electric utilities were quite diverse. The quality of engineering and planning, and the amount of scrutiny and detail paid to reliability design at the feeder level, varied greatly over the industry. Some utilities had tough standards and rigorous engineering methods. Most, however, applied simple formula methods based on topology and standardized

¹ See the IEEE Red Book (1994). *Recommended Practice for Electric Power Distribution for Industrial Plants*.

layouts.²

Weaknesses in Traditional Systems

As was the case at the sub-transmission - substation level discussed in Section 9.2, as budget pressures increased throughout the 1980s and 1990s, utilities responding by cutting back on reinforcement of their feeder system as their loads grew, and accepted higher utilization ratios (peak load/capability ratio) as a consequence. This had an impact at the feeder level that is qualitatively similar to that at the sub-transmission – substation level discussed earlier. Higher utilization ratios eroded contingency margin. Situations where contingency currents were too high, and contingency voltage drops too low, became not uncommon in many utilities, limiting their ability to restore service quickly during storms and when equipment failed, and generally leading to degraded customer availability.

Furthermore, at many utilities where budgets were cut back significantly, planners were forced to use the contingency switching capability of the existing system to accommodate load growth. When the customer demand in one feeder area grew to where it exceeded the feeder's capability, a zone would be transferred to neighboring feeders. At several utilities (e.g., Commonwealth Edison of Chicago) widespread application of this stopgap approach during the 1990s "used up" much of the contingency withstand capability of the feeder system.

Optimizing Engineering and Planning of Feeder Systems for Reliability

As was observed earlier, adequate and even superior customer service level reliability can be assured by properly engineering the switching and capacity margin of the feeder system. This level of reliability can be accomplished by engineering each feeder into an appropriate number and location of switchable zones. A key factor of success is the artful arrangement of the circuit pathways for each feeder and its neighbors so that all have the capacity to support switchable zone loads, and so that there are a sufficient number of strategically placed intersections at which to put switches. This type of engineering can be involved and can require specialized techniques and approaches to work best, but there are proven methods in well- documented sources. See the *Power Distribution Distribution Planning Reference Book – Second Edition* (Willis, 2004), Chapters 13 and 14, for a detailed discussion of switch zone engineering methods and design trade-offs. As an example of the difference such methods can make, Figure

² For example, switches were not located so they would segment the feeder into zones designed for improved reliability, but located based purely on topology, at branching points and intermediate locations along long trunks.




Figure 10.8 Left, the traditional large trunk feeder design, a simple to engineer and operate feeder layout based on 1930s engineering technologies. It has tie points to feeders at both ends of the trunk (Figure 10.7). Right, a multi-branch feeder serving the same area, designed to the same standards, has more switches and switching zones, with ties to similar multi-branch feeders on both sides of it, and at its end points. It is more difficult to engineer. But it but costs less and delivers better reliability. The point is that switching and layout can be engineered to deliver maximum reliability per dollar.

10.8 shows a feeder layout called a short-contingency path, multi-branch feeder. This feeder has a more “dendrillic structure” (repeated hierarchical branching into more and smaller flow paths) as compared to the large trunk feeders shown in Figures 10.5 and 10.6. Large trunk feeders are the most traditional types of feeder design, having only one path that is used in contingency switching (their trunk). Large trunk designs date from the 1930s and are used, mainly due to tradition, as the layout concept for feeder engineering at about half of all utilities in the United States.

By contrast, the multi-branch feeder shown has a more complicated layout consisting of branches and sub-branches. These segment it into more and smaller zones. Basically, the feeder is broken into smaller “chunks” for contingency operation, and those chunks are transferred among several neighboring feeders during a contingency. This means:

A smaller portion of the feeder customer base will be unrestorable due to the failure. The switched zone containing failed equipment can never be transferred. Service in that one

zone can be restored only by repair. Smaller zones mean a smaller portion of the customer base falls into this category during every outage. Consequently, a higher proportion of the customers can be restored through re-switching, cutting average interruption time.

Less contingency margin is needed. Loading levels and voltage drops during contingencies are far lower than with large trunk designs. Such a system is more challenging to design, although with modern electronic feeder CAD systems, this is quite within the grasp of any reasonably trained distribution-planning engineer.

Short pathways mean lower voltage drop. Although not obvious from visual inspection, the increase in MW-mile of flow during contingencies in the multi-branch shown is much less than in the large trunk designs. This means less money must be spent on conductor upgrades for contingency capacity and voltage drop capability (Figure 10.7).

High engineering and operating complexity. A distribution system composed of this type of feeder does involve slightly more work to design and to operate well during contingencies. However, the widespread use of CAD systems for feeder planning and CAM systems for distribution operations makes such capabilities commonplace.

The net impact of the above is that the multi-branch feeder has a 12% better SAIDI and 8% lower capital cost to serve the same load, while requiring roughly 10% more engineering effort and the same overall O&M cost.

The point of this discussion is not to promote the use of multi-branch feeders (there are other types of layout types besides large-trunk and multi-branch, and all have advantages and disadvantages that need to be carefully weighed in each situation). Rather, what is important here is that feeder layout, switching, and reliability can be engineered well, using proven methods. How this is done will be summarized in Chapters 13 and 14.

Strategic Planning – Analysis and Optimization of the Utility Feeder System’s “Strength”

Feeder systems can be engineered to provide only contingency support capability for feeder-level failures or to provide support through the feeder system for substation outages as well. Providing support for feeder outages – designing the feeder system so that it has switching and contingency capability to provide

backup during failures on the feeder system, is certainly a best practice, and practically mandatory for urban and suburban systems.³ Layout, conductor size, and switching can be engineered so that all, or only a portion, of each feeder's load can be transferred to neighboring feeders.

The primary means of supporting a feeder's outage is to transfer its switching zones to adjacent feeders out of the same substation. This keeps the substation load as close as possible to actual, and is generally easier to accomplish in practice (line crews have to drive shorter distances to reach the various switches involved). It is also often considered a more prudent policy to transfer loads only between feeders of one substation, rather than between feeders served by different substations.⁴

Designing a Feeder System to Partly or Fully Support Substation Outages

If a particular feeder can be transferred so that it is completely served through a tie to a feeder emanating from a neighboring substation, then its electric consumers can remain in service during the failure of its substation (or the portion serving it). The ability to transfer the entire feeder load requires great strength among neighboring feeders in terms of feeder trunk capacity and voltage (factors such as discussed in Figures 10.7 – 10.8). A general capability to do this throughout a system – a feeder system designed so all feeders for any one substation can be transferred among its neighbors – results in what is termed a “strong” feeder system.

Clearly, to permit this to work well, in addition to a strong feeder system, the

³ Given that most feeder systems are radial, there is no other way to assure quick termination of customer interruptions caused by feeder failures. In the design of rural distribution systems, there generally is no switching capability, due to the distances and scarcity of the system. But in suburban and urban systems where there are neighboring feeders in all areas, the design of the system to switch around feeder-level outages is a given in almost all utility systems.

⁴ The issue here is mostly a desire, when switching back to the original configuration after the outage is repaired, of avoiding a short-term interruption of service to any customers. To do so, the utility has to use “hot” or “make-before-break” switching. The “to” tie switch to the permanent feed feeder is closed before the “from” switch connecting it to the temporary contingency feeder is opened. For a moment during this process the transferable section has two feeds. This avoids a short interruption to the customer, which occurs if the “from” switch is opened prior to closing the switch to the other source. Hot switching between feeders served by the same substation is generally “safe” in the sense that it will cause no loop flows or overloads. But hot switching on feeders of different substations can create high loop flows, if phase angle differences between the substations caused by circuit flows on the transmission system is large. Some utilities have a firm policy of not doing such switching between substations, and as a result, prefer to do switching only between feeders served by the same substation.

substations themselves must also have contingency withstand capability to accept the additional transferred loads. However this usually is not a design problem. In fact, use of a strong feeder system permits increased utilization ratios (peak load/capability ratio) at the substation level, i.e., lower contingency margin in substation designs. If a substation transformer fails, the remaining transformer(s) at the station do not have to pick up that transformer's load – it can be distributed over many transformers by transferring some of the feeders to neighboring substations.

Often, a utility will design its feeder system to support a limited amount of feeder load transfers to neighboring substations. For example, suppose that the outermost one-third of load on every feeder can be transferred to neighboring feeders. This is not an undue amount of load (one-third that of transferring the full feeders) and more importantly, is the portion of load closest to the neighboring feeders, meaning flow paths and voltage drops are lower, so that it is not typically difficult to arrange.

With this feeder system strength, the loss of one of two identical transformers at a substation can be supported, even if they are both loaded to 100% at peak, with an overload of only 33%. During the contingency, load amounting to 66% of a transformer's capacity is transferred to neighboring substations (the outermost one-third portion of feeders served by *both* transformers), leaving the remaining 133% to be served by the one unit in place. Overloading is equivalent to what occurs without feeder transfers, if the transformers were loaded to only 66%. The ability to transfer 33% of feeder loads results in an ability to increase transformer loading by 50% (from 55% to 100% of rating) with no increase in substation contingency stress levels.

Optimizing the amount of strength in a feeder system can be accomplished by balancing the marginal cost of feeder strength against the marginal savings in transformer capacity. Again, as throughout this book, the keys to success in gaining reliability/dollar are:

- *Systems approach* – in this case the feeder and substation levels need to be considered simultaneously, trading strength at one level against the other, in order to achieve reliability targets at minimum cost.
- *Marginal cost optimization* – as described elsewhere, the planners “buy” and “sell” reliability between levels until there is no bargain at either level.

In any real system, application of this design approach results in areas of the system with stronger and weaker feeders, and weaker and stronger substations respectively. These differences are due to constraints and factors unique to locations throughout the system. At substations where there is a limitation in capacity, a high utilization factor already, and very high marginal cost for any




Figure 10.9 Feeder configuration evolution. By selectively reinforcing segments and changing configuration, an existing large-trunk feeder (left) can be modified, or through gradual changes, evolved, into a multi-branch feeder with greater inter-substation tie capability (right). Engineering such changes so that they both work well, and fit within constraints of route and duct size, etc., and fit budget and work management schedules, generally requires a modern feeder system CAD software package with optimization capabilities. But, this is well within the capability of several existing commercial, feeder design software systems.

additions, reinforcement of reliability by use of a strong feeder system is generally the lowest cost option. Similarly there are areas where the reliability can be bought at the substation level less expensively than at the feeder level. But the fact is that the system if designed like this overall, delivers the performance required at minimum cost.

Increasing Feeder Strength to Augment Aging Infrastructures

Generally, in aging infrastructure situations it is impossible, or very expensive, to augment the substations – sites are small relative to need and tightly packed with equipment. Reinforcement of the feeder system to add strength is a good means of improving reliability overall, as it can usually be achieved within existing constraints and at lower cost. Such design, called configuration evolution, is not practiced by most utilities but is feasible and long proven in numerous cases (Figure 10.9).

Details of configuration evolution can be found in the *Power Distribution*

Planning Reference Book – Second Edition (Willis, 2004). The important point here is that the reliability of a system, aged infrastructure or not, can be managed through artful use of modern reliability-engineering methods.

10.4 SMART DISTRIBUTION SYSTEMS

Although the term “grid” is usually reserved for the high-voltage (wholesale transmission) portion of a power system, the term “smart grid” is widely used to refer to any and all parts of a power system that employ “smart” technology. But as was discussed in Chapter 1, for a system to truly be a *Smart Grid*, it must have two characteristics:

- 1) It must use smart technology: sensors, digital and analog electronic controls, and data communications.
- 2) It must have a revised or updated design of the power system itself and the way it is operated to take full advantage of the capabilities that the smart equipment provides.

These two capabilities are discussed in more detail below.

1 - Variation in something is the key to “smart”

The essence of smart systems is *variation* in some significant factor(s) of system behavior or settings on a context sensitive manner: the system *responds* to different situations it senses by varying its own equipment settings, configurations, or perhaps rates or limitations on demand in order to best fit those situations. A power system with widespread use of “smart technology” just to gather detailed information on loading, customer usage, and equipment status might be omniscient, but if does not do something with the information it gathers it is not truly “smart.” A smart system might increase the price of power during peak periods and lower it during times of lower demand. It might set protection into a “fuse-saving” mode at some times but not at others, based on the expected customer sensitivity to the blinks that fuse-saving protection schemes produce, based on a function of day of week and time of day, system configuration and customer loading. It might re-configure switching automatically in response to failures or equipment loadings, etc.

Such variation *is a requirement* for a truly smart system, and this means equipment installed on the power system must nearly always be able to vary its electrical performance. A capacitor bank must be able to turn itself on or off (perhaps in stages), a voltage regulator vary its tap ratios, a recloser relay vary its reply programming, etc. With this equipment installed, the system can now do something with the data it gathers and the inferences based on that which are developed by analytical engines in the system hub or in local equipment. Capacitors and voltage regulators on a circuit act in concert, cooperating (because

they sense what each other is doing) and working to maintain voltage profile, losses, and power quality at the best points possible regardless of circuit loading and configuration, and so forth.

2 - Re-Designing the distribution system to take advantage of “smart”

The traditional design for primary distribution in North America is to lay out circuits in a “big-trunk” radial configuration and to arrange sectionalization (where and how fuses opened if there were problems) in a cascading protection sequence up and down the circuit.⁵ As was discussed in Section 2.7, the chief advantage of the traditional big-trunk radial feeder approach is simplicity, of design, protection coordination, and operation, to a degree that first the needs of the period during which the industry settled on this approach (first quarter of the 20th century). Circuits can be successfully designed and equipment properly specified with a minimum of information and, quite literally, on the back of the envelope calculations.⁶ The system is straightforward to operate, dispatch, and refer in a safe manner with only maps as databases and common sense procedures used to assure safety. The industry stuck with radial distribution systems throughout the 20th century and into the 21st because they provided satisfactory performance for the time and no change was needed.

Considerable *improvement in every and all aspects* of a power distribution system’s performance can be achieved if a more complicated basic design approach is used as compared to simple radial configurations. System cost, customer service reliability, voltage stability, power quality, electric losses, and safety can all be improved at the same time. The price paid for the simplicity of

⁵ See for example, Chapter 20 of the *Transmission and Distribution Reference Book* (the original Westinghouse “Blue Book,” now published by ABB in a red binding) or Chapter 4 of the *Electrical Engineering Reference Book*, Volume 3 – Distribution Systems also by Westinghouse. Both books have been published in various editions since the 1950s and contain the most widely available discussions of traditional, pre-computer age distribution design available. For a more recent discussion including modern alternatives to traditional designs, see the *Power Distribution Planning Reference Book – Second Edition*, Chapter 14, particularly the portion after page 483 and the section title “Switching and Sectionalization – Managing Complexity and Cost.” Smart systems greatly reduce the cost and increase the efficacy of the high-reliability configurations discussed there.

⁶ This is not to say one cannot do a somewhat better job at assuring good performance and lowered cost using the greater data gathered by modern IT systems and the computerized tools in wide use today. However, with only a pencil and paper, and estimates of number of customers and peak load/customers, a competent engineer can lay out a radial distribution system so it will both worth satisfactorily when all equipment is functioning as intended and fail safe when something goes wrong. In a period when the most advanced tools available to power engineers were a slide rule and a book of logarithms, this was a necessity.




Figure 10.10 Left, a three-branch feeder somewhat similar to that shown on the right in Figure 10.9 provides better switching opportunities, operational flexibility, and potential service reliability than the big truck feeder shown on the left in that figure. Right, a four-zone feeder configuration uses more switches and complexity in transfers.

Arrows indicate the direction that switches would transfer power if the feed were lost at the substation. The four-branch configuration divides the load being restored both more evenly and into smaller portions. This generally means that the restoration can be accommodated more easily and for longer periods by neighboring equipment. For this reason the four-zone configuration will provide slightly better customer service reliability, particularly in a system with high equipment utilization rates. However it will be more complicated to operate, requiring more information on the part of the dispatchers, or the DMS system.

early radial distribution systems is a substantial margin in equipment and system capability that is often there but unused. In the early 20th century, this margin covered all the unknowns due to engineers and operators not having comprehensive load and system metering data at their fingertips, and for the approximations used in back-of-the-envelope calculations, so that circuits designed with linear models of non-linear processes would still work well. Figure 10.10 illustrates these concepts.

A system designed without that margin can be cheaper – perhaps up to 15% lower in overall cost – and more reliable with better quality of service. But engineers will need much more data in order to plan and engineer it. Modern smart data acquisition systems can gather this data, and modern Enterprise IT systems can provide it to the engineers in a way they can efficiently use it. Further, system operators will find the system a challenge to operate compared to the traditional radial system. It will require that they constantly adjust equipment settings in response to changes in load and system conditions, and that from time

to time they recognize the system is in a different operating mode and that they have to re-orient some of its basic scheduling rules and priorities in order to assure that it “fails safe” should something go wrong.

All of this can be accommodated, usually in a completely automatic way, by the proper use of smart grid equipment and designs. The distribution system is no longer entirely radial in configuration or operation, at least as regards all the configurations it can be switching to. It is deliberately designed so that it will have to change its configuration (switching) and equipment settings (voltage tap levels and capacitor switching) in order to function well throughout the course of a year or even perhaps a single day. Automatic equipment can do that. Protection and fail safe capability require that one unit of equipment “know” the status and loading of others nearby: data communication can provide this.

The net result is that a smart distribution system has better performance in every category than a traditional system, including the performance category cost. There is a price to pay, however. Smart systems are *much* more complex from nearly every standpoint, and require advanced design tools and planners, engineers, and operators of exceptional skill and training to plan, design, and operate them.

10.5 CONCLUSION

Distribution systems often provide fertile ground for cost-effective improvement of power system reliability. This is not only due to the fact that many power quality problems start at the distribution level and are therefore most effectively solved there. It is also because the distribution level can be used to provide load transfer capability between substations, thus reinforcing weaknesses in the overall customer service reliability chain due to weaknesses at the sub-transmission and substation levels. A distribution system’s strength – its ability to transfer load between substations – can be tailored as needed at locations throughout the power system.

In aging infrastructure areas, reliability-based augmentation of the distribution system is a particularly attractive option for improvement of overall customer service quality. Basically, the cost of “buying” reliability improvements at the primary feeder level is often far less than at the sub-transmission and substation levels, which are far more constrained. The distribution has a small granularity – it can be upgraded in stages – and essentially evolved into a system designed to bolster sub-transmission/substation level weaknesses. This is not to say that distribution can provide all, or even most, of the “cure” needed in aging infrastructure areas, but it is one of the most effective among the many options which a utility can assemble into a viable overall plan.

Explicit Distribution Reliability Methods

Given the increasing failure rates of its components, a distribution system will benefit as to its electrical and customer service reliability, and the utility owner benefit as regards business results, by some degree of re-design for higher reliability and to take maximum advantage of smart equipment and systems to augment system performance.

Distribution planning and methods traditionally used at most utilities applied standardized concepts, often not completely articulated, but instead institutionalized as tacit knowledge or culture, for how the system should be laid out, tables used, and guideline-driven design methods to simplify effort and engineering. While adequate to meet the traditional expectations for performance of the distribution system, these did not address reliability directly nor get the maximum possible from the distribution system in terms of reliability interaction with other levels, or economy.

For these reasons, modern planning needs are best met using power system design and planning techniques that directly address reliability performance. The most important points are:

1. *Systems approach* – planning of the distribution system must be completely coordinated with the sub-transmission and substation levels. The goal is not to design an optimal distribution level, or substation level, but to design an optimum system. Any one portion of that whole might be decidedly “non-optimum” until its context within the entire system is taken into account.
2. *Balancing of cost versus performance* among multiple levels of the system using adjustment of design based on achieving equivalent marginal cost of reliability. This results in near-optimum use of limited funds to “buy” reliability improvements.
3. *Engineering of configuration and switching* – a wide variety of possibilities in the overall layout of a feeder exist and should be reviewed in order to pick the best approach for the needs in each area. Very often standards at a utility will be so tightly institutionalized that engineers are not aware of neither other options nor the advantages they could obtain with flexibility of configuration and layout.
4. *Reliability-based engineering* – explicit rather than implicit methods for reliability engineering of the distribution system. Applications of these methods to power systems provided a much more dependable design method to achieve operating reliability.

Dependable, practical, methods for applying these concepts to aging infrastructure (and others types of) power systems have been proven in numerous

applications. While these are different than traditional methods, and require more skills and somewhat more labor, they are often essential in order to achieve any type of meaningful performance improvement from the distribution system within a reasonable budget limit. This chapter served only as a summary of modern distribution planning methods. Additional details are found in the references below, particularly in those by Engel (1992) and Willis (2004).

REFERENCES AND BIBLIOGRAPHY

ABB Power T&D Company Inc., *Electric Power Transmission and Distribution Reference Book, Fifth Edition*, Raleigh, 1998.

R. E. Brown, *Electric Power Distribution Reliability – Second Edition*, CRC Press, Boca Raton, 2009.

R. E. Brown, *Business Essentials for Utility Engineers*, CRC Press, Boca Raton, 2009.

M. V. Engel, editor, *Power Distribution Planning Tutorial*, IEEE Power Engineering Society, 1992.

Institute of Electrical and Electronics Engineers, *Recommended Practice for Electric Power Distribution for Industrial Plants* (the Red Book), Institute of Electrical and Electronics Engineers, New York, 1994.

Institute of Electrical and Electronics Engineers, *Recommended Practice for Design of Reliable Industrial and Commercial Power Systems* (the Gold Book), Institute of Electrical and Electronics Engineers, New York, 1991.

H. L. Willis, *Power Distribution Planning Reference Book - Second Edition*, Marcel Dekker, New York, 2004.

United States Department of Energy, *Power Outages and System Trouble (POST) Report*, March 2000, Washington DC.

11

Equipment Condition Assessment

11.1 INTRODUCTION

Assessment of the condition of power system equipment, and decisions about what service and utilization is appropriate for each unit, begins with the inspection, testing, and diagnostic methods covered earlier in this book. These basic tools for *condition assessment* provide the data, which when interpreted, lead to decisions regarding:

1. Whether the unit is suitable for service or must be withdrawn and repaired.
2. Whether more comprehensive, invasive, and expensive tests are justified before making such a decision.
3. The condition of the unit – if and how it deteriorates with age and service and its viability for future service.
4. To be effective, inspection, diagnostics, and testing must be applied in a carefully coordinated program that also uses the results from system operation (trouble reports) to identify overall equipment condition and performance, and that prioritizes preventive maintenance, repair, rebuilding, and replacement. Optimum results are derived by then loading equipment to levels determined based on its assessed condition, the loss of lifetime caused by the usage, and the available economic alternatives. Procedures for organizing and managing such programs will be discussed in later chapters. This chapter focuses on condition assessment inspection, diagnostics, and testing methods themselves.

This chapter discusses the concepts behind condition assessment and condition tracking, and their use, the testing needs and approaches for each major class of power system equipment, and how they are organized into a program of condition assessment. Sections 11.2 and 11.3 discuss condition and state analysis of equipment, and analysis and modeling of them for asset management purposes. Sections 11.4 through 11.8 discuss in turn the testing and diagnostic needs of various categories of equipment, in sections for transformers, breakers and switchgear, underground cables and equipment, overhead lines and equipment, and service transformers. Section 11.9 provides some detail and examples of assessment analysis methods in practice, including using test results to reach conclusions about equipment condition – and draws some conclusions about what testing and diagnostics can and cannot do. Section 11.10 provides a summary of condition tracking, which is typically where most organizations want to be with regard to the use of condition assessment. Section 11.11 concludes with a comparison of needs among the various classes of equipment and a summary of key points.

11.2 WHAT DOES “CONDITION” MEAN?

Condition means pretty much what any person or group using that term wants it to mean. Everyone agrees that “condition” refers in some way to if and how much a unit of equipment or a system is in working order and likely to remain there for a period of time. But one can cite a number of slightly varying dictionary definitions, standards group usages, or books like this one to defend a particular definition or push its use over others, and companies differ in how they use and view condition, beginning with how they define it. Here, it is defined as:

Condition is a meaningful gradation or classification based on all important characteristics of the unit of equipment related to actions required for its *proper management*.

Condition usually has a range of variation from something like *as new* to *very poor – withdraw from service immediately*, with several intermediate grades in between (Table 11.1). Conceptually, the transition from one end of this spectrum to the other is gradual but inevitable: a typical unit might be expected to go from “like new” to “serviceable but of concern” in perhaps five decades of normal use. As it does, aspects of its performance, its cost, and most important, the actions management should take with regard to owning and operating it, will change. When “as new” it is rated to take full summer load and certain substantial overloads during emergencies; when fifty years old those ratings might be reduced or mitigated. When new it is inspected once a decade; at fifty years it is inspected every other year. When new it gets basic service once a decade; at fifty years it gets invasive inspection and fix-as-needed every five years.

Table 11.1 A Condition Code Set Used By One Utility

Condition	Meaning	Actions
A	as new	Service and inspection as prescribed for “as new units” in Company Asset Management Guidelines
B	normal	Service and inspection as prescribed for “normal” units in Company Asset Management Guidelines
C	acceptable	Service and inspection as prescribed for “acceptable” units in the company’s guidelines
D	poor	Needs attention now. Service and inspection as prescribed for “poor” units
E	very poor	Withdraw from service immediately

In many cases condition is measured by an *asset health index* or a score similar in nature to that shown in Table 11.2: perhaps on a scale of 0 – 4, or 5 to 1, or as A, B, C, D, F, or very good/good/sufficient/poor/very poor, etc. The scores are determined from inspection and test results, operating records, and so forth as will be discussed later. In other cases, equipment age (years in service) is used as the condition measure: a utility has determined statistically an equivalent age score that it assigns based on analysis and/or inspection data. This is most accurately referred to as *equivalent age*: a badly deteriorated 35-year old pole might be assigned an equivalent age of 60 years, as compared to nominal deterioration rates, etc. That means it is as deteriorated as the average 60 year old pole. A pole installed 60 years ago might be rated equivalent age of 35 years – it is an outlier among 60-year old poles as regards a lack of deterioration.

Conditions Have Consequences

Regardless of how condition score is tallied, the qualities most often sought in a condition measure are that it is simple and easy to use, and that it relates to:

Ability to do its job: to serve demand, support a certain weight under certain wind loadings, break a certain amount of currently within a certain time, etc.

Durability: the expectation that it can continue in its current role, or not, for a foreseeable time: remaining service life, expected failure rate, etc.

Actions required by the owner/operator: run at full rating, or reduce loading by 20%, or withdraw from service immediately. Inspect every three years or inspect annually. Service this unit immediately.

Conditions have consequences, either with regard to operating capabilities and restrictions, or actions needed to be taken (or not taken) for the equipment (Figure 11.1).




Figure 11.1 Relationship of state, conditions, and consequences and actions as it relates to the management of electric T&D assets.

Condition, State, and Priority

State is a technical description that is unambiguous in context. In theory, state refers to a set of data that comprehensively defines a system so that its response and behavior can be completely understood and predicted. From a practical standpoint, state usually involves several measures that are needed to understand and define the status of a unit unambiguously for modeling, analysis, and prediction.¹ Often these state factors are component related: the “state” of a power transformer involves its core, bushings, oil, case, LTC mechanism, and cooling systems, etc. Failure or poor status of any one could lead to condition assessment of poor – needs attention now.

Condition assessment is ambiguous as to the details. State is not. A condition of “poor – needs attention now” might mean that oil readings are outside of permissible range, or that DGA readings indicate the core is deteriorated, or that the case is rusted and leaking, or that bushing diagnostics indicate a developing problem. Condition ratings are ambiguous in this situation, but the point is that these details are not needed to address the three managerial aspects listed above. By contrast, the state assessment or model will provide all those details: the core is like this, the oil like this, the case like this, and similarly for the bushings, LTC, cooling, and all other equipment associated with the unit. Even for a commodity unit such as a pole, state might give more details behind the “condition.” (This pole is located in an acidic soil area. This pole is in a high risk location.)

Condition is unambiguous as to actions needed; state might be too complicated to immediately determine that. But the value a simple condition score has is that it is unambiguous as to management action required – *attention is needed now*. Condition is the most succinct and unambiguous scoring possible

¹ Model, as used here, is not necessarily a computer program or even a numerical representation. It could be only the concept used to predict behavior: “My boss is a stickler for detail and abrupt in the morning, particularly on days when she has meeting with her boss. She is easier going in the afternoon, particularly when she is tired. She can be distracted from an unpleasant topic by bringing up the Duke women’s basketball team. . . .”




Figure 11.2 Methods that make predictions and analysis based only on condition tend to do much poorer as planning and management tools than models that use state representation. See text for details.

that achieved that one quality, and that is all that should be expected of it. The first step in the attention the unit receives will almost certainly be for someone to gather up the state estimates and the data itself, and examine them in detail, but the condition score has done its job at that point. An important distinction is that modeling and prediction should always be based on state, not condition, as illustrated in Figure 11.2. Condition, by its nature, is not as useful for modeling and prediction as state. State, by its very nature, is not as useful a summary of management issues as condition.

Priority is Used for Optimizing or Trying to Optimize of Results

Priority, or prioritization grade is a rating of importance with respect to spending and budget. Usually priority is based on both condition (often including details depending on state) and a unit’s importance (reliability, grid systems security or




Figure 11.3. Prioritization is based on condition (or state) and importance (potential impact on reliability, system interconnected security, and operating results). This matrix summarizes one utility’s representation of the priority it gives to equipment based on both evaluations, and is used by it to communicate how it manages its assets. Equipment in poor condition is withdrawn from service regardless of criticality. Critical equipment that is of concern is given high priority of attention. A detailed numerical prioritization that also takes into account the cost of actions and the cost of consequences (good and bad) is actually used to make all final decisions on what to do and where to focus resources.

customer service needs). Units that score highly in terms both condition needs and high importance are given priority over those that do not. Figure 11.3 shows a prioritization matrix as used by one utility. Prioritization will be discussed further in Section 11.10.

11.3 ANALYSIS, MODELING, AND PREDICTION

This book is not intended to focus on numerical or logical analysis and modeling methods or to present computer algorithms for such. Organizations and professionals doing asset management have or can gain the skills and knowledge in those areas that they need to properly apply them to aging infrastructure asset management. The bibliography at the end of this chapter provides several of the authors' favorite references on those subjects. The discussion over the next few pages will present two views on how numerical and logical methods are applied to aging asset base analysis and modeling, in order to put the use of these methods in context and make several points the authors think important.

Figure 11.4 shows the flow of *knowledge creation and application* from raw data through to its final product, a model that is supporting asset management decision making.

Data include operating records, inspection, test, and diagnostic results, maintenance and service reports, load readings, failure reports, monitoring system records, etc. The vast majority for data in most asset management databases is *numerical*: measured and recorded as numbers in some range or unit scale, etc. Such data are analyzed and handled using numerical methods. A very small portion may be *logical data*: values that are binary -- yes/no or 1-0. These are handled using logical methods. Some data might be *gradings*: recorded as excellent-good-poor-bad, or high-medium-low-very low, or A, B, C, D, F, etc. These are best handled, at least if rigor and fully defensibility is important, using what are called grade logic methods – what was once called fuzzy logic.

Aging infrastructure analysis and asset management are data intense processes. There is (or should be) *lots* of raw data: equipment purchase and installation records, equipment outage and failure records, operating records, incident reports, test and inspection results, as-found and as-left maintenance reports, etc. Hopefully there are numerous sources. There is also hopefully not only data for recent years but data going back decades. Twenty years of data is more useful than ten, and forty is much more useful than twenty. In a few cases the authors have used data going back a century in their analysis of aging T&D infrastructures.

Less than perfect data is a major issue in nearly all asset management programs. Available data is seldom in “good shape.” Invariably, it has gaps where no data was recorded or retained, “patches” where someone filled in data gaps using a (now) unknown estimation method (but one that leaves obvious traces), inconsistencies (e.g., definitions or categories used in reporting were changed in 1994), incompatibilities (the purchase records list poles only by class, the maintenance records categorize them only by type of wood and preservation), or




Figure 11.4 Analysis converts data into *information*. That is *organized* into knowledge, then used with *structure* in models that guide decision-making and management by predicting the consequences of future actions.

data was measured using a different method for 15 days while monitoring equipment was being repaired, before returning to normal type, and mysteries (there is an unusual jump in values starting 01/01/04).

Beyond this, data often have significant quantification errors. For example, inspection forms might provide places to record surface rust as minor (less than 10%), quite noticeable (up to 33%), serious (up to 66%), or complete (100%). Where does 50% rust fit in this range? Different inspectors will use slightly different interpretations, even if trained as a group. Much of this will be discrete sampled data – values taken every so often – hourly, monthly, annually, which raises a host of issues relating to if and how some numerical methods are legitimate or not. Further, the discrete measurement may be truly discrete sampled data, or period integrated data, a distinction that can make a very large difference in their value for some applications and models.²

Analysis has two purposes from a practical standpoint. The first is to determine and “fix” data quality issues. The second is to find useful information within the data.

- 1) *Data quality* can be determined using various filters and diagnostic methods that detect errors or anomalies in the data. Problems can be “fixed” but exact knowledge of the repair mechanism needs to be

² Truly discrete measurements record the actual value of a factor at the moment it is measured- load at this instant is X, temperature is Y; integrated readings record the total or average for the period: total energy used was Z, total thermal loading was L.

retained. Gaps can be filled using pattern recognition and regression methods that provide a best estimate of missing values. Inconsistencies among definitions can be assessed and adjusted, and incompatibilities can be addressed. It is important to realize that if carried out correctly, these “repairs” do not add any real information, but merely make the data compatible with processes and models to be used later. For example, if three-quarters of quarterly DGA readings are missing from a particular transformer’s records, they can be filled into the records by a simple straight line trending the data from before and after.³ This puts numbers in the data set so that models using them will not “go ballistic” when they see completely unreasonable (zero) values. The data is fixed in a functional manner but not as to the lost information in those missing values.

- 2) *Analysis extracts information* from data. The data may be large streams of historical values. Analysis will find information: “Failure rate increases with age in this way Failure rate increases with the number of through faults within 20% of the design limit in this way” “OH lines that have had tree-trimming within the last two years have 40% fewer outages per mile annually.”

Information – patterns and relationships found in the data – is not necessarily always useful. Technically, information is just non-randomness in the data. From a practical standpoint, it is valuable only if applicable to the interests of the person or organization using it (asset management, in this case). Therefore, analysis is usually *directed*: the analysis does not seek to find all patterns and relationships in the data, but instead focuses on finding patterns and relationships that tie directly to managerial targets and factors under its control.

Generally data analysis is directly at finding applicable information: relationships between data and factors used in or important to the asset planning and management process, so it is relevant to decisionmaking. For example, analysis of data for a utility’s asset management program would be directed at specific desired knowledge such as:

- *Is maintenance effective?* Can one see a reduction in failure and problem rates among equipment that has recently been serviced or treated, as compared to equipment that has not?
- *What leads to failure?* Determine if and how failure rate varies as a function of age, peak loading, climate area, and maintenance levels, etc.

³ A better estimator – one that minimizes expected error, generally can and should be developed using other data such as cumulative loading, temperature, too, but this example provides a simple illustration of the concept.

- *Playing “data detective,”* looking for useful clues or patterns. One of the most useful patterns the authors look for is transformers that have had case rust severe enough to require repair will rust at a noticeably higher rate when put in service, than those that have never been repaired.⁴

Information is identified, or extracted, from data through this analysis, creating information of many types. There are clusters in the data (“This group of 177 transformers are all quite alike with regard to their measurements, and somewhat different than any of the other 2,372 units”), patterns (“This type of unit has no significant operating problems due to rust for the five years immediately after surface rust is first detected”), and relationships (“Cable failure rate is equal to this polynomial function of time-in-service, peak line loading, and per unit bus operating voltage”).

Information found in this way, particularly through analysis directed at asset management application, is almost always *potentially* useful. But it is the conversion into knowledge that makes it applicable, and therefore useful.

Knowledge is information that has been organized and structured so it is useful and applicable. In Figure 11.4, the authors identify this with the term “Organization.”

Classification is the first part of organizing and transforming information into knowledge, grouping assets or situations into classes or categories, or labeling and putting patterns found into a framework. This involves much more than just obvious equipment classes (transformers, breakers, etc.) which actually are based on *a priori* knowledge (asset managers knew this before data analysis began). Here, cluster and pattern results lead to useful groupings or distinctions about assets:

We identified two groups of transformers by their relative measures: wet transformers have both high rust rates and water-in-oil contamination rates; dry transformers have neither. It is quite easy to distinguish members of either sub-class from the greater mass of power transformers, and within each sub-classes, formulae relating data to failure and problem rates and maintenance needs have much tighter standard deviations than for transformers overall.” Defining sub-classes around these cluster definitions leads to more useful application of the data.

⁴ The authors had seen this pattern at several utilities and across several classes of equipment, not just transformers. The pattern and accelerated rate of continued corrosion are not due to poor quality of repairs. New metal may be welded in, and the entire case sanded, primed, and re-painted, yet the unit will rust again sooner and at a higher rate than average. Investigation shows that units that develop rust are just more rust-prone and/or in an area or particular site where rust is more common. This will be discussed in more detail later.




Figure 11.5 The four major aspects of data mining, which is applied with a formalized structure and guided by theory and rigorous method, includes many of the concepts discussed earlier for analysis and knowledge synthesis. A data mining approach often takes the perspective that information, knowledge and model structure are developed jointly and gradually throughout the process.

“Transformers less than 20 years old with HIGH rust problems
HIGHLY LIKELY to have HIGH levels of water contamination
of their oil.”

Rule-identification. The second aspect of the information-to-knowledge transformation is rule determination. Information is culled to identify if-then, this-implies-that or functional relationships that are useful in asset management, and put into a framework where it can be applied.

“Failure rates and problems for this equipment class respectively
best fit the following two functions of the data”

“When surface rust is LOW and peak loading is NOT HIGH, the
consequences of NO INSPECTION are LOW.”

“Condition of equipment in this class and category can change by
32% within six months, so we need to set up our inspection on a
semi-annual basis.”

“This type of unit has no significant operating problems due to
rust for the five years immediately after surface rust is first
detected, so we will schedule inspections and testing only 60
months after the first report of surface rust.”

Classification and rule-identification organize the information so it can be effectively applied to asset management. It is now knowledge.

Models apply knowledge. As stated earlier, a model is not necessarily computerized: human beings aware of the patterns, correlations, and regression results found by data analysis can keep them in mind and apply them in their work. But computerized models apply knowledge consistently and completely (at least if well designed and written), and economically and (again if designed and written well) in a self-diagnosing and self documenting manner. Thus, the classifications, rules, and formulae found in the analysis and organization of the information are embodied into algorithms or logic used in a computer program, an unambiguous definition of the knowledge and its method for application.

One benefit of computer model building is it forces an organization to rigorously define the classes, rules, and methodology of its knowledge. Mental models can be and often are a bit exact or inconsistent in application from one situation to another. Computerization forces rigor and documentation of that throughout.

Prediction. From a practical standpoint, the purpose of models is prediction: they will be applied to anticipate future trends and problems, to study the consequences of plans, and to make decisions about resource and budget allocation, for the future. Models are not used for analysis of past or present data (analytical methods discussed early may be used on an ongoing purpose for this purpose).

Prediction here is used in its largest sense, meaning much more than just projections for future values for the current situation (“Failure rate for this asset class will increase 13% over the next decade”). Prediction includes forecasting how *changes in controllable factors will change results* (“Cutting inspection periods in two will result in identifying 52% of incipient problems before they cause outages”).

A Data Mining Perspective

A modern term for analysis and use of very large data sets and their conversion to useful knowledge is *data mining*. It can and is described in various references in slightly different ways, but regardless is always a process that takes data all the way to useable knowledge, usually in stages. Figure 11.5 shows the authors depiction of it based on several common web and technical references (Wikipedia among them). Four techniques are used to develop knowledge from data, usually in stages. These are similar overall to the analysis-to-knowledge path described above. .

Clustering is finds patterns or groups of assets that are similar to one another, etc.: “This set of transformers has operating and histories that are very similar and distinctly different than average.” “Failure rate of an OH service transformer is 20% higher in the three months after a lightning outage on the line within a quarter mile of its site.”

Classification seeks to identify groups and patterns and put them in context with regard to the framework needed: “We have two groups here: wet transformers have both high rust rates and water-in-oil contamination rates; dry transformers have neither. There are very few transformers not in one group or the other.”

Regression attempts to find a relationship between variables: “Annual failure rate is equal to this function of time in service, annual peak load, climate area code and manufacturer code.” Readers should note that this term does not mean just the classic multiple linear regression method of polynomial curve fitting of data (it is among the simpler and more direct of dozens of methods used); Regression includes all methods that seek to identify a functional relation among variables, and includes grade logic methods: “When all factors are GOOD, the problem rate for the next year will be LOW, MOST of the time.”

Rule Development uses patterns, classifications, and relationships found by clustering, classification, and regression to build applicable rules and guidelines. “Wet transformers over 15 years in service should be inspected for rust annually and the oil tested for contamination bi-annually.”

Data mining is not just a term but a professional practice area with established theory and rigorous methodology that, if followed correctly, generally guarantees

good results, or at least freedom from the effects of many common mistakes and oversights that are often made in data analysis. An example of the methodology is *relational normalization*, which does not refer to numerical normalization (i.e., adjust all values to a scale of 0-100), but to methods that adjust the structure of the database and its rules to avoid duplication, inconsistent updates, conflict among different parts of the mode, and error, and logical completeness which assures that the resulting model covers the range of expected application fully, without gaps.

11.4 POWER TRANSFORMERS

Transformers are one of the basic building blocks of power systems. They alter the voltage-current constitution of alternating current power passing through them, essentially changing the economy of scale of transmission of the power from one side of the transformer to the other.⁵

Power transformers at distribution substations constitute a major category of distribution equipment. Generally, their capacity is summed throughout a system, or at a substation, in order to determine the capability of that system. This is only one of many indications that they are considered *the* key aspect of a distribution system. Regardless, large power transformers are a major concern to any electric utility when it comes to reliability evaluation, because each one feeds large numbers of customers and its replacement would involve a considerable amount of time and expense. Much of the contingency planning of power systems relates to concerns about transformer loading limits, and arrangements for contingency operation when one fails (see Chapters 8, 9, and 13).

Transformer Aging and Deterioration

Transformers “age” in the sense that the strengths of their components deteriorate as a function of chronological time, time in service and loading, and due to severe abnormal events like through-faults and switching surges. The amount of load served has a great deal to do with the rate of deterioration. Chapter 7 discusses in detail aging and loss of life with service time.

Many aspects of a transformer deteriorate with time: 1) the insulation level of its windings; 2) its oil and its bushings; 3) the mechanical strength of its core stack and internal bracing and electrical connections; 4) the desired chemical and physical properties of its materials (anti-oxidants in the oil, corrosion resistance of paints, flexibility of insulation, etc.). However, deterioration of the insulation (of both winding and the oil itself) is the major area of concern, and the predominant area evaluated, to determine a transformer’s “condition.”

There are several reasons for the focus on internal insulation strength. First, a core winding insulation failure will not only prevent the transformer from doing its job, but probably lead to a large fault causing very severe damage if not catastrophic failure. Bushing failures, corrosion, and other types of deterioration can lead to failure of the device, but rarely cause such high levels of damage.

⁵ See the *Power Distribution Planning Reference Book*, Chapter 7.

Secondly, many other defects, when found, can be repaired or replaced rather quickly and inexpensively. Bushings can be replaced, tap changers repaired, and most non-core/winding related equipment serviced while the unit is left in place but de-energized. However, the winding insulation is both the most difficult and the most expensive item to repair in a transformer. The unit must be withdrawn from service for a long period of time and returned to a factory or refurbishment center for what is essentially a complete rebuild.

The insulating oil in a transformer, if significantly degraded, can be reconditioned or replaced. This will improve the oil and have a positive effect on the condition of the winding cellulose insulation. But in some cases the winding insulation is sufficiently deteriorated so that replacing the oil alone will not necessarily make a satisfactory improvement in the transformer as a whole.

For all these reasons, the main focus of power transformer condition assessment is on insulation strength, the various tests focusing on direct (measures of dielectric strength) and indirect (measures of contaminants, dissolved gas analysis) means of determining the winding and oil insulation strength. However, other portions and aspects of the unit should be checked regularly and maintained as needed.

Acceptance Tests

New units are usually tested for insulation resistance, dielectric strength, turns ratio, amount of losses, and run through AC hi-pot and power factor tests. They may be subjected to load cycle (thermal tests). Acceptance tests are not a factor in condition assessment of older equipment and will not be discussed here.

Routine Transformer Inspection

Routine physical inspection of transformers includes : Examining the exterior of the unit for signs of leakage or corrosion of the case; examining radiator joints and grounding straps, etc., for cracked or dirty bushings, for damage (vandalism, weather), deterioration (paint chalking), and loosening of brackets or seals for attached sub-components (radiator fans, pressure sensors). It should also include testing the operation of ancillary devices such as oil pumps, fans, pressure relief valves, and the tap changer mechanism and its control system. This level of inspection should be done annually, if not on a more frequent basis.

Overall, inspection tends to focus on aspects of the transformer other than insulation strength and winding condition. Inspection may include viewing the unit through an infrared scanner to enhance detection of hot-spots caused by loose terminations of incoming lines, leakage current, or other similar developing flaws. Similarly, it can include enhanced “audible noise” analysis using a spectrum analyzer and pattern recognizer programmed to identify acoustic signatures indicative of trouble.

Routine inspection should include a thorough review of all operating records (temperature, pressure, load records) since the last routine inspection, with particular attention paid to abnormal events such as nearby lightning strikes,

through faults, and storm damage.

Depending on the results of these inspections, follow up tests or diagnostics aimed at determining more precisely the condition and any necessary maintenance may be required.

Routine or Periodic Tests

Along with routine inspection, certain tests are done on a periodic basis. These tests have two uses:

1. “Good to go” evaluation. The tests determine if any measured values deviate from their permissible or recommended ranges. This indicates a need for further testing or for maintenance, and caution in using the device at high loadings. Values above a certain event lead to recommendations that the unit be withdrawn from service.
2. Condition monitoring. Periodic values are compared to identify long term trends – a particular value (e.g., dissolved methane in the oil) could still be below the maximum permitted – but significantly higher than its value in the last test, indicating a possible problem just beginning to develop.

Many of the most common periodic tests and diagnostics for a transformer focus on measuring some aspect of insulation strength, or look for the products or signs of insulation and/or oil deterioration. Table 11.2 lists the periodic inspections and tests that are done on power transformers. The frequency of tests shown are the authors’ impressions of typical industry practice (not necessarily recommended practice) with regard to units that are not thought to be high risk. Units that are suspected of having problems or a high potential for failure should be checked much more often – in extreme cases on a daily basis if not continuously (with real time monitoring).

Two guides useful sources of information on transformer tests and their use are *Fifty Years – A Guide to Transformer Maintenance* by Myers, Kelly and Parrish (1981), and *Service Handbook for Power Transformers*, by ABB’s Transformer Remanufacturing and Engineering Services Division (2006).

Four Categories of Risk/Recommended Power Transformer Evaluation

IEEE Standard C57-1004-1991 classifies oil-filled power transformers into four categories or conditions of risk based on the results of dissolved combustible gases tests in their oil. Within each of the four categories, the standard then recommends specific inspection and test intervals that vary depending on the *rate* of change of the gas content as measured from interval to interval. It also contains recommendations for use of units and if and when they should be tested again and if they should be removed from service, based on the gas test results.

Table 11.2 Inspections, Tests, and Diagnostics Performed on Power Transformers

<u>Monitor for . . .</u>	
Liquid level	Continuously
Load current	Continuously
Temperature	Continuously
Voltage	Continuously
 <u>Inspections and Tests</u>	
Exterior for signs of damage, deterioration	Routinely, quarterly
Interior for signs of damage, deterioration	5 – 10 years
Ground connections	Semi-annually
Lightning arresters	Semi-annually
Protective devices and alarms	Semi-annually
Radiators, pumps, valves, and fans	Semi-annually
Tap changer function	Semi-annually
Other exterior ancillary devices	Annually
 <u>Solid Insulation</u>	
Hi-pot (AC)	5 years
Induced voltage	5 years
Insulation resistance	1 - 3 years
Power factor	1 - 3 years
Polarization index and recovery voltage	1 – 3 years
 <u>Insulating Oil</u>	
Acidity	Annually
Color analysis	Annually
Dielectric strength	Annually
Interfacial tension	Annually
Power factor	1 – 3 years
TCGA	Annually
 <u>When Condition Is Suspect</u>	
All inspections and tests above	Immediately
DGA (gas chromatography)	Problem identification
Insulation resistance	If winding faults suspected
TTR	If winding faults suspected

Table 11.3 IEEE Standard C57.104-1991 Power Transformer Risk Conditions

Condition	Meaning
1	Transformer is operating properly.
2	Greater than normal gassing. Suspected higher deterioration rate. Additional monitoring and testing is warranted.
3	High level of decomposition has almost certainly occurred. Internal faults are likely. Periodic gas analysis should be used to establish a trend. The unit may need to be scheduled for repair.
4	Internal faults are very likely. Continued operation may result in failure.

Table 11.4 Summary of Risk Categories and Recommended Actions Based on TCGA Tracking, from IEEE Standard C57.104-1991

Risk Condition	Defined by TDCG (ppm)	Rate of Increase (ppm/day)	Means Operating Recommendations of . . .
1	< 720	< 10	Good for normal operation, and test again annually
		> 10	Good for normal operation, but test again quarterly
2	721 – 1920	< 30	Exercise caution in use and do further tests soon
		> 30	Exercise caution in use of unit and do DGA tests to determine causes.
3	1920 – 4630	< 10	Schedule an outage to check unit, and test monthly until then. DGA tests.
		> 30	Exercise extreme caution in use of unit and test weekly
4	> 4630	< 10	Exercise extreme caution and test weekly, schedule repair soon.
		>30	Consider removal from service now, and test daily until then

Table 11.5 Limits to Amounts of Individual Gases in Each Category

Risk cond.	H ₂	CO	CO ₂	C ₂ H ₂	CH ₄	C ₂ H ₄	C ₂ H ₆	TDCG*
1	below 100	350	2,500	35	120	50	65	720
2	up to 700	570	4,000	50	400	100	100	1,920
3	up to 1,800	1,400	10,000	80	1,000	200	150	4,630
4				above level 3				

The total dissolved combustible gases total does not include CO₂, which is non-combustible.

The full standard gives more detailed breakdowns of recommended testing frequency versus TDCG totals and increase rates than shown here. Table 11.3 summarizes the meaning of the four categories, which are defined by the totals shown in column two of Table 11.4. The table also shows how rate of increase data is interpreted in each of the four categories. Table 11.5 gives the concentrations of various gases (which must be determined by DGA) that are considered tolerable within each category.

Thermal Load Tests

The electrical losses created inside a transformer produce heat, which leads to elevated temperatures unless the heat is moved away from the core with some means of cooling. If the unit has internal problems that create higher than expected amounts of heat, the unit might operate at temperatures noticeably higher than it should for a given load. Sources of excess heat include short circuits, severe sludging problems, damaged radiators or damaged pumps. This could shorten lifetime significantly and bring about premature failure.

Units can be “laboratory tested” to determine their temperature versus load characteristic by running test currents of known amounts through them for known periods of time, while measuring their internal temperature rise. However, such tests are expensive (the unit has to be de-energized, taken to the lab, and left there for a considerable period).

Nearly the same accuracy in determining thermal behavior of a transformer can be carried out in the field. This is done by monitoring load, ambient temperature, and temperature inside the unit on a periodic basis (every 15 minutes) and then calculating the thermal time constant and the maximum rise expected at full load from the resulting data (Figure 11.6). The actual computation is straightforward, although it is generally best to gather much more data than theoretically sufficient – several weeks of readings. Use of an analytical method based on signal processing, rather than time-series analysis methods, is also recommended.

This computed “thermal response” signature of the transformer will reveal whether the unit is performing up to its design specifications or if thermal capability is degraded for some reason. Typically, such field tests will not indicate the cause of the problem, only that a problem exists. For example, Figure 11.6 gives no indication whether the deviation from expected performance is due to a degradation in cooling capability (due perhaps to sludge buildup or blocked cooling ducts), or to creation of greater amounts of heat inside the unit due to a defect such as an internal fault.

The causes of thermal problems have to be traced through other inspections and tests, some of which will be prescribed based on the initial, basic test results (see the aforementioned books). These begin with an inspection and test of the entire cooling system – pumps, fans, radiators. If these are fine, then tests for sludging, such as acidity and interfacial tension tests of the transformer oil, are recommended.




Figure 11.6 A thermal load “test” of a power transformer. Top, electrical load and temperature rise above ambient are monitored over a 24-hour or longer period. Bottom, de-convolution of the two time-series produces a calculated impulse response function (temperature rise and time for an arbitrarily short pulse of loading) for the unit. This can then be used to compute the step response shown, which is the rate of rise and final temperature that would result from having the unit serve a load equal to its rating on a continuous basis. This example reveals an as-tested signature that both have a higher rise time and asymptote (high eventual temperature), indicating some type of internal problem.

11.5 SWITCHGEAR AND CIRCUIT BREAKERS

Circuit Breakers

Circuit breakers are electromechanical devices. They are tested in both mechanical and electrical performance and for signs of deterioration.

Routine inspection

Physical inspection aims to identify signs of past, present, or forthcoming problems, and is carried out routinely (every quarter or semi-annually). The unit is visibly checked to identify any noticeable corrosion, deterioration or damage, or outward signs of internal malfunction (leaking oil around seals). Physical inspection of the device usually includes visual or infrared examination of the exterior of the unit while in service. Inspection of internal parts of a circuit breaker (condition of contacts, etc.) is a big part of that maintenance, but not normally considered part of “inspection” because it is done only after de-energization and disassembly.

Records should be kept on routine inspections, so that gradually deteriorating paint, corrosion problems on a particular unit, etc., and other similar problems can be tracked over time.

Temperature rise test

This can be viewed as a form of inspection and must be done while the unit is in service. The temperature of the breaker is measured accurately and correlated with load. Significant rises in temperature can identify internal electrical problems. Use of an infrared scanner can aid in this test, particularly for determining where in the unit the additional heat is coming from.

Electrical tests

All electrical tests done on breakers require de-energization, but can be carried out in the field if recommended safety precautions are taken and if instructions for the test are followed properly. If they reveal problems, the unit will have to be disassembled and visually checked, or its components tested, to determine the problem and its fix.

Standard electrical tests are performed on circuit breakers. An insulation resistance test applied twice, once with the contacts open and once with them closed. An AC high potential test (or a DC version if AC equipment is not available) can be performed, but generally to only 75% of factory test voltage, so that it does not unduly stress the unit. Manufacturers’ recommended procedures for such tests differ substantially depending on the type of unit. Power factor tests may be done – they are usually done on air circuit breakers.

Contact resistance tests are the only special type of electrical test for circuit breakers, in the sense that these tests are not done on other types of electrical devices like transformers. This is simply a measurement of the resistance from

terminal to terminal across each phase, made with the breaker closed, and is a way of inferring the condition of the breaker contacts without disassembly of the unit. A standard ohmmeter is not used. Instead, a current source (typically about 100 A) is used in order to put some magnitude of current close to normal operating levels through the breaker to simulate actual load. Expected values for medium voltage breakers are in the range 100 – 250 $\mu\Omega$, but each manufacturer provides specific target data for each type of unit that it manufactures. Signs of possible deterioration from this test include either a higher resistance than expected or significant deviations in the measured resistance from one phase to another.

Time and travel tests

Mechanical tests of the circuit breaker are called time-travel analysis. They are carried out with the unit de-energized, but are done on site (in the field). These are tests of the breaker mechanism to determine if it is opening quickly enough, or binding in some part of its analysis or otherwise not performing as expected. Modern test procedures use electronic time-travel analyzers rather than the older style mechanical devices, which were used for testing into the late 1980s. These new test units provide not only an accurate reading of opening and closing times for the unit, but also detailed analysis of mechanical performance. They provide plots of opening and closed speed during each phase of the cycle, and numerical tables that identify quantitatively the breaker's performance in each of several categories. Time-travel tests of a breaker include:

1. The opening and closing time
2. Contact speed and gap distance as a function of time during the cycle
3. Contact bounce
4. Trip operation
5. Time to re-charge or re-energize the trip mechanism

Tests of ancillary and associated equipment

Tests of breakers should include, either as part of the breaker inspection and testing, or as a separate program, inspection, testing, and appropriate diagnosis of the rest of the protection system that includes the breaker. Instrument transformers, protective relays, panels, and wiring should all be routinely checked for signs of deterioration and to see that they function properly.

- *Instrument transformers* in service are tested periodically to determine that their breakdown voltage remains sufficiently high, and for ratio and phase angle performance. Other than that, they are “left alone.” However, if de-energized for any lengthy period (as when the surrounding equipment goes through a long period of maintenance), and particularly if left outdoors, they should be removed and dried.

- *Protective relays* controlling the units need to be tested to determine that their entire system is functioning correctly. Testing of relays involves methods often tailored specifically to each type of relay. However, a highly recommended general procedure is to record the condition and settings of the relay system in “as found” condition at the start of each inspection/test. Many of the problems with relaying are not due to deterioration or degradation of units, but mistakes in setup, calibration, or installation. Some of these can be systematic (due to training errors or poor habits of personnel) which will be found only if accurate tracking of problems of this nature is maintained.
- *Wiring and panel equipment*, including any recorders and indicators, should be checked and assured to be in operating condition. Indicators and records should be checked periodically (e.g. every 5 years) for accuracy, and re-calibrated if need be.
- *A complete system test* (of the relay-breaker combination) should be performed at the end of the inspection/testing. The fact that all of the equipment is individually in good condition does not absolutely assure that it will function as a whole. A final test of the complete system is the only sure way to verify this.

Switchgear

Switchgear other than breakers is generally tested and inspected on the same periodic basis as breakers. Switches and other mechanical devices should be routinely (semi-annually or annually) operated. Unlike the situation with breakers, in many cases visual inspection of contacts and mechanisms is possible with much of this equipment.

- *Routine electrical tests* include only the insulation resistance tests. They are relatively easy and quick to apply in the field and put no particularly high stress on the equipment. If they indicate problems other tests might be called for. These include hi-potential tests and power factor tests.
- *Radio interference* is a very good way to identify problems due to dirty or cracked bushings, post insulators, and corrosion-caused weaknesses in connections. As mentioned earlier, a small AM radio can detect the interference given off by some types of electrical problems.

A detailed discussion of breaker and switchgear inspection and testing procedures, including comprehensive lists of recommended procedures on a step-by-step basis and a method for a detailed RIV (radio interference voltage) test of substations is included in *Electrical Power Equipment Maintenance and Testing*, Chapters 7 and 8 (Gill, 1998).

11.6 UNDERGROUND EQUIPMENT AND CABLES

Underground cables

Underground cable is available in a wide variety of types, including three-phase, single phase, either carrying or not carrying a concentric neutral, and with various types and degrees of “armor” or protective sheathes. Cables are made from a diverse range of materials, including copper, aluminum, and alloys for the conductor. Insulation and construction may be composed of any large set of materials with somewhat differing properties, and different construction methods, including paper, cotton derivatives (varnished cambric), natural or synthetic rubbers, and other materials. Cables also vary in the type or material and construction of their outer sheath.

With all of these types of materials and construction, there is no one type of inspection or testing program that is applicable to all types of cable. Each manufacturer provides information or recommendations. Table 11.6 lists general categories of inspection and testing and the typical periods of their application.

Acceptance tests

Acceptance tests are usually performed on cable when received from the manufacturer. This includes insulation resistance measurement and DC high potential tests. Acceptance tests often include a series of destructive tests on samples of the cable, carried out in a laboratory, to verify its withstand capabilities.

Routine inspection

It is impossible to visually inspect the vast majority of installed underground cable simply because it is underground and inaccessible, inside ducts or directly buried. Despite this, routine annual inspection should be carried out on those portions (the ends of cable sections) that can be inspected, and on all associated terminations and ancillary equipment. This last group of equipment is particularly

Table 11.6 Cable Inspection and Testing

Activity	Purpose	Frequency
Visual inspection	Check for visible deterioration, leaks, corrosion	6 mo.
Indentor test	Track material deterioration	1 – 5 yr.
Insulation resistance	Non-destructive test of insulation quality	1 – 5 yr.
PD test	Detects flaws/incipient failures. High stress test.	Infrequent
Hi-potential (DC)	Detects flaws/incipient failures. Very high stress	Infrequent
Fault location	Identify location of failure	As Needed

important, because a good deal of cable problems occurs at terminations and are due to failures or deterioration of equipment other than the cable itself.

Inspection involves passive evaluation of condition based on visual (including infrared) examination of cable section ends, termination, conduit, and other equipment. Corrosion and deterioration of ground connections can be a particularly important factor for cables and deserves special attention. Cables should be checked for oil leaks and swelling of insulation. Oil circulation systems for oil-filled cables should be checked fully, including full evaluation of the condition and capability of the pumps.

Inspection may include a type of thermal analysis based on comparison of “temperature as measured” versus “temperature expected” using data obtained from monitoring of load and temperature at points on the cable. This is similar in concept to that shown earlier for transformers (Figure 11.6). However, it cannot be done quite as accurately as with transformers because the thermal factors for cables are not known quite as accurately as they are for transformers. Factors such as the thermal transfer capability deep inside the duct, etc., can only be estimated. Therefore these types of tests are useful only for tracking over time so that deterioration – deviations from a previously verified “good” performance level can be identified.

Condition evaluation of installed cable

Underground cable failures are one of the most vexing trouble areas for modern electric utilities. Cables cannot be inspected visually; deterioration rates can be exacerbated by many factors not visible because the equipment is buried. Repair times are lengthy, and often very costly. Many types of cable have very long service lifetimes (e.g., paper insulated, lead covered), but some types of extruded cable, particularly bad batches of it, gained reputations for lifetimes far shorter than needed (i.e., less than 20 years).

Therefore, indirect means of testing cable to determine the condition or state of deterioration of its insulation and hidden or difficult-to-access splices, is a significant part of a utility system reliability improvement program. Electrical tests require de-energizing the cable, but can be done in place. The most frequently performed tests that can be carried out include:

1. *Insulation resistance tests* as described in Section 9.3 are electrical tests that can be done routinely on cable. They are relatively quick and inexpensive and cause no undue stress on the cable, so frequent testing will not degrade cable quality.

Since the measured resistance of a good cable will depend on its length, which is different for nearly every section, these tests are useful only if they are done on a periodic basis (including ideally, when initially installed) with good records kept, so that results can be plotted over time. Unfortunately, small flaws or early signs of

deterioration may not show up on a resistance test. Therefore a partial discharge or DC hi-potential test may be required.

2. *Partial discharge (PD) tests*, discussed in Chapter 6, are available from a variety of commercial suppliers of PD test equipment and services. Interpretation of the results of PD test is still something of a “black art,” but a number of very clever means of analyzing timing and spatial pattern of results have been developed to infer a tremendous amount about cable condition. A particularly good discussion of partial discharge methods including diagrams of different types of partial discharge display patterns and their interpretation, is included in *Electrical Insulation in Power Systems* (Malik et al. 1998).
3. *Indentor tests*, while not strictly an electrical test (they are a type of mechanical procedure) can be done while the cable is de-energized and being otherwise tested for electrical properties. These tests involve pushing a small indentor against the insulation of the cable and measuring the material’s resistance to its penetration. Material will harden during service, so that this test provides a quantitative measure of material deterioration.
4. *High potential tests* can be done to stress test the cable and reveal any developing flaws. As mentioned in Chapter 6, hi-pot tests of UG cable usually use a DC rather than AC voltage approach, to avoid capacitance problems that might occur with AC hi-pot tests. As also discussed in Chapter 6, these tests are well-known for the high level of stress they can put on cable, which can worsen existing weaknesses to the point that the cable fails much sooner than it otherwise would. Therefore, these tests are usually done on a non-routine basis.
5. *Destructive or forensic laboratory testing* is often carried out on samples of cable taken out of service. While this type of examination is frequently done only on the “good portion” of failed cable sections removed after a fault, it is often performed as part of a random sampling process for UG cable. Portions of cable are destructively tested (to the point of insulation failure) in controlled laboratory conditions. Other portions are sliced and analyzed under a microscope for evidence of treeing or other deterioration in the insulation.

Fault Location

When an underground cable fails due to a fault or open circuit, various test procedures can be used to try to locate the fault. The traditional method is a manual sectionalization to isolate the fault, section by section, which is tedious and time consuming, and often abuses equipment.⁶ Other methods include the radar method and arc-reflection method (two forms of time reflectometry) and a resonance method in which electrical signals or pulses are tracked down the cable, pulse return timing or resonant frequency indicating distance to the fault. Several methods inject a signal or pulse that excites some type of electromagnetic signal from the fault location, which can then be found using a portable location (this only works if it is possible to walk above the cable route).

Fault location is not part of “condition assessment” of cable operating in normal fashion, and will not be discussed further here. A very comprehensive discussion is included in Malik et al.(1998).

Repair and replacement

Flaws and problems found by routine inspection can generally be repaired: such things as corroded casings, poor grounds, leaking seals, non operative cathodic protection, and oil pumps not up to specification can and should be repaired immediately. Similarly, poor terminations or equipment at terminal sites that is judged to be bad can be repaired or replaced.

If and when a cable fails, the fault can be repaired with a splice, but it is often advisable to replace the entire section if post-fault analysis indicates the cable had a severe amount of deterioration. Splices have a failure rate higher than equivalent lengths of un-spliced cable. They represent a potential weak-spot in the already weak system (the remaining cable is aged and potentially at risk of failure itself).

But when PD or hi-pot tests indicate that a cable section insulation level is poor, it *must* be replaced. There is no repair for badly deteriorated cable.

11.7 OVERHEAD LINES AND ASSOCIATED EQUIPMENT

Damage, Not Deterioration, Is Often the Major Concern

A majority of power distribution in most utility systems consists of overhead primary (MV) and overhead service, or secondary (LV) lines and equipment. The basic materials for overhead lines include poles, conductor, insulators, and their associated hardware, all of which are incredibly robust equipment. Their expected lifetimes in normal service are 50 years or greater. While deterioration is a concern, the very slow rate of change under most conditions and the very great

⁶ An expedient for finding faults in URD systems at many utilities involves “opening and closing” URD primary sections by pulling elbow connectors with a hot-stick. This is not recommended and most elbow connectors are not designed or rated for such duty. The stress put on the elbows is unknown and cumulatively could lead to early failure.

times to failure for this equipment generally mean that routine monitoring to track the rate of deterioration is not an effective or justifiable O&M procedure.

This is not to say that condition assessment is not an issue in overhead distribution. Far from that, assessment and tracking of condition should be a high priority in many distribution systems throughout the United States, particularly in those where large portions of the system exceed 50 years of age, an age where deterioration of wood, metal, and bell insulators may have reached a point of concern. Condition assessment of overhead equipment in these areas is a critical part of good reliability management.

But one definite fact that colors all approaches to inspection and maintenance of overhead lines is that by its very nature, overhead equipment is quite exposed both to the elements and to external sources of damage such as errant automobile traffic, nearby vegetation and fires, etc. *A majority of overhead outages occur due to damage*, not deterioration, whether that damage is from severe weather, contact with trees, automobiles and construction accidents, forest and building fires, and vandalism. For these reasons, identification of damage, not deterioration, is a large part of inspection and testing for overhead lines.

Inspection of Overhead Lines

Generally, elements of the primary and service voltage systems are not tested with electrical tests like insulation resistance, power factor, hi-potential etc., as are substation equipment and UG cable systems. This is not to say that they could not be, and that such tests have not been done on a special or experimental basis. Such tests aim to identify deterioration of insulation, and as mentioned above, deterioration is typically not a big concern on overhead lines; damage from external causes is.

But while such tests could perhaps identify some types of damage, visual inspection of overhead lines can reveal a great deal more. Visual inspection, done to a high standard and organized with good record keeping will identify damage and in fact can detect signs of deterioration in many elements of the system. Typically, the best procedures utilize written classification criteria (“Less than ten square inches of visible rust per side” “Scaling rust flaking off”). Such routine inspection might be augmented using infrared imaging to identify hot equipment and fittings, and/or the use of developing acoustic/radio frequency detection devices that find traces of leakage currents, etc.

Therefore, inspection is the chief means of determining overhead system condition. The only electrical test routinely done on overhead distribution involves checking the grounding of poles and neutrals in grounded and multi-grounded Y-connected distribution systems.

Poles and Pole Assemblies

Poles are made of natural wood (treated, straight tree trunks), laminates of wood, composite wood (wood epoxy mixes), concrete, fiberglass, or steel. Concrete, fiberglass, and steel poles have very long lifetimes, if their materials are properly manufactured (cured, etc.).

The vast majority of power distribution poles in use are wood. Wood poles have an expected lifetime of from 35 to 85 years, depending on the type of wood, chemical preservative and treatment method, installation, soil and climate conditions, and periodic inspection and treatment given them. Cross-arms have a similar lifetime. Deterioration takes two forms: 1) a general weakening of mechanical strength of the entire pole, due to aging, cracking, etc., and 2) rotting of the pole near ground level due to trapped moisture. They are also subject to damage from automobile accidents, trees falling, fire and all manner of unusual situations (a property owner nearby may use the pole as one corner of a garage or barn being built, etc.). Deterioration, like damage, tends to be very pole-specific. One pole can be near failure due to deterioration; those nearby might have a considerable lifetime left.

Visual inspection is notoriously unreliable in determining the state of deterioration of a wood, although it is very good at identifying damage to poles and their assemblies. Testing for mechanical strength and rotting takes several forms. The simplest, but least exact, is the hammer test. The pole is struck with a large hammer near ground level. The resulting “thump” has a very different sound depending on whether the pole is sound or rotted. This test is very quick (many poles can be checked per person-hour and inexpensive (the equipment, a standard small sledge, costs very little). However it is not particularly quantitative. It tends to identify poles with severe stages of deterioration, but does not measure interim levels of deterioration. In addition, dependability of results has a good deal to do with the skill of the person doing the test.

A type of sonogram approach uses a device that sonically scans a cross-section of the pole (usually near ground level) to determine its density and detect any voids within. The sonogram device is reasonably portable (about the size of a small rug vacuum cleaner) and is taken from pole to pole. This test is more expensive than the hammer test, and requires more operator skill, and a good deal more time for each pole than the hammer test, but it is much more accurate in determining condition on a pole. Coring to obtain a sample of interior wood for analysis is also often used.

A type of indentation-mechanical/electric resistance test can also be used. A portable device about the size of a small suitcase is attached to the pole near ground level, using a web belt to tighten it to the pole. The device then pushes two thin electrical probes into the pole, measuring the mechanical resistance to their insertion, one measure of the condition of the pole. The electrical resistance through the pole is then measured with the probes, using an ammeter. The combined mechanical-electrical values can be correlated with the pole's mechanical strength (different tables are needed for each type of wood, e.g., fir).

Deteriorated poles can be repaired by injection of strengthening and/or drying materials into them near the base or by reinforcement with sheathes or stub-poles. The most effective long-term approach however, is replacement.

Inspection of Cross-arms, Insulators, Conductors and Hardware

Overhead line equipment includes cross-arms, mostly wooden, which can deteriorate much like poles can, insulators and attachment hardware, and conductor and associated hardware like splices and vibration dampers.

- *Careful visual inspection* of a distribution line, often called “driving the line,” can identify damage of cross-arms, hardware, conductor, and deterioration if it has reached advanced stages (near failure). One factor in favor of such inspection is that it is inexpensive, requiring no special equipment. However, it fails to find some incipient problems (in fact, a majority may go undetected), and it is not easy to carry out on sub-transmission or distribution lines that do not parallel roads or highways.
- *Radio interference tests.* As mentioned at the beginning of this chapter, a simple and inexpensive enhancement to visual inspection of electrical equipment exposed to the elements is radio interference monitoring. Routine policy at several utilities is to leave a radio on the AM band but not tuned to any station while driving line routes for visual inspection of equipment. Cracked insulators, corroded brackets, and broken conductor strands create low levels of radio interference, so a sudden increase in radio “static” and noise is a sign of a nearby problem. Generally, inexpensive analog radios work best for this purpose: expensive stereo systems and digital radios have filters and signal enhancement circuits that eliminate all or most static and interference.
- *Infrared scanning.* Two improvements in overhead line inspection are infrared scanning and inspection by helicopter, often combined into a much more effective and rapid means of inspection lines. Infrared scanners are essentially infrared video camera-recorders that display on a CRT an image of equipment as seen in the thermal infrared band. Intensity of items in the image is proportional to temperature, so brackets, splices, or conductor sections that are overheating for whatever reason show up clearly.
- *Inspection by helicopter* provides several advantages. First, a number of important deterioration modes are visible from the air better than from the ground. Cross-arms crack and rot from above where standing water tends to collect after rains. They may look fine from below but can have signs of serious deterioration if viewed from the air. Loose attachments of conductor to insulators are easier to see when the insulator does not block view of the tie point, as it does if viewed from the ground.

Second, in all but very heavily congested urban areas, a helicopter can “fly” line routes while staying close enough to visually inspect in detail. Often, it is possible to approach closely enough that service-level equipment and lines can also be checked.

Third, the helicopter can cover much more line per hour than inspection by foot or automobile. Helicopter surveys often use a video recorder in both visual and infrared bands to record the inspection for later review in detail as needed.

Replacement and Repair of Basic Line Equipment

Damage to or deficiencies in basic overhead line equipment found by inspection can be repaired on an as-identified basis by replacement of the failed or failing components. Little of this equipment is “repairable” in the sense that it can be taken down, reworked or refurbished, and put back up. Instead, replacement with new is the only alternative. The exception is conductor. Birdcaged conductor, broken strands, and similar damage can be repaired by replacement (splicing) of only the damaged part of the conductor.

One recommended exception is annealing. Experience has shown that failures or severe deterioration due to annealing at one part of an overhead line will generally indicate severe deterioration in other parts nearby. Ideally, if a conductor falls down or other problems from annealing are suspected, the entire segment should be replaced or carefully inspected at close range (from a bucket truck) for signs of annealing.

Switches and Cutouts

A portion of manually operating switches in many electric distribution systems are suspect simply because they have not been operated in years. Over time, switches will become inoperable for a number of reasons. Even a minor amount of corrosion on the contacts of closed switches, or in the operating mechanism of open or closed switches, will freeze the switch so it cannot be operated. Water working its way into mechanisms can freeze, warping rotating surfaces and jamming devices. Closing a switch into a current higher than its rating can “weld” its contacts closed, so it will not open next time an attempt is made to operate it. Overloading can reduce the temper in the switch springs, causing it to malfunction. In addition, connections sometimes just weaken or loosen.

Cutouts and fuse connections also suffer from the same types of problems as line switches. However, they are almost never test-operated. There is some likelihood that each will go bad, but given their number, the difficulty in reaching them, and that fact that any quick test would assuredly cause a customer interruption,⁷ few utilities test anything but a handful of cutouts, and only at

⁷ To test without an outage a jumper has to be placed in parallel with the cutout prior to the test, then removed afterward. While certainly within the capabilities of any well-trained line crews, this increases the time, and hence the cost, of such tests.

special sites, each year.

Infrared scanning of both switches and cutouts can reveal a large portion of mechanical and electrical problems. Loose and corroded connections or a partially broken conductor near them shows up as a hot spot. The corrosion that causes frozen contacts often creates higher impedance in the switch contacts, creating a hot spot at the top of the blades.

Fuses

Fuses can be damaged by overheating, nearby lightning strikes, or long periods of high loads. However such problems are very rare, and fuses have very long lifetimes in service, assuming no fault occurs that requires them to blow.

One of the problems with fusing experienced to varying degrees by some utilities is a mismatch between designed and actual protection size of fuses on overhead lines. Mismatches occur because of mistakes made by line crews, expediency (a trouble-man has no 30 amp fuses so installs a 40 amp), and “coppered out” fuses.

The last problem arises due to either laziness on the part of line crews (rare) or to good intentions, overwork, and poor coordination of line repair records (more common than is appreciated). After long outages, particularly in very cold or very hot weather, “cold load pickup” can be a nuisance when restoring service. Starting current and loss of diversity due to the lengthy outage may mean that the “cold loads” are up to three times higher than normal. As a result, on primary feeders with tight fuse coordination, the initial load seen on a branch or lateral that is closed into service after a long outage may exceed the specified fuse rating. The fuse links keep blowing when the line crews try to close the cutouts to put the customers back into service.

The proper way to “fix” this problem is to jumper around the fuse for a period of thirty minutes to one hour, until load diversity returns. However, one expedient solution frequently used by line crews working against a tight deadline (common during storms) is to replace the fuse link with something that will not blow. This might simply be a larger fuse link than specified, but even a larger fuse size (i.e., a 40 amp instead of a 30 amp) may not provide enough capacity for the initial cold load surge. However, a length of #6 solid copper, or a large strand of aluminum from AA conductor, will fit in a fuse barrel, and “cures the problem.” Usually, the intent of the line crews is to return after one or two hours and change the fuse link back to the proper fuse. But due to difficult schedules and forgetfulness, a noticeable portion of such coppered-out fuses probably remain in service. Accumulated over many years, it can be a noticeable problem – engineering records and visual inspection show that fusing is installed on the system, but protection doesn’t work properly.

It is impossible without pulling the fuse element to verify the size of fuse actually installed in a fuse barrel. The “solution” to this inspection problem is a detailed assessment of trouble reports from the field. Statistical analysis of outages in which protection coordination is suspected of not functioning properly

can reveal the degree of the problem – the “advanced TLM” discussed in Chapter 15 is somewhat successful at this. However, nothing short of site by site inspection, will identify all problems.

Regulators, Capacitor Banks, Reclosers, and Sectionalizers

These types of power system equipment are considerably more complicated than basic line equipment, and should be inspected, tested, and maintained in a manner similar to the transformers, tap changers, breakers and station capacitor banks it resembles. Routine inspection of this equipment, including visual and infrared inspection from the ground or helicopter, is typically included with inspection of the entire line.

Unlike the case with the basic overhead line equipment, electrical tests are typically applied to this equipment. Line regulators are tested in the same manner as small medium-voltage transformers/tap changers. Reclosers and sectionalizers fall into a category similar to breakers and relay sets. All of this equipment is routinely tested on something like a 2-5 year cycle for proper operation, and usually with instrumentation to assure proper turns ratios or operating speed and synchronization, as the case may be. Insulation resistance is also measured as a routine course. Reclosers and sectionalizers will have specific inspection, test, and maintenance procedures recommended by their manufacturers.

Shunt and series capacitor banks have unique characteristics (high capacitance) which require special test procedures, generally involving DC potential testing and induced potential tests. Capacitor switches in some systems are particularly failure prone and are tested annually, generally just prior to peak period.

11.8 SERVICE TRANSFORMERS AND SERVICE CIRCUITS

Because they are so numerous and there is no contingency backup for each one, service transformers can be an issue in the reliability of any system. In most cases, restoration of service can be achieved only by replacement since no re-switching option exists. Service transformers and service-level circuits are predominantly “fit it and forget it” devices. At most utilities, they are not inspected, except as part of general inspection of overhead lines. This inspection does catch a good number of problematic service transformers and weak connections of service-level wiring. Many utilities also depend on meter readers and other personnel in the field to “keep their eyes open” for problems. But despite this, deterioration and damage can go undetected. Fortunately, service transformers and low-voltage circuits, whether overhead or underground, are very robust, so lifetimes even in this light-inspection environment are quite good (see Chapter 7).

One very effective “inspection” method for service transformers is a transformer load management (TLM) program. “Program” here means both a computer program and all the procedures and internal organization required collecting the data and using the information produced to manage the service transformer asset base. TLM involves gathering customer-billing data on a

monthly or quarterly basis. This is input into the TLM program, and correlated with transformers (a TLM requires data telling it which customers be served through which transformers). The TLM program then estimates the peak load and peak duration, and hence the total stress on the unit, for each service transformer. Its estimate is based on the number and type of customers and the total kWh sold through each transformer during the analysis period.

The TLM does this calculation for all service transformers in the system. It then produces a sorted list of units for which the calculated stress is higher than recommended – a so called “jeopardy list.” This is sent to the field so that each unit can be checked by line crews (leads can be amped to determine true loading level, the unit can be visually checked for signs of leakage caused by overheating, etc.). Overloaded units are replaced with larger capacity transformers.

A good TLM has among the highest payback ratios of any activity related to maintenance and asset management of distribution equipment. Often, it pays for itself in a matter of months, by permitting overloaded units to be changed out before high loading levels lead to premature failure. Advanced approaches like this can apply much the same concept to evaluation of service quality and lifetime problems with the service-level circuits.

11.9 EVALUATING AND PRIORITIZING EQUIPMENT CONDITION

Theory of Condition and Correlation with Test Results

The theory of equipment aging and deterioration is well understood by the power industry. Three major factors “age” electrical equipment: voltage-frequency stress, thermal stress, and water damage.

1. *Voltage aging.* Given enough time, equipment that is simply energized (e.g., a cable kept in service but carrying no load) will fail due to the accumulated stress of voltage – somewhere, somehow, its insulation will break down. The level of stress, and hence the rate of deterioration in insulation strength, is a function of both voltage level and frequency. Insulation lifetime is exponentially related to voltage level by something between a fifth and ninth order relationship depending on type of material and design. This means doubling voltage will cut lifetime by anywhere between a factor of 2^5 (32) and 2^{25} (33 ½ million) an incredibly wide range, indicating that insulation type and design are key factors in equipment quality.

The voltage-lifetime relationships for typical electrical insulation is in the range of fifth to tenth order. For example, XLPE insulation has a relationship that is ninth order. Doubling voltage means lifetime is cut by a factor of 2^9 (512), meaning the unit loses the equivalent of about 3 weeks of lifetime for every hour it is energized to twice its nominal voltage. Its lifetime will be cut by 35% if it is operated at 1.05 PU rather

than 1.0 voltage.⁸

2. *Frequency* also plays an important role in aging. Time to failure is related to frequency, roughly on a linear basis (doubling frequency will roughly halve lifetime, halving it will roughly double lifetime. Frequency is regulated so tightly – within .02% even in “sloppy” systems – that the equipment installed in any power system will see very nearly the same operating frequency all its lifetime. Basically one can assume that the equipment sees nominal system frequency during all hours of operation.

However, voltage spikes due to lightning strikes, switching surges or other unusual transients have a very high frequency component due to the rapid rise and fall times of their transient pulse. By their very nature these events are higher than normal voltage, and that higher voltage causes added stress in and of itself. But even though brief, these disturbances can cause considerable stress on insulation due to their high frequency content. For example: assuming that lifetime and voltage are inversely related by a ninth order relationship (XLP cable), one can compute that a spike to 1.5 PU voltage, lasting one half millisecond (.0005 sec), would cause roughly 50% more stress than normal due to its higher voltage, multiplied by a factor of 16.6⁹ (almost 150 billion) due to the higher frequency component of the spike. This means that this event lasting half of a millisecond ages the cable by the equivalent of about two years. In fact the impact may be somewhat different than this computation predicts, due to a host of secondary effects that can only be computed with an EMTP analysis.⁹ But the stress will still be great. Brief voltage events put undue stress on insulation.

Secondly, many diagnostic procedures such as hi-potential

⁸ This does not include the effects of any current increase due to the raised voltage. Generally, raising the voltage on a cable or unit of power system equipment will also raise the current passing through it (since it is a constant or very close to constant impedance). That increases the losses and hence the rate of thermal aging. The effects discussed here are due only to voltage – increasing the applied voltage on an *unloaded* section of XLPE cable by 5%, for example, is most likely to cut time until voltage breakdown occurs due to voltage stress alone, by a factor of about 35%. If that also leads to an increase in load, the additional thermal loading could contribute further acceleration of the cable aging.

⁹ This very simple computation neglects a host of electromagnetic transient (EMT) related factors that alter the induced stress, including changes in: 1) the impedance of the cable or insulators as a function of frequency, 2) the non-uniform distribution of stress that occurs as a function of high frequency, and 3) the fact that the exponential order of the relationship is itself a function of frequency.

and partial discharge tests, and some cable fault location methods, use pulses (again, a high frequency content) or variable frequencies higher than normal. They can produce noticeably high stress on insulation.

3. *High temperature* has two effects on electrical insulation. First, some types of insulation (XLPE) soften at high temperatures to the point that plastic deformation occurs in the cable. The concentric conductor falls or “flows” through the insulation; mechanical stress on the cable may result in deformation. Minimum insulation distances can decrease as a result and failure can occur.

Secondly, and more important for the vast majority of equipment which does not suffer plastic deformation, higher temperatures accelerate *all* types of deterioration. The different modes of deterioration and the different deterioration rates of different materials are all affected in varying degrees, but qualitatively all forms of power system equipment lifetime loss are accelerated by higher temperature. A rough rule of thumb is that lifetime of cables is halved for every 9°C increase in temperature. Transformer insulation behaves in a roughly similar manner.

4. *Combined stress of all three factors* is what all-electrical equipment in service sees, the simultaneous effects of voltage, frequency and temperature, over time. Basically, the relative stress levels from these three factors are multiplicative, meaning that each exacerbates any problems due to the others:

$$\begin{aligned} \text{Total amount of stress} = & C \times (\text{voltage stress factor}) \\ & \text{times} \\ & (\text{frequency stress factor}) \\ & \text{times} \\ & (\text{thermal stress factor}) \end{aligned}$$

Thus, if voltage and frequency at the moment are causing high levels of stress, and the insulation is at high temperature – say 9°C higher than normal (which can be expected to roughly halve lifetime), then the accelerated rates of lifetime loss due to the voltage and frequency are doubled.

Condition Assessment from Inspection and Test Data**Go/No-Go versus Condition Measurement**

Table 11.7 reprises data from Chapter 6's Table 6.4 on the types of test and diagnostic methods available for power system equipment. Some of the tests are basically "go/no-go" tests of equipment, such as the high potential test and a transformer turns ratio test. A unit will either pass, or fail this test. Units that pass are put in service. Units that fail should be withdrawn from service.

Other tests, such as the total combustible gas test (TCGA) provide a numerical measure that is useful in assessing condition, and by tracking over time, provides clues as to the gradual deterioration of the unit and its continued capability to provide dependable service. These diagnostic tests lead to classification and fuel decision making on future inspection, testing, and/or replacement of units. For example, TCGA is used to determine classification of transformers by risk category and to recommend inspection and testing intervals (IEEE Standard C57.104-1991).

Some tests can be viewed as both go/no-go and condition assessment. Insulation resistance can be used both as a certification test and diagnostic. Below a certain value of insulation resistance, the unit is not put into service, and the value can be tracked over time so that changes in the unit can be detected. Like it, most of the test procedures discussed in Chapter 6 provide quantitative results – numerical values that can vary over some range. For example, insulation resistance tests measure resistance in megaohms, and can return a value of from 1100 to 1500 ohms when tested on the windings of high voltage power transformer. A power factor test for cross-linked cable can give a value of from just above 0.0 to over 2.0%. Transformer turns ratio tests for a 138/12.47 kV transformer should give 11.06, but can give values below or above that ratio depending on internal faults in the high or low side windings.

Generally, go/no go tests are dependable indicators of the viability of power equipment for service during the next year. If a unit "passes" the appropriate tests, it will usually provide good service in the next year. Leave a unit that has failed such a test in service and it will very likely fail, often catastrophically but nearly always to a point of at least destroying any residual future value in the next year.

Condition assessment's predictive capabilities

Given the wide ranges of numerical values that these tests return, it might be expected that "condition assessment" analysis would be able to translate these test results into a dependable evaluation of the condition of each transformer, cable, or other power system equipment. This would be something akin to an accurate ranking of "equivalent life" – this unit has the equivalent of 35 years (35/40th of normal service lifetime) left, that unit only 11 years, etc. But despite over a century of commercial power system application, electrical engineering technology is still not at this level. This is not to say that many attempts haven't been made to develop such methods of translating test results into accurate status

estimation. But at best they result in “fuzzy” or inexact evaluations. The use of inspection and test data to determine the overall condition of equipment is part science but still part “black” art.

With very few exceptions, the results from any one type of test are relatively straightforward to analyze and apply. The real challenge is combining results from multiple tests into a weighted or “single valued” evaluation that is both easy to use and accurate. For example, exactly what do the results mean when a cable passes an insulation resistance test with perfect scores (equivalent to when new), but PD tests indicate it has a number of weaknesses? Certainly, a qualitative interpretation is that the cable is currently operable but that it has imperfections that will develop into flaws over time. But how does the concern for this particular problem rank with respect to concerns about other cables that had other patterns of test scores? If budget is limited (and when isn’t it?) which cable will get the most attention?

Consider the data shown in Table 11.8 and Figure 11.7, which show total combustible gas accumulation rates, and insulation resistance test results, for a dozen 25 MVA 138-12.47 kV transformers tested by a large utility in the southern United States, in the spring of 1984. Individually, each of these two tests provides a numerical value that is a measure of the “condition” of the unit. Higher gassing rates indicate faster rates of deterioration or internal damage. Lower resistance indicates deteriorated insulation condition.

Table 11.7 Power System Electrical Test Procedures

Test	Type	Where?	In Service?	No-Go	Condition	Stress Lvl
Thermal load tests	AC	Field	Yes	No	Yes	-
Resistance test	DC	Field	No	Yes	Maybe	-
Dielectric absorp.	DC	Field	No	No	Yes	-
Power factor tests	AC	Field	No	No	Yes	-
Polarization recov.	DC	Field	No	No	Yes	x
Hi-potential test	Both	Field	No	Yes	No	X
Induced potential	VF	Field	No	No	Yes	X
Partial discharge	AC	Field	No	No	Yes	x
Transf. turns ratio	AC	Field	Yes	Maybe	Maybe	-
Oil acidity tests	-	Lab/Field	Yes	No	Yes	-
Interfacial tests	-	Lab	Yes	No	Yes	-
TCGA tests	-	Field	Yes	No	Yes	-
DGA tests	-	Lab	Yes	No	Yes	-
Disassembly/inspection	-	Field	No	Yes	Yes	-

Table 11.8 Test Data for Twelve Power Transformers

Unit	Resistance - MΩ	TCGA – ppm/day
1	1520	5
2	1455	21
3	1430	12
4	1405	4
5	1360	26
6	1510	15
7	1480	14
8	1470	25
9	1475	8
10	1480	13
11	1490	3
12	1520	<1




Figure 11.7 Position of transformers 1 – 12 based on their results in two tests.

Table 11.9 Service Lifetime Data for Twelve Power Transformers

Unit	Age at Time of Test	Status 1999
1	5	S.I.S.
2	17	Failed 1990
3	22	Failed 1986
4	31	S.I.S.
5	34	Failed 1994
6	4	Failed 1993
7	11	Rebuilt '86, Failed '97
8	6	Rebuilt 1985
9	13	S.I.S.
10	9	Rebuilt 1996
11	9	S.I.S.
12	New	S.I.S.




Figure 11.8 Remaining lifetime in service of the twelve units, plotted at the location of test results obtained. The rightmost of the two units with the highest data point vertically has an estimated only one year of service remaining. Clearly, units near the lower right hand corner were “in good condition,” but other lifetimes were not completely proportional to test results.

Table 11.9 and Figure 11.8 provide additional data on these twelve transformers, including age at the time of the tests (years in service up to 1984) and status of the unit fifteen years later (in 1999), a “result” not known at the time of the test. Note that Table 11.9’s data verifies that unit 4 is indeed an older unit as was surmised from the data interpretation discussed above. Note also that that unit provided good service for (at least) another 15 years. Similarly bearing out the test results, Table 11.9 shows that unit 8, although only 6 years old at the time, was withdrawn from service within a year of these tests and completely rebuilt.¹⁰

But what do the data mean when taken together, as when plotted in Figure 11.7? What *additional* information on condition can be inferred by using the combination of the two in some analytical manner? Note that, except for unit numbers 8 and 4, there is an identifiable if rather wide inverse relationship exhibited between the two tests – gassing rate goes up as resistance goes down. Roughly speaking then, either one of these tests provides roughly the same assessment as the other (ranking of units by either factor is not that much different), all except for the two outliers, 8 and 4. Of those, unit 4 apparently has considerably degraded insulation, but a low gassing rate. It may be an old unit, with deteriorated insulation due to age, but in otherwise very good condition (inferred because it is producing very little gas). Unit 8, on the other hand, is very likely a failure waiting to happen. Although its insulation tests quite well, the unit is producing gas at a high rate. If deterioration is not the cause (unlikely given the high resistance of the insulation), then the gas is being produced by some other mechanism (e.g. internal arcing) all of which are very undesirable, such as due to internal faults or arcing.

Figure 11.8 plots the years of remaining service given by these transformers at their test score’s position on the diagram.¹¹ Note that either of the tests provide a fair ranking of remaining lifetime versus test results, in that remaining lifetime at the time of the plot clearly correlates, for the most part, with vertical or horizontal position on the plot. However, the correlation in either direction is only fair, not excellent. A combination of both tests helped identify problems somewhat better than either would alone. Certainly insulation resistance tests alone would not have caught unit 8 as a critical unit to investigate further (it was the combination of high resistance and high gassing that indicated a real problem there).

Yet the tests, even combined, did a poor job of getting all the details right. A

¹⁰ The high gassing rate on a relatively good unit with high insulation scores led the utility to do a DGA test (see Chapter 6) on the unit which indicated high levels of acetylene – an indicator of internal arcing. The unit was withdrawn from service, torn down and rebuilt. During that process traces of carbonization were found.

¹¹ These units are part of larger group which included units moved to another location (their positions were upgraded with larger units) and units that failed due to suspected lightning fails and/or vandalism. These units have been removed from the test data, leaving, as far as is known, only a group including units that failed due to “normal” causes.

combination of poor scores in both tests clearly indicated transformer 5 was the “worst scoring” unit of the lot, but it lasted longer in service than units 2, 3 and 7, all of which tested much better in both tests. Unit 7 was a particular failure for the predictive capabilities of both tests. These tests gave no indication of any significant reason to think it did not have a good number of years of life remaining, like those other units nearby on the plot. It might have failed due to something other than age and deterioration (a direct lightning strike) but there was no indication of that.

Greater amounts of data improve condition assessment

The utility that owned these twelve units and actually did the tests shown above in 1984 had more test data than just the two sets of measurements given in Tables 11.8 and 11.9 and shown in Figures 11.7 and 11.8. It also had the total combustible gas counts, not just the rate of increase. Plotted in three dimensions (not shown), this additional data did not help improve the analytical results for these twelve units, except to verify that unit 4 had a low gas total, that thus despite its age really was in good condition.

Thus, a third set of measurements helped resolve one of two anomalies in the test results (units for which the other two test results differed as to evaluation of relative condition to the set as a whole). But it did little to improve the long-term predictive capability of the data set as a whole – it changed the prognosis for units 2, 3, 5 and 7 in no substantive way.

Periodic and Frequency Testing is Required for Efficacy

The data shown here and the results they give in application are typical. Despite careful measurement, estimation of remaining lifetime – long-term condition assessment – is inexact. This test (and others) bears out several lessons:

1. *Multiple tests* help, but mostly because it gives more opportunities to “flunk” a unit on the basis of a single test. Note that unit looked fine from the standpoint of insulation resistance. But it scored so poorly on the TCGA test, that it received almost immediate attention.
2. *Explanations based on other data* help make decisions. Unit 4 was an old unit, something known to the test personnel from their property and maintenance records. That age data provided an explanation for why unit 4 had a lower resistance reading – age means deterioration means lower resistance. Therefore unit 4’s reading was basically expected. Had unit 4 been three years old, the interpretation of that test data would have been far different – there is a problem.
3. *The best of tests are imperfect predictors.* As mentioned above, roughly 4 out of 12 units (2, 3, 5 and 7) in this set ultimately had service lifetimes different than might have been expected based on their test results. These strike the authors as typical results – long-term condition

analysis is dependable only about 50% of the time.

Thus, the overall conclusion one reaches about using test data for condition assessment is that:

1. *Multiple tests* have value because they test for different things or test the same thing (insulation strength) in different ways, and therefore increase the number of ways that flaws or weaknesses can be detected.
2. *Good records* are a critical part of inspection and testing and can be used in conjunction with test data to improve evaluation.
3. *Periodic testing* must be done, because long-term predictive capability of the tests is far from perfect.
4. *Models of equipment behavior* and function are essential, as they permit the test results to be interpreted both for consistency and importance.¹²

This mirrors what was said in Chapters 6 and 7 when comparing modern equipment diagnostic technology to medical diagnostic capabilities. In both fields, the tests are fairly good at identifying a reason (if there is one) for immediate concern due to an impending failure, but poor at predicting how condition will deteriorate, or not, over the long term.

Basically, both the results shown here and experience gained elsewhere all indicate the same thing: that testing and diagnostics, while giving some indication of long-term condition prognosis and remaining lifetime, are basically accurate only at identifying flaws and weaknesses that *currently* or soon will limit capability. Thus, the only viable way for a utility to maintain a good picture of its asset base's condition, is to perform frequent "check ups" on that equipment.

11.10 CONDITION TRACKING

As a carefully considering categorization of equipment that leads to recommendations for service and management actions, condition determination alone provides a great deal of value to a power delivery system owner-operator. But *condition tracking*, which involves recording equipment data and monitoring change over time, provides much greater value. Generally, condition tracking involves periodic gathering of equipment data ("asset data") along with automatic or routine state and condition evaluation, following by evaluation of condition and state changes. The review period used varies greatly among applications the authors have seen, from quarterly to daily in cases of key equipment. In many cases, a variable temporal resolution approach is taken to condition and data archiving. This might involve monthly readings which are archived for a twelve-month period, quarterly readings only for the prior four years and annual readings

¹² Models contribute value by relating test results to how the device functions or reacts and helping identify if the results are merely important or critical. "Model" here does not necessarily mean a computer model. It could be simply the conceptual understanding that test personnel have of how a device like a transformer ages and behaves in use.




Figure 11.9 Condition tracking and the process built around it will require immediate attention any time condition changes. Here, a unit installed in 1992 is tracked over time. It transitions from “as new” to “good” in 1999, at which time the attention given to it consists solely of moving it onto a list for more routine inspection and testing, as for all “good” units. Seventeen years later, roughly on schedule as far as age and its effects go, it moves to “nominal” at which time the utility’s guidelines call for an immediate full inspection and periodic inspections and service afterwards corresponding to that condition categorization. Only eight years later it transitions to “concern,” far ahead of the normal anticipated trend, it moves to condition D. At this point it gets immediate attention as a potential problem.

and data retained prior to that. Data for bad actors, critical units, or those being monitored for R&D or other reasons of special interest, might be retained in more detail. Condition tracking can be done as a stand alone process and application or as part of a integrated asset management and registry. Benefits are:

Incipient problem detection. As in human and veterinarian medicine, change in a key measure of the health index is often the earliest and most dependable warning about impending problems. Particularly in major equipment, a sudden change in gassing or acoustic readings may indicate a problem that must be dealt with now. Thus, *condition change alarms* or notifications are often given (Figure 11.9). An asset tracking system might e-mail particular managers with alerts if condition changes radically in a way that should cause concern, etc.




Figure 11.10 Tracking of the state variables, which could also be considered tracking of the major component conditions of the unit in this case, reveals the cause for the most recent condition change and also a reason for some haste in taking action. Even though the DGA level itself is firmly within mid range, the *rate* of gassing is extreme concern, justifying a high level of concern for the unit's integrity. This transformer's core is apparently not healthy, and the unit will receive immediate attention.

Cost management. Rate of change in health metrics is another factor to be considered in prioritizing units and allocating limited service and repair resources. Of two breakers that are both rated “acceptable, but of concern,” the one deteriorating fastest may be the one most in need of service.

Improved predictive models. Recording and studying the history of condition and how it changes can clearly lead to improved models of how such trends proceed. This applies both to the computer models and processes built around them and the intuition of those using the systems. As one utility manager put it: “Using our system for two years taught me that the category ‘acceptable but of concern’ means we’re never going to have time to look at it and it’s really good enough to go anyway.”

People Often Say Condition Tracking When They Mean State Tracking

Using the distinctions made in the summary, tracking of condition can be useful but not nearly as useful as tracking of state (Figure 11.10). The trigger that brings

a unit experiencing change to management's attention will invariably be a change in condition ("Bayfront Substation's 345/138 auto transformer has gone from B to C in the last week."). But details related to its state and perhaps even the raw data itself may be needed immediately in order to understand the significance of the change and any urgency for action. ("Gassing has increased fourfold since the through-fault on the 19th of April . . . Traces of acetylene have been detected").

Thus, when most people or organizations talk about "condition tracking" they actually mean the periodic archiving of inspection, operating, and test data along with the state estimates and the overall condition determination. Condition is the factor that is *tracked* in terms of alarm criteria and used as the trigger for an alarm, notification, or recommendation for action – i.e., state might change but if the unit's condition remains the same then any recommended actions or consequences of its condition remain the same.

11.11 SUMMARY AND FINAL COMMENTS

Understand the Difference Between Condition Assessment and Failure Prediction

As mentioned in Chapter 6, exact knowledge of when a unit would fail would be worth a great deal to a utility. But even the best condition-assessment methods can not predict exactly the expected remaining lifetime, other than in extreme cases where the unit's failure is imminent. In many such cases monitoring and testing can usually determine the unit is about to fail.

In all but those "near to failure" cases, the best that good condition assessment can do is to rate the *relative likelihood* that a unit will fail. One hundred transformers judged to be in "poor condition" will see far more failures in the next twelve months than one hundred that have been assessed to be in "very good" condition. Just how poor and just how good the units have been assessed will give some indication of failure rate. The information provided is a step in the right direction and worth a good deal toward improving service quality and allowing optimal management of equipment lifetimes. But again, this condition assessment is not a prediction of expected lifetime. Instead, it can and should be used to set priorities and general operating policy.

Good Record Keeping

Good records on the service and maintenance history of major equipment should be a part of any comprehensive condition assessment and improvement program. Records of the annual peak load and duration of peak served, and kWh served through the unit should be maintained. This equipment operating record library should also contain information on the number, type, and time of all through-faults and other types of abnormal events that occurred to the unit. Similarly, records of all inspections and maintenance should be maintained along with any comments on findings during inspections and repair ("Lower northwest corner of case dented").

Record keeping has a good deal of value. Recall Chapter 6's discussion of the value that exact knowledge of future time-of-failure would be worth. Accurate prediction of the exact failure of time is probably not feasible and is certainly not within the capabilities of existing technologies. Nevertheless, an old transformer is worth considerably more to a utility, in terms of being able to estimate its likelihood of failure, if records are maintained so its operating history is available to help determine its condition. Therefore, good record keeping has real value.

The most common failing among utilities in terms of record-keeping is a lack of historical data on through-faults, which many power systems engineers (including the authors) believe is a leading cause of transformer failure.

Inspection, Test, and Diagnosis/Tracking Program

The methods covered in this chapter are worthwhile only if applied in well-managed programs aimed at improving and maintaining service quality at high levels while keeping costs as low as possible. The recommended method of determining priority and of setting the intervals for inspection, testing, and diagnostic evaluation is to use the appropriate approach as described in that chapter.

REFERENCES AND BIBLIOGRAPHY

ABB Transformer Division, *Service Handbook for Power Transformers*, ABB Transformer Remanufacturing and Engineering Services, St. Louis, 2006.

P. Gill, *Electrical Power Equipment Maintenance and Testing*, Marcel Dekker, New York, 1998.

H. H. Malik, A. A. AlArainy, and M. I. Qureshi, *Electrical Insulation in Power Systems*, Marcel Dekker, 1998.

S. D. Myers, J. J. Kelly and R. H. Parrish, *Fifty Years: A Guide to Transformer Maintenance*, Transformer Maintenance Institute, Akron, OH, 1981.

O. C. Seevers, *Management of Transmission and Distribution Systems*, Fairmont Press, Lilburn, GA, 1995.

W. A. Thue, *Electrical Power Cable Engineering*, Marcel Dekker, New York, 1999.

12

Optimization

12.1 INTRODUCTION

In the first edition this chapter was titled *Prioritization Methods for O&M* and discussed methods for reliability centered maintenance (RCM) of a utility delivery system. The new title reflects a slightly expanded scope. RCM is still covered as before. But the basic prioritization approaches applied in the RCM examples can be set up to work toward maximizing something other than reliability, such as increased lifetime or business value – related but slightly different goals. This chapter discusses the concepts that apply to trying to optimize any goal with respect to service, maintenance, refurbishment, and life extension of power delivery equipment, using RCM as its main example.

The goal of virtually every organization is to try to maximize the results it wants and minimize the consequences it does not. Often this is referred to as optimization. That word, along with its verb, adverb and adjective versions, are often horribly abused terms. Few optimization methods, at least as practically applied, truly find an optimal answer or plan. They find good or pretty good solutions instead, often better ones than a project team can find unaided by such tools. Often other products of optimization's use are as valuable as the "optimum answer" it provides, products such as the determination of the true cost of constraints and limitations that must be respected, or the interaction effects of various conflicting business goals, etc.

This chapter focuses on methods utilities and industrial power system owners can use to decide how to maximize the care of the power systems they own. Here, "care" means the expenditure of O&M budgets and resources – non-capital resources – on the system, in order to assure that equipment can and will continue

to be able to do its job dependably and safely. Regulated utilities are permitted to pass legitimate and approved amounts of O&M expenses on to their customers via the rates they charge for power and energy. However, those O&M budgets are limited to what has been approved by regulators. Beyond that, many utilities and industrial power system owners have only limited numbers of skilled employees and contractors available to do such work, and only limited windows of time during the year when they can do it. For all these reason, a practical reality of power system ownership is that there are always more apparently worthwhile O&M activities that *could be done* than can be done. Allocation of O&M boils down to priorities – to favoring certain activities like inspection and tracking over more service, or focusing on equipment of a certain age and equipment rather than all equipment, or on only certain types of equipment and not all.

This chapter discusses both the function of deciding on how much to fund on what equipment care activities, and why, and analytical and procedural methods used by power system owners to do so. Entire books can and have been written about power system equipment maintenance and service, and provide great detail on techniques as well as the ifs, ands, and buts of maintenance and service applied to different equipment and situations (See References and Bibliography). This chapter focuses on breadth rather than depth, providing an overview with a focus on how the function and the techniques available work toward managing an aging power delivery infrastructure.

Section 12.2 begins with some basic concepts about maintenance and service focus and the use of optimization. Sections 12.3 and 12.4 then examine a large (for a textbook) RCM example in a series of evolutions from simplistic to fairly detailed evaluation. Sections 12.5 and 12.6 review several practical issues required to fit technical methodology to a power system's business needs. Section 12.7 then takes up the concept of *optimization*, a decision-making process that addresses more than one goal (e.g., reliability and extended equipment lifetime) and considers practical constraints (Project A has to be done before Project B can be done, only one of Projects C, D, or E can be done). The chapter concludes with a summary of key recommendations in Section 12.8.

12.2 PRIORITIZING INSPECTION, MAINTENANCE, AND SERVICE

One Utility: One Decision-Making System

Any utility that wishes to provide good customer service in spite of the challenges created by an aging infrastructure must prioritize its equipment needs and utilize both its physical assets and its O&M resources as effectively as possible. Doing this involves managing three interrelated activities:

1. *Prioritization of O&M resources* including those for inspection, testing, and diagnostics, along with preventive maintenance, repair, and replacement.
2. *Prioritization of capital spending* to meet new expansion needs and for replacement of equipment that is too old or costs too much to repair.

3. *Utilization* of existing and any new equipment, including the loading level for each unit.

Each of these three activities affects how the distribution company utilizes its assets. Each of them interacts with the others: a change in one might change, good or bad, the situation for one or both of the other two. Thus a utility needs to look at how its total budget should be allocated, beginning with whether O&M or capital buys it more performance, while taking all three, and their interactions, into account.

The first two activities spend different parts of the utility's budget, but each part represents money that is spent. Utility executives and managers must never forget that the third activity – decisions on how to use the equipment – “spends” in a very real way, too. It “uses up” the lifetimes of existing equipment, wearing it out at a faster or slower rate depending on loading, maintenance, and operating policies. Higher loading levels get more value from equipment now, but increase the loss of life rate, “spending” that remaining lifetime faster. Sooner or later, that means both increased O&M costs (as equipment deteriorates from the wear and tear of service, inspection periods must be more frequent, and breakdowns are more common – see Chapters 6 and 11) and higher capital spending (for replacements).

Thus, the power system owner's goal is to balance all three aspects – O&M, capital spending, and usage, in a way that maximizes the business “value” received for ownership of the equipment, howsoever the organization defines that value. Chapter 15 will present six examples of companies with very different business goals and values and show how those lead to very different perspectives and guidelines on use, maintenance, and life cycle management.

Prioritization is Needed Because of Limited Budgets

Resources are always limited with regard to what the power system owner would *like* to do. A utility might have 750 medium voltage breakers in service. It might have data that indicates that inspection and service every three years leads to the fewest problems, best service reliability, and longest lifetimes. But it may have resources, budget, and scheduling windows that permit it to do maintenance and service on no more than 125 years per year. Should it simply cycle through all 750 once every six years? Probably not, if it is trying to minimize customer reliability problems: it can try to focus on units that have an above average effect on customer reliability *and* an above average need for service - those that have operated more often than most units, or are older and more deteriorated, and are on circuits or in places where they affect the service to large numbers of customers.

Regardless, the point is that this question would not arise at all if the utility could do everything and anything it thought necessary: prioritization is needed, and is all about, dispatching or allocating a resource that is scarce – less than all that could be used.

The Need for Prioritization Makes Knowledge of Condition Valuable

A key element of managing an aging infrastructure is condition tracking based on inspection and condition assessment – *knowing* the level of deterioration and the dependability of specific equipment – so that the infrastructure operator can both decide how best to use that equipment and decide if and how to care for it best. Inspection gathers data that is the basis for that knowledge. By and of itself, inspection and condition tracking do nothing to improve condition, prolong lifetime, or increase the value a utility will see from the equipment it owns. The value they provide is purely in the ability it gives the utility to focus the remainder of its resources – those not spent on inspection but on doing something to the equipment – with greater effectiveness. Thus, a point for all aging infrastructure managers to keep in mind, a key concept behind effective management – is the value of condition knowledge. How much is the knowledge inspection provides worth to the power system owner-operator?

Given that a power system owner has an O&M budget it can allocate to equipment inspection and service of a , how much should it spend on inspection? Spending amount b on inspection leaves only $(a-b)$ for service, but presumably that amount will be more effectively spent because of the knowledge inspection provides.




Figure 12.1. Some amount of inspection and condition tracking will lead to optimum results for the infrastructure owner. This diagram shows a hypothetical situation where 20% of budget spent on inspection doubles the effectiveness of the remaining budget, and spending 80% quadruples it, resulting in 175% and 80% of average results respectively. Spending all monies on inspection results in no benefit.

Thus, the value the power system owners derive from their total inspection and service budget is

$$\text{Value } (a,b) = F(b) \times (a-b) \quad (12.1)$$

where:

a = the total budget

b = the amount spent on inspection and condition tracking

$F(b)$ = an effectiveness multiplier for service

Function $F(b)$ is an effectiveness of service multiplier that measures how effective the prioritization or targeting of particular items is. It is defined to have a value of 1.0 at $b = 0$, meaning that the effectiveness at nothing spent on inspection and condition tracking is essentially average: pick the units to maintain at random and over the long haul, the utility will get the average. If an inspection budget b is effective, it might improve results by 10% – $F(b) = 1.10$. If $b = 5\%$ to get that 10% improvement in effectiveness, then overall results improve (total result would be $110\% \times .95$ of the budget remaining after inspection expenses, or 104.5%, a 4.5% improvement in results.

In the example of 750 medium voltage breakers above, with no inspection knowledge the owner will select a batch of 125 breakers “at random.” There will be a few in this group that really need service, some that don’t, and many somewhere in between. $F(0) = 1.00$. The owner gets 125 sectionalizers serviced times 1.0 effectiveness, or 125 units of value from this work.

But suppose the owners spend 20% of the service budget, a , on inspection, so they can bias their selection of units to those that are more in need service. Perhaps the knowledge gained doubles the value they get for every dollar spent on service. But since they spent 20% on inspection they can service only 100 units. Still, they come out a winner: 100 breakers times 2.0 effectiveness = 200 units of value.¹ Spending 20% of a on inspection improved their results by 75%.

If the owners increase the amount they spend on inspection above 20% perhaps they get even more total value. Better inspection leads to more knowledge permits even better dispatching of the service to units where it makes a difference. But at some point, as the monies they devote to service increase, the total value will drop. If they spent all their budget on inspection, it does not matter how effective their service would be because they have no money to spend. Thus, any total-value-derived-versus-amount-spent-on-inspection-curve looks something like that shown in Figure 12.1. It will start off at zero spent on inspection with a value of 1.0. It will rise to some peak value, and then drop to zero and 100%. The owner’s goal is to find that peak value. *Knowledge of the function $F(b)$*

¹ The authors are deliberately not identifying what effectiveness means here, because this example is, at least on a qualitative level, generalize-able to any and all types of results. Effectiveness could mean the decrease in SAIFI affected by the service, or reduction in unexpected work orders required, or increase in lifetime, etc.

is the key to optimizing this. Thus, not only does knowledge of condition have value, but so does knowledge of how changing the focus on service changes that value. The case study covered in Chapter 16 will discuss this and how it played out in practical terms for one utility.

12.3 RELIABILITY-CENTERED MAINTENANCE

A recommended approach to allocation of O&M resources (money, people) for electric utilities faced with aging infrastructure problems is to adopt a form of Reliability-Centered Maintenance (RCM).² What form of RCM is best depends on details of the utility's particularly situation. The basic concept is simple: allocate maintenance resources based on how they contribute to reliability.

Maintenance and Maintenance Management

The traditional view of maintenance's purpose is that it is used:

- To make certain equipment performs its intended function in a satisfactory manner
- To reduce long-term costs by servicing equipment before deterioration causes any avoidable damage
- To avoid unexpected outages by detecting failure in advance

Resources can be allocated by a utility to activities aimed at achieving any and all of these goals (inspection, testing, service, repair). There are several "philosophies" or perspectives on how maintenance can be managed.

1. *By the book* – do maintenance on a periodic basis. The period between service and the maintenance to be done are exactly as prescribed by manufacturer's recommendations or applicable standards.
2. *By the book* (modified) – do maintenance on a periodic basis. To reduce cost, the period between service work and the actual work done each time are respectively the lengthiest period and the leanest amount of service the manufacturer/standards will permit.
3. *Only as needed* – maintain equipment only when it gives operational signs of needing service (e.g., tap changer non-operative).
4. *Reliability-Centered Maintenance*. Allocate maintenance resources among equipment in a manner that will maximize the improvement in "reliability."

² The authors will use "RDM, to denote *results*-driven management, which encompasses the intra-project, marginal benefit-cost methods described here (see pages 179-180) applied to both maintenance and capital spending, and the more commonly seen all-caps designation, RCM, for "Reliability-Centered Management, which will mean here the expanded version of RCM described in this chapter.

5. *Asset Management*. Allocation maintenance resources in an optimized way, taking into account all options and aiming to satisfy multiple corporate goals or targets, and perhaps while also considering the capital budget, as well.

Among utilities in the United States, most traditionally adopted a maintenance program that was somewhere between 1 and 2, performing all the prescribed maintenance on equipment, but often on a slightly less frequent basis than the manufacturers would recommend.

During the 1990s, as electric utilities throughout the United States, faced both increasing competition and uncertainty as to cost-recovery rules in the future, many made every effort to cut spending, including cutting back on maintenance budgets. As a result, there was a shift from (1) above toward (2), with some utilities adopting procedures that were very close to (3), essentially performing maintenance only when equipment failed or gave signs of imminent failure.

Always: A Budget-Constrained Situation

A point worth reiterating is that nearly every modern distribution utility, and certainly every aging infrastructure utility, is permanently in a “budget constrained” situation – there will always be more worthy, and *justifiable*, uses for resources than there are resources. A modern electric utility simply cannot afford to replace all of its aging equipment just because that equipment is old. It does not have enough funds. Therefore, the number of potential projects (capital) or programs (O&M) that are worth funding will always exceed the number that can be approved. The utility will permanently find itself in a “budget-constrained” situation, where it must make decisions between worthy projects on the basis of “bang for the buck.”

Prioritization for utilities in this situation is about much more than deciding on the order in which projects will be done. It is about deciding which projects will be done, and which won’t be done at all, even though those projects may be worthy and, in some sense (perhaps the traditional sense), justifiable. Prioritization in a budget-constrained situation means determining which projects and programs are the *most* worthy – which will do the *most* good – and funding only those with the limited budget.

Techniques to apply that concept to capital projects and to non-maintenance related operating costs are covered in parts of Chapters 5 and given in the case studies covered in later chapters and the Appendices. That same basic concept is used in this chapter with respect to maintenance. All projects, and all alternatives for all projects, are evaluated with respect to how well they help improve customer service quality. Those that provide enough “bang for the buck” are approved. Those that don’t are not.




Figure 12.2 Reliability-centered maintenance does not seek to balance service quality and cost (top), although management practices used to administer it and similar marginal benefit-cost based programs may (see Chapter 15). Instead, RCM seeks to maximize bang for the buck from O&M dollars, which it does by identifying inequalities in what various potential projects and programs will provide for the dollar (bottom).

Reliability-Centered Maintenance

As mentioned above, the real question in allocating resources when faced with service quality challenges *is how much will each activity contribute to the utility's goals, and at what cost?* In RCM, the costs of various inspection, testing, service, repair, and replacement activities that the utility could do are assessed on the basis of how much they will contribute to the utility's goal (good customer service) versus their cost – a merit score. They are then ranked by merit score. Winners are those that score particularly high on a “bang for the buck” basis – those that provide more customer service quality improvement for less money.

Figure 12.2 illustrates this concept. RCM is depicted at the bottom of the figure. Basically, RCM involves weighing the merits of various projects and picking those with the greatest benefit at the least cost. Figure 12.2 also illustrates what RCM isn't. It is not, as shown at the top, an attempt to “balance” or optimize the total amount of maintenance budget against the value of the resulting service – so called Value-Based Maintenance (VBM). VBM has a place in the planning and decision making at an aging infrastructure utility, but that place is as a strategic planning tool to determine overall budget, and to prioritize the total maintenance budget versus capital and other operating budgets.

RCM Prioritizes Candidate Programs and Projects

RCM prioritizes a group of projects or programs based on merit score, permitting the utility to select the most effective candidates to achieve its goals. Four important points are that:

1. RCM effectiveness is usually determined by both the number of customers served and by the improvement in reliability that is expected from the project (customers x reduction in outage) or something similar.
2. Cost is the total cost of the service or maintenance program, including materials, labor, etc. A key element of good RCM programs is that costs are consistently treated among all classes of equipment and types of activities.
3. Cost may include a credit for avoided costs down the road, such as the avoided cost to repair a failure in the future, suitably discounted for time value of money and any lack of certainty.
4. Benefits do not last forever. Good methods recognize that if a breaker is maintained, it will be in “good shape” for only a period of time, after which it will again need service.

Projects are then ranked on the basis of bang for the buck and the most effective one picked. Working down such a list, a utility chooses all it can afford until it reaches a point where its maintenance budget is exhausted. The remaining projects on the list are not funded.

12.4 BASIC RELIABILITY-CENTERED PRIORITIZATION

Table 12.1 lists 35 distribution system maintenance projects or programs that cover a range of typical maintenance activities.³ These include both major maintenance projects on “large equipment” and smaller projects, as well as a major program (pole inspection/replacement). The entire list of 35 projects has a cost of \$3,174,500. However, for this example, the utility has only about 40% of this, or \$1,375,000, to spend. It must select the best \$1,375,000 worth of projects from out of that list.

These projects could be prioritized in any number of ways, but in this example they will be ranked on the basis of impact on customer service quality per dollar. Impact is the decrease in customer-minutes out of service that is expected from each project during the next few years, defined as

$$\text{Impact} = \sum R_t \times P_t \quad (12.2)$$

Where R_t = reduction in customer minutes out of service
per year, in year t , expected due to the activity

And P_t = present worth factor for year t , $t = 0$ (this year),
and years 1, 2, 3 and further into the future

The impact must be evaluated over multiple years, because maintenance projects have a positive impact on equipment reliability that lasts more than one year. This sum usually needs to be evaluated over only the short term, for the next three to seven years. Three years is used in this example, mainly because the tables would become unwieldy if more years of impact improvement were displayed. However, in actual implementation, decisions comparing replacement versus maintenance often need to take a longer perspective, as will be demonstrated with examples given later in this chapter. The concept and method of implementation are identical, however.

Present worth factor is applied to customer outage reductions in future years, in order to adjust all factors onto a present-time basis for the current year. In this case a factor of .9 is used.

Cost in this example is total cost of the proposed activity. In all cases in this example (at least for now), the costs are operating budget items within the present year (year 0), and so no present worth adjustment or consideration of costs beyond the present year needs to be done. Programs or projects that include a continuing cost for several years or a commitment to future expenses, or that make partial payment today with future payment of the remainder for some service or replacement done next year or thereafter, must all be reevaluated with

³ This is literally a “textbook example” created by the authors to illustrate well certain key concepts in RCM. However, all of the data given is both consistent with the authors’ experience in actual utility systems, and with the failures rates, repair costs, and expected improvements given in the various equipment-specific appendices in the IEEE Gold Book (*IEEE Recommended Practices for the Design of Reliable Industrial and Commercial Power Systems*) and its references.

Table 12.1 Example List Of Projects As Evaluated By RCM Ranking

Project Number	Equipment Type	Work Proposed	Cost of of Work	Number of Customers	Peak Load (kVA)
1	40 MVA transf. EC03	Service/Rebuild	\$65,000	10,899	38,000
2	40 MVA transf. EC01	Recond. Oil/Servc.	\$22,000	7,800	37,000
3	40 MVA transf. JT235	Inspect, Test, Follow up	\$11,000	842	31,000
4	40 MVA transf. DT03	Inspect, Test, Follow up	\$11,000	4,612	41,000
5	HV breaker # EC03	Complete Service	\$22,000	10,899	38,000
6	HV breaker # EC01	Complete Service	\$22,000	7,800	37,000
7	HV breaker # 221	Complete Service	\$22,000	634	26,000
8	HV breaker # A0121	Complete Service	\$22,000	11,300	31,500
9	LV breaker #ES06	Complete Service	\$6,500	1,681	6,891
10	LV breaker #EC06	Complete Service	\$6,500	2,200	7,100
11	LV breaker #NB01	Complete Service	\$6,500	580	6,535
12	LV breaker #NB02	Complete Service	\$6,500	650	7,800
13	LV breaker #NB03	Complete Service	\$6,500	1,191	4,883
14	LV breaker #CN1A	Complete Service	\$6,500	1,058	4,338
15	LV breaker #CN2A	Complete Service	\$6,500	891	3,653
16	LV breaker #DT3N	Complete Service	\$6,500	2,091	8,573
17	LV breaker #DT2N	Complete Service	\$6,500	112	10,000
18	UG cables, feeder EC06	Replace All	\$27,000	2,200	9,020
19	UG cables, feeder EC07	Test, Replace	\$18,000	592	4,805
20	UG cables, netw. DT3B	Replace Complete	\$43,000	245	12,600
21	UG cables, netw. DT3A	Test, Replace	\$31,000	340	11,300
22	OH feeder BN06	Test, Replace	\$58,000	950	4,503
23	69 kV circuits 34 & 37	Test, Replace	\$170,000	22,005	103,800
24	OH feeder EC17	Test, Replace	\$58,000	1,578	7,350
25	Poles, central east dist.	Test, Replace	\$290,000	80,134	328,000
26	OH feeder BN04	Inspection, Replace	\$43,000	1,345	5,780
27	OH feeder BN05	Inspection, Replace	\$17,000	1,609	6,108
28	OH feeder CT14	Inspection, Replace	\$56,000	1,671	6,875
29	OH feeder CT05	Inspection, Replace	\$43,000	1,389	7,324
30	OH feeder JM07	Inspection, Replace	\$86,000	1,555	4,120
31	Old Town sub feeders	Test, Replace	\$210,000	22,789	72,007
32	Riverside sub feeders	Test, Replace	\$335,000	28,093	85,000
33	Hillside sub feeders	Test, Replace	\$219,000	14,004	45,070
34	Central sub. feeders	Test, Replace	\$310,000	28,345	110,000
35	Farnton substation	Test, Replace	\$905,000	43,000	124,000
Total			\$3,174, 500		

Table 12.2 Example List Of Projects With Change in Failure Rates for the Next Three Years (%/yr.) Without and With the Proposed Work Being Done

Project Number	Equipment Type	Work Proposed	New	Failure Rate			Failure Rate if Worked		
				Now	Yr +1	Yr +2	Now	Yr +1	Yr +2
1	40 MVA transf. EC03	Service/Rebuild	1.00%	6.0%	6.3%	6.6%	1.8%	1.8%	1.9%
2	40 MVA transf. EC01	Recond. Oil/servc.	1.00%	5.0%	5.3%	5.5%	3.0%	3.1%	3.2%
3	40 MVA transf. JT235	Inspect, Test, Follow up	1.00%	4.7%	4.9%	5.2%	3.2%	3.3%	3.4%
4	40 MVA transf. DT03	Inspect, Test, Follow up	1.00%	6.0%	6.3%	6.6%	4.0%	4.1%	4.2%
5	HV breaker # EC03	Complete Service	0.60%	2.0%	2.1%	2.2%	1.1%	1.1%	1.1%
6	HV breaker # EC01	Complete Service	0.60%	2.0%	2.1%	2.2%	1.1%	1.1%	1.1%
7	HV breaker # 221	Complete Service	0.60%	2.0%	2.1%	2.2%	1.1%	1.1%	1.1%
8	HV breaker # A0121	Complete Service	0.60%	2.0%	2.1%	2.2%	1.1%	1.1%	1.1%
9	LV breaker #ES06	Complete Service	0.92%	2.5%	2.6%	2.8%	1.5%	1.5%	1.5%
10	LV breaker #EC06	Complete Service	0.92%	2.5%	2.6%	2.8%	1.5%	1.5%	1.5%
11	LV breaker #NB01	Complete Service	0.92%	2.5%	2.6%	2.8%	1.5%	1.5%	1.5%
12	LV breaker #NB02	Complete Service	0.92%	2.5%	2.6%	2.8%	1.5%	1.5%	1.5%
13	LV breaker #NB03	Complete Service	0.92%	2.5%	2.6%	2.8%	1.5%	1.5%	1.5%
14	LV breaker #CN1A	Complete Service	0.92%	2.5%	2.6%	2.8%	1.5%	1.5%	1.5%
15	LV breaker #CN2A	Complete Service	0.92%	1.8%	1.9%	2.0%	1.2%	1.2%	1.3%
16	LV breaker #DT3N	Complete Service	0.92%	1.9%	2.0%	2.1%	1.3%	1.3%	1.3%
17	LV breaker #DT2N	Complete Service	0.92%	3.7%	3.9%	4.1%	1.9%	1.9%	2.0%
18	UG cables, feeder EC06	Replace all	3.50%	15.0%	15.8%	16.5%	3.5%	3.6%	3.7%
19	UG cables, feeder EC07	Test, Replace	4.10%	9.0%	9.5%	9.9%	7.0%	7.2%	7.5%
20	UG cables, net DT3B	Replace Complete	2.80%	8.5%	8.9%	9.4%	2.8%	2.9%	3.0%
21	UG cables, net DT3A	Test, Replace	2.60%	6.0%	6.3%	6.6%	4.6%	4.8%	4.9%
22	OH feeder BN06	Test, Replace	4.50%	10.0%	10.5%	11.0%	7.8%	8.0%	8.3%
23	69 kV circuits 34 & 37	Test, Replace	0.09%	0.3%	0.3%	0.3%	0.2%	0.2%	0.2%
24	OH feeder EC17	Test, Replace	2.90%	17.0%	17.9%	18.7%	11.4%	11.6%	12.0%
25	Poles, CE dist.	Test, Replace	0.03%	10.0%	10.5%	11.0%	6.0%	6.2%	6.4%
26	OH feeder BN04	Inspection, Replace	3.40%	10.2%	10.7%	11.2%	7.5%	7.7%	7.9%
27	OH feeder BN05	Inspection, Replace	3.80%	11.4%	12.0%	12.6%	8.4%	8.6%	8.9%
28	OH feeder CT14	Inspection, Replace	1.70%	5.1%	5.4%	5.6%	3.7%	3.8%	4.0%
29	OH feeder CT05	Inspection, Replace	3.40%	10.2%	10.7%	11.2%	7.5%	7.7%	7.9%
30	OH feeder JM07	Inspection, Replace	3.20%	9.60%	10.1%	10.6%	7.0%	7.2%	7.45%
31	Old Town sub feeders	Test, Replace	4.50%	13.5%	14.2%	14.9%	9.9%	10.2%	10.5%
32	Riverside sub feeders	Test, Replace	6.60%	19.8%	20.8%	21.8%	14.5%	14.9%	15.4%
33	Hillside sub feeders	Test, Replace	8.10%	24.3%	25.5%	26.8%	17.8%	18.3%	18.9%
34	Central sub. feeders	Test, Replace	2.80%	7.1%	7.5%	7.8%	4.5%	4.6%	4.8%
35	Farmton substation	Test, Replace	1.70%	3.9%	4.1%	4.3%	2.0%	2.08%	2.15%

a present worth evaluation of cost. Cost is defined on the same present worth basis as benefits, as

$$\text{Cost of project } i = \text{Cost}_i = \sum C_{i,t} \times P_t \quad (12.3)$$

Where $C_{i,t}$ = cost of project i in year t

And P_t = present worth factor for year t

Reduction in failure rate due to maintenance

Table 12.2 shows the 35 projects listed with their new failure rates, and present and expected outage rates for the next three years with and without the proposed project performed. The failure rates used here are within the typical ranges seen for distribution equipment in various stages of aging.

This example illustrates one highly recommended procedure. The improvements made in failure rate for each project are defined using a *formalized method* of estimating improvement (in this case a probabilistic method discussed later in this chapter). It involves classifying maintenance into categories, with each category modeled as having a certain expected impact.⁴ Such an institutionalized set of standardized expectations for improvement assures uniformity of purpose and that no one “cheats” the system by using their own maintenance efficacy figures. Categories used here are the following:

1. *Complete rebuild* returns failure rate to a failure rate equal to “as new condition” plus 5% of the degradation expected over thirty years. Thus, if a new transformer (but one past its infant mortality times) has an annual failure rate of .05%, and a thirty-year-old one has a failure rate of 2.7%, then a recently rebuilt one has an expected failure rate of $.05\% + .5 \times 2.2\% = .61\%$ /year.
2. *Full service done on the unit* cuts the amount of degradation of the equipment by two-thirds. As an example, suppose that a HV breaker has a new failure rate of .92%/year, whereas one badly in need of full servicing has an annual failure rate of 3%. Then, a recently serviced unit has a failure rate of $.92 + .33 \times (3.00 - .92) = 1.61\%$.
3. *Inspection and all repairs found to be needed made* cuts failure rate by 40%. This, if a group of old low-voltage breakers, all in need of servicing, has an average expected annual failure rate of 3.5%, but only .056% when new, then inspection and needed repairs reduces the value to an expected annual failure rate of $.56 + .60 \times (3.5 - .56) = 2.3\%$.

The above are rough rules of thumb used for this example only and not recommended as hard and fast rules the reader should apply to any utility system.

⁴ “Expected” here means in a probabilistic sense. Failure is a random process that is best modeled probabilistically. So is improvement.

They illustrate that repair activities have been *classified* by expected improvement on a category or equipment class basis. No attempt is made to determine what the improvement made by a specific type of service on a specific device would be. Such categorization is highly recommended. Repair and improvements, like failure rates themselves, are probabilistic in nature. Both the concepts of failure rate itself, and failure rate improvement, are probabilistic, and their application works well only when used as probabilistic expectations applied to large sets of equipment.

Obtaining Failure-Rate Reduction Data

A concern for any utility maintenance management team will be “where do we find the data on failure rate reductions for various maintenance functions?” Ideally, these data should be developed by the utility specifically for its situation, or at least based on studies done elsewhere with judgmental adjustment for any unique factors for the utility’s situation.

This is not as difficult as it seems to many. Data is available at most utilities, and if not, then “generic” data based on published and available public domain data, such as that included in various IEEE books and publications, can be used as a starting point. The authors have never seen a situation where the appropriate data cannot be developed to the point where the RCM approach can be applied in a dependable manner.

One strong recommendation is to never use a lack of data as an excuse not to take this approach. In the author’s experience, that is most often used an excuse by people who simply want to continue to do things as they have always been done. If a utility cannot develop a set of reasonable estimates of the impact maintenance makes on its equipment, on a broad category basis as required here, then it is basically admitting it has no idea of the value of the maintenance it is performing. Doing anything whose value cannot be estimated, in the present industry environment, is as extremely imprudent. A utility that does not develop failure rate reduction data has only a very weak basis upon which to justify its maintenance expenses.

Prioritization Based on Ranking

Table 12.3 shows the result of applying the formulas given earlier to the data on the 35 projects shown in Tables 12.1 and 12.2. There, the customer minutes of improvement per year for years zero, one and two is computed and a present worth sum formed. This is then divided by cost to obtain reduction in customer service minutes expected per dollar – “bang for the buck.”

The projects are listed in Table 12.3 in descending order according to their merit score (improvement per dollar), given in the leftmost column. The “bang/buck ratio” for projects on this list varies from a high of 12.84 to a low of 0.01. Figure 12.3 plots the merit score (bang/buck ratio) versus total budget required to buy all projects of that score or better, for this set of 30 projects. For example, the top project on the list (merit score 12.84, has a cost of \$27,000). The

Table 12.3 Projects Evaluated By RCM Ranking, Listed in Descending Order of Merit

Project Number	Equipment Type	Work Proposed	# Cust	Reduced Yr. 0	Cust-Min Next	Out +2 yr.	PW Sum Cust-min	Cost (yr. 0)	Bang/Buck
18	UG cables, feeder EC06	Replace all	2,200	121440	128436	135502	346789	\$27,000	12.84
4	40 MVA transf. DT03	Test, Follow up	4,612	44275	48703	52681	130779	\$11,000	11.89
2	40 MVA transf. EC01	Recond. Oil/servc.	7,800	74880	81432	87461	219012	\$22,000	9.96
1	40 MVA transf. EC03	Rebuild Complete	10,899	222343	235749	249132	636315	\$65,000	9.79
20	UG cables, net DT3B	Replace Complete	245	67032	71207	75341	192144	\$43,000	4.47
25	Poles, CE district	Test, Replace	80,134	575234	625675	672069	1682718	\$610,000	2.76
21	UG cables, net DT3A	Test, Replace	340	22195	25198	27777	67373	\$31,000	2.17
3	40 MVA transf. JT235	Test, Follow up	842	5982	6606	7163	17729	\$11,000	1.61
32	Riverside sub feeders	Test, Replace	28,093	177997	199134	217621	533491	\$335,000	1.59
33	Hillside sub feeders	Test, Replace	14,004	108895	121826	133136	326379	\$219,000	1.49
----- Approval Limit Based on Budget -----									
31	Oldtown sub feeders	Test, Replace	22,789	98448	110139	120364	295068	\$210,000	1.41
34	Central subs. feeders	Selected Replace	28,345	131634	143981	155199	386928	\$310,000	1.25
35	Farmton substation	Complete Rebuild	43,000	385968	415741	443829	1119637	\$905,000	1.24
27	OH feeder BN05	Inspection, Replace	1,609	5870	6567	7176	17592	\$17,000	1.03
19	UG cables, feeder EC07	Test, Replace	592	5570	6348	7014	16964	\$18,000	0.94
8	HV breaker # A0121	Complete Service	11,300	6265	6760	7225	18202	\$22,000	0.83
5	HV breaker # EC03	Complete Service	10,899	6043	6521	6969	17556	\$22,000	0.80
10	LV breaker #EC06	Complete Service	2,200	1376	1493	1602	4018	\$6,500	0.62
6	HV breaker # EC01	Complete Service	7,800	4324	4666	4987	12564	\$22,000	0.57
24	OH feeder EC17	Inspect, Replace	1,578	10680	11752	12714	31555	\$58,000	0.54
9	LV breaker #ES06	Complete Service	1,681	1052	1141	1224	3069	\$6,500	0.47
16	LV breaker #DT3N	Complete Service	2,091	811	891	963	2394	\$6,500	0.37
13	LV breaker #NB03	Complete Service	1,191	745	808	867	2175	\$6,500	0.33
29	OH feeder CT05	Inspection, Replace	1,389	4534	5072	5543	13588	\$43,000	0.32
26	OH feeder BN04	Inspection, Replace	1,345	4390	4911	5367	13158	\$43,000	0.31
14	LV breaker #CN1A	Complete Service	1,058	662	718	770	1932	\$6,500	0.30
23	69 kV circuits 34 & 37	Test, Replace	42,300	5198	5708	6169	15332	\$70,000	0.22
12	LV breaker #NB02	Complete Service	650	407	441	473	1187	\$6,500	0.18
30	OH feeder JM07	Inspection, Replace	1,555	4777	5344	5840	14317	\$86,000	0.17
11	LV breaker #NB01	Complete Service	580	363	394	422	1059	\$6,500	0.16
28	OH feeder CT14	Inspection, Replace	1,671	2727	3051	3334	8174	\$56,000	0.15
15	LV breaker #CN2A	Complete Service	891	310	342	371	919	\$6,500	0.14
22	OH feeder BN06	Inspect, Replace	950	2508	2856	3153	7632	\$58,000	0.13
7	HV breaker # 221	Complete Service	634	351	379	405	1021	\$22,000	0.05
17	LV breaker #DT2N	Complete Service	112	30	34	37	90	\$6,500	0.01
Totals and averages							6,158,862	\$3,174,500	1.81




Figure 12.3 Merit scores of projects versus cumulative budget required to buy down to that level. This set of projects exhibits the “scree pattern” typical of many RCM prioritization results. There are a very few extremely high scoring projects (above 3), a number of “pretty good” (3 – 1.5) projects, and a host of projects with lesser merit.




Figure 12.4 Total “bang” that can be purchased versus budget. By selecting the top ranked projects, the utility can buy over two-thirds of the total benefit available for its budget limit of only 40% of the total cost (dotted line).

utility can buy all projects with scores of 12.84 for \$27,000. The next project on the list has a score of 11.89 and costs \$11,000. The utility can buy all projects with scores of 11.89 and better for \$38,000. And so forth. The first eight projects together total only \$168,000 – just 5% of the total list’s cost – but provide over 25% of the total improvement in reliability from the entire list.

The plot has a “scree” bend, a sharp change in slope between a few very worthy projects and a large number of others.⁵ Here it occurs at a budget value of about \$250,000. This is typical. Most evaluations of this type are dominated by a few very high-scoring projects, amidst a plethora of others.

Figure 12.5 plots total benefit purchased versus budget. It shows how much money is required to buy what total level of improvement. The cost of the first ten projects in Table 12.3 (those with bang/buck values shown in bold) is \$1,374,000, exactly \$1,000 less than the allowed budget total in this example. The utility can spend its budget on these projects – only 43% of that needed for the entire list – and receive 68% of the total benefit in customer-minute reductions of the entire list (4,152,729 customer minutes avoided out of 6,158,862 customer minutes possible from the entire list). The “cutoff value” for funding at the \$1,375,000 budget limit is 1.49-customer service minutes per dollar. No other set of projects it can afford will attain as much improvement.

What This Prioritization Accomplished

RCM prioritized the proposed projects in a way that permitted the utility to rank and select projects that would use budget-dollars most effectively. Among the projects listed, no other subset will provide more value, assuming the utility’s goal is to reduce customer minutes of outage, than those ranked highest on the list. For any budget limit, the utility merely has to begin selected projects from the top, working its way down the list until it exhausts its budget.

RCM, then, is basically a tool to help maximize reliability-related operating activities against a budget limit.

Changing the Figure of Merit

RCM does create one conundrum for a utility. It must choose a figure of merit – a particular performance factor to minimize – in this case the expected customer-minutes of outage. Clearly, that choice will greatly affect just what projects are selected, and what benefits the utility “buys” with its maintenance budget.

Continuing with this same example, what if the utility had selected minimization of kVA-minutes of outage, not customer-minutes of outage, as its

⁵ Scree is a geologic term for the pile of rocks and debris at the foot of a volcanic cinder cone or other tall “monument” rock feature (e.g. Shiprock, NM). As a hard rock cinder cone weathers, bits of rock break away and tumble down to ground level, often rolling some distance from the cone’s near vertical face, and creating the scree – a wide, nearly flat mound surrounding the near-vertical rock monument.

Table 12.4 Projects Evaluated By kVA-Minutes, Listed in Descending Order of Merit

Project Number	Equipment Type	Work Proposed	Load kVA	Reduced Yr. 0	Cust-Min Next	Out +2 yr.	PW Sum Cust-min	Cost (yr. 0)	Bang/ Buck
20	UG cables, net DT3B	Replace Complete	12,600	3447360	3662064	3874675	9881705	\$43,000	229.8
4	40 MVA transf. DT03	Test, Follow up	41,000	393600	432960	468325	1162607	\$11,000	105.7
21	UG cables, net DT3A	Test, Replace	11,300	737664	837466	923193	2239169	\$31,000	72.2
3	40 MVA transf. JT235	Test, Follow up	31,000	220224	243214	263723	652732	\$11,000	59.3
18	UG cables, feeder EC06	Replace all	9,020	497904	526588	555557	1421834	\$27,000	52.7
2	40 MVA transf. EC01	Recond. Oil/serv.	37,000	355200	386280	414878	1038903	\$22,000	47.2
1	40 MVA transf. EC03	Rebuild Complete	38,000	775200	821940	868599	2218511	\$65,000	34.1
25	Poles, CE district	Test, Replace	328,000	2354515	2560978	2750877	6887606	\$610,000	11.3
19	UG cables, feeder EC07	Test, Replace	4,805	45205	51525	56931	137692	\$18,000	7.65
34	Central subs. feeders	Selected Replace	110,000	510840	558756	602288	1501574	\$310,000	4.84
32	Riverside sub feeders	Test, Replace	85,000	538560	602514	658447	1614165	\$335,000	4.82
----- Approval Limit Based on Budget -----									
33	Hillside sub feeders	Test, Replace	45,070	350464	392082	428480	1050407	\$219,000	4.80
31	Old Town sub feeders	Test, Replace	72,007	311070	348010	380316	932335	\$210,000	4.44
27	OH feeder BN05	Inspection, Replace	6,108	22282	24928	27242	66783	\$17,000	3.93
35	Farmton substation	Complete Rebl	124,000	1113024	1198882	1279880	3228720	\$905,000	3.57
5	HV breaker # EC03	Complete Service	38,000	21067	22734	24298	61209	\$22,000	2.78
6	HV breaker # EC01	Complete Service	37,000	20513	22136	23659	59598	\$22,000	2.71
24	OH feeder EC17	Replace as needed	7,350	49745	54737	59220	146976	\$58,000	2.53
8	HV breaker # A0121	Complete Service	31,500	17464	18845	20142	50739	\$22,000	2.31
12	LV breaker #NB02	Complete Service	7,800	4880	5295	5678	14245	\$6,500	2.19
10	LV breaker #EC06	Complete Service	7,100	4442	4820	5169	12967	\$6,500	1.99
9	LV breaker #ES06	Complete Service	6,891	4311	4678	5016	12585	\$6,500	1.94
7	HV breaker # 221	Complete Service	26,000	14414	15555	16625	41880	\$22,000	1.90
11	LV breaker #NB01	Complete Service	6,535	4089	4436	4757	11935	\$6,500	1.84
29	OH feeder CT05	Inspection, Replace	7,324	23906	26744	29227	71649	\$43,000	1.67
16	LV breaker #DT3N	Complete Service	8,573	3327	3655	3950	9815	\$6,500	1.51
13	LV breaker #NB03	Complete Service	4,883	3055	3315	3555	8918	\$6,500	1.37
26	OH feeder BN04	Inspection, Replace	5,780	18866	21106	23066	56545	\$43,000	1.31
17	LV breaker #DT2N	Complete Service	10,000	2693	3000	3270	8042	\$6,500	1.24
14	LV breaker #CN1A	Complete Service	4,338	2714	2945	3158	7922	\$6,500	1.22
22	OH feeder BN06	Inspect, Replace	4,503	11888	13536	14947	36178	\$58,000	0.62
28	OH feeder CT14	Inspection, Replace	6,875	11220	12552	13718	33628	\$56,000	0.60
15	LV breaker #CN2A	Complete Service	3,653	1273	1403	1520	3768	\$6,500	0.58
23	69 kV circuits 34 & 37	Test, Replace	103,800	12755	14008	15137	37623	\$70,000	0.54
30	OH feeder JM07	Test, Replace	4,120	12657	14160	15474	37934	\$86,000	0.44
1b	40 MVA transf. EC03	Complete Service	38,000	601920	644328	685126	1736767	\$38,000	

goal for prioritization? There is considerable merit in using kVA-minutes as the decision-making criterion in RCM. Often, a circuit or transformer will serve only a small number of customers, but those customers each buy a lot of power, and spend a good deal more on power than customers who buy less. One can argue that every kilowatt these larger customers buy should be treated the same with respect to attention and quality control, as those bought by smaller customers. Further, some of these large consumers may be “more important” in terms of societal outage impact.⁶

⁶ Opinion in many parts of the industry holds that interruption of service in the central core of cities and key municipal facilities is much more unfavorable than outages to most other loads, partly due to the impact such outages have on the function of society as a whole.

For example, Table 12.1 showed that project 19 (UG cables, feeder EC07 – test and replace as found necessary) had only 592 customers, about one quarter of the customer count on most feeders, but served 4805 kVA peak load, rather typical for a feeder. Clearly, it serves customers who have a higher demand than average. As a result of the relatively low customer cost, the evaluation shown in Table 12.3 put it mid-pack in the table – it is not to be funded this year. Service to those larger customers will suffer, and for that matter, service to smaller customers on that feeder could suffer in comparison to customers elsewhere.⁷

Table 12.4 lists the 35 projects re-evaluated and prioritized on the basis of kVA-minutes of outage reduction expected per dollar, rather than customer minutes as done earlier. The values are much different since they are based on a different measure. They vary from a high of 229.8 to a low of .44. Again, the ranking exhibits a scree-like shape (not shown). The utility's available budget is reached somewhere in the 12th project (approval of 11 leaves it with about \$225,000 left, approval of project 12 spends about \$100,000 too much. Assuming the first 11 projects are approved and the remaining budget is spent on part of project 12 with proportionate results (it involves many feeders so it would be feasible to only do a portion of it), the total budget of \$1,375,000 buys 22,273,339 MVA minutes of reduction, or 64% of the possible reduction. Project 19, the feeder with few but larger-than-average customers, just barely makes the cut, for although it moves up greatly in relative score, a number of other projects do, too. Other changes are approval of Project 34, which did not make the approval list in Table 12.3, and non-approval of Project 31, which falls off the approval list.

The top projects will be at the top of just about any list

While the change in figure-of-merit altered the composition of the approval list, eight of the top eleven projects stayed the same regardless of the prioritization scheme. Looked at another way, in this example 84% of the budget is spent on the projects that make it into both lists.

This is often the case - the very top rated projects stay top rated even if the exact definition of "merit" is changed. The reason is that projects that tend to score extremely high generally do so because they provide a tremendous improvement in *equipment reliability* per dollar. Whether serving many customers or a high load, that large increase in equipment reliability pays off with a handsome increase in merit score. Careful examination of Tables 12.1 through 12.4 will show that the winners here are all projects where a correction from "very

⁷ In actual fact feeder EC07 (#7 out of East City substation) serves three customers of roughly 1 MVA peak demand each and 499 typical small residential and commercial consumers. Had this feeder, which has very severe reliability problems, served only residential and small commercial customers of the same total peak demand (about 1,700 of them would be typical) the higher customer count would have led to a figure of merit of 2.71, putting the feeder on the approval list. As a result, one can argue that these 499 customers were "screwed" just because they happened to be on a feeder serving several large consumers and therefore having a low customer count.

Table 12.5 Comparison of Top Projects

Ranking Order	As Prioritized by Customer-minutes	As Prioritized by kVA-minutes
1	18	20
2	4	4
3	2	21
4	1	3
5	20	18
6	25	2
7	21	1
8	3	25
9	32	19
10	33	34
11	-	32

poor” to “very good” reliability is made at a good price. Another way to look at it is that prioritizing on the basis of kVA minutes gains 90% of the optimum result from the standpoint of optimizing for customer-minutes.

Both merit indices used above were duration-related. Had a frequency of outage index been used, such as customer interruption count, or kVA-interruptions, the results would have been somewhat less similar to those in Tables 12.3 and 12.4. But even so, more than two-thirds of the projects at the top of the list would remain the same even if a frequency-related merit score (e.g., SAIFI) were used. In general, the use of any reasonable reliability definition seems to gain at least two-thirds of “optimality” as judged by any other figure of merit. Qualitatively, this is a general result:

The use of any reasonable definition of “merit” as the prioritizing index in RCM will provide qualitatively similar overall results.

What “Reliability Definition” Should A Utility Select?

The answer is quite simple: whatever index its regulatory authority has indicated will be used to evaluate performance of the utility. Generally, that is some combination of SAIDI, SAIFI, and a measure of “outliers” – customers who have much worse than average reliability.⁸ Regardless, if law or regulation defines “what is important,” the utility would be imprudent to pick anything else,

⁸ Most utility regulatory commissions have some provision in their regulations requiring attention to customers/feeder who have much worse (e.g., three times) than average duration or frequency of outages, or who are perennially on the list of worse performance every year.

particularly if performance incentives or rate penalties are linked to its performance as assessed by that measure.

This is where performance based rates (PBR) provide a great advantage to the utility. PBR not only provides a completely unambiguous definition of *what* “good performance” means, but it provides a quantitative measure of *how much* it is worth. This permits not only prioritization of projects as shown here, but optimization of the total budget as will be discussed in Chapter 15.

With very few exceptions, there are no significant problems applying the methods discussed above to any type of regulatory definition of “good performance” or PBR rate schedule. The examples given here are a bit simpler to execute and display on textbook pages, than what is generally required in the most actual utility cases, but the concepts are the identical and the execution requires nothing not shown here. Non-linearities in PBR functions (see Chapter 14, Figures 14.6 and 14.8) create some challenges in analyzing individual projects and prioritizing a list to a budget constraint, but the authors have never encountered any system that cannot be optimized in the manner discussed above.

12.5 PRIORITIZATION OF THE TYPE OF MAINTENANCE

Section 12.3’s RCM prioritized projects by picking from a list of proposed maintenance projects. Projects were prioritized based upon:

- The expected reliability improvement,
- The number of customers, the load, or other measures of importance of the equipment’s service, and
- Cost.

Projects were ranked so that the best can be selected against either a goal of using a constrained monetary budget most effectively (as in the example) or against a goal of achieving some target reliability improvement at the lowest cost possible. Applied with good data, focus on objectives, consistency and common sense, that basic method provides good results. However, results can be improved.

The improvement is based on letting RCM determine what and how much maintenance should be done to each unit. In Section 12.3, each item on the list had a specific maintenance activity (e.g., full rebuild, inspect and replace as needed, etc.) *already assigned* to it. Prioritization then picks the best breaker-service activities to fund. But overall “bang-for-the-buck” can be improved beyond this, if the RCM method is used to also decide *what activity* should be done to each breaker, too. Maybe one unit should receive full service, and a second breaker just a quick servicing of only critical components: the money saved there can pay for similar service on another breaker, providing more benefit than if that second breaker were fully serviced.

Table 12.6 Project 8 and Two Alternatives

Project Number	Equipment Type	Work Proposed	# Cust	Reduced Yr. 0	Cust-Min Next	Out +2 yr.	PW Sum Cust-min	Cost (yr. 0)	“Bang/Buck”
8	HV breaker # A0121	Complete Service	11,300	6265	6760	7225	18202	\$22,000	0.83
8b	HV breaker # A0121	Partial Service	11,300	4746	5204	5617	13979	\$9,000	1.55
8c	HV breaker # A0121	Test, Replace	11,300	3797	4231	4612	11340	\$6,750	1.68

Table 12.7 Project 8 and Two Alternatives Listed As Marginal Project Decisions

Project Number	Equipment Type	Work Proposed	# Cust	Reduced Yr. 0	Cust-Min Next	Out +2 yr.	PW Sum Cust-min	Cost (yr. 0)	“Bang/Buck”
-	HV breaker # A0121	Do Nothing	11,300	0	0	0	0	0	-
8c	HV breaker # A0121	Test, Replace	11,300	3797	4231	4612	11340	6750	1.68
8b	HV breaker # A0121	Partial Service	11,300	949	973	1005	2639	2250	1.17
8	HV breaker # A0121	Complete Service	11,300	1519	1557	1608	4222	13000	0.32

Half Measures Sometimes Deliver More than One-Half the Results

Application of “intra-project” alternatives evaluation

Table 12.6 lists Project 8, from Table 12.1, which was the complete service on a HV breaker #AO121, along with two other alternative categories of service that are possible for that unit – *different types of service* on that same breaker. These are option 8b, “partial service,” and option 8c, “test and replace as found needed.” The table shows data equivalent to that shown in Table 12.3 for all three alternatives (i.e., the evaluation done against customer outage minutes). Both Options 8b and 8c provide less reliability improvement than the original project proposed, but they also have a lower cost. They provide more “bang for the buck.” Their figures of merit (right-most column in Table 12.6) indicate that unlike the original project (complete service) both 8b and 8c score high enough to have gained approval in Table 12.3 (the selection limit was 1.49).

Table 12.7 lists these two projects in a slightly different way. First, they are listed in order of increasing cost. Secondly for completeness, “do nothing” is listed as an option. Thus, do nothing (the least expensive) is listed at the top. The next line gives Project 8c, along with its costs and benefits, then 8b is listed, and finally 8.

Most importantly, the costs and benefits of each option have been computed as *marginal differences*. The least cost project (8c) is listed with its costs and reliability improvements as it was in Table 12.6. But project 8b, the next least costly, is listed with marginal costs and benefits. The cost shown is its cost over and above 8c’s. The improvements shown are its improvements minus those that 8c would deliver. Similarly, the original project’s reliability improvements and cost have been converted to the marginal differences that choosing that alternative over 8b would reveal.




Figure 12.5 The three possible projects for Breaker A0121, viewed as marginal decisions. This is *intra-project evaluation* – comparison of project alternatives that all fall within one equipment unit or set of equipment – as opposed to *inter-project comparison* – comparing projects applied to one unit of equipment to those on others.

This *marginal benefit and cost table* has converted the list of possible projects into a *decision chain*. It tells the utility what it gets for what it has to pay when it picks any option over a less expensive option. Table 12.7 lists the three projects and “do nothing” as a serial decision chain. This is illustrated in Figure 12.5. These projects are now decisions among *different categories of service* for the same breaker, a chain of increasingly expensive projects, each providing more reliability improvement than the preceding, but at higher cost.

While viewing a set of possible alternative projects in this manner may seem obtuse, and of no practical value, when dealing with an optimization method as described here, it provides a big improvement. The “chain” of options shown in Figure 12.5 basically reduces a decision with multiple outcomes (which of four alternatives do I choose?) to a series of binary (yes-no) decisions. It simplifies performing the decision-making process with a formal (i.e., rigorous, repeatable, documentable, defensible) procedure that essentially performs a series of binary evaluations (this is better than that). It also illustrates a key point. Doing nothing should *always* be the basis for comparison: all activities should justify themselves against a “zero base.”

Looking at Table 12.7, note that the scores for these three options (last column), the lowest cost option, 8c has the same as it did in Table 12.6 – a score of 1.68 (\$11,340 customer minutes expected reduction over \$6,750). But option 8b has a score of only 1.17, rather than 1.55 as evaluated in Table 12.6. While Project 8b looks good enough when viewed on its own, viewed as a decision to spend \$2,250 more than Option 8c to gain the improvement it gives over 8c, it does not look quite as good – definitely not good enough to gain approval in Table 12.3. That \$2,250 is not nearly as effective in improving reliability as the \$6,750 spent on only buying Option 8c. Similarly, when viewed as a marginal option, Option 8 itself looks very inefficient, with a score of only .32.

Table 12.8 shows the RCM evaluation and ranking of all projects from Table 12.3 with these three projects substituted for Project 8. Note that 8c, with its figure of merit score of 1.68, makes it into the “winners” category. Projects 8b and 8 fall far below the approval limit – their *additional cost* over 8c is definitely not justifiable based on their *additional benefit*. The recommended action, then, is to select to perform only “testing and replacement as found needed” on this breaker.

What got bumped off the list to make room for project 8b? In this case, nothing. Instead, whereas originally the utility could just afford to buy the top ten projects, it can now not quite afford to buy all of the final projects on the approval list (#32, Hillside substation feeders – testing and replacement as found needed). By the time it gets down to that item in Table 12.8, it is \$6,7500 short of having enough to fund this project completely. Prorating that project by assuming that some small portion could be deleted with proportional impact on results, indicates that this change in using alternatives for Project 8 results in a .03% improvement overall.

Using the approval limit’s marginal cost of reliability (\$1.41/customer minute) the additional reliability wrought by this improvement over the results of

Table 12.3 would cost the utility another \$859 to buy if it stuck to the method used in Table 12.3.

This small improvement may hardly seem worth the effort. However, the point was to introduce the concept of extending the RCM evaluation to deciding *how to maintain* each unit of equipment, not just what gets maintained. By evaluating options in this manner – letting the RCM decide what is done to what equipment, the total benefit obtained from the spending is improved. As will be shown later, the resulting savings can be significant.

Table 12.8 Projects Evaluated By RCM Ranking, Including Two Options for Project 8

Project Number	Equipment Type	Work Proposed	# Cust	Reduced Yr. 0	Cust-Min Next	Out +2 yr.	PW Sum Cust-min	Cost (yr. 0)	“Bang/Buck”
18	UG cables, feeder EC06	Replace All	2,200	121440	128436	135502	346789	\$27,000	12.84
4	40 MVA transf. DT03	Test, Follow up	4,612	44275	48703	52681	130779	\$11,000	11.89
2	40 MVA transf. EC01	Recond. Oil/Service	7,800	74880	81432	87461	219012	\$22,000	9.96
1	40 MVA transf. EC03	Rebuild Complete	10,899	222343	235749	249132	636315	\$65,000	9.79
20	UG cables, net DT3B	Replace Complete	245	67032	71207	75341	192144	\$43,000	4.47
25	Poles, CE district	Test, Replace	80,134	575234	625675	672069	1682718	\$610,000	2.76
21	UG cables, net DT3A	Test, Replace	340	22195	25198	27777	67373	\$31,000	2.17
8c	HV breaker # A0121	Test, Replace	11300	3797	4231	4612	11340	\$6,750	1.68
3	40 MVA transf. JT235	Test, Follow up	842	5982	6606	7163	17729	\$11,000	1.61
32	Riverside sub feeders	Test, Replace	28,093	177997	199134	217621	533491	\$335,000	1.59
Approval Limit Based on Budget									
33	Hillside sub feeders	Test, Replace	14,004	108895	121826	133136	326379	\$219,000	1.49
31	Oldtown sub feeders	Test, Replace	22,789	98448	110139	120364	295068	\$210,000	1.41
34	Central sub feeders	Selected Replace	28,345	131634	143981	155199	386928	\$310,000	1.25
35	Farmon substation	Complete Rebuild	43,000	385968	415741	443829	1119637	\$905,000	1.24
8b	HV breaker # A0121	Partial Service	11300	949	973	1005	2639	\$2,250	1.17
27	OH feeder BN05	Inspection, Replace	1,609	5870	6567	7176	17592	\$17,000	1.03
19	UG cables, feeder EC07	Test, Replace	592	5570	6348	7014	16964	\$18,000	0.94
5	HV breaker # EC03	Complete Service	10,899	6043	6521	6969	17556	\$22,000	0.80
10	LV breaker #EC06	Complete Service	2,200	1376	1493	1602	4018	\$6,500	0.62
6	HV breaker # EC01	Complete Service	7,800	4324	4666	4987	12564	\$22,000	0.57
24	OH feeder EC17	Inspect, Replace	1,578	10680	11752	12714	31555	\$58,000	0.54
9	LV breaker #ES06	Complete Service	1,681	1052	1141	1224	3069	\$6,500	0.47
16	LV breaker #DT3N	Complete Service	2,091	811	891	963	2394	\$6,500	0.37
13	LV breaker #NB03	Complete Service	1,191	745	808	867	2175	\$6,500	0.33
8	HV breaker # A0121	Complete Service	11,300	1519	1557	1608	4222	\$13,000	0.32
29	OH feeder CT05	Inspection, Replace	1,389	4534	5072	5543	13588	\$43,000	0.32
26	OH feeder BN04	Inspection, Replace	1,345	4390	4911	5367	13158	\$43,000	0.31
14	LV breaker #CN1A	Complete Service	1,058	662	718	770	1932	\$6,500	0.30
23	69 kV circuits 34 & 37	Test, Replace	42,300	5198	5708	6169	15332	\$70,000	0.22
12	LV breaker #NB02	Complete Service	650	407	441	473	1187	\$6,500	0.18
30	OH feeder JM07	Inspection, Replace	1,555	4777	5344	5840	14317	\$86,000	0.17
11	LV breaker #NB01	Complete Service	580	363	394	422	1059	\$6,500	0.16
28	OH feeder CT14	Inspection, Replace	1,671	2727	3051	3334	8174	\$56,000	0.15
15	LV breaker #CN2A	Complete Service	891	310	342	371	919	\$6,500	0.14
22	OH feeder BN06	Inspect, Replace	950	2508	2856	3153	7632	\$58,000	0.13
7	HV breaker # 221	Complete Service	634	351	379	405	1021	\$22,000	0.05
17	LV breaker #DT2N	Complete Service	112	30	34	37	90	\$6,500	0.01

Table 12.9 Projects 1 and 35 with Multiple Alternatives

Project Number	Equipment Type	Work Proposed	# Cust	Reduced Yr. 0	Cust-Min Next	Min Out +2 yr.	PW Sum Cust-min	Cost (yr. 0)	Bang/Buck
For Project 1									
1d	40 MVA transf. EC03	Test, Replace	10899	104632	115095	124497	309060	\$11,000	28.10
1c	40 MVA transf. EC03	Complete Service	10899	41853	42899	44315	116357	\$10,000	11.64
1b	40 MVA transf. EC03	Recond. Oil/Servc.	10899	26158	26812	27697	72723	\$17,000	4.28
1	40 MVA transf. EC03	Rebuild Complete	10899	49700	50943	52624	138174	\$27,000	5.12
For Project 35									
35d	Farnton substation	Test, Replace	43000	181632	206297	227473	551552	\$195,000	2.83
35c	Farnton substation	Serv. Transf. Only	43000	36326	37235	38463	100993	\$65,000	1.55
35b	Farnton substation	Full Service	43000	81734	83778	86542	227234	\$300,000	0.76
35	Farnton substation	Complete Rebuild	43000	86275	88432	91350	239858	\$345,000	0.70
For Project 9 (example for 9 – 17)									
9b	LV breaker #ES06	Test, Replace	1681	637	716	785	1918	\$1,100	1.74
9	LV breaker #ES06	Complete Service	1681	414	425	439	1152	\$5,400	0.21

Intra- versus inter- project alternatives evaluation

Tables 12.6 – 8 showed one example of letting RCM determine *what should be done* to a unit of equipment. The earlier “optimizations” – those carried out in Section 12.3 – basically decided only among projects where it had already been decided what would be done to each unit. Comparison like that, where decisions about what maintenance will be done has already been made, is called *Inter-Project Alternatives Evaluation*. By contrast, evaluating the possible maintenance options on one unit of equipment to the possible maintenance options on other units is termed *Intra-Project Alternatives Evaluation*. It applies the optimization to a broader range of options and thus provides an improvement.

Comprehensive Example Using Intra-Project Optimization

Table 12.9 lists three alternatives to Project 1 and four for Project 35 – the original and two and three alternatives respectively, for each. Also shown is one example of splitting the nine LV breaker projects (numbers 9 through 17 in Table 12.1) into two alternatives. In the original list (Table 12.1, optimized in Table 12.3), nine LV breakers were listed for “full service” at \$6,500 each. Here, an option – inspect and replace (as found needed) - is an alternative at only \$1,100 in cost but with less effective reliability improvement. All of these options are shown figured as marginal costs and marginal improvements and grouped together as sets of ordered, marginal decisions.

Table 12.10 Projects Evaluated By RCM Ranking Using “Intra-Project” Alternative Evaluation for Twelve Projects (#s 1, 8, 9-17, and 35).

Project Number	Equipment Type	Work Proposed	# Cust	Reduced Yr. 0	Cust-Min Next	Out +2 yr.	PW Sum Cust-min	Cost (yr. 0)	Bang/Buck
1d	40 MVA transf. EC03	Test, Replace	10899	104632	115095	124497	309060	\$11,000	28.10
18	UG cables, feeder EC06	Replace All	2200	121440	128436	135502	346789	\$27,000	12.84
4	40 MVA transf. DT03	Test, Follow up	4612	44275	48703	52681	130779	\$11,000	11.89
1c	40 MVA transf. EC03	Complete Service	10899	41853	42899	44315	116357	\$10,000	11.64
2	40 MVA transf. EC01	Recond. Oil/Service.	7800	74880	81432	87461	219012	\$22,000	9.96
1	40 MVA transf. EC03	Rebuild complete	10899	49700	50943	52624	138174	\$27,000	5.12
20	UG cables, net DT3B	Replace complete	245	67032	71207	75341	192144	\$43,000	4.47
1b	40 MVA transf. EC03	Recond. Oil/Service.	10899	26158	26812	27697	72723	\$17,000	4.28
35d	Farmon substation	Test, Replace	43000	181632	206297	227473	551552	\$195,000	2.83
25	Poles, CE district	Test, Replace	80134	575234	625675	672069	1682718	\$610,000	2.76
10b	LV breaker #EC06	Test, Replace	2200	834	938	1027	2510	\$1,100	2.28
21	UG cables, net DT3A	Test, Replace	340	22195	25198	27777	67373	\$31,000	2.17
9b	LV breaker #ES06	Test, Replace	1681	637	716	785	1918	\$1,100	1.74
8c	HV breaker # A0121	Test, Replace	11300	3797	4231	4612	11340	\$6,750	1.68
3	40 MVA transf. JT235	Test, Follow up	842	5982	6606	7163	17729	\$11,000	1.61
32	Riverside sub feeders	Test, Replace	28093	177997	199134	217621	533491	\$335,000	1.59
35c	Farmon substation	Serv. Transf. Only	43000	36326	37235	38463	100993	\$65,000	1.55
----- Approval Limit Based on Budget -----									
33	Hillside sub feeders	Test, Replace	14004	108895	121826	133136	326379	\$219,000	1.49
31	Oldtown sub feeders	Test, Replace	22789	98448	110139	120364	295068	\$210,000	1.41
16b	LV breaker #DT3N	Test, Replace	2091	492	564	625	1505	\$1,100	1.37
34	Central subst. feeders	Selected Replace	28345	131634	143981	155199	386928	\$310,000	1.25
13b	LV breaker #NB03	Test, Replace	1191	452	508	556	1359	\$1,100	1.24
8b	HV breaker # A0121	Partial Service	11300	949	973	1005	2639	\$2,250	1.17
14b	LV breaker #CN1A	Test, Replace	1058	401	451	494	1207	\$1,100	1.10
27	OH feeder BN05	Inspection, Replace	1609	5870	6567	7176	17592	\$17,000	1.03
19	UG cables, feeder EC07	Test, Replace	592	5570	6348	7014	16964	\$18,000	0.94
5	HV breaker # EC03	Complete Service	10899	6043	6521	6969	17556	\$22,000	0.80
35d	Farmon substation	Full Service	43000	81734	83778	86542	227234	\$300,000	0.76
13	LV breaker #NB03	Complete Service	1191	294	301	311	816	\$5,400	0.74
35	Farmon substation	Complete Rebuild	43000	86275	88432	91350	239858	\$345,000	0.70
12b	LV breaker #NB02	Test, Replace	650	246	277	304	742	\$1,100	0.67
11b	LV breaker #NB01	Test, Replace	580	220	247	271	662	\$1,100	0.60
6	HV breaker # EC01	Complete Service	7800	4324	4666	4987	12564	\$22,000	0.57
24	OH feeder EC17	Inspect, Replace	1578	10680	11752	12714	31555	\$58,000	0.54
8	HV breaker # A0121	Complete Service	11300	1519	1557	1608	4222	\$13,000	0.32
29	OH feeder CT05	Inspection, Replace	1389	4534	5072	5543	13588	\$43,000	0.32
26	OH feeder BN04	Inspection, Replace	1345	4390	4911	5367	13158	\$43,000	0.31
10	LV breaker #EC06	Complete Service	2200	542	556	574	1508	\$5,400	0.28
23	69 kV circuits 34 & 37	Test, Replace	42300	5198	5708	6169	15332	\$70,000	0.22
9	LV breaker #ES06	Complete Service	1681	414	425	439	1152	\$5,400	0.21
30	OH feeder JM07	Inspection, Replace	1555	4777	5344	5840	14317	\$86,000	0.17
16	LV breaker #DT3N	Complete Service	2091	320	328	338	889	\$5,400	0.16
28	OH feeder CT14	Inspection, Replace	1671	2727	3051	3334	8174	\$56,000	0.15
14	LV breaker #CN1A	Complete Service	1058	261	267	276	725	\$5,400	0.13
22	OH feeder BN06	Inspect, Replace	950	2508	2856	3153	7632	\$58,000	0.13
15b	LV breaker #CN2A	Test, Replace	891	188	217	241	579	\$1,100	0.11
12	LV breaker #NB02	Complete Service	650	160	164	170	445	\$5,400	0.08
11	LV breaker #NB01	Complete Service	580	143	147	151	397	\$5,400	0.07
15	LV breaker #CN2A	Complete Service	891	122	125	130	340	\$5,400	0.06
17b	LV breaker #DT2N	Test, Replace	112	18	21	24	57	\$1,100	0.05
7	HV breaker # 221	Complete Service	634	351	379	405	1021	\$22,000	0.05
17	LV breaker #DT2N	Complete Service	112	12	12	13	33	\$5,400	0.01

A different set of decisions

Table 12.10 shows the result of an RCM ranking based on customer-minutes that includes these alternatives added to the projects and alternatives used in Table 12.8 (including Project 8's options). Among the first eight projects at the top of the list are all four of the alternatives to Project 1. Although the RCM method had a list of partial measures it could consider for this unit, it chose to "fully upgrade" Project number 1 through selection of its alternatives 1d, then 1c, then 1b, to 1. Thus, for that transformer it picked, in a serial, incremental, four-decision fashion, the full service as was originally selected in Table 12.3.

But on other projects the prioritization often selected only an interim level of permitted maintenance. For Project 35, alternatives 35d and 35c were selected. Alternatives 35b and 35 did not make the cut. Similarly, one low voltage breaker (9b, breaker ES06) makes the cut with testing and replacement as needed. Cumulatively, these results are considerably different than those obtained in Figure 12.4, even though the same 35 equipment units are being considered, and even though the same reliability index (customer minutes) is being used as the goal. Important differences in the results are given below.

A 6.4% improvement in results has been obtained

For the same budget, an increase of 264,326 customer minutes of interruption expecting to be avoided has been obtained. Gaining this in the original optimization (Table 12.3 would have cost the utility an extra \$187,500 (264,326 customer-minutes at 1.41 customer minutes per dollar).

Approval target is raised

Intra-project prioritization raises the bar for approval, because it gets more "bang" per dollar from some of the projects now being approved. Here, the approval limit increases by nearly 10%, from 1.41 to 1.55. Better projects are being implemented. This increase in the target value (10%) is greater than the increase in overall results (6.4%) because of the scree-shape of the top part of the project mix – the very best projects approved did not change.

More units of equipment receive some improvement.

The number of units seeing some reliability improvement increases by over 30%. Maintenance effort is spread around the system more.

Smaller projects are being done

Since this method tends to "spread around" the money on more projects, that means the limited budget is splintered into more, but smaller, projects. On average, they are of greater average value compared to their cost.

The optimization selected “partial measures”

It is significant that in 40% of the equipment units selected for some maintenance, the RCM decided to maintain it by selecting only an interim measure. More significantly, in all but one case where it had the choice, it chose an interim measure. *“Half measures” often provide more than half the benefits.*

Application to all units would do even better

Table 12.10 applied prioritization using multiple alternatives for only thirteen of thirty-five units of equipment involved. If a multiple alternatives approach is applied to all projects on the list, the improvement in figure of merit obtained grows to about 10% or slightly more depending on details of the reliability index used and how thoroughly project alternatives are represented. Such an example is beyond the scope of what can be given here (it results in a list with over 110 entries). However, the concept of its execution is exactly as used above.

12.6 PRACTICAL ASPECTS FOR IMPLEMENTATION

The foregoing examples, while comprehensive, simplified several of the details needed when RCM is applied in actual practice in order to streamline the examples so that the major points could be made in a succinct manner. When applied to actual utility cases, the authors make the following recommendations for reliability-centered maintenance prioritization.

Cost and Benefits Should Be Treated as Expectations

As discussed in Chapters 6, 11, and elsewhere throughout this book, equipment failure is a random process that cannot be predicted with accuracy on a single-unit or small-lot basis except in a few very special cases. Similarly, both the cost of doing some types of maintenance and the benefits that maintenance provides, is probabilistic. For example, the category of service “Inspect and Make Repairs as Found Necessary” will vary in both cost and benefits from unit to unit depending on what is found when inspection is carried out. This cannot be predicted accurately in advance for any one unit of equipment (a specific pole, a particular power transformer) but it can be predicted en masse over a large group of units.

Reliability-Centered Maintenance should be applied in this context: the failure rates being dealt with, the benefits expected from maintenance, and the costs of that service are all expectations, accurate and dependably forecastable on a large-scale basis but determinable on an individual basis only as probabilistic expectations.

Use Only Historically-Based Cost and Improvement Factors

The failure rates for equipment, as a function of condition as used for “before service” expectations of reliability, the improvement expected from various levels of service applied to equipment of various categories of age/condition, and the cost of various types of service should all be developed from historical data. For example, the entries in a table like Table 12.2 should be based, to the extent

possible, on historical experience specific to the utility, or data from a large set of samples published in a reputable reviewed technical journal. If that is not possible, the data should be developed by consultation with experts who have experience in such work.⁹

Represent Improvements as Having Only Short-Term Impacts on Reliability

Even the most comprehensive service on a transformer, breaker, underground or overhead line will result in improvements that only last for a short time. Left unattended after the preventive maintenance, the device will return to its original “aged” failure rate condition in only a short period of time. The improvements in reliability used as the RCM benefit should reflect this in a realistic sense. Overestimating the length of time improvements make a positive impact will overemphasize the value of maintenance, leading to too much spending and a shortfall of results versus expectations.

Figure 12.6 illustrates the difference between the impact that service might make, and replacement. Here, a 30-year old transformer can either be replaced or re-built. Replacement results in “like new” condition and it will take 30 years for expected failure rate to return to the level of just prior to replacement.

Performing any type of inspection, testing, and preventive maintenance, however, has a different impact. First, it does not result in “like new” condition and failure rate expectation – the improvement is somewhat less. Secondly, the improvement does not last for 30 years, as replacement does, but for a much shorter time, as shown. Generally, noticeable improvements from servicing only older units last for only 3 – 7 years.

Use Category-Based Equipment Condition and Maintenance Types

Expected RCM costs, as well as the benefits (reliability improvement gained) from inspection, testing, and various grades of preventive service and the cost and improvement numbers assigned to each category represent expectations for equipment-service combinations in that category.

All equipment being considered for maintenance, and all maintenance efforts that are being considered to be performed, are classified by these category definitions and the numbers assigned to each category are used in estimating the expected cost-improvement combination for that equipment, in *every* case. “Adjustments” or tailoring of data based on judgment is *not* permitted. As an example, one of several possible ways to implement this would be to:

1. Rate all equipment by condition based on “equivalent age” and lump all equipment into categories by equivalent age group (see column headings in Table 12.11, below).

⁹ Many consultants can develop sets of good-looking numbers. Far fewer seem able to develop numbers that produce dependable results.




Figure 12.6 Top – replacement of equipment after 30 years returns a facility to “like new” condition. It will take 30 years for the equipment to age to the same condition as it was before replacement (assuming all other factors remain the same). Middle – anything less than replacement provides a somewhat lesser improvement in failure rate. It is a mistake to model that improvement as lasting for a long period of time (in this case 30 years). Bottom: a more realistic representation is to model the failure rate as escalating at the same rate as it historically does from the lower rate after service, refurbishment, should be classified into categories for RCM planning and evaluation purposes, as shown.

Table 12.11 Example of Improvement Categorization by “Generic” Categories – Improvement Factor I in Equation 12.3, for Power Transformers

Service Category	Condition Category – Equivalent Years in Service					
	0 – 5	10 – 20	20 – 30	30 – 35	35 – 40	40 – 45
Age Category →	1	2	3	4	5	6
1 No service	0.00	0.00	0.00	0.00	0.00	0.00
2 Inspect/ test but do not follow up	0.02	0.03	0.05	0.10	0.11	.013
3 Inspect and replace as found needed	0.05	0.15	0.25	0.33	0.42	0.50
4 Inspect & test and replace as needed	0.15	0.24	0.33	0.40	0.48	0.55
5 Partial service - 1	0.22	0.30	0.38	0.45	0.52	0.59
6 Partial service - 2	0.29	0.37	0.44	0.50	0.57	0.63
7 Full service	0.52	0.57	0.62	0.66	0.71	0.75
8 "Field rebuild"	0.64	0.68	0.72	0.75	0.78	0.81
9 Remanufacture	0.70	0.75	0.80	0.85	0.85	0.85
10 Replace	1.00	1.00	1.00	1.00	1.00	1.00

2. Model expected failure rate as a function of equivalent age or condition, rather than actual age.
3. Group maintenance operations into several categories.
4. Develop a table of expected equivalent age versus specific category of maintenance, from a table of expected improvement factors (Table 12.11). In this example, the factor R (equation 8) was determined based on such categorization, using the formula

$$F_s = (F_p - F_c) \times M(a,s) \quad (12.4)$$

Where:

F_s = expected annual failure rate after service s is performed

F_p = expected annual failure rate before service s is performed

And $F_s \leq F_n$, the expected failure rate of new equipment

$M(a,s)$ is the entry in Table 12.11, the maintenance improvement factor for age category a and service category s .

What to Do When Marginal Figures of Merit for Project Alternatives are not Strictly Increasing

In some cases, when the hierarchical alternatives chain for a particular project is developed (e.g. Figure 12.5), it will not have strictly decreasing marginal scores as one moves “down” the serial chain of decisions regarding how much maintenance to buy. Everything works well when costs are strictly increasing as one moves down the decision chain. Optimization may encounter problems when that is not the case. Adjustment of the list is required.

This problematic condition was the case with project 1 in the final version of the examples used in Section 12.3. (Tables 12.9 and 12.10). The final option for transformer EC03 (moving to 1 from 1b) had a figure-of-merit of 5.28, whereas the alternative in the chain that was ahead of it in the “chain” (moving to 1b from 1c) had a lesser figure-of-merit, of only 4.28. As a result, in Table 12.9 Option 1 was picked *before* Option 1c. In Table 12.10’s optimization, this did not cause a problem, because Alternative 1c was eventually selected (both 4.28 and 5.28 were nicely above the approval limit of 1.55). But, if the approval limit had fallen between 4.28 and 5.28, it would have meant picking 1d, 1c, and 1, without picking 1b – impossibility – in the real world.

This is a complication that must be dealt with in the real world. There are two ways of adjusting the alternatives in such tables to avoid such computations in the tables, both requiring rather involved “programming” of the spreadsheet or database manipulation system being used. The first involves making the costs and improvements listed for the alternatives “active” functions of those which have already been selected (i.e., their values change depending on what other options have been selected). This necessitates an interactive spreadsheet application.

The second way is to use multiple entries representing both the entries in the chain shown in Figure 12.6, and combinations (e.g., moving from 1c directly to 1). Optimization requires a “post processor” to resolve ambiguities caused by two overlapping alternatives (i.e., both “move to 1b from 1c” and “move to 1 directly from 1c”). Either method works well and both give identical results when implemented correctly.

Real-World Application Requires a Computerized Approach

Implementation of rigorous extended RCM using intra-project alternatives evaluation as shown here does result in a large list of proposed alternatives to consider. The list of 35 projects in the example used in section 12.3 grows to over 110 entries if alternatives are used for all 35 units of equipment being considered. Additional entries, or complications in database manipulation due to having to deal with marginal costs that are not strictly increasing in all cases, makes for a further increase in complexity. Typically, simple “sorting” prioritization methods will not work here: rigorous optimization, using a commercial optimization engine, is strongly recommended.

As a result, for a utility with a large metropolitan service area the list of possible alternatives for all significant distribution equipment in its system will easily run to hundreds if not thousands of possible entries. However, this is exactly what computers are made to do: manage and analyze involved lists. RCM as described here is implementable with nothing more than an electronic spreadsheet and a set of well-designed and implemented organization procedures. The categorization of equipment by type, condition by equivalent age, and maintenance by service category makes it possible to develop procedures that generate the entire alternatives database, as needed for RCM evaluation, from the raw equipment/conditions lists.

Multi-District or Multi-Department Application

One “trick” that helps a great deal in implementation of RCM at a large distribution utility is that the marginal cost of reliability – the “cutoff limit” in the optimization – can be used as both a general target to reduce planning and evaluation work. It can also be used as a tool to apply the RCM concept when O&M decisions are broken into areas or districts due to the utility’s management structure. Figure 12.7 illustrates this concept. Three different regions of a utility each manage their maintenance budgets autonomously, using the same RCM procedure based on the same company-wide reliability index, economic factors and categorizations of equipment condition and maintenance service types.

All that senior management has to do to keep the overall use of maintenance resources “optimal” across all regions is to see that the marginal cost (approval limit) in each region is adjusted until all are the same. Adjusting the budget allocation for each region or department until the approval limit is the same for all regions and departments does this. Figure 12.8 illustrates this function. The utility has three regions that operate on a very autonomous basis. Each applies RCM, producing the prioritized lists that were shown in Figure 12.7. Based on preliminary estimated budgets, each has already picked an approval limit beyond which it cannot afford to spend (e.g. 3.95 for Region 1).

Note that in the top of Figure 12.7, Region 3 has a limit that is 2.98 and Region 1 is at 3.95. This means that the utility can pick up an improvement of .97, without spending anything more, by deleting a dollar from Region 3’s budget where it buys only 2.98 units of improvement, and transferring it to Region 1, where it will buy 3.95 units of improvement. When it makes such a change, the approval limit in Region 3 will rise slightly, and that in Region 1 will gradually drop.¹⁰ It can continue making such transfers, and the approval limits in the regions will rise and drop respectively, until they are equal. At that point spending is adjusted between the two regions: it is “optimal” in the sense that the utility is getting the most “bang for the buck” possible.

The company-wide budget limit determines where on the company-wide curve the utility must place the RCM approval limit. The individual curves then indicate what budget each region should get to fit that limit. Executive management does not even have to re-allocate budget – it simply approves only, and all projects above or equal to the limit – and the budget and its optimal allocation will work out correctly.

Rather than make a lengthy series of incremental adjustments to such budgets, the utility can determine the overall acceptance limit for all regions. As shown in Figure 12.8, the prioritized lists from all regions are used to produce an aggregate score-versus-budget limit curve. The utility’s total budget (\$80,000,000 in this

¹⁰ The approval limits are falling and rising as the budgets in Regions 1 and 3 are respectively rising and falling, since they are respectively moving down and out and up and in on curves similar to that shown in Figure 12.4.




Figure 12.7 Three regions in a utility each autonomously manage their own maintenance budget using basically the same RCM procedure, locally administered by each regional office. Executive management achieves overall optimization by controlling budget allocation using the method shown in Figure 12.8, by using the individual regional projects lists (top) to determine a uniform company-wide approval score (bottom) which it applies to all three regions.




Figure 12.8 Executive management can easily control the overall maintenance budget to see both that maintenance budget is not exceeded and that RCM is optimized over the entire utility, not just within regions, with the procedure shown above. Regions submit their complete project alternatives list (top), each of which is converted to the figure-of-merit versus budget limit plot (similar to Figure 12.4) for each, and the curves summed (alternatively the lists can be added together, sorted, and plotted).

case) is applied to this curve to determine the required approval limit (3.28 in this case). Executive management then uses the score versus budget curve for each region to determine what regional budget corresponds to an approval limit of 3.28. The budgets determined in this manner for the three regions will sum to \$80,000,000. Alternatively, executive management does not have to even re-adjust budgets – it can simply approve all projects with figures of merit above 3.28, but not below, regardless of regional budget limits. Company-wide, this will result in \$80,000,000 in spending, but no more.

12.7 EXTENDING RELIABILITY-CENTERED PRIORITIZATION TO OTHER OPERATIONS PROJECTS

This section will briefly look at how the RCM concept can be extended to include activities and expenses that are not, strictly speaking, equipment maintenance. It will also demonstrate several nuances with respect to setting up comparisons of operations projects that differ in time or scope from one another.

Example of a Much Longer Period of Benefit than Normal Maintenance: Computerized Trouble Call Analysis

Often a utility has options with respect to spending on reliability improvement that are not classified strictly as maintenance. For example, it can implement a computerized Trouble Call Management (TMS) system to facilitate the acceptance, processing, analysis of trouble reports from its customers, and optimize the dispatching of resources and equipment for restoration and repair. Such systems do little to reduce the occurrence of equipment outages and the interruptions they cause (i.e. SAIFI). However, in the authors' experience such systems directly reduce the duration of customer interruptions by about 15 minutes per year, and accommodate other changes (e.g., they permit more complicated and dynamic switching programs to be implemented for restoration) that make a further 5 minutes improvement. Computerized TMS systems are expensive – the advanced systems that are truly effective can cost a large utility as much as \$16 million including hardware, software, training, system maintenance, and labor for operation. However, they provide long term benefits. If updated, maintained and used well, they do not lose their effectiveness at reducing reliability problems within only a few years, as do the preventive maintenance activities discussed earlier.

As an example, suppose Metropolitan Light and Power is considering a system improvement plan that is expected to provide an average 20 minutes reduction in outage time annually. Implementation will take 18 months, at a total cost of \$11,300,000 for hardware, software, facilities, auxiliary equipment, labor, and training, which will be treated here as \$627,778 per month average cost over the 18 month period. During that time the system provides no improvement in reliability. Subsequently, the system has a fixed operating cost of \$1,700,000 per year (facilities rent and utilities, etc.), as well as a variable cost that will average

Table 12.12 Analysis of Long-Term Expenses and Reliability Reduction for a Trouble Call Management System

Year	PWF	Customers	Expenses	Cust.-Min.	PW Cost	PW Impr.
0	1.00	1,137,000	\$7,533,333		\$7,533,333	-
1	0.90	1,150,644	\$6,260,611	23,012,880	\$5,634,550	20,711,592
2	0.81	1,164,452	\$3,074,053	23,289,035	\$2,489,983	18,864,118
3	0.73	1,178,425	\$3,090,542	23,568,503	\$2,253,005	17,181,439
4	0.66	1,192,566	\$3,107,228	23,851,325	\$2,038,652	15,648,854
5	0.59	1,206,877	\$3,124,115	24,137,541	\$1,844,759	14,252,977
6	0.53	1,221,360	\$3,141,204	24,427,191	\$1,669,365	12,981,611
7	0.48	1,236,016	\$3,158,499	24,720,318	\$1,510,700	11,823,651
8	0.43	1,250,848	\$3,176,001	25,016,962	\$1,367,164	10,768,982
9	0.39	1,265,858	\$3,193,713	25,317,165	\$1,237,310	9,808,388
10	0.35	1,281,049	\$3,211,637	25,620,971	\$1,119,829	8,933,480
11	0.31	1,293,859	\$3,226,754	25,877,181	\$1,012,589	8,120,534
12	0.28	1,306,798	\$3,242,021	26,135,953	\$915,643	7,381,565
13	0.25	1,319,866	\$3,257,441	26,397,312	\$827,998	6,709,843
14	0.23	1,333,064	\$3,273,016	26,661,285	\$748,761	6,099,247
15	0.21	1,346,395	\$3,288,746	26,927,898	\$677,124	5,544,215
16	0.19	1,359,859	\$3,304,633	27,197,177	\$612,355	5,039,692
17	0.17	1,373,457	\$3,320,680	27,469,149	\$553,796	4,581,080
18	0.15	1,387,192	\$3,336,887	27,743,840	\$500,849	4,164,202
19	0.14	1,401,064	\$3,353,255	28,021,279	\$452,975	3,785,259
20	0.12	1,415,075	\$3,369,788	28,301,491	\$409,688	3,440,801
21	0.11	1,427,810	\$3,384,816	28,556,205	\$370,363	3,124,591
22	0.10	1,440,661	\$3,399,979	28,813,211	\$334,820	2,837,441
23	0.09	1,453,626	\$3,415,279	29,072,530	\$302,694	2,576,680
24	0.08	1,466,709	\$3,430,717	29,334,182	\$273,656	2,339,883
25	0.07	1,479,910	\$3,446,293	29,598,190	\$247,409	2,124,848
26	0.06	1,493,229	\$3,462,010	29,864,574	\$223,683	1,929,575
27	0.06	1,506,668	\$3,477,868	30,133,355	\$202,237	1,752,247
28	0.05	1,520,228	\$3,493,869	30,404,555	\$182,851	1,591,215
29	0.05	1,533,910	\$3,510,014	30,678,196	\$165,326	1,444,983
30	0.04	1547715	\$3,526,304	30,954,300	\$149,484	1,312,189
					\$37,862,950	216,875,180
Figure-of -merit score = 216,875,180/37,862,950 = 5.73						

an estimated \$1.38/customer/year.^{11, 12}

Table 12.12 lists these costs over a 31 year period, beginning with 12 months in year zero during which the project implementation begins. In year 1 the system is implemented at half year (thus there are half the variable costs and half the expected results in that year. There are no fixed costs in year 0. Customer count is expected to grow as shown (thus variable cost grow each year, too). Both future costs and future benefits are discounted using a present worth factor. As shown, the computed figure of merit is 5.73, enough to justify this TMS system against expenses in Table 12.8 or the example in Figures 12.6 and 12.7.

Figure 12.9 shows the “figure of merit so far” over time for the TMS system. When evaluating a long-term project with significant up-front expenses, it is sometimes useful to look at how good it looks in the short run, not just the long term. Since there are no improvements in year 0, and only one-half a year’s improvements in year 1, the TMS system looks rather poor on a very short-term basis. Note however that Figure 12.9 shows that its figure of merit passes 3.28 (the acceptance limit in Figure 12.8) in only 4 years. Therefore, this looks good enough to gain approval within the example evaluations used in Figure 12.8, anytime it is evaluated over a four-year or longer period.

Example of evaluating very short-term impacts and comparing options for different impact periods: tree-trimming

Tree trimming is clearly a form of “maintenance” and yet not equipment maintenance in the strictest sense. At many utilities it is administered by a separate department or from a separate budget than equipment maintenance. Since it has an identifiable cost and a direct impact on reliability there is no reason it cannot be included in RCM evaluation, with evaluation used to determine if, how much, and where tree-trimming should be used and how much budget should be allocated to it versus other maintenance issues.¹³ Tree trimming generally makes only a relatively short-term improvement in distribution reliability. Some utilities manage vegetation control on a three-year cycle, but two years is considered the most effective by some experts.¹⁴ Regardless, it often has a very good “bang for the buck.”

Table 12.13 shows the evaluation of six vegetation-trimming alternatives for

¹¹ This analysis assumes that none of these costs are to be capitalized. If that were the case, adjustments to the cost basis for this project might have to be done following the guidelines shown in Chapter 15 with respect to comparing capital to operating expenses.

¹² There is no doubt that future improvement in hardware and software design will lower cost and improve the performance of such systems. However, the analysis is done over the long run using the costs and performance of the original system. This is what the utility is committing to, today. Future changes will justify themselves on their own merits.

¹³ Very often tree-trimming has a very high effectiveness per dollar and moves to near the top of the RCM list.

¹⁴ See, for example, O. C. Seevers, *Management of Transmission and Distribution Systems*, Fairmont Press, Liburn GA, 1995, p. 170.




Figure 12.9 Figure of merit for the computerized trouble call system as a function of the time period over which the analysis of benefits and cost is carried out. The system exceeds the 3.28 acceptance limit in Figure 12.8's example in year four, meaning that if that utility is "looking ahead" even just four years, the TMS system's benefit justifies taking money from the three regional maintenance budgets to pay for the system.

the eastern operating district of Region 2 of the utility example used in Figures 12.6 and 12.7. The six alternatives involve two sets of three options, the sets being for programs that will provide respectively two- or three-years of benefit. Within each of these sets, there are three alternatives. These include: 1) trimming back vegetation from all primary and secondary overhead lines; 2) restricting trimming to just primary circuits; or 3) further restricting it to just along primary trunks and major branches (i.e., not along laterals).

As mentioned above, any of these three programs can be carried out with a three-year, or two-year duration program. Two- and three-year programs vary in how far back the trimming is done, and how much off-right-of-way and off-easement trimming is attempted (with customer permission). They differ in cost (the more comprehensive trimming in the three-year program costs more) in a very subtle way, in their impact on public relations.¹⁵ Table 12.13 compares

¹⁵ The customer impacts differ in a way that makes it difficult to identify one program as better in this regard than the other, but on the other hand the difference is significant and should be considered. The three-year program garners a higher number of complaints about aggressive trimming, but it also involves a good deal more proactive customer contact (requests to trim trees on private property which might fall during storms) and prior notification of trimming that, while legal, might concern homeowners. Surveys indicate this leads to a mostly positive public perception of the utility's efforts. (Table 12.13 is based on the experience of a utility in the southeast United States that decided on a three year program due to both its superior economics and because of a simple policy with regard to public relations: "all customer contact, if well-managed, is positive").

Table 12.13 Tree-Trimming Projects Evaluated for Comparison to RCM

Project Number And Area	Type of Work	Years Dur.	Reduced Cust. Yr. 0	Min Out Next	PW Sum +2 yr. Cust-min	Cost (yr. 0)	Bang/ Buck
Tr 1a-2 E. distr. fdrs	Trim all primary/secondary	2	3500987	2345661	5612082	\$1,002,158	5.6
Tr 1b-2 E. distr. fdrs	Trim only primary	2	2625740	1759246	4209062	\$454,542	9.3
Tr 1c-2 E. distr. fdrs	Trim only primary trunks	2	2310651	1548136	3703974	\$314,429	11.8
Tr 1a-3 E. distr. fdrs	Trim all primary/secondary	3	3579759	3161829	2455067	\$1,517,500	5.6
Tr 1b-3 E. distr. fdrs	Trim only primary	3	3209639	2786258	2118362	\$825,000	9.2
Tr1c-3 E. distr. fdrs	Trim only primary trunks	3	3000000	2550000	1980000	\$600,000	11.5

Table 12.14 Tree-Trimming Projects Ranked As Marginal Decision Chain

Project Number	Type of Work	Years Duration	PW Sum of Improvements	Cost (yr. 0)	Bang/ Buck
Tr 1c-2	Trim only primary trunks	2	3,703,974	\$314,429	11.8
Tr 1b-2	Trim only primary	2	505,087	\$140,113	3.6
Tr 1c-3	Trim only primary trunks	3	2,689,738	\$145,458	18.5
Tr 1b-3	Trim only primary	3	639,175	\$225,000	2.8
Tr 1a-2	Trim primary and secondary	2	-1925,893	\$177,158	-10.9
Tr 1a-3	Trim primary and secondary	3	2,941,967	\$515,342	5.7

Table 12. 15 Tree-Trimming Projects Ranked As Marginal Decision Chain with “Dumb” Alternatives Deleted

Project Number	Type of Work	Years Duration	PW Sum of Improvements	Cost (yr. 0)	“ Bang/ Buck”
Tr 1c -2	Trim only primary trunks	2	3,703,974	\$314,429	11.8
~~Tr 1b -2~~	~~Trim only primary~~	~~2~~			
Tr 1c-3	Trim only primary trunks	3	3,194,826	\$285,571	11.2
Tr 1b-3	Trim only primary	3	639,175.	\$225,000	2.8
~~Tr 1a-2~~	~~Trim primary and secondary~~	~~2~~			
Tr 1a-3	Trim primary and secondary	3	1,016,074	\$692,500	1.5

these six alternatives in absolute terms (i.e., it is *not* set up as a set of marginal decisions). On its own merits, any of the six programs exceeds the acceptance limit in Figure 12.8. All look like “good projects.”

Table 12.14 shows these programs ranked in order of increasing cost, with the impacts and benefits computed as marginal values. It is now a serial list of marginal decisions for use in RCM optimization, as described in Section 12.3. The arrows at the right of the table show the serial decision chain, moving from one item to the next in the list, in this case as a series of single jumps down the list. Note that the second and fifth alternatives in Table 12.14 make little practical sense. The second on the list (Tr1b-2) has a lower marginal figure of merit (3.6) than the step after it (18.5), indicating that perhaps it should be deleted from the

serial list and the decision chain aimed to go around it to the next decision on the list. That would mean that the only decision to make is whether or not to take both it and the decision after it at the same time.

Similarly, item five on the list (Tr1a-2) has a negative marginal figure of merit and should not be selected. Although alternative Tr1a-2 looks good in absolute terms (Table 12.15 shows a figure of merit of 5.6), spending an additional \$177,178 on it rather than the next less costly alternative (Tr1b-3) makes no sense because the additional money buys a reduction in improvement.

Table 12.15 shows the marginal decision table adjusted with items 2 and 5 in the original list deleted, and the marginal costs for items 3 and 6 changed to reflect the benefits and costs of a double jump in the table (the arrows indicate the decision sequence). There are now four items in the serial list with strictly decreasing marginal benefit/cost ratios as one moves down the list.

The optimal decision seems to be the second alternative on the list, Tr1c-3 – trimming only primary trunks, but for a three-year benefit period. This option has an absolute figure of merit of 12.5 (Table 12.13) and a marginal figure of 12.2 in this table. The next alternative in Table 12.15 has a marginal ratio of 2.8, below the utility's service territory-wide acceptance limit of 3.38.

Results can be Valid even if Alternatives Cover Different Time Periods

The two foregoing examples bring up an important point. The evaluation method presented here is consistent with respect to different impact periods. For example, one aspect of the difference between the two- and three-year tree-trimming programs that was not addressed above was that in actual implementation, the two-year programs would be applied every two years, and the three-year programs every three years. That does not need to be considered. In the evaluation method both cost and value are discounted over time, so the computed benefit/cost ratios are consistent to shifts in time. If one takes a two-year program and shifts it two years into the future (two years after doing it the first time we will repeat it), its figure of merit be the same. Similarly, shifting a three-year program three years ahead does not change its figure of merit.

As a result, it is not necessary to compare alternative programs of different equipment programs on the basis of the same time period. Readers interested in confirming this can use the data in Table 12.13–12.15 to build tables representing a six-year period (three two-year periods versus two three-year periods) or a thirty year program (fifteen of the former, ten of the latter). The figures of merit remain the same.

12.8 OPTIMIZATION

Using Prioritization for KPIs Other Than Reliability

The prioritization method described in Sections 12.3 – 12.6 can be applied to boost the performance of *any* numerically representable measure, not just reliability. If the projects in Tables such as 12.8 were evaluated for extension of

expected service lifetime instead of reliability, and then evaluated and ranked by what was otherwise the very same method, the resulting prioritized list would give the maximum lifetime extension possible as a function of spending. Similarly, any other numerical measure could be “optimized” by this method.

Optimization and its verb, adverb, and adjective forms are among of the most misused and abused technical terms in the power industry (and a lot of other industries). Strictly speaking, the word involves finding from among many options that which minimizes or maximizes the particular performance index. Also implicit in its application is an important concept: that there is a range or large set of possible solutions. Thus, there is no optimum answer to “What is two plus two?” There is a unique solution – just one. By contrast there are dozens of ways a utility could allocate \$14 million to service, maintenance, and refurbishment of its medium voltage breakers. Every one of those ways is equivalent in terms of its basic qualitative characteristics: it spends \$14 million and it spends it all on breakers. But among them, one plan might be better than another. That is the optimum plan. Thus, “optimization” involves selecting the best from among a range or set of options or alternative possibilities.

In addition, it is necessary to have a formal performance index – a definition of exactly what one considers best. Among the options might be one plan that minimizes the expected number of misoperations, essentially maximizing the reliability of the breakers themselves as a set. Another might maximize customer service reliability, something similar but subtly different. Yet another might maximize the expected service lifetime of the equipment. Several important practical points emerge from this fact:

- One has to know exactly what one wants and be able to define it with regard to the problem. Minimize the number of breaker misoperations. Minimize SAIDI. Maximize expected number of years of service left in the equipment. Best has to have a specific definition.
- Changing the definition of best will change which option is defined as best. The three plans discussed above – minimize misoperations, minimize SAIDI, and maximize service lifetime – are each optimum as regards their particular goal.
- Optimization implies trade-offs. The optimum SAIDI plan will minimize SAIDI, but it might lead to less extension of lifetime than a plan that optimized that (and vice versa).

Prioritization methods as discussed earlier determine the best spending order for maintenance activities with regard to reliability, in a certain sense, and to a certain extent. To that extent, they optimize results. However, they are generally not considered optimization methods because:

Constraints. The real world imposes limitations on and requirements for the plans a utility can make: there are only so many breaker crews, only so much equipment can be withdrawn for service at one time, etc. These are

constraints that limit the possible plans for creating requirements they must meet. In addition, practical issues of precedence, dependency and interaction create constraints in a mathematical sense (defining qualities the solution must have or cannot have). Project A has to be done before Project B can be done. Projects C and D are mutually exclusive: you can do one or the other but not both. Doing Project G reduces the expected benefit of H by 23%.

In many practical situations for both O&M and capital budget optimization, there will be more constraints involved than options to consider.¹⁶

Multi-attributes. Most organizations and executives care about more than one important attribute. Ideally, they hope to minimize SAIDI *and* the number of misoperations *and* expected equipment lifetimes. Executives realize that there are trade-offs, but everyone hopes to find ABE solutions (Almost Best Everywhere) – a plan that, while not strictly optimum anywhere, is *nearly* optimum everywhere (Table 12.16. The ABE solution is not optimum in any respect, but provides KPI measures within 4% in all measures, whereas any optimum solution falls up to 20% off from the optimum results in categories that are not optimized.

Table 12.16 Example of an ABE Solution

Solution Name	Cost \$ x10 ⁶	Δ SAIDI Minutes	Mis Ops %	Lifetime Years
Traditional Standards	\$14	-2.5	.60%	29
Random Selection	\$14	-2.9	.55%	27
Optimum SAIDI	\$14	-4.10	.51%	27
Optimum Mis-ops	\$14	-3.85	.42%	26
Optimum Lifetime	\$14	-3.65	.58%	31
ABE	\$14	-4.05	.43%	30

Risk and Sensitivity. The executives running nearly any business want to know the sensitivity of a particular plan to factors not entirely known or

¹⁶ A medium size utility the authors worked with in the late 1990s used optimization to determine the best service, maintenance, and replacement program for its distribution and substation equipment. Mathematically, this resulted in 934 candidate distribution and substation programs and projects to choose from, and over 1,200 constraints of various types. The optimum plan funded 180 of the projects (about 20%) but accommodated all the constraints.

under their control (how does weather and the variable likelihood of storms play in to the possible annual SAIDI result that would be obtained?)

Prioritization, at least as described in Sections 12.3 – 12.6, cannot handle constraints and multi-attributes well. It cannot work against more than one attribute at a time, and it cannot directly address risk and sensitivity issues. Work-arounds to these limitations exist and are at times partly successful. Procedures that apply prioritization serially and in step-wise fashion can work around small numbers of constraints. A number of organizations have built heuristic procedures that combine multiple attributes into a single aggregated score, to give the appearance of addressing multiple attributes, and to a small extent these help. Risk and sensitivity can be studied using prioritization by slightly changing the input and looking at the results, thus obtaining sensitivities and gradually plotting risk. But in general, the basic prioritization approach – rank all candidates on the basis of merit and then select from the top of the list—cannot optimize in the sense defined above.

To find the best solution to a practical problem in the presence of constraints, when there are multiple attributes, and when risk is involved, one needs true optimization algorithms, which employ a formal numerical or logical search method to identify the optimum answer, or in some cases to find a sub-optimum answer. “Sub-optimum” sounds as if it could be far from preferable and in strict terms of the words, it means any solution less than the best. However, common mathematical parlance uses it to mean the following: “Mathematically one cannot prove this algorithm always finds the optimum solution – but what we are calling the sub-optimum solution is both a really good solution that fully meets all constraints, and the best solution that can be found.”

12.9 CONCLUSION AND RECOMMENDATIONS

It is difficult to argue against RCM, particularly in the case of a utility with significant aging power delivery infrastructure to manage. Depending on how one looks at it, RCM either maximizes the value a utility can obtain from a limited budget or minimizes the cost to obtain any targeted level of reliability improvement through maintenance.

Utilities with aging power infrastructures have need of more than the average amount of *results* from inspection, testing, preventive maintenance, and repair. Often they have fewer budget dollars than required for traditional maintenance approaches. As a result they need all the advantages possible with respect to getting the most “bang from the buck” from maintenance. RCM does precisely this. It is implementable with information that any distribution system owner (industrial or commercial or public utility) has, and uses only straightforward procedures (electronic spreadsheets, good record-keeping), although those procedures do have to handle large amounts of projects and alternatives to be effective. Therefore, it is strongly recommended.

Table 12.17 lists a number of recommendations for maintenance prioritization that generally hold true to all utilities but particularly those with aging infrastructures to manage. The rest of this conclusion discusses each item in the table in more detail.

**Apply Prioritization to All Parts of the System,
Not Just One Level or One Department's Venue**

RCM can work at all levels and for all types of equipment at a utility. There is no reason to limit its scope to only one type of equipment (e.g., power transformers) or to within only one department (Substation Operations). In fact, as will be covered in Chapter 14–15, the basic concept should be extended to all activities, and all spending throughout the utility, not just restricted to maintenance.

Zero-Base All Decisions

All good O&M planning begins with an assumption that nothing is sacred. Any action has to justify itself on the basis of the benefits it makes compared to the costs it requires. All possible O&M actions are in competition with one another for funding. The most effective will win.

Apply Prioritization Using Intra-Project Alternatives' Evaluation

The improvement in “bang for the buck” wrought by applying the optimization to determine both what equipment should be maintained, and how maintenance on each unit should be done can be substantial. The added value is well worth the cost of the more comprehensive and detailed analysis required.

Table 12.17 Guidelines for Maintenance Scheduling on Aging Infrastructure Systems

-
- Apply the RCM concept using intra-project alternatives' evaluation in the optimization
 - Use a reliability definition as close as possible to regulatory mandate or formula
 - Zero-base all decisions
 - Apply intra-project alternatives evaluation for RCM
 - Use the same reliability definition in all regions and for all departments
 - Use generic categories of work and equipment along with pre-defined expectations and costs
 - Do not confuse inspection period with improvement period
 - Rely on inspection and testing to carry a lot of the burden in any maintenance program
-

Minimize the Same Reliability Definition Everywhere

Similarly, unless regulatory mandates decree different definitions of reliability or “good customer service” be applied in different areas of the system, the same definition should be used in all RCM in all regions and all departments.

Use a Reliability Definition as Close as Possible to Regulatory Mandate

As discussed in Section 12.3, while conceptually there may be indices (definitions of reliability measure) that make sense or seem fair from one perspective or another, ultimately a utility should let the regulatory process define “what is wanted.” Once defined, this should be the performance index used in all decisions aimed at “improving reliability.”

Use Generic Categories of Work and Equipment Along with Pre-Defined Expectations and Costs

As discussed in Section 12.3, use of categories for both equipment condition and service classes both simplifies application of RCM on a large scale, and keeps it within in a context of using cost and results factors in only a legitimate probabilistic, manner. Categorization within a rigorous format also discourages adjustment of data by individuals based on claimed “judgment” or special knowledge (i.e., fudging the data).

Do Not Confuse Inspection Period with Improvement Period

The periodicity needed for inspection and the period during which improvements from service are expected to last are not the same. Section 12.3 (Figure 12.6) recommended that the period of improvement due to servicing a unit be modeled as relatively short-term. However, short as it may be, it is often longer than prudent periods for inspection. For example, “major service” might produce an improvement in expected reliability of a transformer that lasts five years. But the utility might decide it is prudent to inspect the unit annually, because problems with something other than what was serviced last time could develop.

Expect Inspection and Testing to Carry a Lot of the Burden

In any optimized program, inspection, testing, and diagnostics carry a good deal of the weight of improvement. In a very real sense, a good RCM plan is its own type of “just in time” system: frequent inspections and testing tend to uncover problems before they occur; a system based on these problems where repairs are only made to items that need to be replaced; and periodic scheduled repair/replacement. On the other hand, something may be replaced that doesn’t seem to really need it, just because it’s time.

REFERENCES AND BIBLIOGRAPHY

- R. E. Brown, *Business Essentials for Utility Engineers*, CRC Press, Boca Raton, 2010.
- J. J. Burke, *Power Distribution Engineering – Fundamentals and Applications*, Marcel Dekker, New York, 1994.
- P. Gill, *Electrical Power Equipment Maintenance and Testing*, Marcel Dekker, New York, 1998.
- Institute of Electrical and Electronics Engineers, *Recommended Practice for Design of Reliable Industrial and Commercial Power Systems*, The Institute of Electrical and Electronics Engineers, Inc., New York, 1990.
- H. H. Malik, A. A. AlArainy, and M. I. Qureshi, *Electrical Insulation in Power Systems*, Marcel Dekker, 1998.
- O. C. Seevers, *Management of Transmission and Distribution Systems*, Fairmont Press, Lilburn 1995.
- W. A. Thue, *Electrical Power Cable Engineering*, Marcel Dekker, New York, 1999.

13

Planning Methods for Aging T&D Infrastructures

13.1 INTRODUCTION

This chapter provides a one-chapter summary of power delivery system planning and the methods and processes used, with a focus on how this planning is made to encompass and focus on aging infrastructure needs. Readers who want more detail than provided here about planning in general can consult Willis (2004). More detail on the methodology for planning of aging infrastructure will be presented in subsequent chapters and in case studies given there and in the Appendices.

Power system planning involves identification of the best equipment to be added, and changes to be made, to the system. It produces specifications and authorization that identify the additions and changes, their locations, manner of interconnection to the system, and schedule of deployment.

Traditionally, system planning was oriented around *expansion planning*. The utility's customer base and load were growing. At some point in the future, the present system would prove inadequate to serve the growing electric demand. When would that be? What additions would need to be made so the system would continue to be adequate? That was the traditional power delivery planning paradigm.

A consideration for aging equipment and infrastructures needs to be added. Equipment that has long been dependable is wearing out. Circuit configurations

and facility locations picked decades ago may not be best for today's and the future's needs. need to be added into the plan to best accommodate these considerations, and mixed with those additions needed to support growth, so that the overall plan is the best – whatever best means. Since cost is an important attribute in power planning, almost invariably one of the planner's chief goals is to minimize overall cost of handling the growth and the aging.

Aging T&D Infrastructures → Reliability Optimization

The primary negative effect created by aging T&D infrastructures is poor customer service reliability. All of the other impacts, which boil down to higher costs in one way or another, stem from the utility's attempts to control the situation with regard to equipment breakdowns, safety, and dependability, or to improve reliability in the face of increasing equipment failure rates. Thus, "Planning for Aging Infrastructures" needs to consider operating costs and risk, and reliability and optimizing reliability gained per dollar, tailoring that reliability of service to consumer needs and utility targets.

Long Term Strategy and Value, But Not Long Term-Planning

It takes decades for an infrastructure to deteriorate to the point that aging effects seriously inhibit performance. The monies required for any viable solution are very large and must be spread over many years. Only programs that are basically continuous in nature can economically address the issues. Thus, a long-term view of the problem and its solutions is needed in order to develop a workable and affordable plan.

This means that at some place in the utility, long-term assessment and planning for aging equipment and infrastructures must take place. A utility must understand its present situation, where aging trends are taking its equipment fleet, and what it can expect from various strategies and approaches it can apply. Ideally, it will develop an overall strategy for when, where, how, and why it will address aging infrastructures and the issues they cause.

But this does not mean that all T&D planning need be done with a long-range perspective. Much of the planning at any utility is short-range – two to five years ahead, oriented, but despite that short-range plans and projects are made to fit within long-term strategy. Traditionally, such short-range planning was coordinated with long-term system and business strategy and needs via criteria and guideline formulae. For example, utilities make minimization of losses cost a long-term goal, specifically aimed at making the lifetime cost minimum. This is accomplished in planning which may look only two to five years ahead, by setting the coefficients of formulae used to compute lifetime losses costs as used in the final evaluation of options from a present worth standpoint.

Conceptually, the same criteria and formulae-based approach is used with respect to aging infrastructures (Figure 13.1).

- Corporate strategic long-range planning of and for aging is based on projection of current trends and corporate goals, eventually leading to an overall strategy and plan. This is generally a business-oriented plan but it provides budgets and metrics (for age, condition, etc.) that can be mapped to system and equipment needs. Examples will be given in subsequent chapters.
- Power delivery planners, working with the corporate planners, develop appropriate criteria, targets, and formula, that position cost and value computations and planning goals in short-range planning toward fitting the aging-related plan.
- These are then applied routinely in the short-range planning process. The result is that short-range plans developed in the two- to five-year ahead process will fall in line with both long-range needs and long-term strategy.




Figure 13.1 A utility's long-range aging-infrastructure priorities are met by the short-range planning process by orienting it to a set of criteria, formulae, and guidelines determined from a long-range plan based upon the utility's aging infrastructure strategy.

At times, the utility may need a special, initial planning effort leading to a program that initiates aging-related programs to get its effort “off the ground” so to speak – something that is basically making up for not having the three steps above for the past decades. This, too, will be discussed in subsequent chapters.

This chapter begins with a look at the basic purposes of planning and the five steps of the planning process (Section 13.2) reviewing the most common pitfalls in each. The purpose of short-range and long-range planning and how they interrelate to aging infrastructures is then discussed in Section 13.3. Section 13.4 then looks at the T&D planning process as implemented in most modern utilities, again emphasizing aspects that particularly interact with the aging infrastructure issues. Section 13.5 stresses the importance of the systems approach in all planning for aging T&D systems. Finally, Section 13.6 summarizes key points with respect to aging infrastructure system planning.

13.2 PLANNING: FINDING THE BEST ALTERNATIVE

What is the Mission?

Planning is a decision-making process that seeks to identify the available options and determine which is the best for the situation at hand. Its purpose is to determine the best schedule of future resources and actions to achieve a company's, an individual's, or an organization's goals with respect to power. Usually, among the major concerns of planning are financial considerations – minimize cost, maximize profit, or some similar goal. In addition, service quality and reliability are almost always considerations. And often, other criteria are important, including environmental impact, public image, and a host of factors that a utility, energy services company, or energy consumer must weigh in making a selection regarding power supply. But regardless, planning is the process of identifying alternatives and selecting the best from among them.

Before beginning the planning step, it is best for a utility to review its mission statement – its overall definition of what guiding purpose and value-system defines its goals. This is particularly true with respect to meeting the challenge of aging infrastructures. Dealing with aging infrastructures, and getting what is desired in the outcome, means having focus and being realistic about constraints, goals, and capabilities. The mission statement is a good place to start. Given the realities of the situation, the aging system, a limited budget, regulatory pressures, is the current mission statement realistic? Making sure that the company's mission statement matches both what it wants to do, now that it has assessed its aging infrastructure problem, and making sure that those principles are realistic and feasible, are big steps toward success.

A company's “Mission Statement” succinctly defines the overall corporate objectives, while its economic goals explain its financial priorities. Usually, mission and financial goals have been incorporated in the economic guidelines

and design procedures that the planner applies in his work. Regardless, while mission statements are often taken lightly by some employees, a review of the company’s mission and financial goals provides the planner with strategic insight into its planning guidelines and economic evaluation formula. Table 13.1 lists the mission statements from four large metropolitan electric distribution utilities in the United States. Planning goals for T&D, for DSM, for IRP, *will* be different among these utilities, because they have different priorities for what their executive management wants to accomplish. How the aging infrastructure problem will be handled in each utility and the priorities the utility puts of various approaches and solutions should be different as a result.

Mission statements are qualitative assertions of the overall philosophies and goals that lead to the quantitative engineering and economic criteria used in distribution planning. By understanding *what* their company is trying to do, the T&D planners have a better understanding of *how* to achieve those goals.

Table 13.1 Mission Statements of Four Electric Utilities

1. “Our company will be the premium regional provider of electric power.”

Recognizing that its financial situation prohibits competing on price, this utility makes quality and service its hallmark. Achieving lowest *possible* cost is not the goal; achieving high service standards is.

2. “Provide economical electric power for the prosperity of the region.”

This municipal utility has a long-standing tradition of low rates, a way of attracting new industry (i.e., growth) to the region. Plans that invest a good deal to improve quality are simply “not with the program.” Marginal quality improvements in a new plan are permissible, only if they lead to lower cost.

3. “. . . maximize return on equity.”

This utility wants to earn the most on what it currently owns. Utilization of *existing* T&D facilities is the key. It is more reluctant to invest by borrowing than a utility that wants to maximize return on investment. Plans that call for massive capital spending may be unacceptable.

4. “. . . maximize return on investment.”

This mission makes no direct distinction between equity and debt – apparently the utility is willing to borrow, if it increases what it can earn on that investment. A plan calling for considerable capital outlay is permissible, if that investment yields a good return.

Planning: A Six-Step Process

Planning is a process of determining the best course of action or schedule of commitments that can be made to achieve one’s goals. The modern power delivery planning process can be segmented into the six steps shown in Figure 13.2. Each one is an important part of the process for planning in any type of industry, but in particularly utility T&D planning.

Any of the six steps, poorly performed, can lead to poor decisions, a poor plan, and ultimately failure to attain those goals. Mistakes and lack of focus in any step will lead planners astray, although the types of mistakes that result will differ, depending which one of the five was the first misstep. Below, each of the five steps is described, along with a brief summary of the goals for this step and pitfalls to be avoided in each. A lengthier discussion, with examples, can be found in Willis (2004).




Figure 13.2 Planning involves the six steps shown. See text for details.

Step 1: Identify the Scope of the Problem

Before beginning their work, every group of planners should identify their goal – exactly what they are trying to do. It is recommended that the planning problem and scope be formally written, as “this project involves planning of short- and long-term electric supply for all 821 homes planned for the new West Oaks subdivision.” This written definition of the scope defines the *extent* of the problem to be solved. All subsequent parts of the planning process should focus on that, and not stray (as some planners are often wont to do) to solve other problems or work in other areas. Trying to resolve a wide range of other problems in the nearby area at the same time this goal is achieved is not part of this project.

Step 2: Identify the Goals

A common mistake made by many planners is to begin their activity without specifying the goals for the particular planning situation at hand, or identifying how these goals might have changed from prior planning situations. It may be one thing to identify the problem, but the goals must explain in detail how to identify what the planners need to accomplish.

Goal definitions are quite different than mission statements. A mission statement lays out an organization’s direction, principles, and values. A goal defines explicitly what is desired to be accomplished in some situation. The former is strategic, the latter tactical. Time devoted to framing the planning situation at hand and explicitly defining the goal for it will be time well spent.

Again, as was the case with scope statement, it is best to produce a formal, written definition. This might read something like: “The goal is to develop a plan, able to be implemented within two years, which identifies all required equipment, labor, permits, and other expenses for electric facilities. This will provide the lowest long-term cost when evaluated on a present worth basis, and provide West Oaks subdivision with expected SAIDI and SAIFI of 2.0 hours and 1.25 events respectively, with all system factors meeting company standards.”

Aging Infrastructure Situations Often Encounter Severe Environmental and “Urban Fit” Constraints

Very often, particularly when dealing with aging infrastructure problems in the interior of older metropolitan areas, there are additional goals unique to a particular plan or situation. For example, it might be important to accommodate the esthetic and land-use preferences of local community leaders, or to use equipment or rights-of-way and sites already in inventory. Or, most often, a key requirement is to *not* require any new sites, rights of way, or visible disruption of an already crowded urban landscape. Some goals are unique to each particular case, created by an unusual situation or set of circumstances. Table 13.2 lists a number of special goals that are not uncommon to encounter when planning for an aging T&D system.

Table 13.2 Examples of Special Planning Goals or Criteria Often Encountered When Dealing with Aging T&D Systems

Fit within a specific budget.

The utility may have a particular amount of money designated for aging infrastructure improvements. This money must be spent in an optimal manner, but further funds will not be available. Budget-constraining techniques (see Chapter 5) are often helpful in this situation.

Substation improvements must be accommodated within existing sites.

The utility either cannot obtain, or management does not want to endure the political hassle, uncertainty, and cost of trying to gain more real estate at existing substation sites. As a result, innovative ways of cramming more capacity and more complex configurations into an existing, already-crowded site, such as use of GIS equipment, compact line designs, and small-footprint transformers, must be found.

No new substation sites and rights of way.

This is occasionally a goal of management – to minimize the number of new substation sites, ideally to zero. New sites are expensive, subject to political and regulatory hurdles that make approval uncertain, and often create adversarial relationships within the community, which management judges it is best to avoid “at any cost.” Lack of new sites does not cause a capacity limitation issue (lack of new room at existing sites does). It causes a problem with economic and capacity reach issues of the feeder system. So called long-reach planning techniques as discussed in *Willis (2004)*.

Expansion plan must include “targeted DG and DSM or renewable energy sources of power.”

Regulator or community leaders may demand that any solutions include the use of modern technologies as a means to cut T&D capacity needs, including load reduction methods and on-site generation. Integrated resource T&D planning and DG planning methods are often needed in these situations. Planners should recognize that while these special resources may provide only a small amount of relief from the reliability of service and capacity challenges of the area, they should be used to their fullest, and that they are required and in a larger sense to address societal expectations and political needs that the company cannot ignore.

High reliability is required to a specific site or customer.

A special customer or site may need extraordinary reliability of service even if the neighborhood around it need only be given the average level of service for the rest of the system. Examples are a professional sports stadium, a hospital, or a police or fire division headquarters located in an older area of the city.

Step 3: Identify the Alternatives

This third step is often the most critical part of planning. This is identifying what *could* be done – all the alternatives open to the planner. What options are available? What variations on these options would be possible? This function encompasses determining the *range* of possibilities for solving the problem.

This is where the majority of great “planning disasters” occurs – the mistakes that lead, weeks later, to statements like, “Why didn’t we think of that?” For, while apparently the least challenging part of planning, a good deal of skill, breadth of thinking, and time, are required to identify the range of possibilities. Time and resources are always limited, and the temptation is strong to assume that one can see all the options just by “a quick look at the problem.” But this is seldom the case. A brilliant solution to an unusual problem most often comes from recognizing that there are also unusual options available to solve it.

The downfall of planners: failing to consider all the alternatives

The discussion of Step 2, above, mentioned that planning power delivery systems in older areas of a system can be a challenge, with locally unique challenges in many unique situations. That makes Step 3 particularly important: often viable alternatives are not obvious. Common fixes used in most situations (run a new circuit) are not possible or preferable. Uncommon approaches might be possible (the neighborhood contains many old but structurally sound buildings: a new substation can be put inside the façade of one, or a new façade that looks old, leaving the neighborhood’s appearance unaltered from the street.¹

This step is often where a surprising number of computerized “automated planning procedures” and tightly prescribed template-based planned methods let their users down. When looking at or reviewing their planning procedure, electric system planners should study if and how it assures them that it *does* consider all possible options. Just as planners need to remind themselves constantly to be open-minded, and not to have pre-conceived notions about where the solution lies, they need to assure that computer programs and processes they use are also not limiting their identification of the full range of options.

In general, it is recommended that planning in older areas in the interior of major metropolitan areas be treated as a special situation, with a team of at least two planners. The team requires both innovation and a lack of adherence to traditional thinking, *and* above average consideration of those things that only come with years of experience. Often, what works best is a team composed of a

¹ The authors used this approach when working together in the late 1990s, on the upgrading of T&D capability in an aging area of a major American city. To this day most people walking past on the sidewalk are unaware they are passing a five-transformer, 200+ MVA power substation that feeds one-fifth of the entire downtown area.

younger planner with an innovative attitude and a very experienced old hand whose personalities mesh well.

A common failing: too short a planning horizon

Beyond a team that will be a bit more innovative than usual and is more savvy than needed in the typically suburban planning situation, another requirement for success in most aging infrastructure areas is a longer planning horizon. Simply put, the system did not get into its current dilemma in a short period of time and a successful plan is going to be one that rights its wrongs quickly and permanently. Managers should require that, in addition to a short-range plan that fixes immediate issues and priorities, the team have an overview of what is required and how the long-term issue will be dealt with. Such a longer range vision often provides insight and help to the short range planning, but will be neglected by planners in the interests of urgent short-term priorities, unless they are explicitly told they have both the time and resources to delve into the longer range

Failing to look far enough ahead often limits the available options by default, contributing to poor planning. A common lament among T&D planners at some utilities is, “We could have done a much better job, but the good sites and land were already taken.” What is usually left out of this complaint is that the planning period used for distribution expansion evaluation was only three to four years ahead. The utility is losing out to other people interested in acquiring sites for other purposes, people who are looking ahead more than three or four years.

By looking only a few years ahead, the utility reduces the number of options available to it. In addition, planners often fail to see the consequences of long term growth and change. If the area has the possible potential to redevelop into significantly more load, then eventually it will redevelop, and the existing system, aged as it may be, will have to be bolstered. The question, then, is not if, but when is that potential redevelopment soon enough to affect elements and economies of today’s short-range plan?²

Having a sufficiently long planning period does not in and of itself assure good planning, but it is a necessary first move in that direction, particularly in difficult situations such as those commonly encountered in long-developed areas with aging infrastructures. For distribution planning, a *minimum* planning period of five to seven years is normally recommended, and in aging areas, a horizon at least 20 years old. Long range plans do not need to be developed with great specificity (nothing 15 years out is going to be built for a long time, and conditions might change). But they do need to be developed to the point that planners can answer the question: “Where are we going over the long run?”

² Anyone who thinks that a blighted, older area of a city can be beyond redemption and will forever remain an older, sleepy area should consider the history of the inner harbor area of Baltimore.

“Do nothing” should always be one of the alternatives considered

One of the most serious mistakes made by planners is to assume that *something* must be done. Occasionally, it is both less expensive and better from a service standpoint to delay any commitment or change to the electric system – to do nothing. Even when this is not the case, “nothing” should always be included in a planning report as one of the reviewed options. This permits the planners to explicitly identify *why* something must be done, by showing that “doing nothing” would violate criteria or requirements, or lead to very poor economy.³ Evaluation of “nothing” is a key element of justifying a plan’s recommendations.

Step 4: Evaluating the Alternatives

All alternatives should be evaluated against a common and comprehensive evaluation standard, one that considers every aspect pertinent to any option, and one that addresses all the goals. For power delivery planning this means evaluating alternatives against criteria and attributes that represent the utility’s requirements, standards and constraints. For DSM programs this means matching customer needs with marketing standards and guidelines.

Very often, the actual planning process will combine the evaluation and selection functions in a process designed to reject alternatives quickly and with minimum effort. Regardless, all alternatives should be evaluated completely with respect to everything that bears on the problem at hand. Planning methods should be examined carefully to make certain that they, too, like the various alternatives being considered, are complete.

Alternatives must be evaluated against both criteria and attributes. *Criteria* are requirements and constraints the plan *must* meet, including: voltage, flicker, other service standards, contingency rules, loading limits, safety and protection standards, and so forth. Such criteria only have to be satisfied: if standards allow voltages down to .95 per unit, then a design alternative is acceptable if it has voltages of .95 per unit, but no lower. On the other hand, attributes are values for which the planners wish to achieve the best performance possible: an *attribute* is a quality that is to be *minimized* (or maximized) while still meeting all criteria. One attribute in nearly all electric supply planning is cost. And while voltage drop, protection coordination, and other requirements are quantities that the planners would *like* to see exceed standards by comfortable margins, they should be treated as the criteria they are: good enough is good enough.

Traditional power delivery planning is *single attribute planning*, in that only the one attribute (cost) is to be minimized. This third step, evaluation, rejects the

³ A common argument against this approach is that evaluation of “do nothing” is not needed because the answer is so obvious that it does not need to be included. But if that is the case, it is simple to do, so why not include it and document its evaluation?

Table 13.3 Some Typical Criteria and Costs Evaluated in Electric System Planning

Voltage standards	Cost of equipment
Noise and emissions limitations	Cost of land and ROW
Safety and clearance standards	Cost of labor and construction
Electrical and fire protection standards	Cost of taxes, insurance, etc.
Pollutant/toxicity material restrictions	Cost of design and engineering
Aesthetic impact and standards	Cost for maintenance and operations
Maintenance access requirements	Cost of losses
Equipment loading standards	Cost of permits, licenses, etc.
Contingency margin requirements	Flexibility in future design
Equipment loading standards	Salvage cost (at end of life)
Construction method requirements	Fuel and fuel delivery cost

alternatives that do not meet all criteria and passed all the others. Among those left, the minimum cost alternative can be selected (in Step 5). Cost is a multi-faceted attribute, *all* aspects of which must be included in the evaluation: equipment, site costs, taxes, operations and maintenance, and losses. Table 13.3 lists some of the more common criteria and cost factors that need to be assessed in electric supply evaluation at the distribution level.

Where mistakes occur in the evaluation step

Mistakes that take place in the evaluation function generally occur because the planners do not check to see if their methodology meets three requirements:

1. *It considers all criteria and factors that are important to the goals.* For example, if a goal is to achieve a certain level of expected reliability, then the planning methods need to be able to assess reliability of all plans on an equivalent basis.
2. *It must evaluate all criteria fully with respect to all resources.* For example, some planning methods evaluate reliability only with respect to outages and connectivity. When evaluating DG and DS, the reliability method must analyze and accommodate lack of fuel availability and potential for failure of the units themselves.
3. *It treats all options equitably.* For example, if planners include the cost of oil spill containment and fire prevention equipment when evaluating T&D alternatives that include oil-filled transformers, they should provide like detail in the assessment of the DG alternatives.

Step 5: Selecting the Best Alternative

In many planning procedures, this function is combined with the evaluation function, and it is difficult to identify where evaluation ends and selection begins. In fact, efficient computerized planning methods that help planning engineers minimize the time and effort required to complete the planning process tend to combine Steps 3 through 5 into one process. But regardless, if the most important point in the identification-evaluation/selection process is the *definition* of “best,” it must match the goals and value system being used. The evaluation/selection method being used must mirror the desires of the organization, and be capable of distinguishing between alternatives in a valid manner in these areas.

A surprising number of planning mistakes, or inefficiencies, occur in the final selection process, because the planners do not assess alternatives on the basis of what is truly important to their goals. A planning method’s selection function should be examined to determine that it:

- *Uses the definition of “best” that matches the planning goals.* The fact that alternatives were evaluated on the basis of a particular attribute in Step 4 does not mean that it is weighted properly in the selection phase. For example, a common failure in T&D planning is to assess the value of electric losses in a distribution plan, but then fail to acknowledge their value in selecting the best design.
- *Can accurately distinguish between alternatives.* What are the limits of accuracy of the planning process? If it is only accurate to within 3%, then selection of a “best” alternative on the basis of its being 1% better than another is questionable. Error range in the evaluation and selection among close cases must be small enough to assure that the determination of “which is best” is accurate.

Many computerized planning procedures produce only approximate evaluations. Planners often fail to consider these limitations and interpret minor differences in computed evaluations as significant, when they are not.

Step 6: Being a Proponent of the Plan

This step is in many ways the key to success in the 21st century utility industry. Plans have to be sold to executive management, to stakeholders including in some cases stockholders, and in all cases, community and regulatory leaders. The planner is often heavily involved in this process and sometimes asked to lead it.

There are several critical points to keep in mind about this step. First, this represents a major change in perspective for the planners. In Steps one through Five, objectivity was a key. Make certain all options are considered

on an equivalent and balanced basis to evaluate and report fully and fairly.

Once management's approval process has selected of an alternative, the planner's job changes. The planner, whose job was to *objectively* present the alternatives and their evaluation to management, now becomes a *proponent* whose job is to help sell that plan to all stakeholders. This can be a challenge, because management may have picked an alternative other than what the planner believed was best. Despite this, the planner's job is to fully support carrying this option forward as the plan, helping sell it to all other stakeholders, and guide it through the regulatory approval process, until it is handed off as a fully authorized project to be built.

This job is easiest if the preceding five steps were well structured within a fully documented and institutionalized framework, used a credible planning method at every step, and was transparently and traceably documented throughout. *Defensibility is often the most important element of a good plan.*

13.3 SHORT- AND LONG-RANGE PLANNING

Lead time – the time it takes to prepare, build, and install facilities – sets the minimum planning period for any endeavor. If it takes up to five years to order materials, obtain permits, survey, build, test, and put into service a new substation in the core of a major metropolitan area, then the power distribution utility has no choice but to plan at least five years ahead in order to make certain it will have all new substations by the time they are needed.

Thus, the minimum possible planning horizon is a function of the lead time required to put solutions into place. Lead time varies depending on the type of equipment, the level of the system, and the local planning approval process for electrical facilities, and other aspects of the situation. Table 13.4 gives some typical lead times by “level of the system.”

Short-Range Planning: Decisions and Commitments

The purpose of electric utility short-range planning is to make certain that the system can continue to serve customer load while meeting all standards and criteria. This stage of planning is driven by the lead time – by the fact that it takes time to get things done and that decisions must be made ahead of this lead time.

Short-range planning must identify the commitments and purchases that must to be made today in order to allow them to be implemented in time to maintain service-within-standards in the future. In the case of the substation example given earlier, it takes five years to make all preparations. Therefore, the utility must plan at least five years ahead, so it knows when to “trigger” the process of creating that substation. Using a shorter lead time means selecting “do nothing” by default.

Table 13.4 Typical Minimum Lead Times at Various Levels of the Electric Utility Power System

Level	Years Ahead
Central Station Generation	10
EHV Transmission (>500 kV)	9
Transmission	8
Sub-transmission (< 140 kV)	7
Primary Distribution Substations	5
Primary Voltage Feeders	3
Rebuilding existing substations	2
Rebuilding existing lines – UG	2
Rebuilding existing lines – OH	1
Lateral primary voltage taps, changes	1
Distributed Generation (>10 MVA)	2
Distributed Generation (10 - 1 MVA)	1
Dispersed generation (< 1 MVA)	.4
Service level	.2

The product of the short-range process is a series of *decisions* – the selection of various alternatives over other ones – for the expansion of the electric service facilities in various areas of the system. It is implemented as a series of projects, often one project per area and major decision, but sometimes with several lumped together into a larger project. Usually, the product of the planning process is a set of detailed project descriptions, each specifying what facilities and equipment are to be bought, where they will be put, how they will be fit into the system, as well as what other commitments and plans need to be enacted. Short-range planning initiates a process which then fills in the design and procurement detail and produces the land acquisition requests, drawings, permit requests, materials lists and construction authorizations, hiring, re-organization, marketing, advertising, customer survey, and other efforts necessary to make it happen.

The “short-range plan”

The “short-range plan” is a schedule of additions, enhancements, reinforcements and changes that are to be made to the system and that have been authorized and committed. Although there is always some recourse in any plan, the short-range plan is composed of decisions which have, in effect, been “locked in.” As such, it tends to be very “project oriented,” as shown in Figure 13.3. The short-range

planning process leads, ultimately, to a series of separately identifiable project authorizations, each committing to a particular addition(s) to the system, with its own scope of work, schedule, budget, and justification. These project authorizations are the hand-off from the planning process to the engineering and construction process.

Long-range Planning: Focus on Reducing Cost

Long-range planning focuses on making certain that the equipment and facilities called for in the short-range plan provide lasting value and the lowest overall cost during their lifetime, unlike short-range planning, which seeks to identify problems and solve them before they occur. The reason to look beyond the lead time, and to plan for the long run, is to assure the decisions being made have the most lasting value.




Figure 13.3 The short-range planning process consists of an initial planning phase in which the capability of existing facilities is compared to short-term (lead time) needs. Where capacity falls short of need, projects are initiated to determine how best to correct the situation. The load forecast is primarily an input to the short-range planning process, i.e., it is used to compare to existing capacity to determine if and where projects need to be initiated.

problems and solve them before they occur. The reason to look beyond the lead time, and to plan for the long run, is to assure the decisions being made have the most lasting value.

How can a planner know if a new distribution line, DG unit, or other electrical addition is a good investment, with a low lifetime cost and a high utilization over that lifetime? That lifetime begins when the new line is completed and put in place, *which happens at the end of the lead time*. Thus, long-range evaluation begins at the lead time and assesses cost, value, and equipment contribution over a lengthy period thereafter. Its purpose is to determine if the short-term decisions will have long-term value. Given typical economic factors, and considering the uncertainties in predicting conditions over the very long term, a ten-year period is generally considered the *minimum* for such economic evaluations.⁴

A good long-range planning process is an end in itself

Long-range T&D planning is one of the few legitimate instances where a process's major goal is itself. The long-range plan, along with its accompanying budget estimates and indicators of economic factors for the future system, *is* the major goal of the long-range planning process. More importantly, it is crucial that this plan's development be ongoing, that it be maintained and continuously updated in order to provide long-range evaluation, guidelines, and direction to T&D planning and investment decisions.

To study a new transmission line, substation, distributed generator, or other planned expansion item and determine its impact over a future period from five to 25 years ahead, the planners need to have a good idea of the future conditions throughout the time the new project will be in service. This includes a forecast of the future load levels it will be called upon to serve, as well as a description of the system and conditions within which it will operate. They need a *long-range plan*, not just for the particular facility they are planning, but as a backdrop to their evaluation of this project.

This long-range plan must lay out the economics (value of losses, etc.) and operating criteria (usually assumed to be the same as today's) that will be applied in future operations. It must specify the load that will be served (long-range load forecast), and it must specify *the other electric additions* that will be made during the period (e.g., other DG units planned for future years). This is so their interaction and influence on the performance and economy of any proposed short-range addition or change can be fully assessed over a good portion of its lifetime.

⁴ When using typical discount rates (e.g., 11%) analysis of a ten-year period of service captures 66% of the present worth of a level investment with a 40-year lifetime, while a 15-year period captures 80% of its present worth. Thus 10- to 15-year periods that begin at the lead time see a majority of the present worth aspects of an investment, but not all.

Table 13.5 Purposes of the Long-range Plan
Definition of “future” for evaluation of short-term planning decisions.
Identification of mismatches between levels – coordination of plans.
Forecast of long-range budget needs (for financial planning).
Identification of long-term direction and strategy.
A basis for evaluation of new ideas or changes in procedure.




Figure 13.4 The coordinated multi-level long-range plan, where the major goal is the maintenance of a long-range plan that identifies needs and development of the system, so that short-range decisions can be judged against long-range goals. Here, the load forecast is a part of the plan, and there is no “project oriented” part of the activities.

The long-range plan does not need great detail

A long-range plan needs only enough detail to permit evaluation of the economics and “fit” of short-range decisions regarding long-range system requirements. A long-range plan needs only a certain amount of detail, sufficient to its purpose. Effort beyond this is wasted.

Other functions of the long-range plan

A good long-range plan provides several other useful functions, as listed below. The first is the long-term economic evaluation discussed above, the reason why a long-range plan *must* exist if good economy of investment is to be assured. But there are other valuable contributions that the long-range plan makes to the planning effort, and to the utility as a whole:

- *Forecast of long-range budget.* The long-range plan identifies capital costs, operations and maintenance costs, and losses’ costs to the point that they can be estimated for budget and financial planning purposes.
- *Identification of long-term direction and strategy.* A long-range plan provides a clear direction and long-term strategy for the electric system’s future. It may not be the best direction or strategy, but good or bad, it shows engineer, manager, operator, marketer, and executive alike, the current vision of how electric power will be supplied.
- *A basis for evaluation of new ideas or changes in procedure.* Proposals for changes in standards or system design are difficult to evaluate unless they can be compared on a fair basis against present standards and future needs. A long-range plan provides a “base case” to which new proposals can be compared, whether they suggest adopting a new primary voltage (23.9 kV versus the current 12.47 kV) or other changes in approach.
- *Coordination of planning among levels of the power system.* At some level, someone benefits if the plans for transmission, substation, distribution and distributed generation levels of the system fit together well. To assure this, their plans must be coordinated. This is best done by looking at the long-term implications of their interactions.

This long-range plan must lay out the economics (value of losses, etc.) and operating criteria (usually assumed to be the same as today’s) that will be applied by the utility in its future operations. It must specify the load (long-range load forecast) the system will be expected to meet, and it must specify how the T&D planners expect to handle that load growth – what new facilities will be installed in the long run to accommodate the future pattern of load growth. Most

importantly, it should specify how the various long-range plans at the sub-transmission, substation, feeder, and customer levels will fit together. The economics of interaction of all four levels, leading to a least-cost system, is assured by evaluation of long-term PW value in a coordinated manner – Figure 13.4 shows the major flow in the long-range distribution planning process, while Table 13.5 summarizes its key purposes.

Handling Uncertainty: Multi-Scenario Planning

An important point to bear in mind is that long-range elements in their plan do not have to be built, committed to, or even decided upon, for years to come. This means that the utility can change its mind, or beyond that, be of two or more minds at the same time. *There can be more than one long-range plan.*

Uncertainty in predicting major events is a major concern in long-range planning. Will a possible new factory (employment center) develop as rumored, causing a large and as yet un-forecasted increment of growth? Will the bond election approve the bonds for a port facility (which would boost growth and increase load growth)? Situations such as these confront nearly every planner. Those of most concern to distribution planners are those factors that could change the location(s) of growth.




Figure 13.5 Multi-scenario planning uses several long-range plans to evaluate the short-range commitments, assuring that the single short-range decision, fits the various long-range situations that might develop. This shows the scenario variations for the case involving two possible future events.

In the presence of uncertainty about the future, utility planners face a dilemma. They want to make no commitment of resources and facilities for load growth that *may* develop, but neither can they ignore the fact that there are lead times required to put facilities in place, and that the events *could happen*. Given the reality of lead times, planners must sometimes commit without certainty that the events they are planning for will, in fact, come to pass.

Ideally, plans can be developed to confront any, or at least the most likely, eventualities, as illustrated in Figure 13.5. *Multiple long-range plans*, all stemming from the same short-range decisions (decisions that must be made because of lead times) cover the various possible events. This type of planning, called multiple-scenario planning involves explicit enumeration of plans to cover the various likely outcomes of future events. It is the only completely valid way to handle uncertainty in T&D forecasting and requirements planning.




Figure 13.6 Four scenarios of load density maps for a large U.S. city. See text for details.

Uncertainty in T&D growth forecasts cannot be addressed by planning for the expectation of load growth

One of the worst mistakes that can be made in T&D planning, integrated or otherwise, is to try to circumvent the need for multi-scenario planning by using a single probability-weighted spatial forecast. Such an approach invariably leads to plans that combine poor performance (as evaluated with respect to both electrical and customer service) with high cost.

As an example, the forecast in Figure 13.6 shows a detailed spatial forecast of electric load growth over a twenty year period for a large city, done in 1993 (intermediate years were forecast, too, but are not shown). At the top are maps of load in 1993 and as projected 20 years later in the base case. This map is very close to what actually developed in the intervening 20 years, although with the 2008 recession, the overall load level and pattern of spatial distribution shown will not be reached until about 2016.

The difference in the two top load density maps represents the challenge that lay ahead of this utility's planners in 1993 - design a system to deliver an additional 1,132 MW in the geographic pattern shown, and add that as needed to the system that was in place in 1993 over the 20 year period.

At the bottom in Figure 13.6 are two alternate forecasts each of which included a possible new "theme park and national historical center" - something like Disney World - that was being considered by an international entertainment mega-company at the time, at one of two sites. The park itself was 35 MW. The real impact on the metro area load would be because it would generate a need for new hotels, restaurants, and other entertainment venues, much as existed in Orlando at the time. Altogether it would cause tremendous secondary and tertiary growth, adding over 12,000 jobs to the region, and leading to an average annual growth rate of about 1.5% more than in the base forecast during the decade following its opening.

Thus, the theme park means an additional 260 MW of growth on top of the already healthy expansion of this city. The maps in Figure 13.6 show where and how much growth is expected under each of three scenarios:

- 1) no park (base forecast, considered 50% likely);
- 2) site A for the theme park (25% likely);
- 3) site B for the theme park (25% likely).

As can be seen, the locations of the additional scenario growth will be very much a function of where the theme park is located. Hence, the three scenarios make a difference in both the amount of equipment and new circuits needed, and where they would be located. T&D equipment will have to be re-located and re-sized to fit the scenario that develops. The problem that the planners faced at that time is




Figure 13.7 The probability-weighted sum of the three forecasts in Figure 13.6. This is perhaps the *worst* forecast the T&D planners could use to develop their plans.

that they had no certainty as to which of the three scenarios would work out.

One thing the planners do *not* want to do is form a single forecast map based on a probability-weighted sum of the various possible forecasts, an “expectation of load growth” map. Such a map is easy to produce: it is just the sum of the three forecast maps, each weighted by its probability (Figure 13.7).

While mathematically valid (from a certain perspective, at least) this forecast map will contribute nothing useful to the planning effort – it represents a future with one-quarter of a theme park at location A, one-quarter at location B and the other half somewhere else. That is one scenario that will *never* happen. Using it as the basis for T&D planning will lead to plans that spread resources too thinly over too much territory. If the base scenario develops, it will have too much capacity in many places not needed. If either theme-park scenario develops, it will have about half the extra capacity needed to serve it in the right places, and half somewhere else. A plan based on a spatial expected load forecast would be, relative to any expected output, a poor plan. Planning with the forecast shown in Figure 13.7 guarantees only one thing: the T&D plan will be inappropriate no matter what happens. In particular, if either of the alternate scenarios occurs, the plan is very non-optimal.

13.4 THE T&D PLANNING PROCESS

Figure 13.8 shows the T&D planning process in overview, as it applies in modern utility cases. Figure 13.9 provides a bit more detail as to the flow of information and process, and will be referred to throughout the rest of this chapter. The exact organization and emphasis of these individual-planning steps will vary from one utility to another. What is important is that all are accomplished in a coordinated manner. Other chapters in this book will focus on the data, procedures, and analytical techniques needed in each step. Here, the emphasis is on what each of these steps is and how it fits into the overall planning process.




Figure 13.8 The basic T&D planning process involves six discrete stages.




Figure 13.9 The T&D planning process is driven by a strategic, corporate-level plan that includes the acknowledgment of aging issues the company has included in its plans, its business goals and criteria it wants to adhere to, and its corporate load forecast. The delivery planning process includes load forecast, and planning functions for transmission, substations, distribution, and customer levels, as well as maintenance of a coordinated long-range plan encompassing all the levels. Its inputs are the corporation's forecast, strategic goals and planning and design criteria. Its outputs are the short-range projects.

Five of the six stages consist of both short- and long-range planning portions of length shown in Table 13.6. The strategic planning is only long-range in context. It might consider the short term, but is really a single comprehensive view of the long-term goals, direction, and consequences for the utility.

For the load forecast, “short-range” planning consists of producing a single “base” forecast – a non multi-scenario “most accurate” forecast for use in all short-range T&D planning. This needs cover no more than a decade ahead, that being the longest short-range planning period for any of the levels shown in Table 13.6. For the other four stages, short-range planning is aimed at achieving recognition of problems and definition of solutions within their lead time requirements.

Long-range planning is focused on evaluating the utilization and investment economics of all short-range projects, using a present-worth perspective or something similar. Short-range projects and decisions are judged against this sounding board of long-range performance. The load forecast can and often will

Table 13.6 Typical Short and Long-range Planning Periods for T&D System Planning

System Level	Planning Period - Years Ahead	
	Short-range	Long-range
Strategic Planning – business & system	none	20+
Large generation* (> 150 MVA)	10	30
Small generation* (< 50 MVA)	7	20
EHV transmission (500 kV and up)	10	30
Transmission (1115 kV to 345 kV)	8	25
Sub-transmission (34 kV to 99 kV)	6	20
Utility Owned Smart and Distributed Resources	5	20
Distribution substations	6	20
“Feeder System” (service area) planning	6	20
Primary three-phase feeders	4	12
Laterals and small feeder segments	1	4
Service transformers and secondary	.5	2
Customer level (DSM, RTP)	5	20

* Generation time periods are for “traditional” planning in a vertically integrated utility planning within a least-cost monopoly-franchise regulatory structure. Generation in a competitive environment might be planned within a profit-payback structure with shorter time periods.

be multi-scenario. Working for that, the long-range planning steps of all four subsequent stages are merged into a single, *long-range T&D plan*, with appropriate and coordinated representation of all system levels.

The short-range planning periods shown in Table 13.6 are recommended for normal utility circumstances. The long-range planning at each level begins at that short-range lead time and extends through a period at least twice and as far as four times farther into the future. *In particular, if planners expect their results to address aging concerns and work toward solving them, planning must include a long-range element. Aging and its impacts are long-term issues and they can only be addressed through a long-term view. A short-term focus inevitably results in deferral of steps needed to reverse, stall, or at least mitigate age-driven trends.*

Spatial Load Forecast

The planning process is driven by a recognition that future customer demands for power, reliability, and services may differ from those at the present. The load forecast projects future customer demand and reliability requirements, which define the minimum capability targets that the future system must meet in order to satisfy the customer demands. Forecasting for T&D planning requires a spatial forecast, as was illustrated in Figure 13.6.

Spatial forecasts: critical to aging infrastructure planning, particularly when on a tight budget

Chapter 8's discussion of the problems traditional contingency-based planning encounters in dealing with modern utility systems demonstrated how one aspect of modern system interaction is a need for broader areas of contingency support. One result of this is that the interaction of capacity/load ratios in one area of a system, with that in other areas, is of greater importance in the overall performance of the power system. When a power system is loaded to modern (high) equipment utilization rates, and composed of aged (failure-prone) equipment, the reliability of service in one area of the system is sensitive to the quality of the planning done in nearby areas. All need each other for support, and a mistake in one area affects not just that area (as it did traditionally) but all its neighbors. The cost of making an error in matching capacity to peak load is much greater in contemporary power systems than it was in traditional power systems.

Spatial forecasts have always been an important element of good power delivery planning. But they take on an even more critical importance because of these aging infrastructure issues. The forecast of location of customer demand is as important in T&D planning as the forecasting of the amount of load growth: both are required in order to determine where facilities need to be located and what capacity they must have.

- *Timing.* Usually, spatial forecasts project annual figures (peak, etc.)

over a period into the future or from three to five times the lead time of the equipment being planned – a spatial forecast for substation planning (five year lead time) would be carried out through 15 to 25 years into the future.

- *Spatial resolution.* Spatial forecasts are accomplished using the small area technique (Figure 13.10), dividing the utility service area into hundreds or perhaps tens of thousands of small areas and forecasting the load in each. As a general rule of thumb, the small areas must be no larger than about 1/10th the size of a typical equipment service area at that level, in order for the forecast to have sufficient “where” information for planning. Thus, spatial resolution needs are less at higher voltage levels – transmission can be planned with “small” areas that are perhaps 25 square miles in size, whereas substation planning must be done on a mile basis and distribution on a 40 acre (1/16 square mile) basis.
- *Quantities forecast.* A T&D forecast projects the peak demand (usually with respect to very specific weather called “design conditions”) at every locality, since these local peaks are what the T&D system must satisfy. In addition, simulation methods forecast customer type (residential, commercial, and industrial) and annual energy requirements (kWh).




Figure 13.10 A load forecast for T&D planning is accomplished by dividing the service territory into a number of small areas – either equipment service areas or the cells of a square grid – and forecasting customer demand in each.

Table 13.7 The Spatial Load Forecast Step

Purpose:	To provide forecast of power delivery requirements for T&D planning
Timing:	Selected years to cover the period from 1 to 25 years ahead, usually 1, 2, 3, 5, 7, 10, 15, and 20 years ahead.
Products:	Small area forecast by year of peak demand, annual kWh, customer types, and other factors (e.g., end-use, reliability)
Coordinated with:	Corporate forecast, customer-level planning (DSM impacts forecast). Marketing plan. Critical input to distribution plan.
Tools used:	Spatial load forecast method (simulation or trending). End-use analysis of load/customer value. Corporate forecast database.
Methods available include:	
• Regression-based extrapolation is good only for short-range forecasting and proven less than ideally accurate if applicable, its chief advantage is low cost and low data requirements. Not recommended. • Pattern-based extrapolation utilized non-algebraic numerical methods to extrapolate future load based on repetition of past trends. Complicated math, but more accurate than regression. • Customer-class based end-use location simulation (land-use based simulation). Very good representation and suitability to T&D planning. Good accuracy if well applied. Requires special skills, data, and is rather data intensive. • Hybrid simulation/trending methods. These combine the accuracy and representation of simulation with lower data and skill requirements of trending.	

Table 13.7 summarizes key points and methods for spatial forecasting. The book *Spatial Electric Load Forecasting – Second Edition* (H. L. Willis, Marcel Dekker, 2002) presents details and practical advice on spatial forecasting, and reviews a number of popular spatial forecasting techniques. Some spatial forecast methods include an end-use based small area analysis, in which usage of electric energy by end-use (heating, water heating, lighting, etc.) is used to provide estimates for marketing and demand-side planning. Advanced methods may use end-use analysis to estimate the customer value of reliability of service.

Cumulatively, these forecast characteristics of load location, amount, customer type and number, reliability, and service requirements set goals that the future system must meet. If the existing system cannot meet them, then additions or changes must be made to upgrade its capability.

Spatial forecast coordination with corporate strategic planning

There are a number of different analytical approaches and computer algorithms that are applied to spatial forecasts for T&D planning. What is best for any given situation depends on planning needs, data, and other details specific to a utility. But by far the most important aspect of the spatial load forecast, and good T&D planning, is that it be consistent with the corporate, or “Rate Department” forecast, which projects sales and revenues for the utility as a whole. The T&D forecast must be based upon the same assumptions and forecast as the corporate plans.

Strategic Planning

No long term goals can be dependably achieved without planning on when, where and how they will be achieved. Utility strategic planning incorporates three aspects that directly interact with and even dictate elements of the T&D plan:

- *Budget.* The utility’s business and corporate planning sets spending limits and constraints. These are passed on to the T&D planning process through the budget approval process and have been a part of traditional T&D planning for generations.
- *Integrated Resource Plans.* The utility’s plan with respect to the use of distributed renewable energy, energy efficiency, demand response, and market-driven incentive plans, along with Smart systems and energy storage heavily influences what is needed from the T&D system.
- *Aging.* Aging is a long term effect and one that will affect not only utility spending and performance, but also the characteristics of that spending and performance (how much is capital, how much is O&M; how SAIFI and SAIF vary from circuit to circuit and their volatility year to year).

T&D planning needs to not only acknowledge and comply with all three elements of the utilities strategy, but ideally be *driven* by the strategic plan. Conversely, the utility must include all three of these elements, with respect to its system, in its corporate strategy. Essentially it must develop a strategy in these three areas of delivery system concern, if it is to build a base for comprehensive and effective planning – and hence effective use of its investment.

In the context discussed here, corporate strategy is not a spatial planning

process. Costs, profits, and risks are not spatial, because money is fungible: a dollar made in the downtown area is just as valuable as a dollar made in rural areas.⁵ As a result the corporate plan and strategy will be unspecific as to locations and interactions with aging areas of the system. Its preparation will need to address the fact that there are aging areas of the system, etc., but locational information is not needed and not used.

This lack of need for spatial detail eases the strategic planning process. As will be shown in examples later in Chapter 16's example and in the Appendices, "big picture" planning has to deal with the fact that the utility has so many old units, etc. Detailed implementation of that strategy *through* the T&D plan, however, has to both deal with specific areas and equipments, and tie the way it deals with them back to those elements of the plan and those required for expansion and other purposes. Lack of coordination so tight it could be called integration of expansion planning and aging-equipment refurbishment and replacement programs, will lead to noticeable inefficiencies in the use of capital.

Transmission Planning

In modern (deregulated) utilities, "transmission" has two very distinct meanings applied to two very different parts of the power system:

Wholesale grid: The regional grid, which interconnects generating plants and key bulk switching stations around the region, and connects the regional power system to other regions, is composed of high-capacity transmission lines, typically of voltages between 230 to 765 kV. This is the portion of the system that creates the wholesale market in electric power, and assures open access to merchant generators and buyers alike. System security, rather than reliability, is really the key factor in assuring "dependability" here.

Sub-transmission: The term sub-transmission has been used for decades to designate those transmission voltage lines whose purpose is to deliver power to distribution substations. This is a mission quite distinction from that of the wholesale grid. Sub-transmission lines are generally of between 35 and 138 kV, although some 230 kV lines exist primarily for this purpose.

The regional wholesale transmission grid is a key element in a healthy and economical regionally electric industry – it creates or enables the open access

⁵ In utilities with service areas that overlap multiple regulatory jurisdictions, such as those that serve customers in several states, strategic and corporate plans do have to take the variations in state regulation and policy into account. In this respect strategic planning can become spatial.

competitive market for electric power. The wholesale grid is usually planned and operated by an independent regional authority somewhat distinct, and often at arms-length from, the distribution utilities. Only part of the criteria and goals at this level relates to delivering power. Without doubt, one purpose of the wholesale grid is to deliver the required power to customers, but in addition, the transmission system is planned to provide:

1. *System Security and Stability* – the ability of the system to maintain interconnected integrity subsequent to any reasonably likely failure or event, and to react to transient events in a smooth manner.
2. *Generation access (deregulation) or generation dispatch (traditional)* – provide sufficient electric tie strength between generation plants and the system to assure all plants can truly access the system – that transmission constraints do not limit choice of what generation can be run.

These two goals mean that a good deal of the transmission system is planned from the perspective of generation: permitting generation access the system, ensuring system security no matter what happens, and providing stability against major disturbances. These aspects of transmission planning are not directly concerned with meeting power delivery needs, but with assuring that the generation-grid level of the system provides a stable bulk transmission capability and open access opportunity for all.

Planning of the “generation-grid” combination, which includes most of the high voltage transmission, is an activity quite apart from what might be termed “sub-transmission” planning – the planning of power transfer capability to the distribution substations, which is the only type of planning relevant to power delivery planning.

In general, sub-transmission planning reacts only slightly with grid planning – the grid is planned and presented as a “given” to the sub-transmission planning process. This is particularly true because of the “Chinese wall” that often has to exist between the wholesale (open access) and retail (local Distribution Company) under deregulation.

Thus, the interaction of the planning processes at the transmission-grid/sub-transmission-connection, has the most limited “coordination” of any of the junctions between adjacent levels of the power system. Sub-transmission is part of the power distribution planning process. It is best done as part of the distribution planning process. The sub-transmission – substation – feeder system ought to be planned as one system.

Figure 13.9 shows the recommended relation of sub-transmission planning to

the other stages of utility system planning. Key aspects of these relationships include:

- *Transmission encompasses two planning functions:* as mentioned above, the “generation-grid” planning which focuses on establishing a strong grid with security and dispatch freedom and meeting all power pool requirements, and the “sub- transmission” planning, which includes all planning activities related to moving power to substations so that it can be distributed to customers. Only “sub-transmission planning” is overtly a power delivery planning function.
- *Forecast of future load for transmission planning* comes through the distribution planning process. Fully half of the statistical variation in future loads due to “load growth” is a function of distribution planning (Willis and Powell, 1985). To the transmission-sub-transmission planning process, the “loads” to be served by the system being planned are the high-side loads of the distribution substations. How much load is collected to a particular substation depends on the distribution planning – whether load growth ends up at one substation, or another, or is split between them is determined by the distribution planning.

This “high side substation load forecast” is best produced by starting with the spatial forecast, assessing the preferred patterns of distribution substation service areas in the distribution planning/substation planning stages, and then passing the target substation loads to the transmission planning process. If particular loads prove difficult to serve, transmission planning can provide feedback to distribution planning, asking if alternate possibilities can be considered.

- *Coordination with substation planning.* The transmission system must reach every substation. Therefore, sub-transmission planning needs to be coordinated with the siting, capacity planning, and scheduling of substations (which of course has to be coordinated with distribution system planning, as shown in Figure 13.10). This coordination involves planning and scheduling of the sub-transmission and substation levels in which the cost impacts among the levels are “traded” back and forth. In many utilities, all these functions are combined into one “sub-transmission level” planning process. Given that this is *coordinated* well with distribution planning, it is a good approach.

Table 13.8 summarizes the important elements of this part of the sub-transmission planning process.

Table 13.8 Sub-Transmission Planning

Purpose:	short and long-range planning of facilities to deliver power to the distribution substations.
Timing:	short-range – 3 to 7 years; long-range – 5 to 20 years.
Products:	short-range – sub-transmission project schedule coordinated with substation and transmission grid plans. Long-range – a base sub-transmission plan.
Coordination with:	“transmission system” planning (EHV, grid, generation), substation-level planning (must deliver power to substations).
Tools used:	transmission network load flow and short circuit, route optimization software, combined sub-transmission-substation optimization programs.

Substation Planning

In many ways, the substation level is the most important level of the system for two reasons. To a great extent this level fashions the character of the whole system. Substations are where the transmission and the primary feeder levels meet, the points to which the sub-transmission must deliver power, and from which the feeders must take it to neighborhoods throughout the system. The substation levels defines how and where cost, contingency, or capacity at one level can be traded or played against the other, and how expansion of both can be coordinated.

The substations’ sites set very definite constraints on the design and routing of the transmission and distribution systems – transmission lines must “end” at the substations and feeders must begin there. Substation location heavily impacts the cost and performance, and particularly the available design freedom, at both T and D levels. Substation capacity limits also heavily influence both transmission (how much incoming capacity is needed) and distribution (what amount of power can be distributed out of a substation).

As a result, regardless of whether substation planning is done as a separate planning function (about 20% of utilities), or as part of transmission planning (60%), or as part of distribution planning (20%), it should be interactively coordinated with both sub-transmission and feeder system planning. This means that communication should be explicit and frequent enough that glaring differences in marginal costs of expansion at the T, S, or D levels will come to the attention of the planning establishment in the utility and be dealt with by appropriate changes to the plan.




Figure 13.11 Residential and commercial growth in the shaded area is expected to add 31 MVA over the next decade. Either Edgewood or Kerr substation could be expanded at about the same cost to serve this growth. However, expanding the feeder system out of Kerr would be expensive – present feeder getaways are saturated, the system is underground and many duct banks are congested, and several geographic barriers make feeder runs quite lengthy. By contrast Edgewood’s all-overhead feeder system would be straightforward to expand, costing \$1,900,000 less. Despite this, the *overall* optimal alternative is to expand Kerr, because transmission limitations at Edgewood would cost more to overcome than their savings at the distribution level.

For example, in Figure 13.11 distribution planners have made plans to expand Edgewood substation from 55 to 90 MVA capacity in order to serve a peak load that is expected to grow from 49 to 80 MVA due to load growth in the area indicated. Both Kerr and Edgewood substations are close enough to the load growth to pick up this new load, and either could be expanded with roughly the same plan and cost – addition of another 27 MVA transformer and associated equipment, at a cost of \$750,000 PW.

However, from the perspective of the distribution system, expanding the feeder system out of Kerr substation instead of Edgewood would cost an extra \$1,900,000 PW. Thus, Distribution Planning selected Edgewood for addition of new feeders and equipment to serve the load growth. But suppose that expanding transmission delivery capability at Edgewood by the required 31 MVA will require extensive reinforcement of towers and conductor, at a cost of \$2,600,000 PW, whereas the changes needed at Kerr to accommodate a similar increase in peak load are minor and would run only \$325,000 PW. Then it is best for the utility to decide to serve the new load out of Kerr, since even though Edgewood is the least expensive distribution option, when overall T-S-D cost is considered, it comes to \$600,000 PW less to serve the load growth from Kerr.

The planning process at many utilities can not accommodate this very simple, common sense approach to cost minimization. The procedural “hand-offs” between transmission, substation, and distribution planning groups does not convey enough information, nor allow for sufficient back-and-forth cooperative communication, so that cost differences like this are noted and the planning process can recognize if and how to adjust to them. At many utilities, the distribution planners will pass no details other than the projected total load for a substation – “We need 80 MVA delivered to Edgewood.” At others the reverse happens, with transmission planners setting limits on substation capacity due to transmission constraints, without information of whether these could be eased if economically justifiable based on distribution costs.

Substation planning is important chiefly because of this requirement, that as the meeting place of transmission and distribution it defines their cost interaction. Planning procedures at this level ought to reflect this relationship and foster the required communication. There are at least three possible avenues to do so.

1. *Planning meetings* between transmission and distribution planning can be scheduled frequently and set up so that a free exchange of planning details, not just planning results, occurs. The authors are well aware of the institutional barriers that exist within many large utilities to this type of communication. In the case of one large investor-owned utility, the managers of the transmission and distribution groups, both of whom had been in their positions for over four years, had never met and had exchanged only a handful of formal “form reports” during their respective tenures. Under deregulation, there could well be formal procedures defining the limits of and requiring such cooperation. Regardless, the results of such communication are significant savings, and worth considerable cultural adjustment. Situations like the Kerr-Edgewood example occur more often than is recognized.
2. *Use of methods and software* for the sub-transmission, substation, and distribution planning functions that accepts marginal or incremental cost data for the other level, or for the substation level that accepts such data for both. For example, some service-area based optimization algorithms applied to substation-distribution system planning can accept transmission system cost curves for incoming capacity, cost which is included in the optimization.
3. *A system of common planning procedures and nomenclature* that sets typical target values for every level of the system and reports the incremental cost of all projects at any level, to all planners at other levels in routine, periodic reports. This would pass on to all the distribution,

Table 13.9 Substation Planning

Purpose:	short and long-range planning of facilities to control, transform and route power from T to D. Consideration of substation located DG, energy storage, and Smart controls.
Timing:	short-range – 3 to 7 years; long-range – 5 to 20 years.
Products:	short-range: substation project schedule coordinated with sub-transmission and distribution plans; long-range - long-range plan.
Coordination with:	sub-transmission planning, distribution planning, sustainability plan, Smart systems/technology strategy.
Tools used:	AM/FM-GIS systems, substation selection optimization applications, sub-transmission-substation-feeder system optimization applications.

substation, and sub-transmission planners information for their use on the costs at other levels of the system. This provides the information, but not the incentive, for cooperation.

Frankly, application of all three of the above measures is best so as to lead to the most coordinated planning, which will minimize the potential for planning errors. Table 13.9 summarizes substation planning.

Feeder Planning

There are two very different aspects of “feeder planning.” The first is feeder system planning (Table 13.10), which involves the planning of the overall distribution feeder system in conjunction with the substations. The second, feeder planning, which involves the detailed layout of feeders and their specification is to a detailed engineering level of itemization and precision (Table 13.11).

Feeder system planning

What might be called “strategic” feeder planning involves determining the overall character of the feeder system associated with each substation, and with the role it plays in assisting in inter-substation contingency support. This includes determining the number of feeders from the substation, their voltage level, the overall service area served by the substation as a whole and the general assigned routes and service areas and peak loads for each. Additionally, operating problems (voltage, reliability, contingency switching) and overall costs may be estimated (budgetary).

Table 13.10 Feeder System Planning

Purpose:	Long-range planning of the feeder system, mainly so that feeder impact of substation-level decisions is assessed for substation planning.
Timing:	Five to 20 years.
Products:	Evaluation of feeder-level cost and performance impact of all substation-level decisions; long-range feeder level cost estimates for budgeting. Integration of Smart technology into local delivery.
Coordination with:	Substation planning, customer-level planning.
Tools used:	Multi-feeder (multi-substation) optimization programs, combined substation-feeder system optimization application for capacity, reliability, economics.

Generally, *feeder system* planning is 6 to 20 years ahead – well beyond any lead time for feeders but with considerably less detail than needed for short-range feeder planning. It is *only* a long-range and strategic planning activity, but it is an essential one. Its purpose is to provide “distribution system” feedback to planning decisions or alternatives being considered at the substation and sub-transmission levels and assess interactions of those levels.

As was mentioned earlier, the interplay of sub-transmission, substation, and feeder costs can be significant. In particular, the value of the primary distribution system associated with any particular substation probably outweighs the value of the substation itself. The 30-year PW of feeder losses on a heavily loaded system may exceed the value of both. Substation planning, particularly siting and sizing, can often be the economic equivalent of the tail wagging the dog. Small changes in location or planned capability can affect the cost of the feeder system dramatically and in some cases, a shift of one mile in the location of a suburban substation can have a PW impact of over \$1,000,000.

Feeder system planning is often called multi-feeder planning to distinguish it from feeder planning which includes functions that can be done on a feeder by feeder basis. To be of maximum impact, multi-feeder planning should cover multiple substations, assessing how the feeder system interacts with substation capacities and costs – optimizing substation service areas while considering both substation and feeder costs (Figure 13.12). Proven techniques for simultaneous substation-feeder system planning exist and are described elsewhere in greater detail (Willis,2002).

Feeder System Planning is a Key Element of Reliability Planning for Aging Infrastructures

Chapter 8 highlighted the need to use the feeder system to provide load transfers and contingency ties between substations. Feeder system planning, aimed at making certain that sets of feeders associated with each substation can accomplish the required target levels for transfers under normal and contingency situations, is essential to assure dependable backup capacity in an integrated T&D system plan. Overall policy decisions on what type of switching support to have both within each substation area, and for transfer of loads between substation areas, is a multi-feeder design aspect. While some single-feeder planning has to be done as a follow-up to the decisions made at this stage, assurance of reliability via good contingency withstand and support capability at the feeder level is done on a system basis.




Figure 13.12 Multi-feeder system planning evaluates the performance of the entire feeder system associated with a set of substations. Its goal is to capture completely the interaction of substation level and feeder level constraints and costs to determine the overall substation-feeder optimum. A big part of this boils down to assigning “optimal service areas” to all substations while acknowledging capital and operating costs and electrical and geographic constraints at both the substation and feeder levels simultaneously, while determining the boundaries (dotted lines) between the substations.

Short-range feeder planning

Short-range feeder planning (Table 13.11) involves determining in detail the routing of each feeder and identifying its equipment to near engineering precision. In many utilities, feeder planning and feeder engineering are merged into one function – “Distribution Planning” produces completed feeder plans replete with detailed equipment specifications and route maps, pole locations, and all construction details. It also provides the final authorization and scheduling for construction.

Generally, feeder planning *can* be done on “a feeder at a time” basis, often using only load-flow and short-circuit analysis tools augmented by engineering judgment to determine alternative designs in a “trial and error” approach to find the best feeder layout. However, many modern CAD systems for feeder planning have load-flow and short-circuit applications that can be run on a number of feeders at once, and augment these with optimization utilities to help refine load balance and switching plan development among feeders. Feeder planning is done differently at every distribution utility, but there are standard recommended procedures that are well documented (Engel, et al. 1992; Burke, 1994; Willis, 2004). Table 13.11 summarizes its essential elements.

Table 13.11 Feeder Planning

Purpose:	Short-range planning of the feeder system, to produce project definitions and authorization for feeder additions and enhancements.
Timing:	One to five years.
Products:	Feeder system project specifications, schedule, and budget, accommodation costs for customer DG, etc.
Coordination with:	Feeder system planning, substation planning, customer planning, DG and DR planning, sustainability plan, and construction.
Tools used:	8,760-hour load models. “Feeder design” or CAD systems, feeder optimization programs (either single or multi-feeder), AM/FM and GIS systems. Reliability “load flows” for the design of optimized switching locations and branch and trunk tie capability. Possibly “shotgun” load flow macros that run 8,760-hour analysis on the year with and without DG & DR interactions.

Customer-Level Planning

Traditional supply-side-only planning, through the first three quarters of the 20th century, included very little if any “customer-level” planning functions. Regulated utility IRP (integrated resource planning) of the type widely practiced in the late 1980s and early 1990s included a great deal of customer side measures to control load and efficiency – various DSM (demand side management) options were optimized on some type of societal resource-value basis, (see Willis and Rackliffe, 1994).

In the 21st century, “DSM” has expanded to include not just load-related adjustments and energy efficiency measures, but customer-owned generation and energy storage. It also includes much more emphasis on free-market driven incentives. Programs and policies in the 20th century envisioned utility control of load reductions. Their equivalents in the 21st century envision customer control and choice responding to real time and other market-driven incentives. Regardless, it appears almost certain that IRP will become the mainstream, expected approach for the industry.

For this reason, “customer-level planning” along with full consideration of customer-side resources and customer operation of them, can be expected to be an aspect in all power delivery planning. Table 13.12 summarizes key points in planning at the customer level. The interaction may be much more than just oblique, because customer actions with regard to load reduction, load scheduling, energy efficiency improvements, renewable generation, and energy storage can have a very substantial impact on T&D, since they affect the demand, the energy use, and the reliability needs the system must meet.

Customer-level planning often has a very long planning period, much longer than distribution planning, and often as long as that of generation. There are four reasons why both its short- and long-range planning periods are longer than almost all other levels of the system:

1. *Lead time.* While DSM programs can often be approved and implemented nearly immediately, it can take years of slow progress to obtain meaningful levels of participation – rates of 50% may take a decade or more to obtain. This “ramp up time” has the same effect as a long lead time.
2. *Customer commitments.* In starting a DSM program, a utility is beginning a very visible customer-interaction. Such programs are not entered into lightly and are often studied over the very long term to make certain the utility can feel confident it will commit to them for the long-term. Starting such a program and then canceling it several years later is not only costly, but leads to customer confusion and lack of confidence in the utility.

3. *Proxy for generation.* Customer side resources are often traded against generation in integrated resource planning. Comparison of the two resources needs to be done over the time span used for planning of generation, which means a present-worth or similar economic comparison over a 20-to-30 year period.
4. *Uncertainty of resource.* Even when the utility has absolute control of major customer loads such as space heating and cooling, there is still a small modicum of uncertainty about the available load reduction at any time, due to the variation in loads that might be on at that time. Taking such DSM measures, and the uncertainty in them fully into account is a normal part of good power delivery and energy resource planning and designing systems to achieve performance goals such as SAIFI targets and voltage regulation criteria.

But there is much more uncertainty about the availability of the customer side resources in an industry where the utility's customers can respond to real-time price signals, and can or cannot operate their generation and storage as they see fit. Customers may or may not maintain their renewable generation over time. They may prefer to operate it differently than the utility expects. At peak periods customers will probably respond to price signals as expected, but perhaps they will not if it is an unusual situation. Even customers who originally buy energy storage systems as a reliability augmentation resource may begin using them, over time, to reduce cost/earn money via daily arbitraging of energy prices, changing the pattern of customer response the utility grew used to seeing on a daily basis.

Planners simply do not know the details of these customer-side response issues with the specificity that they know, say, the expected capacity of a feeder or a utility owned DG unit. That uncertainty needs to be taken into account in their planning of the delivery system's capability and performance. It means achievement of traditional SAIFI targets, voltage regulation, and loading guidelines faces much more volatility. Where to put "a stake in the ground" with respect to the needed amount of power delivery capability is much more difficult.

Here, important aspects of regulatory policy and industry expectations have not yet been fully identified. Since the utility's customers are now resource owners and operators, they can be viewed as partly responsible for performance lapses in the system's ability to satisfy their energy needs. Does a customer who owns an energy storage system capable of powering his home or business for three hours have any right to complain that his lights went out when the local feeder was out of service for two hours? Does a homeowner who fails to maintain her 15 kW of PV panels so they are available during peak demand

periods have a right to complain that voltage at her service entrance dropped to 108 volts? One can argue either side of these questions, and people will, because there is validity in some of the arguments on either side. Different solutions set different precedents and have far different implications on costs and benefits, particularly as to how the cost of the entire utility system is allocated to energy users via the utility's rates. Regardless, eventually, policy and expectations will harden around a set of guidelines and expectations that the industry pretty much follows, but at present they are not yet fully formed. And regardless of what those are and how they are implemented, planning for customer-side performance targets and criteria will need to be done probabilistically, and planning for the attainment of the utility's own business performance goals and targets will be messier, face more volatility, and require more comprehensive risk-mitigation methods than in the past.

Customer-level planning generally begins by adding more detailed, end-use and customer value attributes to the load forecast – forecasts by end-use and appliance category and of reliability value functions. It then tries to optimize PW cost of providing end-uses to the customers by juggling various DSM and T&D options. Generally, optimization methods function by balancing marginal cost of the various customer-side options against marginal cost curves obtained for the supply side. DSM and DG evaluation may also be part of value-based planning in which the utility and customer value functions are jointly optimized.

Table 13.12 Customer-Level Planning

Purpose:	Short- and long-range planning of customer-level resources including DSM, DG, DS, RTP and TOU rates.
Timing:	One to 25 years.
Products:	Schedule of customer-side resource projects/programs; long-range plan of DSM/DG/DS targets.
Coordination with:	All levels of planning, corporate forecast/rate plan, spatial load forecast. Note: this planning is often done by corporate or rate department planners.
Tools used:	End-use load models. Integrated resource (T&D or otherwise) optimization. Customer response models. Value-based planning analysis. Eco-demographic models. Spatial customer/load forecast models.

13.5 THE SYSTEMS APPROACH

Throughout this chapter, discussion of T&D planning has stressed the need to coordinate the planning of the various levels of the power delivery (T&D) system. The various levels – transmission, substation, and distribution – are interconnected, with the distribution, in turn, connected to the customers (and hence a function of the local load density). The best equipment types, sizes, and designs at each level of the system are a function of the types of equipment, sizes, and designs selected for the other levels of the system. In general, the equipment is so interconnected that it is impossible to evaluate any one aspect of design without taking all others into account.

Consider the question of substation spacing – determining how far apart substations should be, on average, for best utilization and economy. Within any service territory, if substations are located farther apart, there will be fewer of them, saving the utility the cost of buying and preparing some substation sites, as well as reducing the cost of building a large number of substations. Thus, overall substation cost is a function of spacing, but how it varies depends on land and equipment costs, which must be balanced carefully. With fewer substations, however, each substation must serve a larger area of the system. Hence, each substation will have a larger load, and thus require a larger capacity, meaning it must have more or larger transformers.

The larger substations will also require a larger amount of power to be brought to each one, which generally calls for a higher sub-transmission voltage. Yet, there will be fewer sub-transmission lines required (because there are fewer substations to which power must be delivered). All three aspects of layout are related – greater substation spacing calls for larger substations with bigger transformers, and a higher transmission voltage, but fewer lines are needed – and all three create better economies of scale as spacing is increased. Thus, transmission costs generally drop as substation spacing is increased.

Nevertheless, there is a limit to this economy. The distribution feeder system is required to distribute each substation's power through its service area, moving power out to the boundary between each substation's service area and that of its neighbors. Moving substations farther apart means that the distribution system must move power, on average, a greater distance. Distributing power over these longer distances requires longer and more heavily loaded feeders. This situation increases voltage drop and can produce higher losses, all of which can increase cost considerably. Employing a higher distribution voltage (such as 23.9 kV instead of 13.8 kV) improves performance and economy, but regardless, it costs more to distribute power from a few larger substations farther apart, than to distribute it from many smaller substations close together. Feeder costs go up rapidly as substation spacing is increased.

The major point of this section is that all four of the above aspects of system

design are interconnected: (1) substation spacing in the system; (2) size and number of substations; (3) transmission voltage and design; and (4) distribution feeder voltage and design. One of these factors cannot be optimized without close evaluation of its interrelationship with the other three. Therefore, determining the most cost-effective design guideline involves evaluating the transmission-substation-feeder system design as a whole against the load pattern, and selecting the best combination of transmission voltage, substation transformer sizes, substation spacing, and feeder system voltage and layout.

This economic equipment sizing and layout determination is based on achieving a balance between two conflicting cost relationships:

1. Higher voltage equipment is nearly always more economical on a per-MW basis.
2. Higher voltage equipment is available only in large sizes (lots of MW).

In cases where the local area demands are modest, higher voltage equipment may be more expensive simply because the minimum size is far above what is required – the utility has to buy more than it needs. How these two cost relationships play against one another depends on several factors: the load, the distances over which power must be delivered, and other factors unique to each power system such as the voltages at which power is delivered from the regional power pool, and whether the system is underground or overhead.

Figure 13.13 illustrates the difference that careful coordination of system design between levels of the power system can have in lowering overall cost. Shown are the overall costs from various combinations of T&D system layout for a large metropolitan utility in the eastern United States, all designed to achieve the same reliability targets. Each line connects a set of cost computations for a system built with the same transmission and distribution voltages (e.g., 161 kV transmission and 13.8 kV distribution) but varying in substation sizes (and hence, implicitly, their spacing).

In all cases, the utility had determined it would build each substation with two equally sized transformers (for reliability), with none over 75 MVA (larger transformers are too difficult to move along normal roads and streets, even on special trailers). Either 161 kV or 69 kV could be used as sub-transmission, either 23.9 kV or 13.8 kV could be used as distribution voltage. Any size transformer, from 15 MVA to 75 MVA, could be used, meaning the substation could vary from 30 MVA to 150 MVA in size. (Peak load of such substations can normally be up to 75% of capacity, for a peak load of from 23 to 100 MW.) Substation spacing itself is implicit and not shown. Given the requirement to cover the system, determining transmission voltage, distribution, and substation size defines the system design guidelines entirely.

Overall, the lowest cost T&D system guidelines are to build a 120 MVA, two-




Figure 13.13 Overall cost of T&D system depends on the balanced design of the sub-transmission, substation, and feeder level, as described in the text. Cost can vary by significant margins depending on how well performance and economy at various levels of the system are coordinated.

transformer (two 60 MVA units) substation fed by 161 kV sub-transmission and distributing power at 23.9 kV. This has a levelized cost (as computed for this utility) of about \$179/kW. In this particular case, a high distribution voltage is perhaps the most important key to good economy – if 13.8 kV is used instead of 23.9 kV as the primary voltage, minimum achievable cost rises to \$193/kW.

The very worst design choices plotted in Figure 13.13, from an economic standpoint, would be to build 25 MVA substations fed by 161 kV sub-transmission and feeding power to 23.9 kV feeders (\$292/kW). This would require many small substations, each below the effective size of both the transmission and distribution voltages, and lead to high costs. Overall, 161 kV and 23.9 kV are the correct choices for economy, but only if used in conjunction with a few, large substations. If substations are to be 25 MVA, then 69 kV and 13.8 kV do a much more economical job (\$228/kW), but still don't achieve anything like the optimum value. The point is that achieving economy in power delivery involves coordinating the interactions, performance, and economies of the multiple system levels.

The Systems Approach Is Important for Aging Infrastructure Planning

The discussion and figure above showed the economic advantages that the systems approach can provide. These are certainly important to all utilities, including those facing aging infrastructure issues. However, there is another, and vital reason to apply the systems approach as a rule in all planning for aging T&D infrastructures:

The systems approach helps work around constraints in siting, size, or topology by drawing on the resources of the other levels in the power system to solve problems present at another.

As mentioned in Chapter 8, very often the owner of the aging T&D system is “stuck” with too few and too small substation sites. Sub-transmission rights of way may be narrow and non-optimal in routing. Expansion and addition may be costly, time consuming, and politically unacceptable.

Regardless, in many cases the utility has little choice but to make the best of these assets, particularly in aging areas of its system where additional land and rights of way are very expensive. The best way to improve performance is to seek out strengths in adjacent levels of the system in order to help support the weakness in another. Techniques discussed in Chapters 10, 11, and 14 can be applied to do just this, as was summarized in Section 13.5. More detailed discussions of the “systems approach” necessary to optimize the synergism among the sub-transmission – substation – feeder system levels to truly maximize the reliability/cost ratio of the entire power system chain are contained in the Willis (2004).

13.6 SUMMARY OF PLANNING IN AN AGING T&D INFRASTRUCTURE

Planning involves the selection of the best alternative from among those at hand. In the power industry, that selection depends to a certain extent on the values of the group making the decision. *All* cost evaluations for power delivery planning are based in part upon economic evaluation that is consistent over the range of all alternatives. Usually, this means that present and future costs must be compared. Present worth analysis, levelized costing, or some similar method of putting present and future value on a comparable basis will be used. While methods and standards vary from one company to another, the important point is that planners apply their company’s method consistently and correctly.

Traditionally, planning at many electric utilities has been something like a

rhinoceros – never looking far ahead but instead focusing on only what is right in front of it, protective of its own turf, capable of surprising speed and powerful enough to get its way much of the time, but frequently content to move slowly and merely stay where it is. Utilities with aging T&D infrastructures need to revise their planning focus, taking the following lessons to heart:

- *Good planning* is more important than ever. A system that has weaknesses or high equipment utilization ratios is especially intolerant of planning mistakes.
- *Optimization of reliability/cost* is a key element of planning for aging T&D infrastructures.
- *Constraints* in siting, size, and routing are usually the major challenge in planning.
- *Think of the long range, but act for the short range.* Money and resources should only be committed on the basis of good short-term benefits and efficacy, but every commitment should fit into a coherent long-range plan.
- *Spatial forecasting* is more important than in normal T&D planning, important as it is in that venue.
- *Feeder system planning*, and coordination of feeder system strength to provide contingency support to the substation level, is a key factor in simultaneously achieving reliability and economy.
- *The systems approach* is a vital element of success in handling constraints in siting, sizing, and routing.

REFERENCES AND BIBLIOGRAPHY

- D. Atanackovic, D. T. McGillis, and F. D. Galiana, “The Application of Multi-Criteria Analysis to a Substation Design,” paper presented at the 1997 IEEE Power Engineering Society Summer Meeting, Berlin.
- T. W. Berrie, *Electricity Economics and Planning*, Peter Peregrinus Ltd., London, 1992.
- J. J. Burke, *Power Distribution Engineering – Fundamentals and Applications*, Marcel Dekker, New York, 1994.
- M. V. Engel, et al., IEEE Tutorial on Distribution Planning, Institute of Electrical and Electronics Engineers, New York, 1992.

- R. Orans et al., "Targeting DSM for Transmission and Distribution Benefits," Electric Power Research Institute, Palo Alto, CA, 1992.
- F. Schweppe et al., *Spot Pricing of Electricity*, Kluwer Press, Cambridge, MA, 1988.
- L. J. Vogt and H. L. Willis, "DSM: Transmission and Distribution Impacts," Volumes 1 and 2, Report CU-6924, Electric Power Research Institute, Palo Alto, CA, 1990.
- H. L. Willis, *Power Distribution Planning Reference Book – Second Edition*, Marcel Dekker, New York, 2004.
- H. L. Willis and R. W. Powell, "Load Forecasting for Transmission Planning," *IEEE Transactions on Power Apparatus and Systems*, August 1985, p. 2550.
- H. L. Willis and G. B. Rackliffe, *Introduction to Integrated Resource T&D Planning*, ABB Guidebooks, Raleigh, NC, 1994.
- World Bank, *Guidelines for Marginal Cost Analysis of Power Systems*, Energy Department paper number 18, 1984.

14

Planning Reliability of an Aging System

14.1 INTRODUCTION

An aging infrastructure will have a detrimental effect on a utility's customer service reliability. Gradually, as equipment deteriorates and equipment failure rates increase, the rate of service interruptions will begin to increase. All measures of customer service reliability – SAIFI, SAIDI, MAIFI and others, will inevitably worsen. An old, partially worn-out power system is inherently less reliable, in terms of customer service quality, than a new one.

It is not possible to add enough Smart equipment and “self healing capability,” enough automated switching, and enough fault location and diagnostic equipment to guarantee this escalation in service interruptions and SAIDI and SAIFI problems can be completely controlled. A utility cannot “fix” the reliability problems created by aging by focusing only on technology alone.

Aging Actually Creates An Opportunity for More Reliability Improvement

As a system becomes inherently less reliable due to aging, the effectiveness of many reliability improvement methods *increases*. In a system where SAIFI is expected to be .66, certain automated switching and fault location programs may not provide a benefit/cost ratio of 1.00 or more. There is not enough opportunity, in such a reliable system, for them to justify their expense. But if the SAIFI is

1.33, they are about twice as cost effective. Therefore, as a system ages and failure rates increase, the cost effectiveness of various reliability improvement methods gradually rises. And as more is spent to renew equipment, the cost effectiveness of reliability improvement methods drops.

What a utility wants is the lowest cost way to deliver satisfactory service reliability to its customers. To the customer, it does not matter if that reliability comes from an a network composed of worn out equipment that is controlled smartly by an advanced monitoring and control system, or from a network that is inherently reliable because equipment condition is maintained at a high level. What is important is that: a) service is reliable, and b) cost is as low as possible. The industrial power owner seeks the same cost “sweet spot.”

Not surprisingly, in every case, the lowest cost, lowest risk plan is always a combination of both approaches, something that basically picks the low hanging fruit from each method in order to provide just enough reliability to satisfy needs. Thus, reliability engineering – planning, specification, and implementation of measures to improve and control customer reliability – is an important aspect of planning for aging infrastructures. The other part is the planning of the best inspection, tracking, condition improvement, life extension, proactive replacement plan. The two interact, creating synergy if used together well and problems if not, so they must be planned in conjunction with one another.

Reliability Engineering of Power Delivery

This chapter summarizes power delivery reliability concepts and methods.

The goal of reliability-based planning, engineering, and operations is to cost effectively maximize reliability of service as seen by the customer, not necessarily to improve reliability of equipment or service time on the system. Customer-level results are what count.

While equipment availability and similar measures are important, the planning and engineering process should focus on customer-level measures of service reliability.

Why is Reliability More Important Today Than It Was Traditionally?

Distribution system reliability is driven by several factors including:

- (1) The increasing sensitivity of customer loads to poor reliability, driven both by the increasing use of digital equipment, and changing lifestyles.
- (2) The importance of distribution systems to customer reliability as the final link to the customer. They, more than anything else,

shape service quality.

- (3) The large costs associated with distribution systems. Distribution is gradually becoming an increasing share of overall power system cost.
- (4) Regulatory implementation of performance based rates, and large-customer contracts that specify rebates for poor reliability, all give the utility a financial interest in improving service quality.

In the past, distribution system reliability was a by-product of standard design practices and reactive solutions to historical problems. In the future, distribution system reliability will be a competitive advantage for a distribution utility, in its competition with other forms of power and energy (DG, alternate fuels, and conservation). Reliability must be planned for, designed for, optimized, and treated with analytical rigor.

Optimal Reliability isn't Necessarily Maximum Reliability

Planners must keep in mind that while reliability is important to all customers, so is cost, and that only a portion of the customer base is willing to pay a premium price for premium levels of reliability. The real challenge for a distribution utility is, within tight cost constraints, to:

- Provide a good basic level of reliability.
- Provide roughly equal levels of reliability throughout its system, with no areas falling far below the norm.
- Provide the ability to implement means to improve reliability at designated localities or individual customer sites where greater service quality is needed and justified.

In Section 14.2, this chapter will introduce an important concept: the engineering of reliability in a distribution system in the same way that other performance aspects such as voltage profile, loading, and power factor is engineered. Section 14.3 will summarize the analytical methods that can be used as the basis for such engineering and discuss their application. Section 14.4 gives an example of reliability-based engineering as applied to a medium-sized power system, designed for overall and site-specific targeted reliability performance. Financial risk, due to performance-based contracts and rates, is a big concern to a 21st century utility. Section 14.5 looks at the analytical methods needed to assess and optimize this, and shows an example analysis. Section 14.6 concludes with a summary of key points.

14.2 RELIABILITY CAN BE ENGINEERED

Figure 14.1 shows a screen from a computer program that can simultaneously assess voltage behavior, reliability, or cost of an arbitrary distribution system design input by the user – what is sometimes called a “reliability load flow.” The program bases its computations on an arc-node structure model of the system, very similar in concept, construction, and data, to that used in a standard load flow program. By using failure rate data on system components, the program can compute the expected number of outages that occur, and the expected total minutes of outage, annually, at each and every point in the network. From these results, indices such as SAIDI, SAIFI, CAIDI and CAIFI, or other reliability indices, can be computed for the network. The particular program shown has several characteristics important to practical application:

- It represents the *impact of equipment failures* on the interconnected system around it.
- It can model the actions of *dynamic changes* associated with a contingency, such as a fuse blowing, a breaker tripping, a rollover switch operating, or a line crew repairing a damaged line section.




Figure 14.1 Screen display from a “reliability load flow” of the network in a small city. See text for details.

- It is *self-calibrating*. Given historical data on outage rates by portion of the system, it computes the failure rates for equipment in each area until its depiction of the base system's reliability performance matches historical reality.
- It has a very *large system capability*, being able to analysis a distribution system composed of hundreds of thousands of nodes.
- Its model of the system itself (the circuit database) is in a *standard* load flow format compatible with most power system modeling.
- It is *easy to use* (graphic user interface).

Results are dependable

The results from this process as dependable from an engineering standpoint as those obtained from a well-applied load flow application. Generally, the reliability assessment has a *better* base of customer data than a load flow. It uses customer count data at the "load end" of the model (information that can be accurately determined) not kW load data (which are most typically estimates based on kWh sales).

Table 14.1 Reliability of Typical Types of Distribution System Components

Equipment Type	<u>Failure Rate Per Year</u>			<u>Mean Time To Repair</u>		
	Low	Typical	High	Low	Typical	High
Overhead Lines (per mile)						
Primary Trunk	0.020	0.100	0.300	2.0	4.0	8.0
Lateral Tap	0.020	0.160	0.300	2.0	4.0	8.0
Secondary & Service Drop	0.020	0.088	0.030	1.5	2.3	8.0
Pole Mounted Transformer	0.004	0.010	0.015	3.0	4.0	10.0
Disconnect Switch		0.014		1.0	1.5	4.0
Fuse Cutout				0.5	1.5	4.0
Line Recloser				3.0	4.0	10.0
Shunt Capacitor				0.5	1.5	4.0
Voltage Regulator				3.0	4.0	10.0
Underground Cable						
Primary Cable	0.003	0.070	0.100	2.5	10.0	30
Secondary Cable	0.005	0.100	0.150	2.5	10.0	30
Elbows & Connectors	6e-5	0.0006	0.001	1.0	4.5	8.0
Cable Splices and Joints	6e-5	0.03	0.159	0.5	2.5	8.0
Padmount Transformers	0.001	0.010	0.049	4.0	6.5	7.8
Padmount Switches	0.001	0.003	0.005	0.8	2.5	5.0

Required component reliability data is not difficult to obtain

One of the primary concerns power system planners have with the use of a reliability-based engineering method is the data requirement. Such methods need data on the failure rates of various types of equipment, and repair and switching times for various situations. Generally, this data is available by culling through utility operating records. However, several sources of generic reliability data are available. These include default data obtained from industry standards (for example, Institute of Electrical and Electronics Engineers, International Conference on Large High Voltage Electric Systems), publications (such as books, journals, conference papers), and tests or surveys conducted by manufacturers. This information provides a good first guess and prepares the model to be calibrated against historical data. Table 14.1 lists typical values of failure rates for most common types of power system equipment.

As mentioned above, the best reliability-assessment programs have a self-calibrating feature: they “fit” their reliability analysis results to actual historical outage rates. These methods start with “generic” failure rate data in internal tables and make various adjustments and re-computations of what failure rates in an area must be, based on the historical data. Self-calibration is the best way to assure that a reliability-assessment model is accurate.

Results are useful even with only generic data

It is important to realize that a reliability-assessment method can help utility planning engineers improve the reliability performance of their system even if it is using approximate failure rate data. Most of the important decisions that are made with respect to reliability have to do with configuration, installation and location of protective and sectionalizing equipment, and switching. Even approximate data, used in accurate analysis of these issues, provides guidance on the wisdom of various decisions. Improvement can be made using only approximate (generic) failure rate and repair-time data.

Using a Reliability Load Flow

Planning engineers use a reliability-assessment program in much the same manner that they use a load flow to design the feeder system. They begin with performance targets - voltages and loadings with certain limits- for load flow planning, reliability within a certain range for reliability engineering. They then enter into the program a candidate system model, a data representation of the system as it is, or as they propose to build it.

The computer analysis then determines the expected performance of the system as represented to it. In a load flow, computing the expected voltages, flows, and power delivery at every point in the system does this. In a reliability assessment, computing the expected frequency and duration of outages at every point in the system does this.

The results are displayed for the user, usually graphically in a way that quickly communicates problems (results that are unacceptably out-of-range).

The planning engineers review the results. If they are less than desirable, they make changes to the system model, using the results as a guide to where improvements or changes should be made. The model is re-run in this manner until a satisfactory result is obtained.

Root Cause and Sensitivity Analysis

Root-cause analysis determines the contribution of each component to poor reliability. For example, if reliability is measured using SAIDI, a root-cause analysis will identify if pole failures, breaker failures, and so forth contribute what impact to the original cause of the SAIDI.

The analysis is considerably more important when part of an aging infrastructure program as well as a reliability augmentation plan. In a way it unites the two endeavors and it definitely provides a starting point to coordinate them well. Both the root-cause analysis, and the analysis of equipment failure trends (whether and how much they are increasing, etc.) are dealing with much the same data. An important point is whether aging of various types of equipment is a root cause. When it is, trends found that show a gradually increasing failure rates should be applied predictively based on the root cause analysis: their contribution to reliability issues will worsen in company with those trends, etc.

In addition to knowing the contribution each component makes to the overall poor reliability of the system, it is desirable to know the impact that improving component reliability or design would have on the system's reliability: again, this links reliability engineering and aging equipment analysis. Will a change make a difference?

For example, what would be the impact of reducing overhead line failures, reducing transformer failures, or reducing cable repair-time? A sensitivity analysis can be used to answer such questions: this simple approach adjusts a component parameter (such as its failure rate) by a small amount and records the corresponding change in system reliability.

Engineering the Design Improvements

In determining how to improve system performance, planning engineers focus on equipment or areas of the system that score high in terms of both root cause (they create a lot of the problems) and sensitivity (changes will make a difference on reliability). After a system has been modeled, calibrated, and examined with root cause and sensitivity analyses, potential design improvements can be identified and modeled. The focus should be on critical components – those with both high root-cause scores and high sensitivity scores.

14.3 Methods for Distribution System Reliability Assessment

Power distribution reliability can be designed just like any other aspect of performance – voltage profile, loading. Doing so in a dependable and efficient manner requires an engineering planning method that can simulate how any particular distribution design will perform in terms of reliability of service. The system reliability will be as targeted by the planners and the effort exerted does not require special skills or undue levels of effort.

Several suitable analytical methods for this type of engineering exist. In the same manner that a power flow model can predict the electrical behavior of a distribution system (such as currents and voltages), a computer program called a *reliability assessment model*, based on one of these methods, can predict the expected reliability performance of any particular distribution system. As reliability becomes more important to electric utilities and electricity consumers, these reliability assessment models will equal or surpass power flow models in importance and usage. Reliability models allow distribution engineers to

- Design new systems to meet explicit reliability targets,
- Identify reliability problems on existing systems,
- Test the effectiveness of reliability improvement projects,
- Determine the reliability impact of system expansion,
- Design systems that can offer different levels of reliability, and
- Design systems that are best suited for performance based rates.

Distribution system reliability assessment is a rapidly developing field. Some of the first work on computer programs for this was done at Westinghouse Advanced Systems Technology, Pittsburgh, in the late 1970s. One of the first commercial distribution system reliability analysis programs, PREL, came from that group in 1981. This software was mainframe based and not widely adopted, largely because utilities in the 1980s did not put the heavy emphasis on reliability as they do 20 years later.

Increasing sensitivity of customer loads and impending deregulation led several utilities to develop their own distribution reliability assessment capabilities in the early and mid-1990s. Demand for distribution reliability assessment software is continuing to grow and commercial packages are now available from several major vendors. Electric utilities that desire to apply such methods have no problem finding several suitable packages from which to choose.

Methodologies for Distribution System Reliability Analysis

Reliability assessment models work from an arc-node type database, very similar to that used in a distribution load flow program. Unlike a load flow, which uses data on each element to compute the interconnected flow of power through the system to the customer locations, the reliability assessment program uses data on each element to compute the interconnected reliability of service to the customer locations.

There are four common methodologies used for distribution reliability assessment: network modeling, Markov modeling, analytical simulation, and Monte Carlo simulation. These differ in their bases – in what they use as the foundation for their analysis of reliability. Each uses a different approach to determining the reliability of the system. A brief description of each is provided below.

Network modeling

Network modeling is based on the topology of the system, which it translates from a physical network into a reliability network based on serial and parallel component connections. Serial components have additive failure rates and each one can disable the entire chain. Parallel components have redundancy (up to a point corresponding to capacity) and can cover outages of other parallel pathways.

Network modeling basically takes the system topology and translates it into a “formula” of serial and parallel reliability characteristics, which it then uses to compute the reliability of the system performance. In computer models this formula is a table-driven set of equation coefficients that implements a reliability computation. This computation computes the likelihood that a continuous interconnection between a source, and the demand point, remains in operation.

Network modeling is simple and straightforward to implement, and produces good results at a basic evaluation of reliability. A major disadvantage for distribution system analysis is that dynamic functions, such as switching, or sequential system responses to contingencies, are outside of its context. Its formula applies to a basic topology, anything such as a switch operation that changes the topology, cannot be modeled. As such, it is not as widely used as other methods.

Markov modeling

Markov modeling is a powerful method for assessment of reliability, as well as for simulation of many other apparently random multi-state processes (such as wind and solar availability for renewable energy systems). Markov modeling is based on analyzing the states that a system could be in: states can be such things as “everything operating normally,” or “component one has failed,” or “component two has failed,” etc. It focuses, though, on analyzing the transitions

between these states: it analyzes the likelihood that the system can move from the condition “everything is operating normally” to “component one has failed,” as well as under what conditions, and how long, it takes to transition back.

Every condition, or state, that the system could be in is identified and enumerated. The analytical method focuses on computing the back and forth transitions between one state and another. It models the conditions and likelihood of these transitions, thereby mapping the amount of time that the system spends in each state. If one state is “everything operating normally” and others represent one failure model or another, then the model can be used to analyze reliability-of-service and its causes.

Markov modeling is excellent for representing systems where details of the transition between states are known or are important. For example, planners may want to study how different repair and replacement part policies would impact the availability of a DG unit. Having a repairman at the DG site, or the parts already there so they do not have to be spent for, would clearly reduce the time to transition from “broken” to “working.” Markov models make it possible to focus in great detail on studying questions about how much such actions would impact reliability, because they focus on the mechanism of moving between states, not on the states themselves.

However, Markov models have some disadvantages for power distribution system applications. The first is that the states modeled are memory less (transition out of a state cannot depend on how the state was reached). This characteristic requires a type of duplication of states when system responses are a function of past events. Thus, if the system is in the state “working” because during the “broken” state a spare part was used (and as a result there no longer is another spare part at the site), the system cannot model this unless a new state is introduced: “working but no spare parts available.” Then, transitions between this state and “working” need to be established in the model (to represent delivery of the new part to the site). In practice, this complicates the application of a Markov model when used to analyze a distributed system such as a power distribution system.

The second limitation is computational. The matrix inversion required by Markov modeling requires that limits to the size of systems that can be represented and/or the complexity that can be represented to systems such as a few DG units. It also limits very simple distribution systems such as only the substation serving a small industrial site and its equipment.

Analytical simulation

Analytical simulation models each system contingency, computes the impact of each contingency, and weights this impact based on the expected frequency of the contingency. At first glance it appears to be very close to a form of contingency analysis, as was described in the Chapter 8 discussion on N-1 contingency

analysis, except that likelihoods are assigned or computed for each contingency case. This is a big difference and not an easy one to add to N-1 analysis. Analytical simulation algorithms are built around the mechanism that computes the probability of each contingency case, rather than an explicit enumeration of contingency states, as in N-1 analysis.

Accurate reliability analysis using analytical simulation means computing the likelihood of each contingency based on failure rates and interconnection of equipment. Generally, dynamic functions of equipment such as fuses, breakers, rollover switches, etc., has to be modeled, too, hence the term “simulation.” A great advantage of this method is that such types of dynamic activity can be modeled.

If the likelihood of each contingency is computed based on a good representation of the failure rates and modes that would lead to it, this type of method can accurately model complex system behavior and dynamically enumerate each possible system state. Aggregation of system states and their probabilities then permits detailed assessment of reliability (all states that lead to success vs. failure).

Monte Carlo Simulation

Monte Carlo Simulation is similar to analytical simulation, but models random contingencies are based on probabilities of occurrence, rather than expected contingencies. This allows component parameters to be modeled with probability distribution functions rather than expected values.

Monte Carlo Simulation can model complex system behavior, non-exclusive events and produces a distribution of possible results rather than expected values (Brown et al., 1997). Disadvantages include computational intensity and imprecision (multiple analyses of the same system will produce slightly different answers). Additionally, Monte Carlo Simulation is not enumerative and may overlook rare but important system states.

Table 14.2 Results Computed By an Analytical Simulation

-
- Expected number of momentary interruptions (per year)
 - Expected number of sustained interruptions (per year)
 - Expected number of interrupted hours (per year)
 - Expected number of protection device operations (per year)
 - Expected number of switching operations (per year)
-

For applications requiring determination of expected values of reliability, analytical simulation is the best method for distribution system assessment. It allows distribution engineers to quantify system reliability (SAIDI, SAIFI) over an entire system, and for individual customer locations, to calibrate models to historical data, to compare design alternatives, perform sensitivity analyses and to run optimization algorithms to maximize the expected results [See Brown, et al., 1998 - 2000]. Monte Carlo Simulation becomes necessary if statistical results other than expected values are required – analysis of such things as the distribution of expected SAIDI from year to year.

14.4 APPLICATION OF ANALYTICAL SIMULATION FOR DETAILED RELIABILITY ASSESSMENT

An analytical simulation method simulates a contingency, determines the impact of this contingency on system reliability, and weights the impact of the contingency by its probability of occurrence. This process is repeated for all possible contingencies, and results in the following information for each component, providing the results shown in Table 14.2.

Simultaneously with the analysis of each contingency and the results it produces, the analytical simulation method evaluates the probability of occurrence of this contingency – how likely is it to occur? Once all contingencies have been analyzed, their results are combined appropriately on a probability-weighted basis to produce aggregate results. In distribution system analysis this aggregation of results is based on system topology, keyed to adjacent equipment, branches, circuits, etc., as well as the entire system.

Modeling Each Contingency

A contingency occurring on a distribution system is followed by a complicated sequence of events. Each contingency may impact many different customers in many different ways. In general, the same fault will result in momentary interruptions for some customers and varying lengths of sustained interruptions for other customers, depending on how the system is switched and how long the fault takes to repair. The key to an analytical simulation is to accurately model the sequence of events after a contingency to capture the different consequences for different customers. A generalized sequence of events as modeled by analytical simulation is:

1. *Contingency:* A fault occurs on the system.
2. *Reclosing:* A reclosing device opens in an attempt to allow the fault to clear. If the fault clears, the reclosing device closes and the system is restored to normal.
3. *Automatic Sectionalizing:* Automatic sectionalizers that see fault

current attempt to isolate the fault by opening when the system is de-energized by a reclosing device.

4. *Lockout*: If the fault persists, time over current protection clears the fault. Lockout could be the same device that performed the reclosing function, or could be a different device that is closer to the fault.
5. *Automated Switching*: Automated switches are used to quickly isolate the fault and restore power to as many customers as possible. This includes both upstream restoration and downstream restoration. In upstream restoration, a sectionalizing point upstream from the fault is opened. This allows the protection device to reset and for restoration of all customers upstream of the sectionalizing point. In downstream restoration, other sections that remain de-energized are isolated from the fault by opening switches. Customers downstream from these points are restored through alternate paths by closing normally-open tie switches.
6. *Manual Switching*: Manual switching restores power to customers that could not be restored by automated switching (certain customers will not be able to be restored by either automated or manual switching). As in automated switching, manual switching has both an upstream restoration component and a downstream restoration component.
7. *Repair*: The fault is repaired and the system is returned to its pre-fault state.

The seven steps outlined above generate a set of system states for each contingency. Switches and protection devices being open or closed characterize these states. For each state occurring with frequency λ and duration δ , the accrued outage frequency of all de-energized components are incremented by λ (if the component was energized in the preceding state) and the accrued outage duration of all de-energized components are incremented by $\lambda \times \delta$.

The simulation becomes more complicated if operational failures are considered. Operational failures occur when a device is supposed to operate, but fails to do so. The probability of such an event is termed *probability of operational failure, POF*. Operational failures cause the simulation sequence to split. One path assumes that the device fails to operate and has a weight of *POF*, the other path assumes that the device operates and has a weight of $1 - POF$. This path splitting is illustrated in Figure 14.2, which shows the steps required when considering a fuse that is supposed to clear a fault.

The result of simulation path splitting is an enumerative consideration of all possible system responses to each contingency (in the context of operational

failures). Enumerative consideration is important since some states may be rare, but have a major impact on the system when they do occur. During restoration, path splitting associated with the enumerative consideration of possible outcomes is important when intended switching fails and customers that would otherwise have been restored are not.




Figure 14.2 Simulation path splitting due to operational failures.

Incremental Studies of Improvement

An analytical simulation will produce identical results if an analysis is performed multiple times. In addition, small changes in input data will cause small changes in results (as opposed to some other methods). This allows the impact of small reliability improvements to be quantified for individual customers and reliability indices by running a series of studies of incremental changes in the system design. For example, the planning engineer can move the location of a sectionalizer or recloser up and down a feeder an eighth of a mile at a time, seeing the results in total reliability, and the distribution of reliability results up and down the feeder, with each change. This capability also allows input parameters to be perturbed and result sensitivities to be computed. However, analytical simulation does not allow the uncertainty of reliability to be quantified. For example, while expected that SAIDI might be 2.0 hours per year, there is some probability that next year

will just be an “unlucky year.” To compute how often such unlucky years could occur, and how extreme SAIDI or other factors would be in those years, requires application of Monte Carlo techniques.

Example Application

An analytical simulation method was applied to the test system shown in Figure 14.1, which is based on an actual U.S. utility distribution system. The system model contains 3 voltage levels, 9 substations, more than 480 miles of feeder, and approximately 8,000 system components. The model was first calibrated to historical outage data, then used to analyze candidate designs as described above. In addition the authors used an optimization module that maximized reliability results versus cost. The diagram shown in Figure 14.1 was originally color coded (not shown here) based on computed outage hours, and shows the system plan that resulted from a combination of engineering study and use of optimization. Individual component reliability results can be easily used to generate a host of reliability indices. For this system, common indices include:

MAIFI (Momentary Average Interruption Frequency Index) = 2.03 /yr

SAIFI (System Average Interruption Frequency Index) = 1.66 /yr

SAIDI (System Average Interruption Duration Index) = 2.81 hr/yr

Targets for the area were 2.00, 1.75, and 2.9 hours (175 minutes) respectively, thus this system essentially meets its target reliability criteria.

Differentiated Reliability Design

Realistically, the expected reliability of service over an entire distribution system cannot be the same at all points. Distribution systems use transshipment of power through serial sets of equipment. Thus there are areas of the system that are more “downstream,” and being downstream of more devices and miles of line, have somewhat higher exposure to outage. Figure 14.1’s example demonstrates this. Although SAIDI is 2.81, expected performance for the best 5% of customers is 1.0 hours and 4.02 for the worst 5%.

However, planning engineers can engineer a system to keep the range of variation in what customers in one area receive versus customers in another, within a reasonable range. This involves the selected use of protective equipment such as fuses and breakers, sectionalizers, reclosers, and contingency switching plans, and in some cases, distribution resources such as on-site generation and energy storage, and demand-side management. Artful use of these measures, combined with application of sound distribution layout and design principles, should result in a system that provides roughly equitable service quality to all customers.

However, there are situations where a utility may wish to promise higher than standard reliability to a specific customer, or to all customers in an area such as a special industrial park. These situations, too, can be engineered. In fact, generally it is far easier to engineer the reliability of service to a specific location than it is to engineer an entire system to a specific target.

Performance-Based Industrial Contracts

Generally, such situations arise in the negotiation for service to medium and large industrial customers. Many of these customers have special requirements that make continuity of service a special consideration. All of them pay enough to the utility to demand special attention – at least to the point of having the utility sit down with them to discuss what it can do in terms of performance and price.




Figure 14.3 Service to three industrial customers in the example system were designed to higher-than-normal reliability performance targets.

The actual engineering of the “solution” to an industrial customer’s needs is similar to that discussed above for systems as a whole. Planning engineers evaluate the expected reliability performance of the existing service to the customer, and to identify weaknesses in that system. Candidate improvement plans are tried in the model, and a plan evolves through this process until a satisfactory solution is obtained. The system design in Figure 14.1 includes three such customer locations, as shown in Figure 14.3. Special attention to configuration, fusing and protection coordination on other parts of the circuit, contingency (backup) paths and quick rollover switching, provided about 80% of the improvement in each case.

The other 20% came from extending the reliability assessment analysis to the customer’s site. In each case, the customer’s substation, all privately owned, was analyzed and suggested improvements in equipment or configuration were recommended to his plant engineer. In one case the analysis extended to primary (4 kV) circuits inside the plant. Table 14.3 gives the results for each. For site A, the target has been exceeded in practice to date because there have been no outages in three years of operation. However eventually performance will probably average slightly better than the target, no more.

In the case of site B, the customer’s main concern was that the length of any interruption be less than a quarter hour. Thus, the target is not aimed at total duration time, but on each event, something quite easy to engineer. Performance has greatly beaten this target – means to assure short outages such as automatic rollover and remote control restore power in much less time than required.

For customer C, any interruption of power resulted in loss of an entire shift’s (eight hours’) production. Thus, the major item of interest there was total count of events. Initial evaluation showed little likelihood of any interruption lasting that long. Reliability engineering focused on frequency only.

Table 14.3 Reliability Targets for Three Special Customers

Customer	Industry		Annual Target	Three-Year	
Operating Avg.	Type		Events	Total hours	Events
Total hours					
System	Entire system	1.61	2.81	1.67	2.65
A	Tractor manufacturer	.33	.02	0	0
B	Synthetic fiber	1.	.5/event	.66	.085/event
C	Carpet manufacture	.25	-	.33	.42


In all three cases, the performance-based contracts call for a price to be paid that can be considered a premium service price. It is typical for industrial customers in this class to negotiate price and conditions with the utility, but in all three cases the price included an identified increment for the improvement in service. Unlike some performance-based rate schedules, there is no “reward” or extra payment for good service – the premium payment is payment for the good service.

But performance penalties are assessed, not on an *annual* basis depending on annual performance, but on an *event* basis. If reliability of service is perfect, the utility has earned all the premium monies paid. Any event (any, not just those exceeding the limits) garners a rebate according to a certain formula. For example, customer B receives a \$30,000 rebate for any interruption of power, but the rebate rises to \$300,000 if the interruption lasts more than 30 minutes. (Customer B’s average monthly bill is about \$175,000.)

14.5 USE OF A HYBRID ANALYTICAL SOLUTION

If one hundred identical distribution systems are built, the expected reliability of each will be identical. In a given year, some of these systems might be lucky and experience nearly perfect reliability. Others may experience near expected reliability, and others may be unlucky and experience far worse than expected reliability. This variation is natural and is vital to understand when negotiating reliability-based contracts.

Variations in reliability can be examined using techniques referred to as risk assessment methods. A risk assessment model identifies all possible outcomes.



1. Break-in period (infant mortality) 2. Useful lifetime 3. Wear out period

Figure 14.4 The traditional “bathtub” failure rate curve, as discussed in the text. Here, the device has a useful lifetime of about 30 – 35 years.

This commonly takes the form of a probability of each outcome occurring. When possible, this is done through analytical methods such as function convolution. However, often this is not feasible and “Sequential Monte Carlo Simulation” is used where time slices are modeled in succession and each component is tested for random failures during each time slice. This type of simulation is very flexible and can handle complex models, but is computationally slow and data intensive.

Many of the problems associated with a sequential Monte Carlo Simulation can be overcome if components are assumed to have a constant failure rate over the course of one year. This is a reasonable assumption and can be visualized by “bathtub failure rate functions.” These functions show that typical electrical components will have a high failure rate when they are first installed (due to manufacturing defects, damage during shipping, and improper installation). The failure rate will reduce after this “infant mortality” period and remain at a fairly constant level over the useful life of the component. The failure rate will gradually rise as the equipment wears out at the end of its useful life. A bathtub failure rate function is shown in Figure 14.4.

Components with constant failure rates follow a Poisson process. This allows the probability of the component failing a specific number of times in a year to be easily computed. If a component has a constant failure rate of λ times per year, the probability of it failing x times in a year is:

$$\text{Prob. of failing } x \text{ times} = \frac{\lambda^x e^{-\lambda}}{x!} \quad (14.1)$$

If constant failure rates are assumed, the analytical simulation described in the previous section can be modified to simulate a *random* year rather than an *expected* year. This is done by determining the number of times each component will fail *a priori*. For each component, a random number between zero and one is generated. If this number is less than $e^{-\lambda}$, no failure will occur in the year being simulated. If the random number is greater, the number of times that the component will fail is determined by Equation 1.

Once the number of times each component will fail is determined, an analytical simulation is performed that substitutes component failure rates with the number of times that they will fail in the random year being simulated. Using this process, many years can be simulated, a list of outcomes can be recorded, and distribution statistics can be computed. This methodology is referred to as an Analytical/Monte Carlo Hybrid Simulation. This particular hybrid simulation has two desirable features:

1. It requires no additional data beyond the requirements of an analytical model.

2. If there is confidence in the expected values generated by the analytical simulation, and there is confidence that component failure rates are constant over a single year; there is equal confidence in the results of the analytical/Monte Carlo hybrid simulation.

Analytical/Monte Carlo Hybrid Simulation is applied to the system being analyzed (i.e., the system in Figures 14.1 and 14.3) by performing 1,000 random simulations on the same system that was used to demonstrate the analytical simulation. The results are shown in Figure 14.5 as both a histogram (the bar chart) and an integral of the histogram (the continuous line). The statistical results of the simulation are (in hours per year): mean value = 2.81, standard deviation = .79, minimum value = .95, maximum value = 6.37.

The Analytical/Monte Carlo Hybrid Simulation is computationally intensive, but provides statistical results not obtainable by using purely analytical methods. This statistical information is vital when assessing technical and financial risk associated with reliability based contracts. The remainder of this chapter demonstrates this by applying Analytical/Monte Carlo Hybrid Simulations to distribution systems subject to performance based rates.




Figure 14.5 SAIDI Results for 1,000 random years (hr/yr). Bars show a histogram of the expected annual results, which, while having a mean of less than 3 hours, have a good probability of being over 5 hours. Solid fits a log-normal curve to the data.

Analyzing the Risk from Performance-Based Rates

Deregulated utilities are reducing costs by deferring capital projects, reducing in-house expertise, and increasing maintenance intervals. As a direct consequence, the reliability on these systems is starting to deteriorate. Since these systems have been designed and maintained to high standards, this deterioration does not manifest itself immediately. System reliability will seem fine for several years, but will then begin to deteriorate rapidly. When reliability problems become evident, utilities often lack the necessary resources to address the problem.

Regulatory agencies are well aware that deregulation might have a negative impact on system reliability. In a perfect free market, this would not be a concern. Customers would simply select an electricity provider based on a balance between price and reliability. In reality, customers are connected to a unique distribution system that largely determines system reliability. These customers are captive, and cannot switch distribution systems if reliability becomes unacceptable. For this reason, more and more utilities are finding themselves subject to performance-based rates (PBRs).

A PBR is a contract that rewards a utility for providing good reliability and/or penalizes a utility for providing poor reliability. This can either be at a system level based on reliability indices, or can be with individual customers. Performance is usually based on average customer interruption information. This typically takes the form of the reliability indices SAIFI and SAIDI.




Figure 14.6 Performance based rate structure for the example in this section. There is no penalty or incentive payment for annual SAIDI performance that lies between 2 and 3 hours. Above that the utility pays a penalty of \$1,666,666 per minute, up to 60 minutes. Similarly, it receives a reward of the same magnitude for each minute below the dead zone.

A common method of implementing a PBR is to have a “dead zone” where neither a penalty nor a bonus will be assessed. If reliability is worse than the dead zone boundary, a penalty is assessed. Penalties increase as performance worsens and are capped when a maximum penalty is reached. Rewards for good reliability can be implemented in a similar manner. If reliability is better than the dead zone boundary, a bonus is given. The bonus grows as reliability improves and is capped at a maximum value. All PBRs will have a penalty structure, and some will have both a penalty structure and a bonus structure. A graph of a PBR based on SAIDI is shown in Figure 14.6.

Most initial PBRs will be negotiated by the utility to be “easy to meet.” This means that business as usual will put the utility in the dead zone. This does not mean that a utility should do business as usual. It may want to spend less on reliability until marginal savings are equal to marginal penalties. Similarly, it may want to spend more money on reliability until marginal costs are equal to the marginal rewards. In either case, a utility needs a representative reliability assessment and risk model to determine the impact of reliability improvement and cost reduction strategies.

In order to make intelligent decisions about PBRs based on average system reliability, a probability distribution of financial outcomes is needed. This requires a PBR structure (such as the one shown in Figure 14.6) and a distribution of relevant reliability outcomes (like the histogram shown in Figure 14.5). It is also desirable to describe the reliability histogram with a continuous mathematical function. A good function to use for reliability index histograms is the log-normal distribution, represented by $f(x)$:

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left[-\frac{(\ln x - \mu)^2}{2\sigma^2}\right]; \quad x \geq 0 \quad (14.2)$$

The parameters in this equation can be determined by a host of curve-fitting methods, but a reasonably good fit can be derived directly from the mean and variance. If a mean value, $\bar{x}$, and a variance are computed, log-normal parameters are:

$$\sigma = \sqrt{\ln(\text{variance} + e^{2\ln \bar{x}}) - 2\ln \bar{x}} \quad (14.3)$$

$$\mu = \ln \bar{x} - \frac{1}{2}\sigma^2 \quad (14.4)$$

The log-normal curve corresponding to the test system is shown along with the histogram in Figure 14.5. The curve accurately captures the features of the histogram including the mean, median, mode, standard deviation, and shape.

Assume that a system analyzed earlier, characterized by the reliability shown

in Figure 14.5, is subject to the PBR formula shown in Figure 14.6. Using these two functions, a financial risk analysis can be easily determined. For example, the expected penalty will be equal to:

$$Penalty = \int_0^{\infty} PBR(SAIDI) \cdot f(SAIDI) dSAIDI \quad (14.5)$$

The probability of certain financial outcomes can also be computed. For example, the probability of landing in the dead zone will be equal to the probability of have a SAIDI between 2 hours and 3 hours. This is mathematically represented by:

$$\% \text{ in Dead Zone} = 100\% \cdot \int_2^3 f(SAIDI) dSAIDI \quad (14.6)$$

The expected outcome of the test system is \$14 million in penalties per year. Thus, despite the fact that the expected annual SAIDI (2.81 hours) lies in the dead zone (no penalty – no reward) the utility will on average have to pay a fine amounting to about 8.4 minutes per year. Penalties will occur 32% of the time, bonuses will occur 12% of the time, and the utility will be in the Dead Zone 57% of the time. The distribution of all possible outcomes and their associated probabilities is shown in Figure 14.7.

For a cash-strapped utility, Figure 14.7 represents an unwelcome situation. And it is risky, because even though the average penalty of \$14 million may be acceptable, that outcome will rarely happen. The \$14 million average occurs because a penalty of \$50 million or more will occur once every 7 years and the maximum penalty of \$100 million will occur once every 12 years. Faced with this situation, a utility would be wise to negotiate a less risky PBR. Possibilities for adjusting the PBR to mitigate risk include:

1. Make the reward and penalty slopes less steep,
2. Widen the dead zone boundaries,
3. Move the dead zone to the right, and
4. Reduce bonus and penalty caps. Each of these options can be used alone, or a combination can be negotiated.




Figure 14.7 Distribution of financial outcomes for the example PBR problem. Although it is most likely the penalty/reward is zero (performance lies in the dead zone) the utility has a significant probability of getting hit with the maximum penalty (about once in every 12 years), while no chance of getting the maximum reward.

A suggested change to the PBR rules to mitigate the risk for this utility system is shown in Figure 14.8. This change stretches out the bonus and penalty transition zones from one hour to two hours. The impact of this changed PBR on the utility's expected financial exposure is shown in Figure 14.9. On the surface, it may seem like not much has changed. Penalties will still occur 32% of the time, bonuses will still occur 12%, but total risk is reduced greatly, by \$5 million. Utility companies can use predictive reliability assessment tools when planning *and* designing their systems and their responses to regulatory rules. This will allow the reliability of existing systems to be assessed, the reliability impact of projects to be quantified, and the most cost-effective strategies to be implemented. In most cases, an analytical model (Chapter 14) will be sufficient. This will allow utilities to compare the expected performance of various systems and various design options. Analytical/Monte Carlo hybrid techniques become necessary when technical and financial risk needs to be examined. This chapter has shown methodologies to serve both of these need categories.




Figure 14.8 Changed PBR schedule that reduces risk. The schedule has a milder rise in both penalty and reward sides and is a “fair” proposal for the utility to make to regulators because it treats both penalty and reward sides equally.




Figure 14.9 Financial outcome with the revised PBR formula. The utility’s expected average is still a penalty, but only \$9 million, an improvement of \$5 million. More important, the possibility of it being hit with a large (over \$50 million) penalty has decreased greatly.

14.6 CONCLUSION AND KEY POINTS

The reliability of service of a power distribution system can be designed to meet specific goals with respect to reliability criteria such as SAIDI and SAIFI. It can also be designed to minimize financial risk whether that is measured in terms of utility and/or customer costs, or includes performance-based rates and contracts as well. Of four possible methods, all capable of computing distribution reliability, the authors believe the analytical simulation method is most applicable to reliability-based planning.

Analytical simulation computes the expected reliability of each component based on system topology and component reliability data. It does so by simulating the system's response to each contingency, determining the impact of this contingency to each component, and weighting the impact of each contingency by its probability of occurrence. The result for each component is the expected number of momentary interruptions per year, the expected number of sustained interruptions per year, and the expected number of outage hours per year. Results can be displayed in tables, visualized spatially, or aggregated into reliability indices.

To assess technical and financial risk, the analytical simulation method is expanded to an Analytical/Monte Carlo Hybrid Simulation. Such a method assumes that component failure rates are constant over a specific year and models many random years rather than a single expected year. The result is a probability distribution of possible results that can be used to obtain information such as variance, median, modality, shape, and extremes. It also allows rigorous risk analyses to be performed on various types of performance-based rates and customer contracts.

After negotiating a performance-based rate, a distribution utility will find itself in a radically new paradigm. An explicit dollar value has been placed on reliability, and systems can be designed to balance the cost of reliability with the cost of infrastructure. If reliability measures are based on indices such as SAIFI and SAIDI, it may be beneficial to the utility to allow the reliability of a few customers in areas where it would be expensive to "fix" problems to decline to levels far below average.

Every power system will have areas where reliability can be improved relatively inexpensively, where for whatever reason the marginal cost of SAIDI and SAIFI improvement is low. Focusing on these areas alone improves overall utility-wide statistics, but increases the variance of customer experience. Eventually, this global "gaming" of the performance criteria will be viewed as unacceptable by customers and regulators alike, sparking more complicated PBR formulas and regulations designed to prevent such situations. Ironically, "deregulation" may therefore result in regulation of a previously unregulated area. It is best to focus on comprehensive programs that keep customer service

everywhere within rather tight bounds. Regardless, *reliability can be engineered*. Key aspects of reliability-based engineering for distribution systems are:

- *Focus on a customer-oriented metric.* What really counts is the reliability of service the customer sees not reliability as measured on the power system itself. Reliability-based engineering should target customer-level reliability performance.
- *Use a good reliability-assessment method.* Such methods apply a rigorous analytical method, compute expected reliability on a detailed (nodal) basis, and use a method that provides consistent evaluation of results so that the effect of incremental design changes can be seen.
- *Use differentiated reliability where necessary* to best meet customer reliability needs and price sensitivity. Premium areas (higher reliability and price) can be designed to specific targets for those areas.
- *Calibrate the base model* to historical outage data as seen in actual operation of the system.
- *Engineer reliability by improving design* in response to weaknesses identified by the reliability assessment method, in existing and proposed system designs.
- *Minimize financial risk*, using a method appropriate to the evaluation of the probability distribution of expected results

REFERENCES AND BIBLIOGRAPHY

- P. F. Albrecht and H. E. Campbell, "Reliability Analysis of Distribution Equipment Failure Data," EEI T&D Committee Meeting, January 20, 1972.
- R. N. Allan, et al., "A Reliability Test System for Educational Purposes – Basic Distribution System Data and Results," *IEEE Transactions on Power Systems*, Vol. 6, No. 2, May 1991, pp. 813-821.
- D. Atanackovic, D. T. McGillis, and F. D. Galiana, "The Application of Multi-Criteria Analysis to a Substation Design," paper presented at the 1997 IEEE Power Engineering Society Summer Meeting, Berlin.
- R. Billinton, and J. E. Billinton, "Distribution System Reliability Indices," *IEEE Transactions on Power Delivery*, Vol. 4, No. 1, January 1989, pp. 561-568.

- R. Billinton R., and R. Goel, "An Analytical Approach to Evaluate Probability Distributions Associated with the Reliability Indices of Electric Distribution Systems," *IEEE Transactions on Power Delivery*, PWRD-1, No. 3, March 1986, pp. 245-251.
- R. E. Brown, S. S. Gupta, R. D. Christie, and S. S. Venkata, "A Genetic Algorithm for Reliable Distribution System Design," *International Conference on Intelligent Systems Applications to Power Systems*, Orlando, FL, January 1996, pp. 29-33.
- R.E. Brown, S. Gupta, S.S. Venkata, R.D. Christie, and R. Fletcher, 'Distribution System Reliability Assessment Using Hierarchical Markov Modeling,' *IEEE Transactions on Power Delivery*, Vol. 11, No. 4, Oct., 1996, pp. 1929-1934.
- R.E. Brown, S. Gupta, S.S. Venkata, R.D. Christie, and R. Fletcher, 'Automated Primary Distribution System Design: Reliability and Cost Optimization,' *IEEE Transactions on Power Delivery*, Vol. 12, No. 2, April 1997, pp. 1017-1022.
- R.E. Brown, S. Gupta, S.S. Venkata, R.D. Christie, and R. Fletcher, 'Distribution System Reliability Assessment: Momentary Interruptions and Storms,' *IEEE Transactions on Power Delivery*, Vol. 12, No. 4, Oct 1997, pp. 1569-1575.
- R.E. Brown, J.R. Ochoa, 'Distribution System Reliability: Default Data and Model Validation,' *IEEE Transactions on Power Systems*, Vol. 13, No. 2, May 1998, pp. 704-709.
- R. E. Brown, S. S. Venkata, and R. D. Christie, "Hybrid Reliability Optimization Methods for Electric Power Distribution Systems," *International Conference on Intelligent Systems Applications to Power Systems*, Seoul, Korea, IEEE, July 1997.
- J. B. Bunch, H.I. Stalder, and J.T. Tengdin, "Reliability Considerations for Distribution Automation Equipment," *IEEE Transactions on Power Apparatus and Systems*, PAS-102, November 1983, pp. 2656 - 2664.
- "Guide for Reliability Measurement and Data Collection," EEI Transmission and Distribution Committee, October 1971, Edison Electric Institute, New York.
- S. R. Gilligan, 'A Method for Estimating the Reliability of Distribution Circuits,' *IEEE Transactions on Power Delivery*, Vol. 7, No. 2, April 1992, pp. 694-698.
- W. F. Horton, et al., "A Cost-Benefit Analysis in Feeder Reliability Studies," *IEEE Transactions on Power Delivery*, Vol. 4, No. 1, January 1989, pp. 446 - 451.
- Institute of Electrical and Electronics Engineers, *Recommended Practice for Design of Reliable Industrial and Commercial Power Systems*, The Institute of Electrical and Electronics Engineers, Inc., New York, 1990.
- G. Kjølle and Kjell Sand, 'REL RAD - An Analytical Approach for Distribution System

- Reliability Assessment,' *IEEE Transactions on Power Delivery*, Vol. 7, No. 2, April 1992, pp. 809-814.
- A. D. Patton, "Determination and Analysis of Data for Reliability Studies," *IEEE Transactions on Power Apparatus and Systems*, PAS-87, January 1968.
- N. S. Rau, "Probabilistic Methods Applied to Value-Based Planning," *IEEE Transactions on Power Systems*, November 1994, pp. 4082 - 4088.
- A. J. Walker, "The Degradation of the Reliability of Transmission and Distribution Systems During Construction Outages," Int. Conf. on Power Supply Systems. IEEE Conf. Publ. 225, January 1983, pp. 112 - 118.
- H. B. White, "A Practical Approach to Reliability Design," *IEEE Transactions on Power Apparatus and Systems*, PAS-104, November 1985, pp. 2739 - 2747.
- W. Zhang and R. Billinton, "Application of Adequacy Equivalent Method in Bulk Power System Reliability Evaluation," paper presented at the 1997 IEEE Power Engineering Society Summer Meeting, Berlin.

15

T&D EQUIPMENT MANAGEMENT PHILOSOPHIES

15.1 INTRODUCTION

This chapter discusses ownership “philosophies” – the best word the authors can find to cover the very palpable differences in thinking and attitude that color electrical equipment management values and decisions at different companies. Readers interested only in technical aspects of equipment aging, their impacts and interactions with power systems, and the methods of analysis and mitigation that companies can use to control aging’s effects can move past this chapter, but it provides an interesting and useful discussion of why companies sometimes take a particular approach to the management of their aging infrastructures.

The authors have never encountered two companies – be they utilities or industrial or institutional owners of power systems – that have exactly the same perspective, needs, priorities, and decision-making approach to the management of their electrical equipment and systems. Many electric delivery utilities share the same broad perspective, shaped by the fact that they all face fairly similar regulatory rules and business environments, and are hardened by decades of being in the same industry. Still, utilities vary somewhat in their priorities and perspectives. Beyond utilities, other companies, some of which own billions of dollars worth of electrical equipment and to whom its performance is critical to success, have very different philosophies of ownership, leading to priorities and management decisions that would seem bizarre to long-time employees of utilities. Regardless of industry, those who recognize there are other perspectives will have a better base from which to say, “This approach is right for us.”

Often a company's interpretation of what will work best for it is seen through the lens of its own distinctly different view of *why* it owns the equipment and *how* it should care for it. That viewpoint is supported by and even a product of a culture that feeds on and reinforces that perspective – “That is the “what way things work here,” etc. Differing viewpoints about that lead to different philosophies – value systems – about to how the equipment should be tended and how age and condition play into decisions with regard to ownership and management. They lead to very different views about “what is correct” and what is not with regard to caring for electrical equipment. They drive greatly differing expectations about how equipment should “age” and what the owner should do about equipment as it deteriorates from the wear and tear of service.

The point of this chapter's discussion is not to identify which approach is right or which is wrong, but to present the different perspectives and show how and why they differ. A person who has long worked under any one of these philosophies is likely to view the values, priorities, and decisions made in the others as “wrong” and the decisions and recommendations flowing out of it as “incorrect” or even silly, but people on the other side of the perspective will have the opposite view.

This chapter begins with an overview of four very different ownership perspectives and philosophies the authors have encountered in their work. The discussion looks at how these different philosophies affect a company's culture – the attitudes and tacit “right and wrong” habits and institutionalized instincts that is perhaps best summed up as “this is how we do things here” within any group activity. Often a particular cultural attitude and value system dominates a company or a large division within a company, and makes a substantial difference in how it reacts or what its values are as compared to other portions of the organization. Very few companies are purely one or the other of the four approaches discussed in Section 15.1, but most are dominated by one philosophy more than the others.

Six example presented along with a discussion of how their differing needs and cultures drive their management of electrical facilities. Names have been changed and in some cases unimportant aspects of the example company's description altered so that the authors can talk frankly about their customers and former customers without revealing confidential information or embarrassing anyone (including the authors).

15.2 FOUR PHILOSOPHIES OF OWNERSHIP

Why does your company own electric equipment, and why do you and your coworkers take care of it? To most people long involved with the operation and maintenance of electrical equipment on a daily basis, these two questions seem pointlessly unnecessary, and the answers obvious. But if one surveyed a cross-section of electric utilities, municipal public works departments, and industrial and institutional owners of power equipment, one would get a wide range of very different answers to these two questions. Moreover, the answers

Table 15.1 Four Ownership Perspectives

Perspective	Equipment is . . .	Service is . . .	Aging is . . .
Buy-Burn-Bash	expendable	often ineffective	inevitable
Standards-Driven	what's important	obligatory	bad, very bad
Reliability-Centered	success enabler	manageable	a reason to use RCM
Asset Management	a business tool	a business expense	part of doing business

would not be consistent among apparently similar businesses. There would be a good deal of difference among electric utilities, although perhaps not as wide a variance as between them and large industrial firms. And among industrial power system owners, there would be a considerable difference, although again perhaps not as great as the difference between the average industrial “owner” (say a refinery plant manager) and the average large institutional owner (e.g., a university) owner of electric facilities.

Table 15.1 summarizes these four perspectives have with respect to equipment, service, and aging of equipment (this limited to how they is aging of equipment in general, not specifically “problems of aging infrastructures,” which will be discussed in more detail later in this chapter). These four different perspectives are:

Buy, Burn and Bash – the authors’ alliterative nomenclature for a practice of benign neglect of electric equipment: “We buy it. We put it in service. We burn it up and just throw it away and replace it when it fails.”

Standards-Driven Maintenance – “We do things by the book, and we write the book.”

Reliability-Centered Maintenance (RCM) – “We pay close attention to equipment and put our attention and effort where it will do the most good.”

Asset Management – “We’re running a business. Electrical equipment, like everything else we own, is here to serve the business.”

Buy-Burn-Bash and the Standards-Driven Approach are both quite old and proven approaches. One can trace either back to over one hundred years – to the very beginning of electrification. There are companies that have used and held onto either approach for decades, and that have a culture that strongly believes in

the approach and very forcefully rejects other perspectives and approaches. Reliability Centered Maintenance and Asset Management are newer methods, enabled by new technologies and requiring more involved analysis and management methods, but fitting ways of thinking that evolved out of business management institutions and academia in the late 20th century.

Buy, Burn and Bash

The equipment is going to wear out anyway. This is the core concept driving ownership and operating decisions at any organization taking this approach; a company can do preventive and service maintenance on some parts of a transformer, breaker, or other electrical component, but there are other characteristics critical to its operation in which deterioration cannot be well controlled. Those failure modes will lead to its failure anyway, so why bother with servicing those other things?

For example, electrical bushings can be renewed and metal cases checked and repaired for rust, etc., but the electrical insulation in a transformer's core will deteriorate with use and nothing short of completely rebuilding it will slow or stop that. In some types of breakers, contacts and mechanical parts can be renewed, but the tight, perfectly round and aligned metal castings supporting its trip mechanism will loosen, warp, and fatigue with every use; it is only a matter of time until it gives problems as a result.

Thus, this perspective holds, in both cases money spent on "proactive" replacement and renewal of components that can be repaired is largely wasted: modes of failure that cannot really be addressed will lead to the equipment's eventual demise anyway. Better just to buy equipment that is reasonably durable, take measures to make certain nothing in its use contributes to really premature failure, and "burn it up" – use it until it fails – and then dispose of it and install another.

In some cases, this perspective is quite correct based on the facts, as will be discussed below, but in some cases a company takes this approach because it views the cost and effort required for proactive maintenance and care, and even routine service, as a distraction from its "core mission," and/or because it has used this approach long enough that it knows no other.

What Buy-Burn-Bash does that other own-operate approaches don't.

Mostly, it is what this approach does not do that most people see as the big difference: the equipment owners perform almost no proactive service and as little "required" service and inspection as possible. Inspections, treatments, and services mandated by law (as for example, those required of some natural gas powered equipment) are done to the extent required. Safety is sometimes a big priority and there are inspections and monies allocated for that. But beyond those, inspections might be done annually only, as one of the authors refers to it, "to make sure the equipment is still there." Anything more is not done. Even manufacturer's service recommendations, including those required to maintain

any warranty in effect, are ignored.

But beyond those differences-of-absence that persons working under the other three paradigms would first notice, B-B-B organizations do practice positive habits unique to their approach. They set themselves up to tolerate a “zero-care” equipment environment. In many cases careful thought and intense attention are devoted to system and process designs that tolerate the loss of equipment without loss of service or production, and designs generally allow for very quick change out when a failure occurs. Owners often put considerable care into evaluating and specifying equipment for their purposes, seeking electric machinery that is durable but not necessarily easily maintainable. Available spares might be kept on site and arrangements made for access to contractors for near immediate replacement when a failure does occur.

Where a “Buy, Burn, Bash” approach is popular

Heavy industrial applications. Many large industrial plants such as refineries have high utilization rates for the electric system: equipment loading stays at or near rated capacity – and thus at or near the equipment’s designed maximum temperatures – for most of the 8,760 hours each year, with no overnight cooling cycle as in utility systems. Such use accelerates electrical-mechanical insulation deterioration and some types of metal fatigue as compared to utilization that peaks each day and then tapers off with a distinct cooling period overnight. For example, a power transformer might lose 50% of its core insulation strength in only 10 years in a high-utilization rate industrial application, whereas it would not reach that state in 50 years in a typical electric utility application.

Additionally, in many industrial venues all of the equipment is not just physically, but also electrically “close”: fault duties are quite high and through-faults cause near maximum mechanical stress at every event. Equipment is mechanically, not just electrically, stressed at high levels over time and so it wears heavily in other ways, again taking perhaps less than a decade to reach points of deterioration and wear that other types of service would bring only after a half a century or more of use.

In the decade it takes for these types of deterioration to take place, metal cases are unlikely to rust through, seals to dry out, crack and leak, oil to reach contaminate levels above guidelines, or bushings to fail. So, since those major potential causes of eventual deterioration cannot be addressed by maintenance, and other failure modes are unlikely to occur, why bother to service and maintain the units at all? In such cases, “don’t service it, just burn it up” may not be the optimum approach, but it is certainly not wildly incorrect, and it is by far the approach to “equipment care.”

Beyond this, many industrial plants have “downtime” costs that are astronomical compared to the cost of electrical equipment: some large industrial facilities cost their owners \$250,000 per hour when they are out of service. As a result, a plant electrical design that is completely failure tolerant is mandatory, as are spares and resources to switch out units quickly: there is no business impact

when a unit fails without warning, and thus no value seen in the reduction in failure rates that comprehensive inspection and service might provide.

Finally, the authors are convinced from working with a number of organizations that “management bandwidth” is a motivating factor for some organizations that take this approach. Part of this philosophy’s popularity is that it not only requires less talent and labor than the other approaches, but it requires less management effort, too. There are fewer processes to set up, maintain, and overview: all advantages from the standpoint of the industrial plant manager who wants to keep her focus only on the plant’s core mission of producing tires, or shipping rolled pipe, etc.

What’s Good about Buy-Burn-Bash

It is simple to implement and can work well in situations where equipment is not maintainable, has short “natural lifetimes,” or where the consequences of failure have been mitigated in system and plant design.

What’s Bad

In many situations and in many ways this ownership policy and the service and maintenance approach it provides is decidedly non-optimal. Particularly if misused or poorly managed, this ownership-operating policy can throw away significant portions of useable equipment lifetimes.

The Standards-Driven Ownership Policy

Ultimately you can’t tell in advance exactly which units need service and which don’t. This is the key concept at the heart of a standards-based approach: a dozen medium voltage circuit breakers might look exactly the same from the outside; but one cannot tell until each unit is disassembled which will need service and which will not. Therefore, if owners expect dependable service they will have to service all twelve even if only four really are in need of that service.

This ownership approach manages operation, inspection, testing, service, maintenance, and refurbishment by sets of rules. Breakers of a certain type must be applied within a set of tightly prescribed operating guidelines and limits. Inspection will be done every XX years and those results followed up according to written rules carefully defining what is to be done or not done with regard to what is found. Every MM number of operations, or YY years, these breakers will be disassembled and serviced. After ZZ years, or when FF number of breakdowns or failures to operate have occurred, these tests will be done. If results are within a certain range, the unit will be sent for a factory rebuild. And so forth.

For example, based on the engineering and design studies that went into a high-voltage particular breaker, a manufacturer might recommend that the breaker contacts and trip mechanism be inspected and serviced after every six maximum fault-clearing operations. Over time, one particular owner might determine that for the set of two hundred such breakers that it owns, that service

can be done after every ten cycles. Perhaps its system has lower fault duties and thus operational stress is lower, or perhaps the manufacturer's recommendations were very conservative. Regardless, based on its experience, and its study and evaluation of past "as-found," it has concluded that the longer inspection cycle will not cause problems, and over time that has proved to be the case. The rule becomes "inspection and fix as required after ten operations." On the other hand, experience and "bad luck" may have led another company using this approach to determine that service after only five operations is required if it is to have the dependability of operation it expects from the set it owns.

Both companies manage maintenance of these breakers via a set of rules that specify when, where, and how that service will be done; the rules just differ slightly. Each company will have similar of standards for all other routine and periodic ownership, operating, and service decisions it must make with respect to breakers and allied equipment (relays, wiring conduits, etc.), each determined so that, if followed, the desired results, such as dependable operation, etc., are highly likely to occur.

Similarly, rules are established in a similar manner for the service of transformers, and for capacitor banks, and line regulators, and all other equipment. Electrical standards are established that set load and voltage limits on a particular type of transformer so that "premature" failures due to overheating and loss of life or overvoltage are extremely unlikely to occur. Voltage withstand levels for electrical cable will be set high enough that even as that insulation steadily degrades over time in service, it remains sufficient to avoid flashovers and failure for many years in service, and inspection of those cables will be specified to the extent determined to be necessary to keep failure rates below "unsatisfactory" levels. Intrusive inspections (disassembly, etc.) will be called on major equipment at frequent enough intervals to catch problems before experience has indicated they will get out of hand. Case and tank inspection and maintenance standards will call for inspection and re-painting on a certain schedule, all to avoid "premature" rust-through. And so forth.

These rules – standards -- are developed over time and may evolve as conditions change. When a technology or equipment type or situation is new, a combination of manufacturer's recommendations, engineering assessment, and past experience thought relevant might be used to set rules. The standards might evolve rapidly as experience is gained and the organization "learns by doing." A company may share and compare the standards it develops with other similar companies, participating in "standards groups" where its personnel meet with like personnel from other companies to discuss best standards with regard to use, service, replacement, etc. These groups might issue "industry standards" which reflect the consensus of all participants on what is recommended in general, but the company will modify them as it thinks fits its needs best.

Companies practicing this approach manage the results they get from their equipment by managing their standards. Generally, only when the ownership experience is not satisfactory will the standards be changed; if the equipment's

lifetime appears to be too short (e.g., MV cable is failing in large amounts after only 12-15 years in the ground), if breakdowns occur too often (breakers stick), if durability and dependability are inadequate (poles appear to be deteriorating too quickly), if safety becomes an issue (lines are failing on the ground more than expected) – the organization will study the causes and change the rules for either the specification of equipment, its utilization, or its maintenance, until it gets what it wants. Thus, over time, the organization's standards evolve until the expected level of results – equipment lifetimes, failure rates, and operating costs, are satisfactory. Thereafter, as long as there are no unsatisfactory results, the standards do not change.

Condition-based standards-driven approach. Many modern practitioners of the standards-driven approach partition otherwise similar equipment into sub-categories based on diagnostics (e.g., DGA tests for transformers) or other critical factors (years in service) and vary the rules for operation and service in each sub-category (older units are not permitted to be as highly stressed, and/or are inspected more often).

This sub-categorization-by-condition-or-age approach is to a great extent driven by the aging infrastructures in the electric utility industry. As large numbers of equipment reach an age where the wear and tear of long service caused them to need more and different types of maintenance and refurbishment, the benefits of testing to determine condition are seen as part of the solution to a situations where results (failures, breakdowns) were getting out of hand: change the rules but only for a sub-set found to need the revised procedures.

What a standards-driven approach does that other approaches don't

Relatively little condition assessment and tracking. In its purest form, this approach sees little value in diagnostic methods, testing and inspections, etc., aimed at estimating condition, and therefore does little or none; the “fact” that such testing/tracking methods are undependable is a core concept – corporate psychologists would say a “central cultural belief” behind the whole standards-driven approach. And the truth is that traditionally testing and condition tracking offered only marginal improvement in results, other than in the cost of the service, because when a standards-driven approach is applied well failure rates are kept at low levels, breakdowns and mis-operations are similarly limited, safety is as expected, and, as near as can be determined, equipment lifetimes appear to be under control. Thus, this approach usually does little testing, recordkeeping and condition tracking, etc.

As stated in the introduction, very few organizations operate purely within one or another of the four approaches discussed here. In the 20th century, most electric utilities and some large industrial power owners did operate in essentially this manner alone. However, some now do testing and use those tests to sub-categorize major equipment such as breakers into sets that require different periods or types of service, as discussed a few paragraphs above.

Still, organizations dominated by this approach to equipment ownership and

the operating approaches it leads to do less testing, diagnostic analysis, and condition tracking than any other but the B-B-B. And getting a purely standards-driven organization to try diagnostics and testing to estimate and use condition in maintenance scheduling, unless there is a problem with something other than cost, is a hard sell. What they do that others don't is:

Service some equipment known to not need service. Even when applied with sub-categories based on diagnostics and condition-related criteria, this approach will still apply "one size fits all" rules to all equipment in a category or sub-category, with an aim of servicing all in order to assure the service of the most needful.

Where a standards-driven approach is popular

Traditional electric utilities and for power systems at military facilities and some large institutions. Traditionally – meaning for most of the first 100 years of the electric power era – all public electric utilities followed this approach. Each company had a standards department that maintained its set of rules for every needed application and worked to see that they were available to and could be followed by all who needed them. In every department and function of the company, work processes were defined and guided by the standards: do the maintenance like this – "The book says so. Don't argue with the book."

Why the standards-driven approach is popular

A standards-driven equipment ownership and maintenance policy became nearly universal within the electric power industry because it worked in a utility environment and was manageable with the technology, human resources, and management structures that utilities had to work with during the middle part of the 20th century.

The simpler "buy, burn and dispose" approach never worked for utilities, mainly due to the different loading characteristics of equipment on a utility system as compared to heavy industry, where B-B-B is popular. Unlike industrial equipment, utility equipment may see its peak rated loading only a few hundred hours a year, and go through high-load/low load cycles on a day by day basis that limit periods of high temperatures to only a dozen hours a year or fewer. Beyond this, fault duties over much of the system are often far, far less than rated breaker duties, meaning many equipments seldom see the maximum mechanical stresses they are designed for during "abnormal events." As a result, insulation deterioration rates and the incidence of high levels of mechanical stress events are far lower than for industrial plant equipment: expected lifetime is decades.

Further, the same conditions that lead to multi-decade expected lifetimes of insulation, etc., often exacerbate other, preventable deterioration modes. For example, a transformer at a large industrial plant that is loaded near its rated capacity 8,760 hours a year will be hot all the time. Water will not remain mixed in its oil, condensation will not develop on its case, and any moisture that seeps

into cracks or crevices will be evaporated quickly. But for a similar transformer that is loaded to peak only a few hours a year, water retention and the problems it causes, including rust, is a likely failure mode. Thus, a utility will see value in routine inspection and service to avoid rust, whereas the industrial owner would not. The utility will perform routine inspections, periodic service and re-painting with epoxy paint, etc. How routine, and how comprehensive that service should be will be defined by standards developed over long practice and honed to work best in the utility's climate and situation.

While the other ownership philosophies discussed here are not uncommon, and several to be discussed later are growing rapidly in use, the standards-driven approach is still the predominate approach for many aspects of electric utility and some industrial application.

Good Aspects of the Standards-Driven Approach

It works: the lights stay on, for the most part equipment lifetimes are reasonable, and O&M and safety issues are controlled to acceptable levels.

The standards-driven approach was developed in the 1920s and 1930s. It is difficult to imagine any way of managing maintenance so the worst cases are caught and fixed, other than this approach, given the limitations of those times. That was before line trucks and dispatch centers had radio communication, before most modern diagnostic technologies were invented, at a time when the national labor pool of technically educated employees was very limited, when records were handwritten and kept in file cabinets, when engineering calculations were done with slide rule and pencil, when equipment diagrams were hand-drafted, and before there were photocopiers, mapping systems, on-line call center systems, non-intrusive and non-destructive scanners and sensors, or condition-tracking data bases. The standards-driven approach needs routine engineering evaluation that can be done with a slide rule, a book of algebraic formulae and a set of log tables (more involved calculations might improve results but it can get by with only that). Inspection, service and repair rules could be written so they can be carried out by people taken from the available labor force in the early 1930s: the utility could teach them the rules and make sure they follow them and things will work. That work required no more than basic tools and service equipment and management resources of the period: these processes could be managed with paper records, filing cabinets, and back and forth communication "to the field" no faster than once a day.

Today, of course, all the technological and organizational improvements mentioned above do exist. In many cases companies using this approach have modified their standards and processes to take advantage. But no matter how obsolete this approach may seem, the fact is that organizations who implement this approach well get the results they expect: very long equipment lifetimes and very few operating surprises. Furthermore, this approach was compatible with and has become a part of the whole regulated-utility paradigm as it is institutionalized within the power industry; it was one piece of methodology and

habits the industry developed which, when taken as a whole, still work for society as a whole. As the standards-driven approach would say itself, if the results are satisfactory, why change anything?

Bad Aspects for the Standards Driven Approach

The fact that this approach works does not mean that it works as well as possible: perhaps other approaches can increase lifetimes while also driving down costs, reduce failure rates, etc.

Standards – rules – are by their nature rather inflexible, one-size-fits-all rules: “we own 680 15kV class breakers and every one will be serviced every five years, whether it needs it or not,” etc. Ideally, as technology evolves, standards would evolve too, to take advantage of new capabilities: as diagnostics are developed that can identify units that need service and those that do not. Thus, one would expect to see that over time, standards evolve to include ifs, ands, and buts with regard to inspection and diagnostic or monitoring results, etc. But while some of this does occur, standards-based organizations are often prone to resist change. The institutionalized habits of a standard-based organization tend to oppose change if “nothing is wrong.” If breakers seem to be failing too often, the organization will investigate why, identify the problem, and modify its standards or create a new one. But the standards-driven approach is “programmed” to change itself only when the results it has been set up to provide are not being produced; it is not set up to seek improvement in its processes and method. Thus, changes for any other reason beyond a failure to obtain the traditional desired level of results the standards are designed to produce are discouraged.

Reluctance for improvements in method. Where its procedure has always called for a test meter to be used, the adoption of new type of test meter that is easier to use and more accurate might be enthusiastically accepted by the organization: its standard calls for a test meter so the new model fits into the existing process.

But where the existing standard does not call for the use a test meter, the organization will be much less likely to adopt a change to take advantage of the improvement its use might bring. The culture and value system are oriented to avoid any deviation from the results that are wanted. Improvement is not really part of the desired results, and even if it were, the possibility that the change could lead to negative results makes the organization rather slow to make such changes. The bigger the change required in procedure and standards application, the more the organization will resist the change.

Inability to evolve standards with changing usage. A standards-driven approach may not evolve as the organization’s operating policies or use of equipment evolve, until problems develop.

For example, over a period of fifteen years, daily load curve shapes on feeders at one utility gradually worsened in load factor. Meanwhile,

following an almost universal trend in the power industry to cut spending, the utility let utilization rates (ratio of peak demand to capacity) increase across its primary feeder system. And in company with all that, it implemented a comprehensive system of peak load control to shave peak demand during summer.

These three trends all resulted in a great increase in the daily variation in uncontrolled voltage drop during summer months, leading to much higher frequency of operation of both voltage regulators and capacitor switches. Not surprisingly, breakdown rates and failures of both types of equipment increase, eventually to a rather dramatic degree. The standards for inspection, test, and service for these equipments were not modified until the higher usage rates led to failures which led to an investigation into why, which eventually led to revision of standards.

Something similar occurred at some utilities when they moved to a “fuse saving” protection scheme. Under such a protection approach, distribution reclosers and breakers will see heightened use, and wear, as a result. Yet many utilities failed to adequately adapt to this, even though their engineering and management anticipated the higher rate of wear and its consequent need for greater service. The standards process eventually adapted but did not keep pace, meaning the utility went through a period of poor MV protection results until standards were revised in response to “unacceptable results.”

These and other similar instances spring from a common characteristic in the standards-driven approach as implemented at many companies: it is reactive, not proactive. At many companies it responds only to deficiencies in the results it is expected to obtain; it is not meant to and will not lead to improvement in a situation that is “good enough.” The authors expect that those electric utilities using this approach will be challenged in this regard with respect to traditional distribution system needs during the first quarter of the new century. The widespread use of electric vehicles will change loading levels, utilization rates, and operating conditions for many types of distribution equipment in ways that may not be easy to anticipate or that will lead to a need for very great changes in maintenance practices.

Reliability-Centered Maintenance (RCM)

It is possible to identify what equipment needs service, and to focus attention on those units where service and repair will do the most good. This is the central tenet around which RCM is centered. “Reliability-Centered Maintenance” actually includes several slightly different maintenance prioritization methods that the authors have lumped together because they are based upon this central tenet. But all base their method on this fundamental assumption (usually true) that condition prioritization works. There are three variations that merely differ in *how* the organization executes the details of selecting the equipment for attention.

Condition-Triggered Maintenance

In this first approach of the three, maintenance is done on equipment based on measured equipment and performance related factors as compared to trigger points or thresholds. As examples, with respect to circuit breakers, this might be:

Number of operations: breakers are not maintained on a routine schedule every few years, but always after they have performed PPP fault clearing operations, whether that is only six months or three years.

Results of a timing test: the service and maintenance a breaker receives depends on the results of a timing test. If the unit fails badly it is rebuilt or full serviced. If it fails only slightly it is lubed, adjusted, etc., until it passes the test. Breakers that pass a timing test receive no service.

Similar rules would apply to other switches and mechanical equipment such as reclosers, load tap changers, capacitor switches, etc. For transformers, the rules might be based upon dissolved gas test results, or upon records of load (units loaded to emergency rating for more than two hours are inspected and serviced, etc.). For structures, there would be other rules, and so forth.

The types of data used to determine if maintenance is done, and the rules that trigger decisions to do or not do maintenance, vary a good deal due to the great variety of equipment types in use in any large electric system and the different companies making decisions as to how to implement RCM. Regardless, the common elements of condition-triggered maintenance are:

- 1) *Performance or recent operating results*, or test results for present performance or condition of the device are used to make the decision.
- 2) *Comparison to a threshold or rule* that determines if the unit will be serviced.

In the authors' view this approach is not a full RCM or CCM approach – it straddles the boundary between standards-driven and the full RCM methods to be discussed below.

Standards triggered. This method depends on rules to determine maintenance. They are condition-based, but they are still rules, applied equivalently to all equipment of a type. Their contrast with other trigger methods will be highlighted later in this section when those methods are presented.

Condition-based. The method distinguishes equipment based on actual data related to each unit. In this sense it represents a tremendous change from the traditional standards-driven approach, in that it recognizes, acknowledges, and acts upon the fact that *individual* equipment of the same type (e.g, MD-80 breakers) and situation (old, heavily used) could

nonetheless have very different characteristics and maintenance needs. Over time, this change alone can and will have a measurable effect on the “culture” of a maintenance department (See Section 15.3): this shift in group instinct and thinking is often as important as the actual improvements brought about by use of this method.

“Before” condition only. Decisions on service and maintenance are triggered only on the basis of condition prior to maintenance: “This breaker has operated six times so it will be serviced,” or “The transformer is gassing above the threshold rate so it will be torn down for inspection and rebuild,” etc.

Condition Models and Prioritization

By contrast, other categories of RCM include three aspects not a part of condition-triggered maintenance approaches. These are:

“Before and After.” Decisions on service and maintenance consider the impact maintenance will make on the condition or reliability by comparing status before to that expected after : “Based on its condition now the unit will contribute an estimated 129.8 customer minutes SAIDI next year, but if we do the service that drops to 417.6 minutes. The difference justifies the service cost,” or, “This life-extension service typically extends expected lifetime of cable in this situation from another 4.5 years to 13.0 years. The PW value and the lessened likelihood of an outage affecting SAIDI justify the cost.”

Model based. Condition-based and reliability-based maintenance methods typically use a model (although the organization might not call it that) of condition or reliability, to combine *several* measured data or tracking factors into a *single-valued* measure or classification of benefit or condition for use in determining service and maintenance decisions. Examples of this metric, used as the benefit in B/C computations, are:

Condition class. Dissolved gas test results, unit age, measured levels of audio and RF noise generation and the effects of possible service are combined into a measure of transformer condition: Near-new, Good, Fair, Poor.

Availability or unit reliability. Dissolved gas test results, unit age, measured levels of audio and RF noise generation, expected loading levels, and the effects of possible service are combined into a measure of expected availability or lack of it.

Service reliability or impact. Dissolved gas test results, unit age, measured levels of audio and RF noise generation, expected loading levels, number of customers served and factors describing contingency capability of the system in and around the unit, and

the effects of the possible service are combined to estimate the customer service reliability for the next year.

Prioritization: Finally, true CCM and RCM methods prioritize potential or candidate projects and programs on the basis of benefit/cost, ranking them from most effective to least. Chapter 14 showed more fully how RCM methods work with detailed examples of both the models used to estimate before and after condition or reliability and the prioritization procedures used.

Once projects are ranked by benefit cost, the organization can approve and fund projects in order of effectiveness: if it has a very constrained maintenance budget this year, it picks only the “cherries” – those projects with very high benefit/cost. It thus gets as much “bang” for its limited bucks as it can afford. If it has a larger budget it picks projects farther down on the list – those that are best in the sense of what has not been selected so far – until the planned spending reaches the new, higher budget level. Prioritization does three critical things for the organization.

- 1) It maximizes the benefit it gets from the maintenance effort and spending it makes. Condition-trigger methods offer improvement over standards-driven results but do not maximize benefits.
- 2) It applies budget constraints in an orderly and effective manner so that the organization can spend what money it has most effectively.
- 3) It provides a means to evaluate the cost-effectiveness of the entire maintenance program and budget, which is why it is one part of the larger Asset Management approach, as will be discussed later.

Condition-Based Maintenance

Within this sub-category of RCM there are two further variations on the RCM theme, differing with respect to exactly how maintenance and service resources and monies are then allocated in one of two ways:

Maintain all units within a band of “good condition,” servicing them when they are estimated to be at the bottom of that range with a goal of bringing them up to the top of the range (but not perhaps to a level “as good as new”).

Maximize effectiveness of maintenance and service. Here, resources are added to keep units above some minimum threshold (“good enough”) but not necessarily to bring them all the way to some “top condition” level. Instead, the goal is to maximize the cost effectiveness of the entire program: if \$1000 of service, re-painting, etc., will “buy” three years of further durability in operation for a particular breaker, at a cost of \$333/year, and \$2,000 would buy five years (a cost of \$400/year), the organization will opt for the more frequent but lower cost program, etc.

In both cases, the organization uses a “model” to consolidate and combine *multiple* data on equipment condition or status into a classification, numerical condition measure, or other factor that is used to determine the merits of potential service and maintenance. For each potential project (i.e., do full service on breaker EA-36M) this is used to compute the “bang for the buck.”

$$\text{Merit} = (\text{Condition measure before} - \text{Condition measure after}) / \text{Cost}$$

This is computed for every candidate project being considered. The projects are then sorted highest to lowest. By picking the project at the top and then working down the list, approving the next, and the next, etc., until the permitted budget is exhausted, management can select the combination of projects that maximizes condition improvement.

Constraints: The ability to handle constraints is important, such as “We only have crews to do 16 full breaker rebuilds this year,” or “If we do Project P, we must then also do Project Q.” “Spreadsheet” CBM ranking methods can handle a few constraints, but the best CBM methods use formal optimization algorithms to handle situations where the number of constraints outnumber the candidates involved.¹

Aging infrastructures are explicitly addressed with this method by making age one of the factors going into the condition model. Doing so effectively means having a statistical base of analysis to develop formulae for [age+other factors]→ condition estimates. This will be discussed in more detail in Chapter 16’s examples. This approach is sometimes called *reliability-centered maintenance*: the condition measure (figure of merit in the equation above) being the expected availability (reliability) of the equipment during the next year or several years. However, the authors reserve that term for the method described next.

Reliability Centered Maintenance

This method is basically like the CBM method explained above, except that the “reliability” being optimized is reliability of service as seen by the utility customer or the industrial plant process. To do so, RCM uses additional data not used in the CBM process, information on equipment’s place in the system, the number of customers it serves and its criticality to their service, and the restoration time for a failure, etc. For example, suppose a utility can do maintenance on only one breaker but has two candidates for service. Both are in equivalent condition. Unit A serves 500 homes, is on a line that historically has had 3 faults every 5 years, and has an expected restoration time of 3 hours. Unit B serves 800 homes, on a line that is expected to have events 7 times in every five years, but where restoration time is only 90 minutes. A unit that has not been serviced has a 5% chance of misoperation, whereas a unit that has been serviced is evaluated as having only a 2% likelihood.

¹ One recent CBM project in which the authors provided assistance involved slightly less than 940 potential maintenance and service project, but over 2,700 constraints.

Reliability gain from servicing A:

$$500 \text{ homes} \times 3 \text{ events/5 years} \times 3.0 \text{ hours} \times (5\% - 2\%) = \\ 27.0 \text{ cust-hrs/year gained}$$

Reliability gain from servicing B:

$$800 \text{ homes} \times 4 \text{ events/5 years} \times 1.5 \text{ hours} \times (5\% - 2\%) = \\ 28.8 \text{ cust-hrs/year gained}$$

Unit B will get the service because it will yield a 7% greater impact in expected reliability improvement. All available resources for breaker maintenance and service will be allocated among breakers in a similar manner.

Furthermore, the organization will perform such “reliability gained” optimization to determine what it should service, and how comprehensively while comparing the cost-effectiveness of servicing breakers in this way to that of inspecting and pro-actively replacing cable, etc. Money and resources will be allocated between these two service areas, and to the inspection, testing, service and refurbishment of other equipment types, based on this type of analysis. Beyond that, the total budget for maintenance and service of all electrical equipment might be determined from this analysis, as the minimum needed to maintain reliability at a certain level. This type of optimization is often referred to as Pareto optimization and is discussed in much more detail elsewhere (see Chapters 5 and 12).

What CBM and RCM Do That Other Approaches Don't

First, an organization following this approach believes that it is possible to track and estimate the condition of equipment, and that it can determine in advance how continued service and various types of service and maintenance will affect condition in meaningful ways; it has confidence in the analysis behind RCM.

RCM organizations do more inspection and more testing than those following buy-burn-bash or standards-driven paradigms, maintain equipment logs and records, and estimate and track condition. They also “optimize” allocation of resources (the aforementioned Pareto optimization).

Where CBM and RCM are Popular

Among utilities and organizations that want to maximize the reliability of limited budgets. An organization practicing Buy-Burn-Bash and having no issues with affordability of the approach will seldom even consider shifting to RCM: its value system and culture sees no value in optimizing something (service of equipment) that is of little value anyway.

Similarly, an organization satisfied with the results it is getting from the use of its standards-driven approach will see no any reason to adopt CBM/RCM either. It is getting the results it wants: nothing is broken so don't fix anything. However,

if challenged by a limited O&M budget, it may consider moving to RCM as a way of getting its traditional level of satisfactory customer (if a utility) or plant (if an industry) reliability at a lower overall cost. CBM/CM also appeals to companies that want predictable and equal budgets from year to year. Often, a budget driven by standards will vary from year to year as maintenance needs change (in addition to just natural variations, more severe weather one year might do more long-term damage and require more repair the following year, etc.). Regardless, a standards-driven company has no way from within the standards-driven paradigm to reduce spending if needed. With CBM/RCM, a company's financial planners can reserve a certain amount of money for maintenance each year and CBM/RCM can spend that "well," getting the most from it. Furthermore, it can estimate in advance what impact the budget reduction will have on reliability, telling decision makers the consequences of their action.

What is Good about the CBM/RCM Approach

CBM/RCM maximizes reliability gained within any budget or resource constraint such as a company having only a certain number of training repair personnel. When applied to a standards-driven organization's budget, it will improve the reliability-related results of that organization's budget (but perhaps not other types of results). As mentioned above, standards one-size-fits-all rules but at many companies there are enough ifs, ands, or buts written into them that they are fairly flexible and tend to put maintenance dollars where they are effective. Still, a shift to an CBM/RCM approach will spend maintenance dollars as well as possible for reliability, not just well. A good RCM program will typically give a 5 – 10% improvement in reliability gained per dollar.

What is Bad about the CBM/RCM Approach

First, CBM/RCM might be "optimizing" the wrong thing, or more typically, something that is not quite what should be optimized. For example, a utility might have decided that SAIDI (the average time a utility customer is without power in a year) is a large factor contributing to the RCM measure or "bang" when it maximizes bang for the buck. RCM will give it good SAIDI per dollar, but there are other reliability-related aspects of equipment management it does not address well. For example, durability – something conceptually close to reliability but not quite the same thing. An RCM program aimed at reducing SAIDI will spend a lot on vegetation management (traditionally called tree trimming), as that is generally among the most effective means of controlling SAIDI. As a result it will spend far less on properly painting equipment that is otherwise going to rust, and on inspecting and replacing seals, etc: money spent there might avoid reliability problems someday, but probably not soon.

As a result, durability-related issues such as painting and maintaining equipment integrity may suffer. RCM proponents will point out that this is not really a problem: over time the rust will become worse and seals will loosen, and no doubt at some point threaten reliability. Then and only then RCM will fund

projects to correct those problems.

This is correct – *eventually* RCM will address the problem. The problem is that in many cases RCM is not set up to balance long-term versus short-term issues well or to recognize the benefits of continuous year to year programs over specific “now” projects: to do so increases the difficulty of nearly every phase of its analysis, evaluation, and management. For example, maybe \$13,000 per year, will inspect, sand, prime, and paint all the breakers in the system, keeping those aspects under control, as it were. However, waiting until rust and corrosion are serious enough to jeopardize reliability could mean the cost will be \$1,500,000. At typical discount rates used by utilities and businesses, it would be far better to spend the \$13,000 per year. However, most RCM methods, as implemented, would see little gain in reliability in the near term due to this approach: they are trying to buy reliability now, not eventual savings.

Second, in many cases CBM/RCM is eventually used to rationalize cuts in maintenance spending *and results*. In almost all cases CBM/RCM is initially adopted to either obtain better results from the same spending as a former standards-driven approach, or the same results with slightly reduced spending. But over time, the temptation is high to make further budget cuts. “We’re facing a tough year – what if we re-optimize reliability again?” Further cuts are made and reliability falls below traditional, and perhaps necessary, levels: what is being spent is cost effective, but not enough is being spent.

These two potential problems – not optimizing long- versus short-term considerations, and using CBM/RCM to make cuts that are too deep in spending at the expense of an attribute (reliability) that is important to the organization, are both examples of a bigger issue: CBM/RCM methods optimize only one aspect of the organization’s interest without considering others or the impact on the organization as a whole.

This and other similar problems an organization might have with the CBM/RCM can be “fixed.” The concept of evaluating options on the basis of “bang for the buck” can be extended to factors other than just reliability. Ranking based on these evaluations can be done to identify which projects and programs are best. All of this can be put into a context that understands whether it is best to spend money now, or later, and how much, and what is best for the business, and not just to maximize a particular attribute like reliability. There are complications that arise when this is done, but they can be mastered, and an “all encompassing, business-based optimization” can be made to work. The resulting ownership-operation approach is called Asset Management. Figure 15.1 compares the three basic approaches within RCM.




Figure 15.1 Comparison of the three types of "Reliability" Based Maintenance.

ASSET MANAGEMENT

Our electrical equipment exists to support our business in attaining its goals. A true asset management organization operates as if all of its assets, including electrical equipment, are business tools whose sole purpose is to further the business goals of the organization. Electrical equipment does not exist to provide power, reliability, and voltage within standards. Those products of equipment ownership, and their quality or lack of it, are important only to the extent they help the organization in obtaining its fundamental goals. Electrical equipment exists to make money, if the organization is profit-focused, promote growth, if the organization wants to grow in one manner or another, or to help maximize the number of orphans who can be housed, fed, and educated, if it is a charity focused on that goal.

How Does Asset Management Work?

Basically, the concept behind CBM/RCM – evaluate options on the basis of benefit per dollar and buy the most cost effective ones – is extended to other important desired results of the organization, and not applied just to reliability. The optimization is applied to all the many disparate goals of the organization – reliability, safety, durability, electric performance, work force morale, a desire to save not spend money, etc. – are somehow combined and coordinated in a context that makes it all work. That, in a nutshell, is asset management. The extension to other qualities beyond reliability uses basically the same approach as for CBM/RCM.

1. The organization must define a metric, or metrics, that measures the desired goal: SAIDI for reliability, maybe lost-time-injuries for safety, number of low voltage and sag complaints for electrical performance, etc.
2. It must be able to evaluate any and every project or program it could spend money on with regard to its impact on every one of these metrics, or have confidence that it has no impact on those metrics it cannot measure.
3. Pareto optimization – ranking the options by “dollars per unit of desired result” – and selection of the highest rated opportunities in each category can again be applied to identify which projects are best.

Only there are complications when more than one result (reliability in CBM/RCM) is optimized simultaneously. These complications can be mastered but they require considerable expansion of the analysis and optimization process, and a lot more work.

Dealing with More Than One Attribute

An attribute is a quality that an organization wants to maximize (or minimize).² Typical attributes sought with regard to electrical equipment are service lifetime, reliability, safety, power quality, and cost. The important point here, and the weakness identified with CBM/RCM in the previous section, is that whenever electrical equipment is concerned, there is always more than one attribute. Reliability might be *most* important to one organization, but equipment lifetime and low maintenance costs are not unimportant. Predictable and steady cost streams – with no unexpected expense surprises – might be the key requirement to another.

When there is only one attribute, evaluation and ranking are simple: measure each option with regard to the attribute's performance metric (for reliability this might be SAIDI), determine its cost, and compute its cost-effectiveness – the ratio of the two values. Sort candidate options – alternatives the organization *could* do – from most cost effective to least and fund projects by starting at the top of the list and working down, until one runs out of money or buys sufficient performance to satisfy need. But when there are two attributes, the option that is most effective for one attribute (e.g., reliability) is not necessarily the most effective for the other (safety). The list of candidate projects can be sorted best-to-worst for either attribute, but not both at once. The situation only gets worse when there are three, four, or more attributes.

Multi-Attribute Optimization

“Multi-attribute” optimization approaches address the following situation.

- There are several attributes – each a “result” the organization wants as much of as it can get or “bad results” that it wants to minimize or control.
- Metrics – measures of the result sought – can be defined and applied for each attribute. They may be very different types of measurement for the various attributes: customer-minutes lost for reliability, lost-time injuries for reliability, voltage variation for power quality, Miller Alienation Workplace Satisfaction scores for employee morale, and so forth.
- Any project or program can be evaluated with regard to all metrics.

² Attributes differ from criteria in that the organization wants as much of an attribute as it can get but only a certain value of a criterion. A *criterion* is a limit or target that must be met but for which little or no additional value is given if it is exceeded. Voltage must be above 110 volts to provide good service. Paint must be 7 mils thick to have enough adhesion and wear resistance. A *result* is a *criterion* if more than that desired limit is not needed. By contrast, more reliability is always a good thing, better safety always to be sought after. They are *attributes*.

Most projects will score zero in some categories but many will have some score in several: a project might impact reliability, have some influence of safety, and affect future O&M costs, etc.

- The “best” budget is sought: a combination of projects that obtains more results across this set of attributes, as the organization sees it, than any other.

This chapter is not intended to discuss how one solves this problem mathematically – single and multi-attribute optimization methodology is discussed in greater detail in Chapters 5 and 12 and given as examples in Chapter 16’s case study. What is important here is the basic ownership philosophies and understanding that lead an organization to *want* to apply that asset management approach, and how those approaches interact with and shape those philosophies.

So the important point here is that this problem can be solved, after a fashion. “After a fashion,” means this is done mainly by “structuring the problem” so that it is both mathematically tractable, and the answers obtained are practical (i.e., the answer is useful). And that is done via the steps discussed below. Not all are used in exactly the same way in all cases – there are slightly different ways to do asset management. However, in general, the following steps are taken anytime “Asset Management” is used by an organization:

1. **Money is the “Attribute of Attributes.”** The asset management approach is centered around minimizing cost or maximizing profit or optimizing some other monetary metric such as risk of failing to meet business goals. There is no doubt that this is done in many organizations because ultimately money is as or more important than any other result. But more fundamentally this is done because money – cost – is the one common thread relating all the candidate projects and programs that could be done. Some projects affect only safety, some just reliability, some reliability and electric performance and nothing else, some only interact with workforce morale and do nothing for reliability. But each and every one of them has a cost.

Strictly speaking, money is an attribute: the authors have never seen an organization or person that does not want more money if it can get it or more savings if that can be achieved. The reader confused about the fact that “single attribute” optimization such as CBM/RCM involves two attributes, reliability and money, needs to keep in mind that terminology is not always scrupulously correct and that money is always treated as distinctly different. The term “single attribute” optimization almost invariably refers to optimizing some desired attribute, be it reliability, or safety, or public image, against the “attribute of attributes,” money. Asset Management always includes money, too, and in some sense, it is also more important than any of the others, at least to the process and the decision making,

within which it is the basis for how evaluations and decisions are made.

2. **Include Both Capital and O&M Costs.** Multi-attribute optimization can be applied to only O&M spending, or to only capital spending. However, in so doing it will tend to optimize that budget at the expense of the other. Multi-attribute optimization of a utility company's O&M budget will maximize the effectiveness of O&M monies spent, but it might do so in a way that has negative consequences for capital. For example, an O&M spending focus on reliability and safety might neglect to spend on durability, so a utility gains good customer service and safety but spends little on paint and service to extend equipment lifetime: future capital costs for replacement are increased. The authors consider that true "asset management" involves the simultaneous optimization of both capital and O&M spending, considering the interactions and effects of one to the other.
3. **Study the Trade-offs Involved.** The organization wants as much of every attribute as it can get: zero-accident safety, world-class reliability, infrastructure durability better than a medieval cathedral's, etc. But an asset management approach will quickly show it that it cannot have everything it wants, but that it usually can have anything it wants: there are *trade-offs* involved. One either spends a lot of money, or accepts less of one results in order to get more of another. As one optimizes capital spending in one way, it has consequences on O&M spending in another way. And so forth.

Asset Management is favored by executive management in many industries, including the power industry, largely because it can quantitatively identify trade-off options the company faces in this manner. True, it helps the organization "optimize" results – it leads to spending the final O&M budget that it decided upon as well as possible. But executive management is often more interested in *understanding* its choices than in honing its operating plan to the last decimal place. In fact, in the authors' experience, organizations often do not choose the mathematically optimum plan: executives may think that the math does not quite address all the intangible aspects of their goals. But understanding the options the company has, and how tradeoffs affect performance and cost, capital, and O&M, one desired result versus another, is the key to the executives making their decisions well.

So, for the investor owned utility, what *will* it cost to improve SAIDI from 78 to 60 minutes if spending is optimized to do so at the least possible cost? If the company does not have the additional money to afford that, what reductions in durability, in electrical

performance, in safety, etc., would it have to accept to free up enough money to fund that reliability gain? What if it is unwilling to reduce safety and wants to “take it from” the remaining areas? Can that be done? What will the resulting impacts look like? When will they set in? How long will they last?

Similarly, for the charity organization focused on orphan care in third world countries, will buying twenty high-quality 150kW field generators and maintaining them “by the book” enable it to care for more orphans than buying 25 lower cost units and skimping on maintenance, or not? For how long? With what risk of falling short of the expected goal?

This is the type of information executives in an asset management organizations want and use to make decisions with regard to their business, including their business’s involvement with electrical equipment, be that incidental or key to the company’s core mission.

4. **Select a Plan that Maximizes Business Results.** All of the above can be considered tactical evaluation and planning – analysis of how to maximize implementation and the options the company has in that implementation. The strategic part of asset management involves how the company decides what combination of trade-offs to accept, and why. There is seldom one clear best answer. A strategy that emphasizes durability (long-term lifetime) and “even business results” (steady year to year spending, low risk of failing to make profit targets) may emphasize different programs and spend money differently than one focused mainly on customer service reliability and system renewable, etc. All four aspects – durability, business results, reliability, and and system renewal – are important, and the organization can have more of one or two by trading away some of the others. It can have more of all four only by giving up more of the “attribute of attributes” – money.

Decisions at this level are always made by the people running the organization, and are basically strategy and policy and are made by the executive management of a company. Fundamentally, any organization, be it an electric utility or the charity, values electrical reliability, low O&M cost, safety, employee morale, and other results because they contribute to its success. A particular plan is selected because those making the decisions – the company’s executive management – ultimately decide it has the greatest likelihood that it will result in the organization’s success, whatever that means to them. For an investor owned utility these might be operating profit, stock price, or executive bonus level. For a children’s charity this might be number of orphans cared for and public image leading to increased contributions next year.

Many companies, particularly large technically-inclined businesses, use the “output” of the electrical equipment asset management analysis as the input to a corporate model that those making policy and strategic decisions use to chart of the overall course of the organization. In many cases this “model” is a computerized numerical model that juggles investment, O&M costs and its emphasis, market share, pricing strategy, workforce size and skills, and investment in all these areas, against financial factors like cash flow, debt ratios, contingency needs, etc. In other cases the model is just the instinct and intuition of top executives, honed by years of experience and applied according to their judgment.

Either way, such models are used by the executives running organizations to evaluate their entire business, and answer essentially the same questions as asset management did for the electrical equipment: “What can I do and what can’t I do: what are the options and tradeoffs I have?” The information on the electric system and the asset management options that an organization has is only one set of inputs to its “corporate model” but often a very key one if electrical equipment is an important aspect of the organization’s function of needs.

Regardless, working from the information provided, the company executives pick a plan. That plan in turn rests on certain options from the electrical equipment asset management, which are now essentially “set in concrete” as part of the overall plan. For an electric utility, the selected plan might be something like “We will spend \$511 million, for which we expect a SAIDI of 75 minutes, safety below .8 lost-time injuries per 100,000 hours labor, power quality leading to fewer than 25 complaints per 1,000,000 customers per year, a major equipment durability index of 47 years expected lifetime as we measure it, employee morale greater than 85 on the Miller Alienation workplace satisfaction scale. Details of how we get each of those and how they hang together were determined in the asset management analysis given to us. That plan and its recommendations are approved, and we expect the results promised to follow from that spending and effort.”

5. **Set and Manage to Targets.** The approved electric system asset management plan identified how much it would spend to be spent on each area such as equipment service, replacement, training, etc., and it identified each of the results that would be generated such as reliability, employee morale, etc.

This plan now permits the organization to pretty much de-couple the management of various aspects of its operations into single-attribute situations within each division or department. For

example, suppose the asset management plan that has been approved called for \$13 million for MV breaker maintenance. That budget can be assigned to Substation Maintenance for MV breaker maintenance, and within that, it has only to focus on doing projects that are most cost effective – using the money well. This is a bit of an oversimplification (breaker maintenance involves following safety practices in order to achieve the companies goals), but illustrates the point: the utility manages its overall combination of results it wants by studying the trade-offs and then assigning “vertical” responsibility among various groups in a way that makes each of their assignments pretty close to a single attribute management focus.

And so ultimately, the business sets its tactics based on the strategy it selects. But the process it took to select a strategy took into account an evaluation of the options it had with respect to tactics. An example case later in this chapter will discuss how this works and give diagrams of the process used at one asset management company.

Regardless, good executive management will define and select a strategy it considers is best for the organization. This will vary greatly from company to company: an investor owned business seeks to maximize profit - owning electric equipment doesn’t change that. A non-profit organization might define its goals very differently, as something along the lines of “We want to maximize the number of orphaned children we feed, keep healthy, and school.”

Typically, an organization’s strategy will not be to maximize that results it *might* achieve, but to *minimize the risk of failing* to meet targets for those results that it promises to meet. Executives running an investor owned utility might reject a plan that promises them a good chance of making 12% profit, in favor of a plan that gives them a near certainty of making 11.5%, etc. “Taking into account factors we cannot control like the weather and economy, we have set our profit goal at \$290 million, a figure we are confident we will make. We have set reliability, employee morale, and other goals in a way that is most compatible with both achieving this goal and in being flexible to help us respond if conditions change, so we still attain that goal.”

15.3 COMPARING THE FOUR OWNERSHIP PHILOSOPHIES

Is there a “right way” to operate?

Perhaps the most fundamental contrast among these four operating approaches is with regard to whether each considers that there is a “right way” to practice equipment ownership and operation. As different as buy-burn-bash and the standards-driven approaches are, when implemented in an organization, either leads to a definite view of how equipment should be used and what service maintenance should be done: *there is a right answer*. For buy-burn-bash, this is “use it up and don’t take care of it.” The standards-driven approach holds that

the right way to do things is “by the book” and views anything more as unnecessary and anything less as “lowering our standards.” Over time, organizations that practice one or the other of these two approaches develop a culture that has great confidence in the approach, that thinks and acts from this viewpoint instinctively, and that tends to regard deviations from “the right way” with cynicism, often as if any such proposal is near heresy.

By contrast, both CBM/RCM and the Asset Management approach never imply or lead to a conclusion that there is a necessarily *absolutely* correct way to do things in all situations, although they will identify a best way to do things in any one situation. Use of these approaches tends to put the fact that there are always compromises and tradeoffs into stark view, forcing an organization and its members to never forget that anything gained comes at the expense of something else given up. This is a fundamental difference that eventually leads to very different cultures inside companies practicing these different ownership approaches.

Buy-burn-bash and standards-driven organizations are in many cases strongly convinced they are not only “doing things right” but that there is no reason to consider improvement or change: “It isn’t broken, so don’t fix it.” There is often little internal willingness to review and refine, to improve, or to innovate, and these organizations often have a difficult time adjusting to external changes. Such organizations tend to operate very efficiently, and to make few mistakes with regard to their approach. But they can become complacent.

CBM/RCM and Asset Management organizations are confronted daily with the fact that there are different “ways to do things” and a range of tradeoffs, and that any decision has potential advantages and disadvantages versus the others.

Over time, these organizations develop a culture of constantly evaluating, even challenging the way they do things. Such organizations tend to be slightly unsatisfied with the results in the best of times. In the authors’ opinion, while the buy-burn-bash and standards-driven organizations might be much more confident and comfortable places to work, CBM/RCM and Asset Management based organizations can be much more effective, particularly in the face of change.

How the four philosophies view aging infrastructures

Table 15.2 summarizes how aging power equipment infrastructures are viewed by the four ownership perspectives. Again, readers should keep in mind that this discussion refers to each in its “purest form” and few organizations are purely one or the other but tend to be dominated by one approach while using elements of the others.

Table 15.2 How the Four Ownership Perspectives View Aging Equipment

Perspective	Aging Power Equipment Infrastructures	Can you handle it?
Buy-Burn-Bash	Never an issue. Equipment is expected to age and die. This is how we operate and it is managed completely already.	We already do.
Standards-Driven	A catastrophe. Equipment will no longer behave in the manner expected by this paradigm: we take care of it like we always have, but it won't give us the results it always did.	Not really.
Reliability-Centered	A challenge to maintaining good reliability and a possible concern from the standpoint of cost and resources needed to maintain our targets.	Yes. We relish the challenge to do so.
Asset Management	A challenge to maintaining good business performance and a possible concern from the standpoint of keeping costs in line.	Yes, it's just one more issue we have to balance in running our business.

Buy-Burn-Bash doesn't lead to aging infrastructures. From within this perspective and for true practitioners of buy-burn-bash, "aging infrastructures" don't exist. The authors have met organizations in which the entire electric department staff is skeptical if "aging electrical infrastructures" are an issue anywhere. In the experience of these organizations, equipment is always run to failure and over time replacement is a continuing part of business: the "infrastructure" at any plant or institution is a mixture of old, new, and everything in between.

If these organizations had long institutional memories (and a few have records going back 60 years or more) they would recognize that there have been "infrastructure" changes over time that significantly affected them. Equipment manufacturers change designs, materials, and technologies. As a result, breakers, switches and equipment that varies its behavior generally performs (lasts) longer than in the past, often with service lifetimes that are twice as long as four decades ago. However, equipment that relies on wound coils and insulation generally has shorter lifetimes today than many decades ago. Today, transformers and motors are designed with the aid of intricate computer models to very cost-conscious

standards with little if any margin, whereas fifty years ago that equipment was robustly “overbuilt” and so gave longer lifetimes in some cases. A few organizations in this category recognize this but view it only as interesting and of no particular importance.

The Other Three Approaches Don’t Like Aging Infrastructures

One way or another, aging infrastructures mean rising costs and potential reliability management issues. Either older equipment is replaced, which is very expensive, or it is refurbished or otherwise has its life extended, which is expensive, or one lives with it as is, which impacts reliability and operational costs. Meanwhile, condition assessment, which is costly in its own right, and difficult and imperfect at best, is required to manage the effects as best as they can be. None of the organizations using the Standards-Driven, Reliability-Centered Maintenance, or Asset Management approaches to equipment ownership and operation like this situation. But their reaction to the challenge, and how they address it, is quite different.

Standards-Driven. As discussed in Section 15.2, a standards-driven approach treats categories or equipment, or sub-categories in more modern variations on its theme, with “one-size-fits-all” rules. But aging in service changes what is a set of identical equipment units such as MV breakers, into a set of individual units, each noticeably and some distinctly different than the others and each really requiring individual attention and consideration to determine how to get the most value from it. Aging, and the *occasional* unit that reaches end of its service life and fails is expected by the standards-driven approach; such units cause the low but accepted rates of equipment failure this approach tolerates and, after years of experience, expects.

Aging infrastructures and the equipment management issues create a mass of equipment that is no longer generalizeable by category: the organization cannot guarantee that it will get good results if it follows any type of one-size fits all rules, even applying by sub-sub-categories. The basic problem is that the situation created by aging infrastructures lies outside of the mechanism that a standards-driven approach can handle.

As a result, the reaction to an aging infrastructure among organizations that are dominated by this paradigm is something close to panic. The problem just “doesn’t fit” and nothing the approach can do within its own methodologies will address the problem well enough. As a result, aging is often seen as a much bigger, and more unmanageable, problem than it is. Simply put, a *pure, traditional* standards-based approach can’t handle an aging infrastructure well. After five decades in service, once identical units aren’t. One quite workable rule, even when applied via sub-category and sub-sub-category, does not give the expected results all the time, or at least enough of the time. Condition must lead to varied service attention, and that means the approach is not purely a traditional standards-based approach.

What is needed is an approach that can treat each piece of equipment as a uniquely different unit, to be measured and tracked and considered carefully in order to make the right decisions at the right time with regard to getting what remaining life can be squeezed out of it.

Reliability-Centered Maintenance. This approach can deal with an aging infrastructure, at least in a tactical sense. It can apply inspection, testing, and diagnostic methods, analyze and evaluate condition, and prioritize the spending of available monies and the commitment of limited labor resources, where needed to optimize the condition improvement obtained or the reliability gained. How well it can do this will of course depend on how accurate the testing and diagnosis is, and how well the processes implementing it can compute condition from that and prioritize maintenance.

But the important point here is that approach has a way to handle aging infrastructures, to both deal with those well and to acknowledge and take into account imperfections and uncertainty in testing methods, data, and analysis, and to do the best possible when the organization may not be able to afford all of the budget the method says it should spend. Optimization methods suitable for application here to aged infrastructures, aging equipment, and other aspects of planning and budgeting, are covered in Chapters 5 and 12.

Asset Management. This approach can deal with an aging infrastructure in both a tactical and strategic sense. Like CBM/RCM, it can apply inspection, testing, and diagnostic methods available to it, analyze and evaluate condition, and prioritize the spending of available monies and commitment of limited resources to optimize reliability. As was the case with CBM/RCM, how well it can do this will of course depend on how accurate the testing and diagnosis is, and how well it can compute condition from that and prioritize maintenance. But even more than CBM/RCM, asset management has ways to deal with imperfections and uncertainty in testing methods and to do the best possible with limited funds. Beyond that, asset management can fit the aging infrastructure issue into the overall business case. Executive management might not like the scenarios it sees or the costs it generates and constraints it puts on the business. But within the limits of the asset management's testing, diagnosis, and condition assessment, "the aging infrastructure problem" and its impacts are completely transparent to the organization, and potentially manageable.

15.4 ORGANIZATIONAL CULTURE

Organizations, like people, have personalities – often referred to as corporate cultures – which embody the values and instincts that determine how they react to situations, make decisions about challenges and opportunities they face, and interact with other organizations and society in general. "Culture" includes tacit knowledge – ideas and rules not written down but nonetheless known to and shared by all. It includes long-standing habits and expectations for "how things are done," implicit and explicit protocols of who and what takes precedence and

why, and the values and norms shared by everyone and the organization as a whole on everything from ethics to safety. It affects how decisions are made, how they are carried out, and particularly, how an organization reacts to external changes or challenges outside the norm.

Much has been written about corporate culture, which readers interested in pursuing can find easily via the internet. There are volumes written about it and at least eight distinct theories or perspectives, probably more, about what it is, how it works, and how one characterizes culture and works to change or reinforce it. Some of what is written on this subject – the authors think a great deal of it – is piffle. However, there is no denying that companies have cultures and that seemingly similar companies can have very different cultures. For example two electric utilities can be of roughly the same, be located in the same state, with each focused on service to a large metropolitan area and its suburbs, and yet have very different cultures. The perspective one sees from executive management, the questions that are asked first by middle management and the way they want to frame a new work initiative, and the look in the eye of workforces about to be taught new skills can be completely different at these two companies. And to the point of this book: these two companies could very well differ markedly in how they perceive an aging electric equipment infrastructure. They could have completely opposite views on whether it is a potential catastrophe or even a problem, on if and how they report their experience to others, or even if they do, and in how they propose to manage their infrastructure and the effects of equipment aging, and why.

Missions and Cultures

An organization's culture grows out of the shared experiences of its members, the type of work it does and how it does it, feedback from its successes and failures, and the leadership and influence of its management. It is fashioned by the routine work procedures and priorities established in the workplace, and the belief people and teams invariably develop that their work is important, that they are doing it the right way. Thus, cultures and basic work philosophies/approaches tend to support and strengthen each other in a type of feedback: given that it sees itself as successful, the culture will reinforce its own commitment to the use of the particular set of values, beliefs and methodology, and conversely, the use of that "successful" methodology will strengthen the culture's belief in it and itself.

Consider for example two hypothetical organizations, both responsible for the maintenance and care of a T&D system and its equipment.

- 1) The first is an organization that has long used a pure standards-driven approach. Over time, given that the organization is successful (as it defines that) this led to a culture comfortable with, believing in, and adapted to the standards-driven approach: do things "by the book." The culture will expect to have written rules and procedures to follow at every turn, and consider that adherence to standards and

completion of work to prescribed requirements constitutes in some sense the end of its obligations.

This organization's culture will de-value methods or approaches that analyze and tailor solutions to individual situations or equipment units: after many years with this approach the organization not only does not have the tools to take that approach, its members believe such an approach will not work technically or practically: that it is wrong. Instead, it firmly believes that performance issues, be they equipment operation or high failure rates or reliability issues, are solved looking at *aggregate* results with detail and then changing the rules or procedures to bring results back to an overall acceptable level. In short, it will believe that there is a right way to do things, period, and that that is prescribed by its standards.

When confronted with a sudden problem or unexpected surprise the organization's instinctive reaction is to search within the organization's standards-driven approach for a solution: what do the rules say? It will not think of or trust other approaches.

- 2) By contrast, suppose the second organization has for a long time applied a pure Reliability-Centered Maintenance approach to do what is essentially the same job as the first. Again, assuming it has been successful as it defines success, its members will grow experienced with, respect, and trust the CBM/RCM method. They will believe in methods that track inspection and test results on, that estimate condition of, and that prioritize work activities for *individual* equipment units or projects. Performance issues, be they equipment operation or high failure rates or reliability issues, are solved by gathering data on individual units and events, by looking at and classifying individual results, and then ranking or prioritizing or otherwise selecting items for action.

This culture "knows" that there is never one right way; that decisions on maintenance, and in a way everything and anything else, depend on background and other factors, that *aggregate* performance results can always be obtained by different sets of *individual* decisions and projects, and that all projects are in a type of competition with one another for priority. When confronted with a sudden problem or unexpected surprise the organization's instinctive reaction is to first measure the problem, evaluate its implications (usually with a focus on reliability), and then structure a decision based on maximizing value.

These organizations might have the same purpose, but they are fundamentally very different otherwise. First, regardless of culture, they have very different capabilities at their disposal; neither group *could* tackle a challenge by applying

the other group's methodology, even if it wanted to – it does not have the tools, methods, or skills.

But beyond that, neither group would instinctively or even intellectually consider using the other's approach. People in a standards-driven organization believe in methods that look, perhaps in great detail, at everything in aggregate and in the big picture. They seek to identify and apply broad rules and generalizations that will work in the management of groups of equipment: if 19 out of 93 MV breakers have suddenly experienced a certain problem, the solution is to find a procedure or solution that can be applied as a rule to all 93 – because the problem is probably just as likely to occur to the other 74.

People in the CBM/RCM organization believe in very nearly the exact opposite: look at and treat every unit of equipment as potentially different and apply resources in a focused, not generalized way. If 19 out of 93 MV breakers have had a problem, find out how to inspect and test MV breakers in order to identify this condition anomaly, then evaluate and prioritize its seriousness so that resources can be applied as needed to units that appear to have the problem, and not to the others.

Mixed Approaches

The preceding examples were hypothetical, each “pure” cultures of one type or the other. In actuality, the authors have not seen a purely standards-driven utility maintenance department in over 15 years (but at one time this was the norm in the power industry). Similarly, even among utilities that have long practiced CBM/RCM or Asset Management, there are elements of system and electrical maintenance that are addressed through what is essentially a standards-driven approach. But as mentioned in Section 15.2, invariably one ownership philosophy/approach always dominates.

Still, even if dominated by one approach, keeping a foot in another camp can lead to a more mature, contemplative culture, one that is more flexible in the face of problems and challenges, that can step outside of either paradigm and think “what would work best here?” On the other hand, the authors have occasionally seen “split paradigms” lead to one or the other approach being given an orphaned status: “We do CBM/RCM here, and where the method is ineffective or unaffordable for a particular class of equipment we use a standards-based approach. But since we know that that approach doesn't really fit here, or work well, in effect we do little if anything about that equipment.”

Culture Can Change

Cultures will change and evolve over time. This can happen gradually, and often haphazardly, when something in or around them is changed: at many utilities a change in methodology from standards-driven to CBM/RCM is first resisted by the existing culture, then warily accepted, and finally becomes the core of a new culture. This can take some time, and workplace results and business performance are not always “pretty” during the transition.

Cultural change can also be driven and directed through effort and focus: often when a new CEO comes into an organization, she or he will start by saying that they want to “change the culture.” They mean that the present organization is not capable of reacting and performing as they require, etc. The authors have seen culture change effected at several companies with remarkable results as far as improvements in productivity, system reliability, etc. In the authors’ view, the fact that cultures can be changed leading to palpable changes in performance is the very best proof that they exist.

Cultures Within an Organization

Large companies, such as major electric utilities or global industrial companies may have very different cultures or personalities within different departments and divisions, particularly if those divisions are functionally isolated from other parts of the company. Nearly every large electric utility has three distinct “cultures” operating within it. These are:

Stockholder steward – This culture “knows” that an electric utility is “first and foremost a business,” that it exists to make money for its investors or, if it is a municipal utility, to meet its budget commitments. Financial metrics such as stock price, budget balance, and profitability are most important to this culture and must be met. Members of this culture will look for ways to cut spending, to reduce risk, etc.

The people at the “top” of a utility, and the staffs and departments that routinely work with and support them, are almost always nearly exclusively of this culture, focused on running it as a business, and a successful business at that.

Equipment steward – This culture focuses on the utility’s physical assets – its T&D system. At its worst, this culture believes that equipment and facilities are the be-all of its existence and must be cared for and preserved in good condition regardless of consequences: operators subscribing to this culture’s view may (and in fact have) dropped customer load rather than load equipment up to approved emergency ratings. People in this culture vehemently oppose spending cutbacks even if driven by desperate corporate circumstances.³ Although often not acknowledging it, people imbued with this culture’s perspective think and act as if electrical

³ One of the authors (Willis) sat through an employee presentation a number of years ago as the SVP for a large investor owned utility announced and explained a major cutback in T&D maintenance spending by saying: “I will not put the company in jeopardy to avoid putting the system in jeopardy.” A manager from T&D maintenance in the next chair over muttered quietly “That’s the stupidest thing I ever heard of.” Later, when asked if he realized that a company in jeopardy might go under and might not be able to pay salaries, he replied, “There are more important things than that.” His commitment is commendable but his perspective was definitely distorted.

equipment that is well cared for and not abused will last forever, and believe that in fact it should: any major equipment failure is a failure of the organization to do its job.

Customer, or public, steward – This culture focuses on the utility’s role in the community and on “doing the right thing” for customers. The customer is not only always right, but he always comes first, so the utility’s first obligation is to accommodate its customers’ wants and needs. If a homeowner wants a slight addition to his service extension or a more expensive re-routing of his service drops, he should get it at no charge; if a change in a facility design would please residents of one neighborhood, that ought to be done; if a local community wants a line moved, it should be moved at company expense and as soon as possible.

Ideally, a large company will be aware of and control and channel cultural perspectives and values in order to maintain an balance of all three, using each where and how it contributes toward the utility’s mission (Figure 15.2). Usually, this means that leadership at the top is either stockholder (typical in investor owned utilities) or community-centered (municipal utilities are often this) and that strategic direction and major priorities are set on that basis. By contrast, portions of the company that deal with customers are allowed (encouraged) to have a mostly customer culture, but one disciplined by priorities and guidelines driven from the top (and the business basis) and with a reasonable regard to the constraints and limitations of the electric system. Service and operations personnel and departments will tend to have an equipment focus, truly caring about the system and equipment they maintain and control, but understanding that the T&D system is there to serve the customers and must be operated and maintained within budget limitations.

All successful utilities have an active upper management that deliberately “uses” all three cultures and that makes certain each is aware of and responsive to the needs and values of the other two. The key to a balance among these three cultures is open and frequent internal communication, to establish an understanding of the “bigger picture” within each culture, and to build a common sense of purpose. By contrast, poor performing utilities permit these three cultures to each become uncommunicative and parochial (locally entrenched and blind to the larger picture): executives may think they are driving their company from a business basis, but the customer culture still spends on the basis of customers first/stockholders last and the equipment culture still operates with more regard for equipment than the stockholders, etc.

Expected long-term shifts in utility cultural mix

Figure 15.3 shows mixtures of the three cultures – stockholder, customer, and equipment – representative of typical electric utilities in the past (1975), the present (2012), and the future (≈ 2025), and indicates the nature of the long-term

shift in attitudes and priorities that is occurring in most electric and gas utility organizations. Overall, the relative weighting of the organizations' focus on their customers stays roughly as it always has been: modern utilities have a more quantitative, technically enhanced focus on their customers and service quality, but the same technological and functional improvements also apply to their business and equipment operations, too. So, despite a shift to a world where customer-level reliability is explicitly tracked, engineered, and managed, a utility's overall *relative emphasis* on its customers is little different than it ever was.

Thus, the net change in organizational priority at most utilities is a gradual move to a purely "business driven" perspective about their assets. Generally, "business cultures" welcome this change and "equipment cultures" fight it (but it is a losing battle for them). This shift in organizational emphasis is a response to the long term trends discussed below and is inevitable given the evolution taking place in the power industry.




Figure 15.2 Every utility is a mixture of three "cultures," different senses of purpose and values that dominate various groups within the organization. The best utilities keep these disparate institutional focuses balanced and operating coherently. If one dominates the other two it will lead to poor performance in every category, including those aspects most dear to the dominate culture. This drawing shows stockholder culture at the top, because it usually dominates upper management, and equipment and customer cultures at the bottom, because they are distributed geographically in division offices close to equipment and customers.




Figure 15.3 Change in the mix of cultures and their focus within a typical US utility over time reflects the larger forces at work on the business, market, and regulatory sides of the electric industry.

15.5 SIX DIFFERENT COMPANIES: SIX DIFFERENT PHILOSOPHIES

This section discusses six companies that own and depend upon electrical power systems for their success, and looks at the ownership and operating approaches of each. All have come to one or more of the authors during the last decade for assistance in power systems planning and operations. In all six examples here, names have been changed and unimportant aspects of their description altered so that the authors can talk frankly about the organizations and their electrical systems experience without fear of embarrassing their clients. Data and results presented are real. Two are utilities. Two are institutional. Two are commercial businesses whose focus is not electricity. None of the six operates entirely within any one of the four “ownership philosophies” discussed earlier, but two follow mostly a buy-burn-bash, two are essentially standards-driven organizations, one follows CBM/RCM, and one uses an asset management approach.

15.5.1 Big States Electric – Asset Management

Company Profile: Big States Electric is an investor owned electric utility, the result of a series of merger-acquisitions over the preceding fifteen years in which Big State Electric absorbed four neighboring utilities. All five companies are able to trace their company’s heritage back to the beginning of the 20th century. The company operates now as six “independent” operating companies, each within a single state, serving a total of 4.9 million connected meters with a coincident peak hourly demand of 24 GW.

Electrical Equipment and System: Each of the four companies it has recently absorbed, and Big State Electric itself, were the result of previous mergers, some going back into the 1920s. As a result the company is a conglomerate of different voltage levels, design types and historical operating methodologies and history. It owns 9,902 miles of overhead and underground transmission operating at 500, 345, 230, 161, 138, 115, 69, 46 and 34.5 kV and 57,800 miles of distribution line

comprising 4,734 feeders operating at voltages of 34.5, 25, 23, 22, 19, 13.8, 117.5, 11, 8, and 4.2 kV. BSE owns 87 purely transmission voltage switching stations and 918 distribution substations with 2690 transformers with an aggregate nameplate rating of 58 GW. It owns 1,134,007 service transformers with a total installed capacity of 45GW.

Ownerships and Operating Approach: BSE uses an Asset Management approach for both capital and O&M budgeting and prioritization of spending, having made this transition after the most recent merger. All five companies that merged into BSE were standards-driven organizations, and initially there was talk of leaving each company free to continue to work from its own traditional standards even as they became part of BSE. The rationalization for this was the fact that since each now operated under the jurisdiction of a different state regulator there was no problem if they each had different practices and standards. A senior manager observed that while this undoubtedly true, the main reason was that no one wanted to take on a messy and potentially contentious job of integrating standards and standards-departments across all the companies.

But in order to gain approval for the latest merger-acquisition, BSE had to promise synergies and savings amounting to nearly \$60 million annually (\$1 per month per customer). The use of common planning, operating, and maintenance practices, and the integration of those departments into centrally administered groups, were among a long list of savings areas BSE gave stockholders and regulators who had to approve the merger-acquisition, in order to get approval.

The company waited for three years after the merger to begin this task. The resulting effort to integrate standards and guidelines coincided with the appointment of a new President of BSE's Delivery Company, to whom the six operating company Presidents reported. She drove the effort toward Asset Management, insisting that the resulting merged organization and all six operating companies use a uniform, integrated asset management approach linked to BSE's corporate model and strategy. This was a rather clever move in one way, senior managers observed, for it meant no one company had to accept any other company's standards – they all had to make a transition to something else.

BSE's asset management is administered by an Asset Management Coordination group reporting directly to the SVP of Delivery Operations. They create and maintain the asset management tools and process and coordinate and support its use within the six operating companies. Figure 15.4 shows how the system is implemented. It is hierarchical, both top down and bottom up. Each operating company inputs all capital projects and its maintenance and service programs and projects into templates for the company's PEPT (Performance and Planning Tool). These are analyzed in aggregate and Pareto and other analysis results output which basically given the needs and options each company has: "We can do SAIDI = 75 minutes and such and such performance in other areas for \$XX, or SAIDI = 80 minutes and so and so performance in other areas for \$YY, or," etc.




Figure 15.4 Structure of BSE’s asset management implementation.




Figure 15.5 Age distribution of BSE’s 1,215 power transformers.




Figure 15.6: Top: Normalized PW annual cost for the next 29 years of three asset management scenarios for management of BSE’s aging transformer fleet compared to the predicted cost of “business as usual” under its former standards-driven approach (St-D). Bottom: SAIDI contribution of the four transformer management programs. See text for details.

Asset Management results and the options from the six companies are then analyzed as a whole, in aggregate providing the “options” that BSE executive management looks at to set its strategy and policy, and to assign T&D capital and operating budgets for each company. A good deal of the effort at this level deals with constraints created by six different operating jurisdictions: the company has different promised reliability targets and rates in each state, and requirements to be able to show it is within agreed upon regulatory limits in how it operates and uses ratepayer money with each. Yet it must operate as one company and achieve the overall synergies and savings it promised. Executives pick a plan that fits BSE overall needs and also fits within legal and policy decisions for each state operating company: segments of this plan are basically the “marching orders” for the six individual operating companies. Each company takes its part of the overall plan and is then free to operate as it sees fit, although it is expected to provide the performance (attributes and spending limit) promised in its submitted asset management information: its executives and management compensation is based on achieving that goal.

BSE upper managers and particularly the staff of the Asset Management Coordination Group are very frank about the fact that initially the asset management initiative did not go well. Inevitably there was resistance to the change within each operating company, including many people who felt the company was “lowering its standards.” Over time, these issues subsided and by the third year of operation under the new approach had mostly disappeared.

But BSE also discovered that while its asset management data, tools and process worked they were not dependable in one sense – the process is accurate and useful only when making *incremental* changes. BSE discovered what many asset management theoreticians, consultants, and internal coordination groups do not want to accept, that in practice the method is only accurate incrementally. For example, suppose it has been spending 15,800 hours labor annually on medium voltage breaker maintenance and is seeing a mis-operation rate of .71%. A detailed analysis that determines that if it increases hours by 5% to 16,600 annually, misoperations will be reduced to .67% is probably fairly accurate. Analysis trying to determine what would happen if hours were increased or decreased by 25% is very likely not to be dependable.

Aging T&D Equipment Infrastructure: Figure 15.5 shows the distribution of ages of power transformer MVA in the BSE system. The newest unit is less than a year in service, the average is 37 years in service, and the oldest unit is thought to have been in service for 72 years (records going back to its acquisition do not exist). BSE has investigated aging and deterioration and their expected effects and created three scenarios that it is using to manage decisions about service, refurbishment, and replacement. No one, including the Asset Management Coordination group that worked with consultants and vendors to produce the study, and upper management, is confident the analysis is completely accurate. However, the analysis is transparent, all the data and conclusions traceable, and

the methodology makes sense and seems credible. Company executives have decided to use it as the basis for their decisions in this regard, and have received no pushback from regulators.

Figure 15.6 shows BSE's analysis results for how it should best manage its aging power transformer fleet. The top diagram shows the annual cost of replacements, refurbishment, repair, inspections, and testing for four scenarios. The bottom number shows the project SAIFI contribution of power transformer failures and operating problems. "St-D" is business as usual under its traditional standards-driven approach. Alt1 is a plan optimized with a constraint that it would never spend more than the standards-driven approach in any year. It reduces PW of future expenses by 2.1% and the time-weighted value of SAIFI over the next 29 years by 3.2%. BSE's staff did it to see if and how "optimization" could improve the situation even if constrained by budgets. Alt 2 is a plan unconstrained as to when money is spent but constrained to the same total PW cost as Alt 1 –a study done to determine if "early spending" on life extension, etc., actually pays off with regard to PW results. It works down SAIFI to 9.1% less than the St-D case, so the answer is "yes." Alt 3 is an optimum plan as far as BSE's Asset Management Coordination group can determine. It reduces PW cost over the period by 4.2% and SAIFI by 13% as compared to the standards-driven case. Again, the staff, BSE's management, and the executive making the decision were aware that these results were only approximate, but the company went with transformer management alternative 3.

15.5.2 Mid State Electric Company: Reliability-Centered Maintenance

Company Profile: MIDSEC, as it refers to itself, is an electric cooperative serving a 4,800 square mile service territory that was once rural but is now the suburban periphery of a large and growing metropolitan area. It sells power to 198,000 connected meters with a peak hourly demand of 1,087 MW.

Electric Equipment and System: MIDSEC owns 391 miles of 230kV, 138kV, and 69 kV transmission lines, 131 power transformers, 197 high voltage circuit breakers, 312 medium voltage breakers and 75 MVAR of high voltage capacitors located at 48 substations, in addition to 4129 miles of 117.47 kV feeders, 175 line voltage regulators, 412 capacitor banks on feeders, and 31092 service transformers in service. The average age of its transformers is 22 years; over half of its customer base has been added in the last two decades due to the high growth of the metropolis nearby. An aged infrastructure is not a problem for MIDSEC. However it has a portion of equipment that is old including two substations and five power transformers dating to the late 1940s.

Ownership and Operating Approach: MIDSEC practices a combination of reliability-centered maintenance and "modern" standards-driven maintenance in which it sub-categorizes some classes of equipment into "bins" based on age and diagnostics, applying different rules to each bin. It applies RCM to vegetation

management, to the inspection, treatment and replacement of wooden structures on transmission lines, to all its 131 power transformers and to 105 of its service transformers (all those over 1000 kVA). It has a Transformer Load Management program that addresses the other 3,987 service transformers it owns, which it considers to be an RC process.

All other equipment is operated, inspected, tested, and serviced “by the book,” according to a set of standards evolved from REA operating and maintenance documents but modified through many years of operational fine tuning to MIDSEC’s needs. The company “sub-categorizes” breakers by age and loading and number of customers into groups that receive more or less frequent and more or less comprehensive service. Similarly, equipment, structures, and conductor on lines in heavily forested areas with high customer densities are inspected on a rolling five-year cycle rather than ten-year cycles as in other areas.

MIDSEC’s move to RCM, as at many utilities taking that approach, was initiated from within the Engineering and Operating division and not by top executive and financial management as is the case with many Asset Management initiatives. MIDSEC is one of the few electric cooperatives to participate significantly in IEEE equipment standards groups as well as NRECA bodies. Its VP of Engineering learned of RCM while working in such meetings earlier in her/his career. A new VP of Operations brought on in 1998 had a background in the nuclear navy and so was familiar with RCM approaches. As they implemented it at MIDSEC, RCM was mostly a technical initiative aimed at improving reliability of service. Beginning in 2001, RCM at MIDSEC focused on reliability—as opposed to condition-based maintenance. In fact, RCM began with application to vegetation management.

When RCM was implemented, the company also revised loading standards. Normal and emergency loading limits for transformers and cable were increased after a very comprehensive “thermal” analysis of equipment operating stress and loss of life that showed there would be no impact on reliability. MIDSEC’s Engineering and Operations department personnel consider this re-rating part of the company’s RCM plan even though, strictly speaking, it did not address maintenance per se. The loading analysis categorized 15 older transformers in a sub-category of reduced allowed loading. The average unit in this “low-load” category had an expected 11.6 year remaining lifetime based on IEEE standard Arrhenius formula calculations of expected insulation degradation. The reduced loading limits are aimed at an average 20 year remaining lifetime for these units. Load reductions were arranged by re-switching the feeder system to shift loading to newer units in the same or nearby substations.

Figure 15.7 shows the results from MIDSEC’s RCM program. Eventually spending on equipment T&D was reduced by over 11%, while SAIDI has improved by nearly 10%: MIDSEC’s standards-driven methodology had been delivering good results, but seems to have been spending much more than necessary.

MIDSEC’s executives are pleased with the program but view their current




Figure 15.7 Annual T&D system service and maintenance costs normalized to system size and corrected for inflation, etc. (solid dots), and SAIDI (open dots). Anomaly in reliability and spending in 1989-1994 was caused by a design flaw in primary voltage lightning arresters that were replaced in '94-95. Begun in 2001, RCM permitted MIDSEC to reduce spending on equipment maintenance by over 11%. Within four years SAIDI results were clearly dropping from an average of 88 minutes to about 80 minutes. Dotted lines show MIDSEC projections of future performance (see text for details).




Figure 15.8 Data for ten years of MV breaker maintenance at five electric cooperatives. Breakers were put into “bins” by years in service 0-5, 5-10, etc. Top: annual cost per unit in each bin was computed from maintenance records. Dots show the average for all 10 years of the survey for each age group, line shows the range of the data. MIDSEC was surprised that there appeared to be an increase in costs after only 25 years but not that cost began to escalate at 40+. Broad line shows the curve MIDSEC used to represent the cost-versus-age relationship in subsequent studies. This operating and maintenance approach is an interim step in a gradual evolution toward a full asset management method. As technology, time, and resources permit, they expect to add more sub-categories to those equipments they manage with their sub-categorizing standards approach, and eventually to extend RCM to a wider range of equipment types, and then gradually add more attributes to the optimization.

Aging Infrastructures: As mentioned earlier, MIDSEC's growth over the past 15-20 years means that it has a relatively new system, at least when viewed as a whole. The average transformer MVA in its system is 21 years old. However, 15% of its system equipment is over 40 years of age. In an effort to analyze aging effects quantitatively, MIDSEC shared data and analysis effort for a statistical analysis of equipment maintenance costs, condition, and failures on its system with five other cooperatives in its region. One conclusion of this study was that breakdowns were more frequent and maintenance cost slightly more expensive as units got older (Figure 15.8).

Qualitatively the escalating cost with age was not entirely unexpected, although MIDSEC was surprised that costs began to rise at age 25. It had expected a later but then faster escalation as a function of age. Regardless, as its breaker fleet ages, MIDSEC expects maintenance needs and costs for it to rise. For planning purposes it uses the broad line in Figure 15.6 as a guideline: it can estimate costs for, say, five years ahead by simply assuming all its current breakers are five years older and applying the cost escalation given by the curve. Similarly, studies on other equipment types lead to a like analysis of future costs. In this way, MIDSEC projects its RCM-based spending will gradually rise over time along the dotting line projection shown in Figure 15.5.

An unexpected consequence of RCM that concerns MIDSEC management is that, as its system equipment ages, RCM will respond by spending less on service and maintenance in some categories. Basically, RCM seeks to spend money where it does the most good from the standpoint of improving reliability – *where can I buy the most SAIDI improvement per dollar?* As equipment, for example the breakers depicted in Figure 15.6 age, they become less reliable, or at least less reliable unless more is spent on them. “Shopping for reliability” by buying breaker maintenance becomes less attractive than, say, buying it through vegetation management or some other program. Perhaps today the RCM optimization balances breaker maintenance and tree trimming in a 50/50 allocation of funds. Ten years from now, that might be 55/45 in favor of the tree trimming.⁴

This non-intuitive response of RCM to equipment aging is the main reason MIDSEC wants to move to an Asset Management approach. Basically, “what is broken” with its current RCM approach is the system's inability to manage

⁴ The interactions between aging, reliability, and RCM-type methods are quite complicated and will be covered in much greater detail in Chapters 5 and 12. RCM responds well to equipment aging *in a tactical sense*, by shifting spending focus toward older equipment where reliability improvement is needed, but its overall strategic response is as discussed here – as equipment ages and costs more to maintain, it “loses interest” in doing so.. All single attribute optimizations have this flaw in some measure: if a systemic change in the cost of the attribute they are optimizing occurs, they may not be able to deal with it well. This is one reason multi-attribute approaches like Asset Management are favored for strategic business planning.

anything more than short-term reliability improvement. It can't manage durability (long-term equipment value and reliability), too, and it can't relate all this to the business overall. For this reason, MIDSEC is committed to eventually moving to Asset Management. The relatively new condition of its current system equipment, and the slow progression of aging and its effects means there is no crisis and it can do so in an orderly manner.

15.5.3 International Bulk Petroleum Fluids – Buy-Burn-Bash

Company Profile: IBPF is a long-established “vertical” oil company into exploration, drilling, recovery, refining, wholesale and retail sales of aviation fuels, gasoline, diesel, and fuel oil, greases and lubricants, and various plastic bulk stock under a variety of brands. It operates 91 major petroleum refineries and 297 distribution centers and sub-refineries around the globe. Refining capacity is a major bottleneck in its operation and it puts considerable effort and focus into operating all at 100% output, 24 hours a day, 365 days per year.

Electrical Equipment and System: IBPF owns plant power systems at all of its sites with peak hourly demands that vary from 12 MVA to 397 MVA. In most cases it owns the site substation and buys power at transmission voltage. IBPF obtains equipment for each of its sites based on the electrical standards and power industry practices in the part of the world where it is located: ANSI in North American, IEC in Europe, and JIN in Asia. It takes this approach both because it sources equipment locally and because personnel hired for the plants will be more familiar with what is typical in their continent.

Typical load factor for the larger IBPF refineries is over 95%. Most of its larger sites do a “plant shut down” or schedule shut downs of parallel portions of each plant for one or two weeks per year, so it can do equipment updating and maintenance of refinery equipment and systems. The rest of the time the plant is running at close to full capacity and loading is nearly constant. Downtime due to power outages can exceed \$250,000 hour and even short outages can create situations that require a full day or more to return to full production. In locales where power supply is an issue its plants will have their own generators. Regardless, IBPF plant electrical systems are designed to be “bulletproof” – with redundant systems anywhere a failure would interfere with plant operation and with redundant equipment anywhere individual unit failures would shut down a system.

Ownership and Operating Approach. As regards the electric power systems throughout its plants, IBPF mostly follows a Buy-Burn-Bash approach, with only a few exceptions. Given the variety of sites it owns, equipment types it buys, and local laws and regulations it must obey, IBPF has to permit a certain amount of flexibility in electric operations at each plant. It has a standard framework for electric operations across all its companies and sites, a set of standards that pretty much mandates a buy-burn-bash approach to everything except protective

equipment, which is maintained fastidiously and for which the company pays premium prices for advanced technology including protective monitoring and replaying systems directly tied into its plant control systems so production can adjust to power system issues and maintain maximum output. Transformers, capacitors, inductors, and similar equipment not connected with protection are all treated as “consumable equipment,” a term used by several IBPF plant managers, one of whom said, “They’re like flashlight batteries – buy them, use them up, and just pop a new one in when the old one is used up.”

Transformers at IBPF plants are sized very close to nameplate rating; typically a unit will have a continuous loading at or just under its nameplate rating. Under these conditions they last an average of 11.3 years at major IBPF plants (Figure 15.9), a figure roughly in keeping with expectations based on IEEE Standard




Figure 15.9 Data on transformer failures versus time in service for three IBPF plants built in the late 1970s and early 1980s shows cycles with a periodicity of the average expected lifetime of the transformers in the type of service IFBP uses them to provide. Dots indicate actual data while the dotted line shows the statistical fitted trend over time. Most of these units are loaded to an average of 98% of nameplate for 50 weeks per year and have a calculated expected lifetime of 11.3 years. By year 25 many of the first series of replacements have also failed and been replaced. Over time the oscillations shown will gradually even out due to random variations in when individual units actually fail.

insulation deterioration rates.⁵ IBPF has done studies on buying larger size units and loading them to lower utilization rates (e.g., buy a 40 MVA instead of 32 MVA unit to serve a 32 MVA load). Its studies show this is not cost effective, partly because to get an assured longer lifetime for the larger units it would have to perform maintenance and service on the units.

Breakers and their control systems at IBPF plants are its one exception to the company's buy-burn-bash operating philosophy. They are serviced regularly and fastidiously and are replaced for one reason or another (often due to plant upgrades or reasons other than age) long before they are "old" – an average of every 22 years.

Within IBPF, MCCs (Motor Control Centers) and other local controls and operating systems, as well as the plant process control systems themselves, are the responsibility of and maintained by a separate group: Process Control and Operations which does not necessarily follow a buy-burn-bash approach with that equipment. Regardless, this equipment is usually updated on average every six to eight years as part of ongoing technology improvement of plant systems and thus never "age."

Aging Infrastructure Issues. IBPF has no problems caused by aging or aged electrical infrastructure. Its plant managers and the executives in charge of its electric systems believe that it never will.⁶ The company recently decided not to pursue a new electric design-operation-management paradigm that it went to a great deal of effort to develop. Working with both electrical and plant design consultants and a software system integrator, it came up with an innovative new plant design that eliminated much of the redundant systems and equipment used in its previous plants, but assured equivalent production reliability by use of real-time incipient failure detection and "self-healing" control systems. This approach was extended completely through the electric system, resulting in a 61% reduction in its cost.⁷ Although the new plant design saved a good deal of money, the complexity and untested nature of the new system was ultimately too much risk for IBPF's management.⁸

⁵ Fully loaded to its "nameplate" rating for all 8,760 hours per year, some power transformers deteriorate to 50% of insulation mechanical strength in less than 7.5 years.

⁶ The company does have issues with structural and chemical processing related infrastructures related to some of its older refineries.

⁷ Elimination of nearly one-half of the equipment reduced price to roughly half. But elimination of double buses meant tie breakers were eliminated, too. Similarly, elimination of one of each pair of redundant transformers reduced fault duties significantly, meaning lower capacity and lower cost breakers could be used, etc.

⁸ One of the authors helped with the risk assessment. Some IBPF plants pay for their capital cost in 17 months. While a savings of say, 10% is significant, it represents just 7 weeks' production. A less than perfect "fail-safe" system could easily create many more weeks of lost production over the next 30 years the plant might be in operation.

15.5.4 Big State University System – Standards Based

Company Profile: Big State University System operates six campuses with a total of 396 buildings, an oceanographic research center on the coast, an inland agricultural research facility, and a full-service, 100 room hotel off campus for the School of Hotel Management. Total full- and part-time enrollment is 52,000. The school has 19,000 staff.

Electrical Equipment and System: Each of the BSU system sites is served by the utility in their area through utility-owned substations, with the university system taking possession of the power at primary voltage at the point of metering as it leaves the utility property. Campus power systems are 2kV, 4.2 kV, and 117.5kV systems, and mostly underground cable in steam tunnels and ducts and spot networks in basement vaults. Most buildings use three-phase 208 volt systems. The two research centers use some overhead 120/240 and 480 volt systems. The main campus chiller plant has two direct three-phase feeds at 117.5KV to power two 3,500 HP motors that run water chillers for the campus cooling system. Overall the BSU system owns 594 network or service transformers, 341 network protectors or equivalent, and over 40 miles of UG primary serving, 417 service points with peak demands from 20 kW to 4500 kW. Sum of non-coincident peaks at 417 delivery points is estimated at 156 MVA (some points are not metered).

The BSU system's Electric Facilities Department is headquartered on its main campus in Capitol City. Its top management and engineering staff are located there, but the rest of its 171 personnel are distributed across all sites roughly in proportion to enrollment at each. The Director of Electric Facilities reports to the BSU Associate Vice President of Facilities and Services who reports to the BSU System Provost.

The Electric Facilities Department is responsible for specification, acquisition, installation, operation, and maintenance of all electric equipment at all eight BSU system sites, including building wiring and electrical equipment which requires slightly more than one-half the department's budget and labor. In addition the Electric Facilities Department specifies and contracts for electric design and construction of facilities at new campuses and for new buildings on existing campuses.

Ownership and Operating Approach: BSU Electric Facilities Department is very nearly a pure standards-driven shop. The standards-driven approach has always worked well for it (and at many other universities), in part, the Director of Electric Facilities has observed, because it fits the mindset of the rest of the institution. Educational institutions are procedural and rule-based, using lesson plans, study guidebooks, periodic tests evaluated against written standards, etc., to specify and channel learning activities and structure their operations. Few types of organizations are more committed to doing things "by the book." As a result, when the Electric Department manager meets with the Provost and other University system executives they understand, even feel comfortable with, the

way the Electric Facilities Department operates and administers its responsibilities via a standards-based approach, and therefore have more confidence in and acceptance of its needs and recommendations.

The Electric Facilities Department's Standards group maintains a set of guidelines for loading, for operational procedures like isolation of equipment for maintenance, switching and balancing load, and for maintenance of all equipment. These standards were based in part and follow almost completely those recommended by the American University Facilities Association committee on electrical standards, and are well proven as far as the Department is concerned. They have and will continue to slowly evolve over time. For example new OSHA Arc-Flash requirements implemented in late 2009 led to changes in BSU's internal guidelines for inspection, testing, and maintenance of equipment.

BSU's electrical facilities standards call for *annual* inspections of all equipment – a periodicity that many organizations, standards-driven or otherwise, would consider too frequent. In this case the motivation is safety – campuses are “schoolhouses” and a regard for the safety of the students and faculty is paramount. These inspections are primarily focused on safety but uncover a large “punch list” of correction items each year, providing all the routine examination of equipment BSU's Electrical Facilities Department deems necessary.

One part of the BSU system is done with a completely different approach. In 1990-2000 the main campus was augmented by the addition of a 141-acre “Millennial Campus.” This provided expansion room for new engineering and chemistry graduate labs, a new physics department building and laboratories, but it was mainly a co-operative industry-academia site, where companies would build new R&D facilities to share resources and work. To attract high-tech companies to its shared resources Millennial Park, it is billed both as a next-generation IT environment wired with fiber and high tech systems throughout, and as a “power quality park” with near-perfect reliability and voltage stability of power delivery, and near-zero harmonics and voltage transients. The system is a modified distributed network using European ring-main buses and power electronics and storage throughout, along with a good deal of automation, all to maintain continuous power flow, stable voltage, unity power factor, and harmonic content near zero. This system is operated in a purely RCM manner, something that the director does not believe is necessary from an operating standpoint, but views as a marketing strategy to give further distinction to the area's advantages to power quality conscious companies.

Aging Infrastructures While some of the BSU campuses are relatively new, having been built or greatly expanded in the last two decades, three date back to the early 20th century. BSU has a good deal of old electrical equipment of all types on these campuses. As discussed above, all equipment is inspected frequently. On the older campus underground cable runs through steam tunnels between buildings and the central plant and is therefore relatively easy to inspect along with equipment in vaults and building basements. BSU technicians visually

scan all cable, splices, and junctions (three-wire splices) and scan with an infrared detector (inspections are typically done in summer when “steam” pipes are carrying chilled water and the tunnel are quite cool). This work catches a good deal of deterioration-related problems before they cause failures or cable vault fires, etc. BSU makes what it calls pro-active replacements before expected failure occurs, but still has considerable numbers of failures – electrical service is seldom interrupted because most of the systems are spot- and distributed-network secondary design.

Figure 15.10 shows failure-replacement rates for cable and splices and junctions (three way splices) for primary and secondary cable on BSU’s three oldest campuses over a thirty year period. There is no dependable data on equipment ages, etc. The data show failure-replacement rates have escalated rather linearly and by a total of 50% over three decades. BSU’s Director of Electric Facilities considers the cable failure escalation to be age related, but attributes part of the escalation in splices and junction failures to the poorer quality of materials and labor of the last twenty years. The data do show a higher rate of increase in failures since 1990, and a number of splices have been repaired several times, but there is no conclusive evidence proving this.

BSU’s Electric Facilities Department is most concerned about the age-related deterioration of the electric system civil facilities on these three older campuses. Some concrete and brick vaults, manholes, and tunnels date from just after WWI and while in good condition for their age, are deteriorated to the point they are a concern. A number of these rated worst condition by visual inspect were “cored” or otherwise sampled and will be evaluated by the Civil Engineering Department’s materials labs.




Figure 15.10 BSU’s 30-year record of underground electrical cable section failures and proactive replacements (solid dots) and splice and junction failures and replacements (unfilled dots) at its three oldest campuses.

Overall, BSUS's Director of Electric Facilities and his staff acknowledge that at least some of their campus systems have aged infrastructures and aging-related failure and maintenance cost escalation problems. Overall maintenance costs have increased slightly faster than inflation and campus growth can explain for the past fifteen years, but budgets addressing this need have been approved every year, largely, the Director believes, because his department's standards-driven approach is credible and respected, the Provost and board are familiar with the aging infrastructure problem, and his department has well-documented data on failures and maintenance costs to back up his budget requests. For this reason, BSU's Electric Facilities Department plans no major changes in how it operates the system it owns.

15.5.5 Third-World Orphan Rescue Foundation: Uniquely Standards Based

Company Profile: The Third-World Orphan Rescue Foundation (TWORF) is a secular charity with a publicly stated goal to house, feed, and educate as many orphaned children in undeveloped countries as possible. The organization is headquartered in New York City and funded purely by support of philanthropic foundations and by individual donations. TWORF believes that by educating orphans it is not only improving their lives but indirectly the lives of others in those nations, by providing leaders and infrastructure development experts to help those countries develop. Its most recent annual report said it was caring for just over 19,000 orphans.

TWORF operates 43 "schools" – orphanages – in 21 countries. Almost all are located in rural areas and are basically small self-supporting villages, each with a population including orphans, TWORF personnel and indigenous staff, of about 600. Most orphanages have a "production" side; many farm the surrounding land for both their own food and to sell in local markets, others have a small factory or facility, such as manufacturing natural-fiber rope and netting, or making furniture – something that both earns money for the orphanage and teaches a trade to the older children.

Electric power is more central to TWORF's success, and power systems equipment is more of a concern to the charity's board and executive management, than an outsider might expect. Electricity provides many functions that TWORF considers essential to orphanage success: it powers water well pumps and conditioning equipment, provides refrigeration for food in its kitchens and medicines in its infirmaries, powers communications equipment including navigation beacons at landing strips in rural sites, runs televisions and computers for both education and recreation, powers equipment in its machine shops and the manufacturing equipment if any, at each site, and, of course, provides lighting and other conveniences. In addition, the teaching of electrical equipment service and repair as a vocational subject is a focal area for TWORF: electric systems are crucial to economic development and by training future electric system technicians it is enabling the countries where its orphanages are located to grow.

All forty-three TWORF orphanages have facilities to produce their own electric power and run what is essentially a private power system inside the orphanage grounds. Seven orphanages do have a connection to the local power grid, but 37 have no source of power but their own. As one TWORF executive put it, “Pretty much by definition, if local power supply was available and dependable these would not be third-world areas.”

Electric Equipment and Systems: TWORF owns a wide variety of electrical equipment types and systems with little uniformity in anything among its 43 orphanages. Peak power needs at its sites vary from just over 80 kW to 300 kW, with an average of about 210W per resident. It has both “European” (50-Hv 250 volt three-phase) and “American” (60-Hz, 120-1Ph/208 3-Phase) systems, and just about all types of equipment possible. A majority of generation is medium speed diesel but there is no uniformity in the types or equipment used, etc. Four sites use small steam generation fired by local fuels, including “agricultural refuse” at one. Some orphanages use “renewable energy” such as windmills for water pumping and passive solar heating for water heating, etc., where possible, but there has been little renewable energy until recently. This is changing as the organization has begun soliciting contributions to help it obtain wind and solar power.

TWORF’s records indicate that it owns 228 diesel or fuel-oil powered combustion (piston) power generators of from 15 to 125 kW output, four steam-powered generators (two piston powered, one turbine, one sterling cycle), four small hydro plants powered by water wheels or “run of pipe” water turbines totaling just under 400 kW output, 87 wind turbines varying in size for 2 to 50 kW, and 23 PV “generator sets” averaging about 8 kW each. It currently has no micro or small gas turbines as the organization does not believe it has the skills to service them.

TWORF has three unbendable rules about power systems and equipment at its sites. First, the “power system” at each site and all electrical equipment operates only at utilization voltages below 260 V: no primary is allowed on site except in a fenced substation owned by the local utility, if there is one. Similarly by policy it does not deal with 440, 480, or 600 volt circuits and equipment either. Second, all circuits are overhead, and uniformly made with covered conductor insulated to 600 V. Third, for reasons of both safety and simplicity, there are no circuit breakers, only fuses.⁹ A fourth *guideline* calls for all circuits to be built and operated as three-phase open loops with isolation switching at each feed point, as is the fairly standard European LV design practice, but this is not mandatory, merely recommended.

Unique Ownership and Operating Approach: Management of electric systems at

⁹ Building wiring systems might have breakers in their service panel boxes, etc., but incoming utility service, and all lines and equipment used to distribute power through each orphanage campus, are fused.

TWORF's 43 orphanages is a "flexible and pragmatic" standards-based approach that unique as far as the authors are aware in one way: at every one of TWORF's 43 orphanages, the power system is a vocational learning tool and the "Manager of Power Supply" is an instructor whose job description is to both run and maintain the site's power equipment and train orphans over 15 who show mechanical aptitude to maintain and service electrical equipment. The two functions of system operation and instruction are heavily intertwined, not just in the manager's assigned duties, but in TWORF's design and administration of its electrical facilities and standards and in its maintenance and documentation processes organization-wide. The organization considers that "electric skills instruction" is one key area for educational development that particularly helps boost local economic growth and prosperity and thus focuses on driving this.

Ownership: TWORF's New York based headquarters staff focuses on obtaining electrical equipment through donations or by purchase of used (often very well used) equipment from supporting foundations. It maintains records on what type of system (European or American) is at each site and lists of the equipment and requested needs from each. They try to address requests by school directors ("We need another 50 kW generator") but take what they can get by donation and ship it to where it seems it will be most useful. TWORF headquarters buys equipment only as a last resort, preferring to let local School Directors, or regional Superintendents, buy equipment locally. Each School Director is pretty free to make any decisions on administration of the equipment, even to the point of selling or trading equipment that is obtained by headquarters and shipped to the orphanage, as he or she sees fit. Headquarters maintains an extensive library of service manuals and guidebooks on all types of electrical equipment, which it provides as needed to each site, trying to support what its management wants to do with respect to ownership.

Operation and maintenance, on the other hand, is tightly prescribed, with the three inviolate rules given earlier: power systems are exclusively utilization voltage with no primary permitted (except for incoming feed from the utility), all circuits are uniformly composed of covered conductor, and only fuses – no breakers – are permitted within the power system, although some individual building wiring systems have breakers in their service panels.

Beyond those rules, TWORF has very tightly documented standards for electric operation, service methods and procedures, and particularly for safety and work practices. It has a mandatory training program for power equipment operation and maintenance. This program and documentation is for all its maintenance and operations personnel, but is also used for instruction in its vocational programs in electric equipment service and repair. Many of its operations and maintenance personnel are the students themselves, who do such work as part of their training.

Annual inspections are made at each site by TWORF's Superintendent of Water and Power or one of three assistants. Standards and inspections focus

mostly on safety, completeness of required documentation, and the quality of curriculum, instruction and results in the vocational areas of water and power equipment and systems. Reliability of power supply, equipment condition, and even operating costs are factors not addressed at all, unless the School Director has issues with them.

With respect to safety, “Rule #1” is that *all* maintenance, repair, and service is done with the equipment shut down, de-energized, and grounded. There are no exceptions. Safety, including grounding and work practices, is a focus of much of the electric instruction program in the first year. Maintenance and service of equipment itself is incredibly resourceful: some sites make replacement parts for generators and re-wind transformers. Several perform such services for surrounding businesses and institutions. One site is in the process of “spinning off” an electric service company staffed completely by orphanage school graduates.

Aging Infrastructures: TWORF is unique among those the authors have worked with in not regarding aging power equipment as any kind of problem. The organization’s culture seems to value old but serviceable equipment as both a symbol that it does not waste resources and an advantage from the standpoint of on-the-job training: TWORF has a lot of student labor available in its vocational classes and so can handle a good volume of repair and service. Its records indicate that the average generator it owns was 20 years old when donated or obtained and is now 32 years. As mentioned earlier, its Water and Power Managers and their students are quite resourceful and fewer than 4% of its generators have been retired as unserviceable – once obtained they are kept in service “perpetually.”

Figure 15.11 shows analysis results for data from thirteen years of TWORF operating records covering a total of 318 motor generator sets in the 25 – 125 kW range. The data shown are used by TWORF’s central headquarters staff and local school administrators to estimate electric system budgets, etc., but the organization’s primary purpose in gathering it is to plan backwards from the hours of labor needed for this and similar electric-system related activities to determine the number of vocational students that can be given useful work and on-the-job training in this area of skill.

Only data on units owned and operated by the Foundation for two years or more have been used (many units donated to the school need major refurbishment and give problems in their first two years). The dashed line is the authors’ best third order statistical fit to data on the number of units found to be un-repairable as a function of age.¹⁰ The majority of failures that result in TWORF scrapping a unit for spare parts are mostly due to failure of major castings such as engine

¹⁰ For both values plotted, generators were grouped into “bins” of 5 years of age: from 0-5 years, from 5-10 years, and so forth. Results were only significant for the range of ages from 15 – 65 years. Trends shown for newer and older generators were estimated by the authors, since TWORF has few generators less than 15 or older than 65 years.

blocks, generator casings, crankshafts, etc. Non-availability of parts is not an issue: except for major castings the schools can make, or scrounge, most any part needed.

The shaded area shows expected annual labor requirements for service per motor-generator and includes all systems in the unit, mechanical, electrical, and electronic or electromechanical controls. The plotted results are shaded to indicate something of the statistical range of variation in service requirements as units age. Hours of service required are measured in “apprentice” hours.^{11,12} The most interesting point from the author’s perspective is the variation in hours/year of service required as units get older. Intuitively most managers and engineers associated with electric equipment will expect that as units get older they will vary more in maintenance and refurbishment needs, but this actual data that shows the very wide spread of “unpredictability” in service requirements that develops as units age, in addition to a steady increase in the average need. The standard deviation of service needs by age 50 is over 50% of the mean.

As a general rule, conductor, cable, switches and fused cutouts are bought new. Poles and structures are locally manufactured, usually by the school’s carpentry shops. The average age of this equipment is estimated at 16 years. Failures are not a reported problem at any of its sites. Documentation is quite good so it was possible to determine that over the whole organization, failure rate for circuits due to all causes other than external damage (truck backing into a pole, etc.) is .2 failures or repairs per year per mile. Similarly records indicate SAIDI is equivalent to 98%, with the majority of that being the occasional late fuel delivery causing cutbacks in power generation.

TWORF does think it has a “new technology infrastructure” problem, with power renewable generation it has tried to add in the last decade. Much of this is a result of unfamiliarity with the specification and purchase of equipment so that much of what was obtained is incompatible with the skills and service capabilities of its schools. All but 8 of its 87 wind turbines were new when installed and the average age of units now in place is only six years. All ran well for two to three years, but a *majority* now give problems within five years of installation: blade angle mechanicals and blade-to-generator gearing breakdown, and DC-AC power electronics often need service; both require tools and skills the schools do not have.

¹¹ Apprentices are students within one year of graduation from the vocational school and the students leading service and maintenance work under the tutelage of an instructor.

¹² Data used in this analysis did not take into account the size of units –a 125 kW unit typically takes somewhat longer to service than a 25 kW unit. However, the authors do not consider that to be an issue in these data plots’ validity. Data were checked and there is no strong correlation between age and size in the data, and the point of this Figure is only to demonstrate the effects of aging from actual data.




Figure 15.11 Data from TWORF records on motor-generators used at its orphanages around the world show the effects of age on annual labor required and % deemed unusable each year. See text for details.

TWORF has recently decided to use only singly-fed induction generators on its wind turbines to avoid these problems. Similarly, 30% of PV panels and inverters fail within three years. Diagnosis leads to no identified causes or systemic issues, but local shops cannot service the electrics or repair panels when PV cell laminates split, etc.

Overall, TWORF is committed to its present system and plans no changes other than to retire all its PV generation and all wind turbines that do not use induction generators, and shift to more use of sustainable energy in non-electric form.

15.5.6 Huey Longwaites International Airport – Buy, Burn, Bash

Company Profile: Huey Longwaites International Airport (LIA) is one of the nation's newest airports, with a runway and ramp layout designed to minimize taxi time to and from terminals and promote fluid air traffic in and out of the airport at peak periods. Opened in October of 1994, the airport features “Smart” everything including automated people movers, baggage systems, and traffic light and control systems into and out of the airport. LIA is a 24 hour a day field with arrivals and departures internationally every hour of every day of the year.

Electric Equipment and System: The airport has a peak annual hourly demand of 72 MVA, with an 84% annual load factor, an indication of just how 24/7 its operations truly are. It owns the power distribution system within the airport territory boundaries, which is served by two “split ownership” substations, located at opposite ends of the airfield. Both are served by 230 kV transmission looped through the substation. Big State Electric owns the high-side equipment, including three 230/25 kV transformers at each site (two operating and one on hot

standby), and the 25 kV buses and supporting equipment. LIA owns everything from the 25 kV feeder breakers on out to its system. In the first 15 years of operation, there was power outage to the airport two times for a total of five hours, both incidents caused by failure of the utility feed due to a systemic problem traced to protection coordination issues that occurred under unusual operation conditions of a nearby independent power provider.

LIA's internal airport delivery system consists of 39 25kV/600-480-208V "substations" each a spot network, fed from 42 miles of mostly UG 25 kV primary feeder operated as closed loops with automated switching. It owns and operates a total of 138 25kV/utilization voltage transformers, with 131 in service and 7 kept in storage as spares. Within the airport system service reliability has been perfect at the terminal level: not one of the five airport terminal buildings or the control building have lost all power since the airport opened. Transformer annual peak demands are on average 96% of nameplate. Average annual load factor is 93%. Expected time to insulation half-strength calculated based on ambient loading and the Arrhenius formula is 16.4 years. However, this figure is considered optimistic by LIA's electric department engineers, because it does not consider the rather high levels of harmonics in the system due to the heavy use of digital equipment and fluorescent lighting throughout the airport.

Ownership and Operating Approach: LIA practices something very close to a pure Buy-Burn-Bash approach to operation and maintenance, and its service staff and management firmly believe that all electrical equipment is expendable and that maintenance practices ultimately will make little difference in service lifetimes. By intention no inspections or maintenance other than routine annual visual checks for obvious leaks and safety hazards are done to any of the electrical equipment other than the breakers and protection systems. These are maintained under contract to an outside firm and are inspected annually and tested and serviced every three years.

Aging Infrastructure. Other than two early failures attributed to flaws in equipment as delivered, during the first 10 years of operation there was only one transformer failure and two 25 kV cable section failures. In early 2002 (8 years after opening) a full "security" inspection was carried out due to terrorism concerns after 9/11, at which point LIA installed equipment to monitor access to all facilities and track temperatures and conditions in each vault and duct, and "hardened" all data-com and control systems, etc. Since the additional cost was minor, visual inspection of all equipment and civil facilities for condition was added to this security inspection. This produced a set of over 600 items that could use attention. Electric department engineers "prioritized" this to a very short list that addressed several bad cases of case rust (including fixing the reason water was on the floor of two vaults), and replaced two sections of cable in an under-runway tunnel that had what looked like damaged sheathes. They also re-did several splices that infrared scanning indicated were hot. LIA records indicate that in the eight years since, the other equipment on this list has




Figure 15.12 Failures of transformers (top) and primary voltage cable sections (bottom) during the first 15 years of operation of the LIA airport electrical system.

had failure rates no higher than other equipment at the airport.

In the seven years since this set of inspections, transformer, breaker, and cable failures have occurred more frequently and indicate expected average lifetimes of 14.8 years for transformers, 47 years for breakers, and 19.3 years for cable. Figure 15.12 shows the failure history and projected trend: these trends clearly indicate uneven and increasing expenses in some years but LIA management is not concerned. That has been anticipated and planned. LIA management and electric department personnel are satisfied with their electric system and its operation and maintenance practices and plan no changes.

15.6 SUMMARY AND CONCLUSIONS

Different companies can have very different perspectives about why they own electric equipment and why and how they should care for it. These “operating philosophies” might be based on sound reasons dictated by the companies’ situations, or on prejudice and non-objective biases unique to the company owners or sustained by the company culture. But in every case, the way a company “thinks” largely determines how it operates and maintains its electrical equipment. Conversely, how and why it operates and maintains its electrical

equipment contributes to its *culture*: the values, priorities, and beliefs that control how it thinks and reacts to challenges and opportunities.

This chapter presented four distinctly different operating philosophies or approaches to equipment operation and care. It discussed how they differ in function, but more importantly, in terms of whether the implications they have for culture reinforcement – whether they imply a “right way to do things” or force organizations to see and deal with compromises and imperfections.

It discussed how and why cultures develop with distinctly different beliefs in (as well as methodologies for) electrical equipment maintenance priorities. Examples were given to highlight and discuss differences.

The main takeaways from all this discussion are that all companies owning and maintaining electrical equipment do not “see” the values, obligations, duties and even necessities of that ownership in the same way, and that from the standpoint of some others they do not “think rationally” - and vice versa. Thus, a person looking at aging electrical infrastructures, or any electrical ownership issue for that matter, cannot draw entirely valid conclusions from reports, data, activities, and discussions given by various companies and organizations, without taking into account these “human” factors. Furthermore, what constitutes “success” in handling an aging electrical infrastructure, and if and how methods are applied to successfully achieve it, can depend greatly on developing and using an understanding of these human issues.

Asset Management, if Done Well, Always Gets the Right Answer

In closing, the authors wish to note that this chapter, although focused on non-technical issues, revealed many of the limitations and issues that have led them to heavily favor the Asset Management approach. From an analytical and technical perspective, Asset Management, assuming it is well executed, will always determine what ownership and operating policies are best for an organization and particular situation. If buy-burn-bash is best, it will determine that, as well as any slight variations or variations or mixtures of that with other methods that will work best. If a standards-based approach works best, Asset Management will conclude and recommend that, if set up to assess that question. It will “dispatch” budgets in CBM or RCM approaches optimally. And of course, it will fulfill its role if executed as a full Asset Management program of gaining the best results possible with any budget, showing management its options, and computing the overall cost-effectiveness of the company’s spending plan.

That said, one must keep in mind that the data collection, clean up, filtering and aggregation needed, the condition tracking and analysis required, and the analysis and process management involved in Asset Management can be a considerable burden to any company owning significant amounts of electrical equipment. If the optimal answer is “buy-burn-bash” then that effort is largely wasted – the overriding reason that organizations using that or less-expensive equipment management approaches do no such analysis.

Technically, that argument is unassailable, except for this: the culture that Asset Management develops and creates and supports around it, of knowing there is no “right answer,” of seeing all decisions as compromises that gain and give at the same time, of using a common business basis to drive all evaluation and decisions, is far and away the best culture for any electric system owner. Organizations that “think” like this make fewer mistakes, come closer to optimizing results even when viewed in hindsight, and react far better and quicker to unexpected surprises and change.

16

CASE STUDY: WOOD DISTRIBUTION POLE MANAGEMENT

16.1 INTRODUCTION

This chapter presents a detailed case study involving the management of an old wood distribution pole infrastructure. This is a particularly long chapter because the study is presented in considerable detail, with tables of data and explanations of each step in the process from initial study to final plan.

Like all examples in this book, the utility name has been changed, and corporate, financial and cost data disguised via normalization to different totals in order to preserve that actual company's anonymity. Otherwise the case study is presented as it occurred. Concentration here is on aspects likely to be unfamiliar to power system engineers who have not dealt with aging infrastructures: data sources, collection, aggregation, and vetting of data for analysis are covered in detail, as this will be new to many readers, but how to compute present worth (PW) factors and Benefit/Cost ratios is not, as experienced power engineers know how their company does that.

Why Wood Poles?

Because they make a good case for learning about aging infrastructure management. Some classes of large equipment can be managed as individual units. For example, a utility may own several dozen 400+ MVA

autotransformers. Given their importance and cost, it can essentially manage each as a separate unit: monitoring, testing, analyzing condition, prescribing maintenance and service, and even de-rating the unit if need be, all on an individual, specific and very focused unit by unit basis.

But the sheer number of poles, and their commodity nature, means a utility cannot do the same with them. And yet poles are critical: a utility's business is literally riding on them as they carry its power to its customers. They must be managed as an infrastructure and by using the less-than-perfect data that seems to be all there ever is when single units do not justify individualized attention. Finally, the authors used wooden poles as an example because they had results from several excellent utility pole studies, so they were confident the results presented here would be representative of results throughout the industry.

The presentation is given in as linear a fashion as possible, following the progress of the actual project from beginning to end, in the order that work was done, results obtained, and decisions made. At points this presentation will delve into technical issues related to how a particular data form, analytical step or decision was set-up or performed, what pitfalls were encountered, and how those missteps can be avoided. At the end, lessons from this example are used to identify rules that apply to all aging electric power infrastructure projects.

16.2 AN INAUSPICIOUS BEGINNING

City Light and Power

City Light and Power was founded prior to World War I by Charles Brush, an early and ultimately unsuccessful competitor of Westinghouse and Edison in the development of electrical technology and the franchising of electric utilities. It has remained an investor-owned utility, and over the eleven decades since its founding has gone through a considerable amount of change and growth, until today it serves as many customers outside the municipal boundaries of the city whose name it bears as it does within. At the time this example begins, CL&P served 825,000 delivery points in an extended metropolitan area of roughly two million people, through its T&D system (Table 16.1). At this time, it owned 437,226 wood distribution poles.

In late summer 2005, CL&P's board of directors appointed a new CEO, who in the way new CEOs typically proceed, took several months to get a feel for the company before implementing the board's instructions about tightening management and focusing onto a core strategy for business performance. Six months after taking over, in a top-down manner, she began to bring in some new upper management, people with fresh perspectives and new ideas.

A new SVP of Power Delivery took over operations of the retail delivery division in the spring on 2006. His attention had been drawn to operating repair and replacement costs that were apparently rising faster than the asset base size or demand or revenue. Data were not complete, so no one had high confidence in the analysis. But the trends seemed to be there, and a popular theory was that it was being caused by equipment aging.

Table 16.1 City Light and Power T&D Facilities and Equipment (Jan 1, 2009)

Category	Number
Company founded in (original date of oldest of sinc- merged companies)	1896
Total number of companies merged into present day CL&P, 1897 – 2010	9
Peak system load (normalized to 90% extreme summer peak weather)	3,785 MVA
Annual energy input (normalized to mean annual weather)	17,242 GWhr
345 kV transmission – miles	all network 411 miles
138 kV transmission – miles	network and loop 927 miles
69 kV transmission – miles	loop and radial 196 miles
46 kV sub-transmission – miles	loop and radial 95 miles
Transmission substations	345/138 kV 5
Transmission substations	138/69 kV 32
Sub-transmission substations	69kV/46kV 8
Distribution substations	most are two transformer 125
Power or auto transformers	503
Primary feeders (12.5 and 25 kV)	7,824
Miles ducted UG	48
Miles OH	13,286
Miles UG	1,828
Miles URD	7,390
Service transformers Network vault	112
OH on poles or racks	212,304
UG or other padmount	137,560
Distribution poles – wood	437,226
– other	3,081
Transmission voltage (69 & 46 kV) delivery points and facilities	11
Distribution voltage (25 & 12.5 kV) delivery points and facilities	23
Service level circuits and drops	825,039
Retail customers	825,039
Retail delivery employees	1,430
Field force and stores centers	3

The new SVP created a temporary asset management team to gather data and produce better information by the end of the year. One of this book's authors was brought in as an advisor to the CL&P team assigned to do that work.

Distribution Poles

The new team took a look at the entire range of CL&P's T&D equipment. Three categories stood out as most in need of attention: lattice transmission towers, underground cable, particularly for URD applications, and wooden distribution poles. This narrative case will follow only the latter, but many of the tools and




Figure 16.1 Data on pole failures over the previous thirty years, gathered at CL&P in early 2007. Lower extremity of the plotted line is interpreted as “normal weather.” Spikes include pole replacements normally excluded in storms categorized as extreme for reliability reporting. A particularly violent series of excludable storms in 2004 resulted in over 4,700 pole failures (dotted line) which had been removed from the data being analyzed because that storm was categorized as a less than one in a century event.




Figure 16.2 Two fitted trends that CL&P thought to bracket the problem. Dashed line – the least exponential fit, predicted the “storms excluded” failure rate would double in the next twenty years. Solid line shows a projection including storm-related failures and shows a faster escalation rate. Not only was CL&P’s pole failure rate increasing rapidly, it appeared its overhead distribution system’s *susceptibility* to storm damage was worsening, too.

methods developed and reported here were used for those other classes, and others that were less urgent but still important.

The annual number of “normal weather pole failures” – failures not associated with storms or excludable reliability events – had more than doubled in the last two decades (Figure 16.1). Some attempt was made to fit the data with formulae for extrapolation of the trend into the future, but a number of different formulae could be fit to what data were available, all producing different results (Figure 16.2). Regardless, every projection showed a continuing escalation. They only differed as to when aging became a crisis.

16.3 OVERVIEW OF APPROACH

Figure 16.3 shows the project team’s initial concept of their pole management approach’s framework. Its three main parts were:

A Pole database would include data from three main sources:

- a. company pole records (purchase, installation, etc.),
- b. operating records (failures and replacements, etc.),
- c. inspection and testing results.

A Planning Model would apply regression to:

- a. Analyze data and identify relationships among variables and factors (e.g., average condition versus age, age of system, best time to replace a pole, etc.).
- b. Extrapolate failure and cost trends into the future to provide estimates of “how bad it will get, when,” providing a base case to which options for improving pole management could be compared.
- c. Study the effects of proposed changes in replacement or maintenance policies, so as to permit CL&P to study the effects for various steps it could take to mitigate symptoms of failures or replace poles earlier.

A Condition tracking function would:

- a. Retain data on past equipment test results and condition estimates, etc., and whatever other data was useful.
- b. Provide the best estimated condition and status for each pole: particularly an estimate of when it should be replaced.
- c. Produce lists of “worst poles” and “poles that should be replaced, etc.” and similar reports.




Figure 16.3 Initial concept for CL&P’s wooden pole management system involved a simple stand alone application built around electronic spreadsheets and three main modules. See text for details. The project team expected to be involved in assembling a database and running the planning model when completed.

Figure 16.3 shows this structure as the project team originally envisioned it, an approach that proved unequal to the task. Figure 16.4 shows the ultimate product, which is both more complicated and requires more work. Its development, and the lessons learned with it, are covered this chapter. Readers wishing to go straight to “the final answer” can skip to Section 16.12.

16.4 THE DATA CHASE

Contrary to expectations, CL&P’s project team discovered that there were a lot of potential sources of data which might be relevant (Table 16.2).

Pole Inspection Records. Since 1988, when CL&P and its state utility commission reached agreement on revised safety management standards, CL&P had done a rolling ten-year pole inspection. The first ten percent of poles had been inspected in 1989, 1999, and 2009, the second ten percent in 1990, 2000, and 2010, and so forth through the whole ten-year repeatable cycle. These records were originally in a computerized flat file database: a line for each pole each time it was inspected: pole number, circuit, address, date of inspection, inspector’s ID number, visual evaluations conforming to tables of condition choices on the inspection form, results of “hammer thump” tests at the base of the poles, and a check list of safety or hazard issues at each.




Figure 16.4 Final pole management system depended on a large amount of analytical work done to extract the maximum information possible for the multiple data sources and minimize estimation error from its use. This “cleaned up the data,” provided formulae for the planning model and condition estimation formula. The system depended on the corporate GIS system: ultimately all pole data used by the team was retained there. Up-to-date condition data from the corporate asset tracking system was transferred to the GIS on a monthly basis.

Table 16.2 Data for Distribution Wood Pole Aging Management Study

Database	Paper	Computer	Comments
Pole Insp. Records	None	1988 - 2010	10% of units per year are inspected and visually evaluated as to condition. Flat file record for each pole inspection.
Birthmark Survey	Archived	Where found	Where the birthmark or “brand” of a pole can be found and read, basic data had been collected during inspections.
GIS Pole Records	None	Recent	GIS system has records for every pole currently in the system including those in Stores, including location, type of pole, attachments, information from the birthmark survey (see above, etc). Age or latest installation date is not among the items available for poles installed prior to 2007.
Regulatory Reports	1933, 1989	None	Reports and information gathered for the initiative that led to the inspection program. Useful mostly as a snapshot of 1988 timeframe condition estimates.
Failure Reports	none	1989 onward	A list of all failures of equipment was maintained starting in 1990, this change made at the same time as the pole inspection program was begun.
Pole Work Orders	1947-1989	1990-2010	Records of pole installations including codes for new, replacement, upgrades, other, etc. Paper records are not consistent in format or confirmation content over their period. Known to be incomplete prior to 1975, thought to be complete and dependable since computerization in 1991.
Stores Inventory	1938-1993	1993-2010	Records of poles bought and poles disbursed to the field for work, replacement, etc. Sketchy in places but some long past years are thought to be complete.
GIS Customer and Spatial data	Fairly complete since 2008		GIS data used for spatial customer and demand forecasting, and spatial information gleaned from municipal property records gave fairly good estimates of the initial build out period (first construction) in every area of the system. (See Appendices for discussion.)

“Brand” or “Birthmark” Survey. Since at least the 1930s, most utility poles have carried a “branding” or “birthmark” burned into the pole at about eye level when installed (Figure 16.5). Information required of modern brands is covered by ANSI Standard 05.1.2008, but older brands may provide less data. Brands are more difficult to read the older a pole gets and the less preservative it retains. In the CL&P system only about half the digits on poles over 60 years of age can be dependably discerned. Some CL&P poles also have attached metal tags or metal identification numbers attached.

Geographic Information System (GIS). CL&P had a GIS as part of its Enterprise IT initiative. The GIS system had just been linked to pole records, and while not all inconsistency issues had been resolved, the system had a record of location, attachments, and characteristics for 445,098 poles in the GIS system – almost all of the distribution poles in the system, wood and otherwise.

From the aging distribution pole project team’s perspective, the timing of its completion could not have been better. Beginning with the inspection cycle in year 2000, data for that year’s 10% of the distribution poles were logged for GIS input each year: GPS location, pole characteristics and attachments, inspection results. The final two 10% portions of the system would be completed in 2008 and 2009. Unfortunately, age or age-related data (the installation date) were *not* among the items carried in this system for poles installed prior to 2007.

Distribution Failure Reports. Since 1989 CL&P had maintained computerized records of every equipment failure on its distribution system: poles, cross-arms, conductor, brackets, service transformers, etc. Initially this was kept in a simple electronic spreadsheet format by month, but eventually moved into an SQL database on the enterprise IT system. There was concern about the quality of this data: the pole inspection program led to replacement of only an average of 19 poles a year for safety or hazard concerns (about 190 per full ten year inspection cycle, or .04%). But beyond that, the data was merely retained and used, so its quality was largely unknown. This is a legitimate concern, where data are collected for years but never used, and will be brought up again later. In 1989, in company with the inspections, CL&P had begun performing a simple spray application of anti-fungal treatment in company with each inspection. CL&P had never been certain how effective such treatment was, but inspectors had to visit every pole anyway and the marginal cost of the treatment was minor. Records indicate if each pole was treated and the date.

Regulatory, Legal and Vendor Reports and Documents. Information gathered for the initial hearings in 1988 on the pole inspection program, provided very little information beyond what was otherwise known. There were annual filings to the public utilities commission on the results of the inspection program, but those were merely summaries of the inspection data and provided no additional information. But a report filed in 1989 by a pole services vendor hired to do the pole inspections, a well-known and experienced supplier of poles and services, did prove useful as a snapshot of estimated system condition in that year.



Figure 16.5 Left, brand following the guidelines of ANSI standard 05.1.2008 gives from top to bottom: manufacturer (MANU), which plant the pole (C) came from, month (6) and year of manufacture (07), code for the wood type (SP-Southern Pine), code for the type of preservative (C-creosote), the amount of preservative infused in the pole (.75 lbs per cubic foot), the class (5) and length (40), and the owner (CL&P). Middle, some of this information as contained in pre-WWII and early postwar brands found on the CL&P system. Right, a depiction of the deterioration and lack of full readability typical of pre-WWII brands.

In addition, CL&P's predecessor company, Metro Illuminating Company, had gone through a business re-organization during the Great Depression. In 1933, a court-ordered audit prior to re-settlement of the company's debts accounted for all the assets the company had. This included a full count of "poles installed" and "poles in hand for spares" (inventory). A complete set of these audit books and the final report for the bankruptcy was available in the company records. Finally, upon a return to a fully civilian economy at the end of World War II, many companies, including CL&P, audited both assets and resources carefully in order to account for costs that could be used as tax credits or deductions due to the wartime regulations and practices of the Wartime Mobilization Commission and Office of Price Control, etc. Summary (but not detail) records from this audit, completed in 1947, were in the company's business records and, again, included actual counts of poles owned. All of this data was old, but there was no reason to think it was not dependable.

Pole Installation Work Order Records. CL&P had been replacing poles due to accidents, storms, failure, and NCR (not-condition related – mostly relocations) reasons for as long as it had been in business. Only records since 1990 were computerized but there were cabinets full of older paper records. The farther back one went, the dustier the filing cabinets became, and the more uncertain people were of terminology and short-hand notations on work order cards, etc.

But it appeared that nothing had ever been thrown away since a “revision” of work practices and records carried out at the end of WWII. In fact, some records were found for periods between the two world wars. Old records were valuable in their own way. This was a decades long process and problem – decades old data might help. There was never any serious thought given to computerizing all these records. Their value would be for spot checks of periods in the past to confirm long-term analysis and model-building later in the project.

Of more use were Monthly Summaries of crew activity and work completed, that included counts of poles installed, miles of line strung, transformers hung, etc. Some months were missing, but in the period 1947 to 1980 (416 months) only 42 monthly summary reports were missing.

Stores Department Reports and Inventory Tallies. CL&P’s pole-yard records had been computerized for only about two decades, but as with work orders and storm reports, paper records going back to near the end of WWII were kept in several old filing cabinets in the back of what had been the Central Meter Shop. These Stores Department records proved to be a treasure from the project team’s perspective. Monthly summaries – not just for poles but for cross-arms, conductor, cable, and many other pieces of hardware and equipment types -- were found for all but five months in the period 1947-1989, and for most of the time during WWII. In fact, somewhat less comprehensive summary reports were found, 77 of 108 monthly reports for the period 1933-1941.

Review and comparison of pole installations work orders monthly summaries from Operations records, and these Stores inventory reports, showed that these two sources did not always tally well on a week to week basis but generally rose and fell and corresponded nicely on a month to month and particularly a yearly basis. Regardless, basic data from these records for the period 1933 – 1989 were entered into spreadsheets by the project team.^{1, 2}

Weather Data on an hourly basis for the last forty years was available as part of weather normalization work the T&D planners did for load forecasts and equipment ratings. They also had data on major storms back to before WWI.

GIS System Customer and Spatial Property Data. GIS data used for spatial demand forecasting, and spatial information gleaned from municipal property

¹ For all equipment types covered by the records. The team rightly concluded that, eventually other equipment types might be of interest, too.

² Forty-two years of monthly paper Stores Inventory Summaries came to slightly more than 500 months. The average summary report was six pages which made a stack of paper a bit over three feet high. Each report covered about one hundred twenty five classes of equipment or materials – distinguishing poles by class and length, etc.

Data from one monthly form took about two hours to enter and verify. Co-ops working two at a time at \$12/hour took roughly six months to complete the transcription. Four weeks were required to transfer merge this with records of the last twenty years and assemble all into a coherent file covering a period of 60 years for a variety of equipment types and classes, at a total cost of about \$20,000: a bargain.

records gave fairly good estimates of the initial build out period (first construction) in every area of the system.

But no direct age data on a pole specific basis

CL&P was able to assemble a count of the poles it owned going back as far as 1933 (Table 16.3). A majority of poles in the system had a readable date in their pole branding, but almost half were not entirely readable.³ Missing from a majority was the one thing that the team most wanted, the year or date that pole was *installed* so that its service age could be established, and so that statistical analysis of the entire set, and trends within that set, could be performed.

Since 1989 the installation date for every pole had been entered in computer files as and when the poles were installed. But such data had not been recorded in any useable form prior to 1980 and perhaps not in all cases until computerization in 1990.⁴ Thus, as of 1989, for poles older than twenty years, the pole brand data taken down in the first round of inspections were fairly dependable for poles new enough that it could be read dependably. There was no concern here that that date was of manufacture and not installation: CL&P typically kept poles in its pole yard no more than perhaps two years and that was not an amount of error thought to matter. However, the older the pole, the more likely it was that a pole brand did not exist or could not be found, or that the characters in its date code could not be discerned clearly.⁵

Thus, CL&P had fairly dependable date of installation, and age, on 75,250 poles (17.2% of the total) – those installed since 1989 – the newest and thus the poles of least concern from an aging standpoint. About 70% had a date taken from reading the pole brand: How accurate those were for older poles was undetermined at the time. But in total CL&P had original date data, of some level of accuracy on just over 85% of the poles in its system.

³ “Not entirely readable”: one of the two digits was uncertain as to of its value.

⁴ Some paper records were maintained in two divisions beginning in 1980 through until the computerized system was installed in 1989. These were known to be incomplete and not used in the initial work described here. By the time they were re-examined several years later, the data proved unnecessary as existing age estimates were good.

⁵ Later, CL&P performed a very detailed statistical analysis of *readings* of the dates on pole brands. Some numbers are more likely to be misread as others (an 8 as a 3, for example) whereas others were less likely (a 1 is unlikely to be misread as a 2). Statistical analysis of mis-reading likelihood was used to refine the data. This led to some improvement in age estimates. However, at the time CL&P first went into this study it did not have this information.

**Table 16.3 Count of Poles
Installed at CL&P by Year**

Year	Pole Count
1933	26,555
1947	42,457
1980	357,852
1981	367,356
1982	374,222
1983	377,715
1984	378,810
1985	378,815
1986	380,476
1987	382,323
1988	384,652
1989	387,263
1990	389,988
1991	392,144
1992	395,426
1993	400,963
1994	406,589
1995	412,162
1996	416,170
1997	419,254
1998	422,327
1999	424,943
2000	426,621
2001	427,521
2002	428,915
2003	430,209
2004	431,219
2005	433,391
2006	434,304
2007	436,044
2008	437,226

The ideal pole management database would have included exact knowledge of the age of every individual pole in the system. From that, one could locate poles of any age or type, etc., and build any and all statistical tables needed, etc. But such data were not available, and so the first order of business was to assemble what data there were, try to resolve inconsistencies and estimate missing elements, to do the best that could be done.

Summarizing its data situation then:

- CL&P had comprehensive, pole-specific data going back to 1989 and slightly less accurate data going back another 9 years (to 1981).
- It had less dependable and comprehensive data from its branding surveys and other sources going back to the end of WWII.
- It had data of even less dependable nature going back a further 14 years (to 1933).
- Among data taken from records more than 30 years old, only overall pole counts for 1933 and 1947, were considered useful.

16.5 BUILDING THE HISTORICAL MODEL OF POLE OWNERSHIP

Here “model” does not refer to an analytical, algorithmic model. Rather, the model is the set of historical data on year-by-year counts of poles owned, poles added, etc., which the project team was going to treat as “correct” with respect to its study of the pole asset base. This database/model would be used for two main analytical/functional purposes.

- 1) *Analysis, prediction, and planning*: to determine the current characteristics (age, condition, remaining lifetime, expected current failure rate, expected change in those factors to change in operations and management practices, etc.). To estimate trends in the scale and nature of future problems, to study if and how trends would impact customer service quality, operations, and expenses, and to study how to best accommodate them in CL&P’s future operating plans.
- 2) *Execution of plans and programs*. There is a big difference between statistical infrastructure analysis and specific pole targeting. Analysis, prediction, and planning might determine, statistically, that 18,000 poles out of its 435,000 poles should to be replaced. But if CL&P could not find the older poles in need of replacement, the plan is useless. CL&P had to be able to locate poles of different ages and conditions dependably enough that its program would be effective. This proved to be the most difficult challenge it faced.

Since CL&P had age data on only a fraction of its poles, it would have to estimate – infer – the rest. The first step was to assemble the data and build up as

“complete” a historical data model as possible. There were two alternative approaches that could have been taken in assembling that dataset.

- 1) In its study of aging, failures, and deterioration, CL&P could restrict itself solely to analysis and inference of trends over only the preceding 28 years (since 1980), the period for which data was fairly complete and considered fairly dependable. Basically, it could build an accurate model of what, in terms of pole lifetimes, was a short period (less than half nominal service lifetime).
- 2) The project team could do a good deal of extra work trying to sort out the older sources of data so as to assemble them into some semblance of a reasonable picture of CL&P’s pole base prior to 1980.

Here, the cost and time required for this additional work was not the major concern of the project team. The key issue was that the older “data” in such a model of historical pole ownership might be less dependable than the data for the last 30 years and thereafter might contribute errors to all subsequent analysis.

After considerable amount of discussion, study and evaluation of its options, the project team took the second approach. It decided that models built on “fuzzy” data of a longer historical duration would be better than models built with the dependable but shorter data histories of the last three decades.

Three facts weighed in this decision. First, the vast majority of wooden poles last far longer than 30 years. In fact deterioration does not really become an issue in most cases until after this period.

Second, models developed from only 28 years of dependable data would *imply* a certain history *before* that historical data period. Even if only 28 years of data were used, it would only be prudent to at least study and compare model results to long-term historical data: basically one way or the other the team would be backing into a full historical analysis anyway, so it might as well do it overtly.

And third, the project team knew that this was only an initial data model, that further analysis would help refine the model, and that the team could always decide *not* to use the older data if the results indicated it was not contributing useful information.

Using an Input – Output Approach

Annual data was assembled from the Work Order Records and Stores’ Department’s inventory monthly counts from 1933 to 2008 (Table 16.4). The

Table 16.4 Estimates of Poles Added to the CL&P System, 1938 – 2008

Year	Added	Replaced Failed	NCR	Total Installed	Stores Bought	Data Disp'd	Surplus
1938				803	726		4
1939				1011	1060		4
1940				608	730		4
1941					863		5
1942				912	997		5
1943	1083	50		1133	1125		6
1944	1109	44		1154	1269		5
1945	1163	46		1210	1334		6
1946	2086	50		2136	2415	2471	6
1947	3052	74		3125	3231	3115	8
1948	3628	105		3733	3957	3880	8
1949	3510	60		3570	4348	4297	9
1950	4256	64		4320	4503	4615	8
1951	4176	68		4244	4309	4388	8
1952	4833	74		4908	5207	5298	9
1953	6425	78		6503	7013	7553	9
1954	7512	95		7608	7840	7977	10
1955	10590	89		10679	9684	9870	10
1956	10380	94		10474	10881	11085	11
1957	10700	100		10801	11163	11538	11
1958	11778	111		11889	13607	14653	12
1959	12093	120		12213	12951	13313	12
1960	14459	140		14599	14860	16276	13
1961	16026	143		16169	16100	16212	14
1962	16897	152		17050	16995	17323	15
1963	10501	169		10670	11992	11895	17
1964	9774	172		9945	10975	11051	16
1965	11514	199		11713	12255	12486	17
1966	10897	232		11128	11141	11941	18
1967	11683	210		11893	14224	13962	19
1968	14179	337		14516	13504	13802	22
1969	14642	247		14889	15869	16050	22
1970	11973	269		12242	13210	13705	23
1971	13678	278		13956	15742	15127	25
1972	15495	320		15816	18002	17941	32

Table 16.4 continued

Year	Added	Replaced		Total Installed	Stores Data		Surplus
		Failed	NCR		Bought	Disp'd	
1973	8423		401	8824	9536	9706	30
1974	4692		352	5044	5771	6010	31
1975	10305				11144	11302	
1976	5914				6366	6712	
1977	4201				4833	5209	
1978	4448				5457	5546	
1979	6479				7674	7985	
1980	8784				11267	10829	
1981	9297	362	151	9810	10805	11334	53
1982	6622	355	165	7142	8205	8597	54
1983	4007	357	162	4526	4686	4916	58
1984	1234	580	177	1991	2266	2374	67
1985	67	391	175	634	2683	2806	76
1986	1681	469	194	2344	3164	3102	69
1987	1829	418	196	2443	3176	3330	72
1988	2288	432	197	2917	3717	3898	76
1989	2618	567	210	3396	4124	4319	81
1990	2755	1062	238	4055	4287	4363	86
1991	2194	540	220	2953	3758	4018	91
1992	3225	534	240	3999	5039	5195	94
1993	5362	534	235	6131	7447	7708	98
1994	6199	798	273	7270	7526	7659	114
1995	5412	588	267	6267	7604	7921	110
1996	3978	646	259	4883	5932	6046	114
1997	3395	1253	294	4941	5081	5373	125
1998	3061	794	280	4135	5205	5223	123
1999	2640	1514	286	4440	4636	5102	127
2000	1726	710	293	2729	3692	3757	134
2001	997	787	325	2109	3025	3094	139
2002	1489	802	313	2604	3485	3692	162
2003	1508	1246	357	3111	3434	3577	158
2004	1116	877	333	2326	3190	3286	157
2005	2229	886	338	3453	4479	4734	166
2006	1064	1183	1176	3423	3186	3244	979
2007	1832	926	2652	5410	6608	4240	2488
2008	1403	1235	4502	7140	8988	3639	4312

specific meaning of the data in various columns needs some explanation in order to be understood in context. Column two is the total number of poles installed in each year: for example, in 2004 CL&P installed 1560 *new* poles on its distribution system in order to serve new customers or add to system capability.

Columns three and four (aggregated prior to 1981) are the poles *replaced* for failure and for NCR (non-condition related, such as relocations) reasons: for 2004 these are 1,074 and 948 respectively. Adding these figures produces column five's figure of 3,582 poles installed. This figure is considerably higher than the net increase in poles for the year 2004, which was only 1,002 poles (431,040 – 430,038 poles). Columns six and seven are the poles bought and dispersed to the field each year and generally match each other closely, and match the total number of poles installed. Finally, column eight is the count of poles which were retired. Thus, the net increase in pole count for 2004 is 3,582 – 1074 – 948 = 1002 poles. The alert reader will note a significant uptick in NCR poles for 2005 – 2008: CL&P embarked on a major “system hardening” program in which selected portions of the system were upgraded to withstand storm stresses more. Nearly 13,000 poles, 3% of the OH distribution system, were reinforced (replaced with a heavier class of pole) in this program.

Table 16.5 shows the “bottom line” for CL&P. As in Table 16.4, column one is the year. Column two gives the actual customer count for the year, with Column three showing the percentage growth rate computed from the numbers shown. Column four is actual pole count (where available, regardless of what accuracy it was thought to have). Columns five through seven show poles added for growth, for failures and NCR respectively, again, actual data only (failures and NCR are available only a sum prior to 1980). Finally, column eight is the initial estimate, formed from all the data, of the actual pole count for that year. These estimates were formed in slightly different ways for different periods of the history:

Estimated pole counts for years between 1933 and 1947 were produced by allocating the net increase in poles during the period (15,901 poles) proportional to customer count growth during the period. From 1933 through 1947, CL&P's customer base grew by 53% while the pole count grew by 59%. (The reasons why poles grew faster than customers will be discussed later.) Net annual increases in poles during the period were thus estimated as:

$$\text{Pole growth rate}(\text{year}) = 59/53 \times \text{Customer growth rate}(\text{year}) \quad [16.1]$$

Pole counts for years 1948 through 1979 were estimated by adding the recorded value for poles added that year (column five) to the pole count for that year, to give the pole count upon entering the next year. The only exceptions are for those years for which the recorded value of poles added has an asterisk – monthly records being incomplete in those years. In such cases the data available was used to fit to the trend smoothly, in company with the preceding and following year data values.

Table 16.5 Key Data and Results for Estimation of Pole Count by Year 1933 - 2008

Year	Customer		Poles	Data, Changes			Total
historical	Count	Growth Rate	(data)	Added	Failed	NCR Retr'd	pole count
1933	75,291	1.40%	26,555				26,555
1934	77,307	2.70%					27,352
1935	80,115	3.60%					28,466
1936	83,451	4.20%					29,794
1937	84,887	1.70%					30,369
1938	86,392	1.80%					30,972
1939	88,560	2.50%					31,843
1940	90,218	1.90%					32,511
1941	92,347	2.40%					33,371
1942	94,773	2.60%					34,354
1943	97,629	3.00%		1,083	50		35,514
1944	100,709	3.20%		1,109	44		36,770
1945	103,786	3.10%		1,163	46		38,029
1946	109,055	5.10%		2,086	50	6	40,193
1947	115,514	5.90%	42,457	3,052	74	8	42,457
1948	121,840	5.50%		3,628*	105*	8*	45,217
1949	127,744	4.80%		3,510*	608*	9*	48,998
1950	133,932	4.80%		4,256*	64*	8*	53,080
1951	139,069	3.80%		4,176*	68*	8*	57,043
1952	144,529	3.90%		4,833	74	9	61,849
1953	150,718	4.30%		6,425	78	9	68,272
1954	158,139	4.90%		7,512*	95*	10*	75,491
1955	167,690	6.00%		10,590	89	10	84,493
1956	179,108	6.80%		10,380*	94*	11*	94,660
1957	191,737	7.10%		10,700*	100*	11*	105,078
1958	206,367	7.60%		11,778*	111*	12*	117,674
1959	221,846	7.50%		12,093	120	12	129,763
1960	240,122	8.20%		14,459	140	13	143,617
1961	259,963	8.30%		16,026	143	14	158,434
1962	282,186	8.50%		16,897	152	15	174,307
1963	297,044	5.30%		10,501	169	17	184,981
1964	311,612	4.90%		9,774	172	16	194,838
1965	328,003	5.30%		11,514	199	17	206,108
1966	342,879	4.50%		10,897*	232*	18*	216,227
1967	360,061	5.00%		11,683	210	19	228,409
1968	378,235	5.00%		14,179	337	22	240,808
1969	398,870	5.50%		14,642	247	22	255,100
1970	416,981	4.50%		11,973	269	23	267,090

Table 16.5 continued

Year historical	Customer		Poles	Data, Changes				Total
	Count	Growth Rate	(data)	Added	Failed	NCR	Retr'd	pole count
1971	436,572	4.70%		13,678*	278*	25*		280,447
1972	460,105	5.40%		15,495	320	32		296,041
1973	473,055	2.80%		8,423	401	30		304,486
1974	480,423	1.60%		4,692	352	31		309,198
1975	495,506	3.10%		10,305				319,011
1976	503,965	1.70%		5,914				324,242
1977	510,290	1.30%		4,201				328,060
1978	517,530	1.40%		4,448				332,399
1979	528,666	2.20%		6,479				338,849
1980	543,846	2.90%	357,852	8,784				357,852
1981	560,636	3.10%	367,356	9,297	362	151	53	367,356
1982	573,409	2.30%	374,222	6,622	355	165	54	374,222
1983	580,653	1.30%	377,715	4,007	357	162	58	377,715
1984	583,179	0.40%	378,810	1,234	580	177	67	378,810
1985	584,788	0.30%	378,815	67	391	175	76	378,815
1986	586,397	0.30%	380,476	1,681	469	194	69	380,476
1987	592,262	1.00%	382,323	1,829	418	196	72	382,323
1988	600,527	1.40%	384,652	2,288	432	197	76	384,652
1989	610,701	1.70%	387,263	2,618	567	210	81	387,263
1990	621,425	1.80%	389,988	2,755	1,062	238	86	389,988
1991	629,995	1.40%	392,144	2,194	540	220	91	392,144
1992	643,157	2.10%	395,426	3,225	534	240	94	395,426
1993	665,655	3.50%	400,963	5,362	534	235	98	400,963
1994	688,986	3.50%	406,589	6,199	798	273	114	406,589
1995	712,576	3.40%	412,162	5,412	588	267	110	412,162
1996	729,882	2.40%	416,170	3,978	646	259	114	416,170
1997	743,379	1.80%	419,254	3,395	1,253	294	125	419,254
1998	756,976	1.80%	422,327	3,061	794	280	123	422,327
1999	768,673	1.50%	424,943	2,640	1,514	286	127	424,943
2000	776,236	1.00%	426,621	1,726	710	293	134	426,621
2001	780,302	0.50%	427,521	997	787	325	139	427,521
2002	786,635	0.80%	428,915	1,489	802	313	162	428,915
2003	792,539	0.80%	430,209	1,508	1,246	357	158	430,209
2004	797,156	0.60%	431,219	1,116	877	333	157	431,219
2005	807,165	1.30%	433,391	2,229	886	338	166	433,391
2006	811,383	0.50%	434,304	1,064	1,183	1,176	979	434,304
2007	819,481	1.00%	436,044	1,832	926	2,652	2,488	436,044
2008	825,000	0.70%	437,226	1,403	1,235	4,502	4,312	437,226

Data for 1980 through 2008 was treated as valid. Values for 1980 through 1988 were thought to be less accurate than for later years, but any error was treated as “normal” data error and no adjustments were made in the raw figures. The result is an actual or estimated pole count for each year, 1933-2008. Figure 16.6 shows the pole count trend along with CL&P’s customer count data for the period.

How Accurate are these estimates?

If one treats the 1933 pole count taken from the court-ordered audit of Metro Illuminating Company Assets done in 1933 and the 1947 post-war pole audit as accurate (as CL&P did) then the long-term trend of pole ownership had to be something close to that shown: the trend has to go through the 1933, 1947, and 1980 to 2008 points.

Overview of Long-Term Trends

Figures 16.6 through 16.9 show several long-term trends including structural changes in system approach typical of any utility over so long a period. Customer count increased by a factor of 11 from 1933 to 2008, pole count by a factor of 16.5. Figure 16.6 shows the total number of customers and poles in the CL&P system from 1913, when the company began operating as Metropolitan Lighting through 2008, as estimated by the project team. Figure 16.7 shows the number of poles known or estimated to have been added each year. Figure 16.8 shows the ratio of customers/pole, and Figure 16.9 the ratio of poles added per customer added, by year.

Prior to and into WWII, CL&P had roughly one pole for every three customers and was adding about .4 pole per new customer, both rather low numbers compared to later values, and both due to two factors in play at the time:

Design practices – distribution primary circuits were very low voltage compared to today’s systems (some only 1 kV) and there was little attention given to aesthetics or to the reliability impact of common mode failures: as a result some poles carried as many as six circuits.

Demographics – prior to WWII, a much higher ratio of residences and businesses were located in relatively dense urban multi-family apartments and buildings located in cities.

New design practices. Beginning in the late 1930s and accelerating up to the beginning of WWII, CL&P, like many utilities, moved to OH distribution design




Figure 16.6 Cumulative connected customers and poles installed on the CL&P system for years 1913-2008, as developed in the data aggregation/estimation. Values for 1913-1932 were estimated by straight-lining to 1933's numbers from zero in 1913, the year the company was founded.




Figure 16.7 Number of poles added to the CL&P system each year from 1913 through 2008. Slightly more than 510,000 have been installed on the CL&P system. Yet only 437,226 remain. The others are retired, upgraded for NCR reasons, or have failed and been replaced.




Figure 16.8 Number of poles per customer in the CL&P system, 1933 - 2008, has changes attributable to shifts in both design policy and demographics of growth in the service territory.




Figure 16.9 Ratio of poles added per new customer gives a more volatile view of changes over the period. Spikes in the 1970s are either data error (there is some indication that totals were "trued up" only every few years) but could represent valid intense building programs dictated by growth spurts and budget issues.




Figure 16.10 Initially as the city grew outward it had mostly flat, developable land around it on all sides. By the 1970s and 1980s, however, development was having to grow up and into canyons between hills and mountains (shaded areas). The net result was less land for development as a function of distance from the city center. This contributed to a slight diminishment of the downward trends in poles/customer that began in the 1970s.

practices that would be much more recognizable to modern utility engineers and managers. It standardized on a new, higher distribution voltage,⁶ and settled on guidelines that called for one, or on rare occasions no more than two, three-phase circuits per pole. This led to a modest increase in poles per customer.

Suburbanization: In the 1940s and 1950s, CL&P's service area like nearly all US metro areas started to develop significant peripheral suburbanization. A greater portion of new homes were built singly, and on larger lots, in more outlying areas. Shopping centers were built outside of the dense downtown area. The widespread "suburban sprawl" resulted in lower customer density and hence a need for more and longer feeder runs to reach customers – more poles per customer.

Underground distribution began to become popular in the CL&P system's suburban areas in the 1970s. While three-phase circuit mains were (and still are)

⁶ Modern utility engineers might find it difficult to believe, but there was a time when 4.16 kV distribution was considered new and high compared to current standards.

largely overhead, laterals were increasingly built as direct-buried (sheathed) underground in residential subdivisions (URD – underground residential distribution). The number of poles per customer began to drop in the 1970s due to this design practice.

“Ruralization” As the metro area grew outward, the topology of land for metropolitan expansion changed (Figure 16.10). The city had originally been founded on flat lands on both sides of a river not too far from the coast and grew uniformly in all directions. But as the city continued to expand, it ran out of such land on its periphery and began moving into hill country valleys to the west and north of the city in narrow (only several miles wide) strips running out from the central hub: the amount of developed land compared to distance outward from the city center by half. Longer “rural” feeders were needed to reach out into these valleys. This tended to push poles/customer ratio up slightly.

“Hardening” A particularly violent storm hit CL&P in 2004, requiring replacement of over 7400 poles and causing multi-week service interruptions in some areas. After consultation with and approval by its Public Utility Commission, beginning in 2006 and continuing through 2008, CL&P implemented a new set of overhead circuit guidelines that called for more, and stronger, poles along key commercial corridors and in key places in its system.. Certain circuits were proactively reinforced with additional poles (shorter spans). About 8,000 poles were replaced and about 3,500 added to the system during this work. This tended to also increase overall pole count/customer ratio.

16.6 BUILDING THE PREDICTIVE MODEL

CL&P needed a predictive model, a tool that would essentially say to them: “Although today’s distribution pole failure rate is X% and the total costs and reliability impacts it produces are Y and Z respectively, in ten years failure rate will be A and cost and SAID contribution will have risen to B and C respectively.” This section discusses how the analytical sustainable-point model (Appendices) was developed from the data discussed in the previous section.

Equivalent age would be used as the measure of condition

Condition – the measure of deterioration for a pole – would be *equivalent age*, defined as the average amount of deterioration in a pole with that number of service years. For example, the level of deterioration associated with 60 years of service is the *average* deterioration for *all* poles of that have been in service 60 years. That for 30 years is similarly the average of all 30-year-old poles, etc.

Since average deterioration as a function of equivalent age was identical to average deterioration as a function of chronological age, given a large enough population of poles, actual chronological age could be used to develop the

overall failure rate curve as a function of equivalent age.⁷ Determining how failure rate varied with age was fairly straightforward to develop from the historical data (Table 16.4), although the project team reminded itself constantly that this did not mean the developed model was highly accurate, due to uncertainties in the “data” as discussed in Section 16.4.

Details Not Included in the Initial Analysis

This initial analysis ignored or glossed over several potentially important complications. First, like most utilities, CL&P had poles of *many different types of wood*. The majority of its recent poles were second-growth southern yellow pine, but CL&P had significant numbers of Douglas fir and other types of wood in its system. Similarly, *preservatives* ran the gamut from creosote to CCA to Penta to other less common treatments. The initial analysis did not distinguish these differences mainly because data on what type of wood or preservative was used in each pole was not available.

Stationary failure and NCR functions. The deterioration rates and failure rate functions (curves) might not be a stationary function in time, i.e., they might have gradually changed over time, as wood quality, pressure treatment techniques, installation practices, and post-installation treatment procedures changed. Regardless, the project team decided to initially assume one model would work for all decades, even six or seven in the past. It would check this afterward, by looking at the fit at ten-year intervals from 1960 to 2010.⁸

A complication that *was* considered was the effects of the proactive danger-pole replacement program and the anti-fungal treatment and replacement programs begun in 1988. As mentioned previously, inspection results were used only to replace a small number of safety-related and egregiously damaged poles each year. It was assumed, therefore, that all such poles would have failed within three years, and adjustments were made in the data base for the model to reflect that.

As to the anti-fungal treatment program, it was begun in 1989, being applied to 10% of the poles in the subsequent years, and re-applied on ten-year increments. The impact of the treatments would be gradual, and poles installed long before 1988 had gone through much of their expected service lifetime before receiving any such treatment and probably would not be that affected. The manufacturers of the treatments and the service providers that performed the work had reports and statistics that showed the effect they predicted on poles, but the project team was not certain that they applied to CL&P’s situation. This factor was not looked at until later in the program.

⁷ Since the relationship is non-linear, the use of averaging here might have slightly biased results. It had a small but negligible effect that will not be discussed here.

⁸ Slight differences seen were well within the expected fitting error. There was no steady trend in the character of the change (i.e., the curve did not gradually grow more exponential each decade, or have any other major characteristic that showed a trend over time, etc.).

The Historical and Predictive Model

The initial work by the project team, and all examples given here, were done with Microsoft Office Excel® and Access® without the use of add-on features or engines, but included the full statistical and “solver” features available with Excel. The model was subsequently transferred to a SAS platform.

The project team developed a historical and predictive model of wood-pole age distribution, NCR and failure replacements, treatment and inspection costs, as a function of year, that is essentially identical to the models used for the examples in the Appendices.

What “Fitting” A Sustainable Point Model Means

Appendix A describes the sustainable point spreadsheet model in considerable detail. This discussion will only highlight key points so the reader can follow what “fitting” means and can understand all it accomplished.

Given a failure rate as a function of age, and a distribution of ages for pole is year N (i.e., 1,831 poles are one year old, 1,763 two years old . . . 340 are 65 years old, etc.), Appendix A’s spreadsheet model computes the distribution of ages for year $N+1$. Suppose that in year N , 4,00 poles are shown to be 62 years old and that the failure rate curve shows that at that age annual failure rate is .92%. Then the count of poles that are going to be 63 years old in year $N+1$ is $4700 \times (1-.92\%) = 4,657$. A total of 43 of those 62 year old poles are predicted to fail in year N . The model similarly computes the number of poles of other ages that survive to fill in row $N+1$. It adds up the total failures, 43 for age 62 plus those for all poles of other ages. Since failed poles are replaced, it makes the total of failures the number of new poles (age < 1 year) in year $N+1$. Hence it has the year $N+1$ distribution.

It then computes the year $N+2$ distribution in the same way.

In doing this, the model shows the effects of NCR and retirement. In addition, the model can add new poles for growth: there were 437,226 poles needed in 2008, but system growth needed quite a few more to be added in 2009, etc. At the conclusion of a row’s computation, the model has therefore determined how many poles of what ages will survive into the following year, how many new were added, for year $N+1$, due to all reasons. The model then computes year $N+2$ from $N+1$, etc., all the way through to the final year. In this particular case, the model was set up with a starting value for 1912 of zero. The number of poles known or estimated to have been added in each year (fifth column in Table 16.5) is a “known” and entered as the total additions that year.

Solve for the Failure, NCR, and Retirement Functions of Age

CL&P had no poles in 1912. It had 26,555 in 1933, and 437,226 wood poles in what was the base year for this study (2008). Tables 16.4 and 16.5 give mostly reliable and mostly complete data on pole counts, additions, NCR, etc., for various years in between. *Solving, or fitting, this model, means finding the*

failure rate, NCR, and retirement rate functions of age that make the model best replicate, or match, the data in Tables 16.4 and 16.5. There was a statistically large sample set: according to the data developed, CL&P had installed 510,221 poles on its system since 1913. Only 437,226 remained. The other 73,204 had either failed (20,027), been replaced for NCR reasons (29,578) or been retired (33,192). The failure rate, NCR, and retirement functions wanted were those that would best fit this data. The reader interested in important if secondary details should note that of the three causes of a pole leaving the system, *failures* had the lowest count.

Years versus Age

In the following discussion of exactly how the model was built and fitted, it is important to distinguish between age and year used throughout.

“Age” is the time a pole or group of poles has been in service.

“Year” is, without exception, the chronological year.

Thus, a pole installed in year 1913 has an age of 75 in 2008. There were 437,226 poles going in to year 2009, but only 8,468 (see Table 16.4) had an age of zero (new that year).

The actual process of fitting the model was straightforward if intricate in certain places. As mentioned earlier, a number of basic types of formulae were tried for each, as shown in Table 16.6. For each, for each of the failure, NCR, and retirement trends, coefficients that best fit the data were found.

Table 16.6 Types of Formulae Tried for Modeling Poles

Type	Formula
Constant	Rate (year) = A
Linear	Rate (year) = A + B*year
Multiplicative	Rate (year) = A^(1+B)
Add and Multiply	Rate (0) =A, Rate (year) = Rate (year-1)*(1+B)
Logarithmic	Rate (year) =A*2.718281828^(B*year)
Binomial	Rate (year) = A+B*year+C*year ²

RMS minimization seeks to minimize the root-mean square of the residuals (modeling mismatches), penalizing the fitting function more than proportionally for large mismatches as compared to small (a residual that is twice as much is penalized four times as much – the fitting method can improve the solution by reducing it a small amount even if it means a much smaller residual for another

data point increases by slightly more. RMS is by far the most popular error measure used in basic statistical modeling, but its popularity has more to do with the fact that standard solution methods (regression) minimize RMS rather than firm evidence that it always provides the most useful models.

AAV methods minimize the sum of the absolute values of the mismatches: a residual that is twice as big is penalized twice as much, and the fitting algorithm can only make an improvement if an adjustment to the formula reduces the mismatch of one point more than it might increase another's. Maximum deviation minimization methods seek to minimize only the largest mismatches. They can improve the result if the adjustment to a formulae results in the greatest mismatch being smaller.

Three types of error measures were used, as outlined in Table 16.7, in order to see if, and why, results might vary depending on which was used. In addition, each was fitted on both an absolute and percentage error basis. All these variations were done without any expectation that the project team would discover one was much better than another, and mostly because the work required only a slight amount of additional time.

Table 16.7 Error Metrics Tried in Determining Best Fit

Type	Formula
RMS	Root mean square of the residuals
AAV	Average absolute value of the residuals
Maximum	Single largest residual

The basic sustainable point analysis model was created for CL&P's pole base (see Appendix A). The last column of Table 16.5 (estimated total pole count) was set up as the totals by the year column in the model. The basic formulae (exponential, simple increasing, or whatever was being tried) were set up for its failure, NCR, and retirement rates as functions of age.

When coefficients (A, B, C) are given for each of the three functions, failure, NCR, and retirement, the spreadsheet computes:

- The amount of failures, NCR replacements, and retirements that would have taken place, each year.
- The amount of poles added as replacements given those failures and replacements and removals, in order for the total at the end of the year to match the pole total for the year (Column 9, rightmost, in Table 16.5).
- The distribution of ages of the pole asset base in each historical year that would have resulted from this.

Each function – failure, NCR, and retirement – had up to three coefficients (A, B, C) depending on the function being used. The specific values of A, B, and C for each would determine how it “predicted” failures, NCR, or retirement rate with age.

Simultaneous Solution for All Three Functions. MS Excel’s Solver feature was used to find the coefficients for each of these functions in a series of cases using each of the formula (Table 16.6), with each of the error metric (20.6) – a total of eighteen possible cases. Theoretically, one cannot solve for these functions individually: all three affect the same population with the same effect (poles disappear over time). In practice the project team did, in stages, focus on each individually. However, it began with a simultaneous solution for all three functions for all eighteen cases. Designating the coefficients for each function as

A_f, B_f, C_f for the failure rate curve

A_n, B_n, C_n for the NCR rate curve

A_r, B_r, C_r for the retirement rate curve

the same type of formula and the same error metric was used for all three functions, failure, NCR, and retirement, i.e., a particular case would involve modeling failure, NCR, and retirement rates as binomial equations, minimizing RMS error.

In every case in this first attempt at solving for model formulae, Excel’s Solver was used to find the values of $A_f, B_f, C_f, A_n, B_n, C_n, A_r, B_r, C_r$ that would minimize error in fitting to:

The total of poles installed each year	– weighted 50.00%
The total of poles that failed each year	– weighted 16.67%
The total of poles that were NCR each year	– weighted 16.66%
The total of poles that were retired each year	– weighted 16.67%

The number of coefficients being determined and the difficulty in getting a solution varied from case to case. But there were never more than nine variables (A, B, C, for each of the three functions) and 75 years of data, so the problem was quite tractable. In this first iteration, the work to determine the various fitted models required about two weeks by one person full time with others assisting at various points.

Refining the Model. Thirty six cases were developed. The project team studied the results and identified which it thought were best. Some models were clearly superior to others, but in several cases there appeared to be no significant difference in the overall quality of one formula versus another.

Various exploratory cases were run to further study results, try various ideas, and refine the model. A variation on the fitting approach was tried, a form of stepwise fitting: a single function was first used to explain failures and NCR,

then functions to disaggregate it into two trends were determined. It yielded no better results. Different formulae types were used for different functions (e.g., linear for retirements and logarithmic for failures, etc.). In other studies, two of the formulae were held constant and only the coefficients of the third allowed to vary. In some cases, it was determined that particular modeling aspects made no difference at all (The failure rate curve described below could also be exactly duplicated with a logarithmic function, etc.).

Eventually, the project team settled on the following models:

$$\text{Failure rate (age)} = .035\% * (1 + e^{(.00095 \times \text{age})}) \quad [16.2]$$

$$\text{NCR rate (new through five years of age)} = \text{zero}$$

$$\text{NCR rate (6)} = .00188\% \quad [16.3]$$

$$\text{NCR rate (age + 1)} = \text{NCR rate (age)} \times 1.022 + .001\%$$

$$\text{Retirement rate (new through ten years)} = \text{zero}$$

$$\text{Retirement rate (11)} = .0088\% \quad [16.4]$$

$$\text{Retirement rate (age + 1)} = \text{Retirement rate (age)} \times 1.02 + .00095\%$$

Figure 16.11 shows these curves plotted for ages from new to one hundred years in service. Table 16.8 gives the fitting statistics for the various data, Table 16.9 various values at twenty-five year intervals. Figures 16.12 and 16.13 show actual versus modeled results of poles installed and failures.

Table 16.8 Error Levels in Model Fit to CL&P Data

Data Trend	RMS - %	AAV - %
Total Poles Installed	8	6
Total Poles Installed normalized storms	7	5
Failures & NCR (1938 - 2008) normalized storms	16	14
Failures & NCR (1938 - 2008)	12	10
Failures (1981 - 2008)	22	19
Failures (1981 - 2008) normalized storms	10	9
NCR (1981 - 2008)	12	9
Retirement	16	10

Table 16.9 Annual Rates as a Function of Age

Trend	Age in years →	New	25	50	75	100
Failure		.05%	.05%	.18%	1.2%	21%
NCR		0	.06%	.11%	.21%	.25%
Retirement		0	.02%	.05%	.12%	.25%




Figure 16.11 Annual rates for failures, NCR replacement and retirement as a function of age, as determined by fitting the sustainable point model to CL&P’s historical data.




Figure 16.12 Modeled (dotted line) versus CL&P data (solid line) for the number of poles installed each year for all causes (growth, failure, NCR reasons).




Figure 16.13 Modeled (dotted line) versus CL&P data (solid line) for the number of poles installed each year for failure and NCR reasons only. Data for the period shown as a gap was unavailable.

Overview of Fitted Trends and Functions

All three functions of age - failures, NCR, and retirements - increase with increasing age (Figure 16.12). Among them, failure rate (including storms) escalates most rapidly as age increases, reaching .5% at age 50, 1% at age 70 and over 20% at age 100.

For a pole in its first few decades of service, the likelihood is slightly higher that it will be removed for NCR or retirement reasons than for failure. In the decades between 30 and 50 years of service, failure rate escalates to slightly more than the sum of the NCR and retirement rates, but a pole in “middle age” is still nearly as likely to be removed (whether replaced or not) for reasons other than failure. But while NCR and retirement rates increase with age their trends are not nearly as exponential as failure rate: by age 75, a pole is nearly four times as likely to fail as to be replaced for NCR reasons or be retired.

Assessing the Quality of the Resulting Model

Qualitatively, nothing with respect to the three functions determined by the fitting surprised, or the age distribution to be examined in a few paragraphs, surprised or concerned the project team.

The overall fitting statistics (Table 16.8) looked reasonable. The single most important assessment of the quality of the model was the fit to failures (Figure 16.11), because this was, ultimately, the one trend that mattered most. The other three reasons that poles have to be installed - NCR, retirement, or for growth itself - do not cause reliability or safety problems, represent work that was both somewhat discretionary and that could be scheduled at the company's convenience, and are not directly related to equipment deterioration.⁹

Since data distinguishing between NCR and failure counts did not exist until 1981, the long-term match had to be assessed on the sum of failures and NCR: AAV and RMS are respectively 16 and 14%. For the period when distinct failure counts are available (1981 to 2008), the AAV and RMS of 19% and 21% respectively are not really a good fit. But as Figure 16.13 makes clear, the major source of mismatch is the spikes in actual failures (storms), which are to some extent random and unpredictable.

However, when the actual data is “normalized” for constant annual storm damage, the fit is much improved: AAV and RMS are 11.7% and 10.4% for the entire history, and 9.9% and 8.6% since 1981: the fitted sustainable point model was judged a reasonably good statistical fit and a useable model for planning.

⁹ Figure 16.10 shows that NCR and retirement rates are both clearly a function of age, but this is not due to equipment deterioration. Instead, they are directly a function of time: after a sufficient period, customers or governments grow and need changes made (NCR) or growth and redevelopment occur and lead to retirement of old equipment. NCR and retirement need to be modeled because they are significant sources of changes in poles (up to 40%) to the extent that if ignored, the fit to failures will be poor.

16.7 WHAT THE MODEL SAYS ABOUT THE FUTURE

Why the Age Distribution Curve Has the Shape It Has

Figure 16.14 (top) shows the Asset Base Model's estimated distribution of ages for CL&P's poles in 2008, with the middle and lower diagrams showing predictions for what that distribution would look like after ten and twenty further years of service. There, and throughout this entire section, the analysis is restricted to only the 437,226 poles that were in place as of January 1, 2008. CL&P would continue to add poles as it expanded its system to accommodate new customers (it would add over 3,250 poles for expansion reasons in 2008 - 2009), but typically an aging infrastructure study "cuts off" additions, as done here throughout.

Failures, NCR, and retirements have had a significant effect on the age distribution of CL&P's poles. Over 70,000 poles, or 14% of all poles ever installed on the CL&P system, are no longer in service, so failures and removals do influence the shape of the distribution. *But 86% of the original poles are still there.* Thus, the original installation date dominates the shape of this distribution: it is high in the middle – representing the period about 40 years earlier – because that is the era when CL&P was growing at a high rate due to a robust local economy. The high values (spikes) in Figure 16.14 correspond to the spikes in the Figure 16.12 (number of poles installed by year). The maximum number of poles ever installed in a year was 16,880, in 1962. There were two high-growth years shortly thereafter, when more than 15,000 poles were installed, but there was no subsequent year close to those figures. As a result, even though in the 46 years from 1962 to 2008, an estimated 1309, or 7.8% of those 16,880 poles, had failed, been replaced for NCR reasons, or been retired, there were more poles of age 46 (15,571) than of any other age. The Asset Base Model estimated the average age of a pole in the CL&P system in 2008 as 37.4 years. Fully half are between 37 and 57 years of age, but 45% are spread among those ages less than 37 years, and only 5% are estimated to be over 60 years of age, so the average is quite far from the median of 47.5 years.

The Asset Base Model's estimated expected failure rate for 2008 is .27% (an expected 1,188 failures), which is equal to the failure rate curve's value for a pole of age 47 years. The CL&P wooden distribution pole asset base might have an average age of 37.4 years, but it is failing at a rate equivalent to all poles in the system being 47 years of age, due to the non-linearity of the failure rate curve as a function of age in combination with the distribution of age. There are a few, *but enough*, really old poles failing at a much higher rate.

Failures were slightly *less* than half of the total number of poles withdrawn from the system in 2008. Counting NCR and retirements, the model predicts a total of 2,516 poles removed from the system, whether replaced or not, with only 1,188 being failures. All these values are the Asset Base Model's estimates, but close to the values actually recorded in 2008 – certainly well within variations due to weather, economy, etc., that play on those trends.




Figure 16.14 Age distribution for poles as predicted for 2008, 2018, and 2028, showing the expected evolution of the current pole population, without growth. As the “age hump” moves to the left, number of failures nearly doubles, from 1,189 in 2008 to 2,275 in 2028.

Sustainable age, if considering only failures, not NCR and retirements, is 51 years. Sustainable failure rate – the failure rate the system will gradually trend toward as aging continues, but then stay at permanently, is estimated at 1.21% - about four and one-half times the 2008 failure rate – or 5,290 per year. Thus, whatever CL&P's management thinks of the situation in 2008 with 1,188 failures predicted out of 437,226 poles, compared to what it can expect in the long run, the situation cannot be considered really bad.

Sustainable age, if considering all causes failures, NCR, and retirement, is 48.7 years – with a total withdrawal rate of 1.19%. The sustainable failure rate with and without NCR and retirements is essentially the same, so it is clear that the reason all the trends will worsen is *entirely* due to deteriorating condition of and increasing failure rates for poles. This is not a surprise, of course, but it has been confirmed by analysis.

The Age Distribution's Evolution Over Time and Its Implications

Looking at the future of only the 437,226 poles currently in the CL&P system at the end of 2008 (i.e., neglecting any it adds after 2008) the age distribution will evolve as shown in Figure 16.13. The hump of what are middle-aged poles in 2008 will gradually move to the right as they age, shrinking very slightly as more and more of this group fails or are withdrawn from service. Due to accelerating failure rate over time, the spike of poles installed in 1962 will shrink at a faster rate than the other spikes (1969, 1972, etc.). After 2019 it will no longer be the highest spike. So eventually, the shape of the age distribution will change. But these changes will not materially affect the situation within the next 20 years: the distribution for 2028 is still recognizably a version of 2008's.

As the mass of poles added in the 1960s and 1970s – the hump – ages and moves to the right, it eventually moves into the range of ages where failure rate begins to climb steeply. Annual failure count increases significantly as a result. The Asset Base Model predicts that annual pole failures will more than double in twenty years, from 1,188 in 2008, to 1,789 (an increase of 50%) by 2018, and to 2,614 by 2028 (an increase of over 46% over 2018's value and 2.2 times that for 2008). Note that the percentage increase is actually dropping every decade – it is 50% from 2008 to 2018 but only 46% from 2018 to 2028. It will continue to decrease gradually. This is the asymptotic trend toward the sustainable point, although the total number of failures will continue to increase until it reaches the sustainable value of an expected 5,290 each year, the percentage increase each year will decrease, until at the sustainable point, it is zero. This, however, will take nearly a century and is of little practical importance to CL&P. The next twenty years are the concern, during which the severity of “pole problem” will more than double.

To the project team the overall trend boils down to a simple “elevator speech” they can take to upper management: *for the foreseeable future, unless action is taken, distribution pole failures and the cost and reliability consequences they drive will increase by about 4% annually.*

A Key Historical Point: the Move to URD Design: 1968 to 1977

In company with utilities everywhere, CL&P originally built the majority of its distribution system using bare conductor on overhead lines. It had a very small amount of underground distribution in its downtown area: ducts under streets for cable, underground vaults for “submersible” transformers, etc. But the majority of its territory, all suburban and residential areas, was served entirely by overhead distribution.

In the late 1960s CL&P was one of the first utilities to take up underground residential distribution (URD) systems in suburban areas. Both traditional all-overhead distributions systems and URD systems use overhead distribution for three-phase portions of primary circuits. But in a traditional all-overhead distribution system the single-phase laterals are also carried overhead. In a URD system that function is replaced by buried single-phase primary conductor: laterals are underground, not overhead. *The URD system has fewer than half the total number of poles, and almost no poles carrying single phase lines.*

This means that the mixture of poles in a traditional all-overhead system differs markedly from that in a URD system. A URD system has fewer poles per customer than an all-OH system. More important here, the average pole is more important: any pole failure in a URD system is likely to be on a three-phase main circuit. Only a portion will be in an OH system.

CL&P began its transition to URD in 1968. By 1978 the company was building nearly all URD in suburban areas -- where the vast majority of its new




Figure 16.15 Distribution of ages for all poles (solid line) and poles for just the three-phase portions of the CL&P system (dotted line). A majority of poles installed prior to 1968 were lighter poles on single-phase circuit sections, but after 1978 the majority of poles were the heavier three-phase type.

construction was located. As a result, the CL&P pole base has a significantly different mixture of pole types and sizes in those portions of the system built before 1968 as compared to those portions built after 1978. The older portions are about 66% lateral poles, which are generally lighter-duty poles carrying lower capacity single phase lines and equipment. Only about 33% of poles in these areas are the heavier three-phase poles.

In the newer portions of the system three-phase poles make up about 85% of all poles installed. There are very few light single-phase poles. In the period 1968 to 1978 the mixture gradually changed from the one figure to the other (Figure 16.15).

Figure 16.16 shows why the shift to URD is quite important to the aging infrastructure analysis done here. The distributions shown there differ from those of 16.13 given earlier in that these distributions show the distribution of ages for poles that are failing (the early distributions were for all poles in service). The area under Figure 16.15's top curve (that for 2008) is 1,188 – the number of poles failing in that year – and the diagram shows the distribution of ages of those 1,188 failed poles. The distribution shows a hump of failures in the 70+ age range. That indicates a lot of failures are due to very old poles failing, in fact since there were never a lot installed (the utility was much, much smaller 70 years earlier) and many have already failed (about 20% or more), this is a very high portion of 70 year old poles failing each year.

But despite this, the greatest number of poles failing – the higher part of the distribution – is in the 35 to 50 year range: while failure rate for this age range is far less than for poles 75 to 90 years old, there are many, many more poles of this age, so that failure rate times number of opportunities to fail nets out to a higher number overall.

The shaded portions of Figure 16.16 show poles that were installed after the shift to URD. A large majority of these are three-phase poles, whereas only about a third of those in the unshaded part of the distribution are. The middle and bottom diagrams show the predicted failure distributions for 2018 and 2028. The most important change is the increase in area under the curve – the total number of failures more than doubles by 2028, by as seen, the portion of post URD poles that are failing increases dramatically: by 2028, a far greater number of major poles are failing than in 2008. This has major implications for the impacts on cost and reliability which will be discussed later in this chapter.

To distinguish between the two types of poles – lateral and three-phase poles, the CL&P project team accounted for this issue in its spreadsheet Asset Base Model by splitting its model into two portions: lateral poles and three-phase poles, splitting the count in the database 33/66 prior to 1968 and 80/20 thereafter. As its models and experience evolved, as it moved to a more specific database (this will be covered in more detail later) and particularly when it moved to the SAS-based rather than spreadsheet based model, the database was modeled to represent a wider range of pole types (by class), but the results given here were not developed with those models.




Figure 16.16 Distribution of ages of pole failures on the CL&P system for 1998, 2008, and 2018 as computed by the calibrated Asset Base Model. Shaded areas are poles installed after 1970, when URD practices were already well established. URD system poles go from a minority to a near majority of failures during the period 1998 to 2018.

16.8 PROJECTIONS OF POLE-RELATED COSTS AND PERFORMANCE

With the completion of the Asset Base Model and its database, the project team could now look at aging, failure, and cost trends in more detail. It evaluated them with regard to safety, cost and resource needs, and reliability of service, and for overall cost to the corporate bottom line.

Safety Impact Was Considered Under Control

Safety issues with poles and other equipment in easements had always been a high priority at CL&P, particularly since a pole accident with injuries led to the 1989 PUC regulation mandating the ten-year pole inspection cycles. Several years earlier a pole had collapsed near a school bus stop, slightly injuring several children. For the past two decades, CL&P had done inspections and had had a program of replacing poles assessed as safety concerns. The increase in the number of “hazard poles” found annually had roughly tracked the annual failure rate, nearly doubling over the past two decades. If one assumed that this would continue then the potential safety issues would be doubling, too. However, CL&P considered that its present program was effective at finding at-risk poles that posed major safety hazards, and so felt that the program would continue to do so. Thus, safety, *per se*, was not a special concern.

However, there were claims filed by persons and businesses for damages caused by fallen poles – those that hit a car or crushed a garage roof, etc. Many were probably legitimate and a few perhaps not, but the company’s legal department’s claims group had to deal with them all. Perhaps damage claims were not a safety issue but it was a public issue.

Furthermore, executive management wanted safety explicitly represented “in a business-like way” in any analysis for pole management: it simply would not do for any pole management program to not include safety as a priority. The project team analyzed data on all injury, damage, and other costs and claims related to pole failures received since 1980, from within the company, the public and other businesses. After correction for inflation and other factors, this resulted in an expected average cost of \$191 per pole failure for “injury and damage” claims.

Cost for Failure Replacement

CL&P had determined that the cost to replace the average failed pole at an average of \$3,262 (2010), with the range varying from \$1,900 to over \$10,800 depending on the size and type of pole and equipment on it, the characteristics of its location (on the street, off street) and the timing of the work required (normal business hours, weekend evenings, storm restoration, etc.). Well over half of pole replacement labor involved overtime. By contrast, records of NCR replacements, work that also involved removing an old pole and installing a new one, averaged only \$1,940 (2008). This difference in cost was attributed to the

Table 16.10 Materials and Labor Costs for Replacing Failed Poles (2010 dollars)

Type	Min	Avg	Max
Single-phase	\$1,900	\$2,400	\$5,600
Three-phase	\$2,400	\$2,800	\$9,100
Trunk	\$3,800	\$5,100	\$10,800

Table 16.11 Projected Average Pole Failure Replacement Cost

Year	Cost (2010 \$)
1990	\$2,660
2000	\$2,917
2010	\$3,262
2020	\$3,533
2030	\$3,624




Figure 16.17 From 2010 to 2030, when pole failure rate is expected to double, pole replacement costs are expected to increase by a factor of 2.35.

unexpected nature of most failures - the work could not be scheduled or deferred so much of it was done on overtime and without planning to fit it most efficiently into the field forces' schedule. In addition, failures did involve more slightly more labor than NCR (cleanup of wrecked components, etc.). In NCR replacements, any equipment on a pole - service transformers, cap banks, etc., is re-installed on the new pole. In failures, the old equipment is scrapped. A portion of the cost of the new equipment is attributed to and included in the cost of a "pole failure."¹⁰ Table 16.10 summarizes CL&P's table for pole failure replacement costs.

The Trend of Future Failure-Replacement Cost

CL&P replaced 1,235 failed poles in 2008 at a total materials and labor cost of \$3.9 million.¹¹ From the simplest perspective, since failure count was expected to grow by 4% per year, this replacement cost of about \$4 million per year could be expected to grow by 4% annually, basically doubling during the next 20 years. But there was the "URD effect" to take into account, and its effect was significant.

As discussed earlier, the shift to URD distribution starting in the late 1960s eliminated the vast majority of the lighter, less costly lateral poles from new distribution construction. From 1913 to 1967 lateral poles made up about 2/3 of all poles installed. By 1980 they constituted only a small minority of poles added each year. Thus, traditionally and into the 1980s about 2/3 of CL&P poles failing each year were lateral poles, and even into the final decade of the 20th century the URD portions of the CL&P system were still relatively new, with noticeably fewer failures than in older areas. Laterals constituted more than 60% of pole failures even if the proportion was starting to drop slightly from that of past decades.

By 2010, the oldest URD circuits are more than 40 years old - an age where failure rates start to rise—a growing portion of the system is composed of three phase failures. The long term trend is toward a gradually changing mix of lateral failures and heavier, three-phase pole failures, the ratio changing from about 2:1 in 1985 to about 1:1 by 2025.

As the mixture of small and heavier poles changes, the cost per pole will slowly rise, too (Table 16.11). This compounds the escalation caused by an annually increasing number of failures. The overall cost of pole replacement

¹⁰ This equipment replacement cost was initially determined using accounting methods, rules designed to assure compliance with generally accepted accounting practices. Several years into the program, CL&P did look at this aspect of its costs from an engineering standpoint, and made some slight adjustments that slightly altered policy with respect to planning replacement equipment but did not materially affect the overall results.

¹¹ The Asset Base Model's estimates were slightly lower - 1,188 and \$3.7 million. The model assumes average weather whereas 2008 had slightly more storms than an average year.

increases at an even a faster rate. Material and labor costs for pole failure replacements were predicted to go from an estimated \$4.2 million for 2010 to \$6.3 million in 2020 and \$8.7 million in 2030 (all figures in 2010 dollars). Failure count would increase by a factor of 2.20 in the twenty years from 2010 to 2030, but failure replacement cost (in constant dollars) would increase by a factor of 2.35.

Reliability: The Worrisome Trend

In 2008, overhead distribution line outages in which poles were a component of the failure accounted for 31 of CL&P’s 96 minutes of aiSAIDI, and a quarter of its reportable SAIDI (20 minutes out of 78 minutes), as determined using data from its Distribution Management System and Work Order Systems. CL&P’s fitted, predictive Asset Base Model estimated that figure at 27/20 minutes (Table 16.12) out of an average 93/75 minute SAIDI: that model assumed average weather in making predictions: CL&P experienced just slightly more storms than usual in 2008.

The fact that poles’ contribution are about one-third of the corporate total aiSAIDI and one quarter of reportable SAIDI makes distribution pole failures a “big deal” at CL&P, with an importance to the company far beyond the cost of replacement. Distribution wood pole failures contribute so much to SAIDI because they are associated with the vast majority of distribution outages and almost all widespread ones. A distribution pole’s failure is almost certain to interrupt service to at least a few customers. How many customers are interrupted, and the length of time that it takes to restore service, depend on the pole’s precise location in the system and the operating circumstances at the time of its failure. If the pole is located at the end of lateral its collapse might interrupt service to only two customers. But if the pole were the first on a main feeder trunk immediately out of the substation, its failure might cut power flow to 1,500 customers.

Table 16.12 Predictive Determination of aiSAIDI Impacts from Distribution Pole Failures on the CL&P System for 2008

1,188	poles fell
81	customers per pole on average
48%	restored by switching, automatic or otherwise
18.00	minutes average restoration time with switching
52%	restored only after repairs are made
422.00	minutes required on average to replace a failed pole
21,947,682	customer minutes (1,235x81x48%x18+52%x422)
824,670	customers in the CL&P system, average for the year
26.61	aiSAIDI contribution (21,947,682 /824,760)

The times required to both restore service through switching where that can be done, and to repair or replace a pole and restore service where switching cannot be done, will vary depending on a number of factors including the difficulty of work at the location, and how busy CL&P's field forces are at the time. If the pole failure occurred during a storm, there could be a backlog of poles that have to be repaired, and it could be many hours before a particular pole is repaired. On the other hand, if the pole failed on a clear-weather day, repair might be affected in only three hours. CL&P's average performance in this regard, over the five years prior to 2008, for repairing pole failures (*including* during storms) is 422 minutes, which was considered very good performance.

Restoration through switching, as opposed to restoration via repair, varies too. Some switches are automated and roll over automatically, others are remote control and switched within about five minutes. But other switches are manual and take up to thirty minutes to change in urban and suburban areas when crews are available. Times are somewhat longer in rural areas and may be much longer during a storm when a backlog of work might prevent field crews from immediately responding. Average is switching time 18 minutes. Still, many customers with service interrupted by a pole failure would have to await completion of those repairs before having their service restored.

URD: again, a trend within a trend, but more complicated this time

The same issue with before-URD versus after-URD that affected cost also affects reliability considerations, but it is both more dramatic and a bit more complicated to analyze. The primary effect is that for areas of the system built prior to 1968, between two-thirds and three-quarters of pole failures are usually lateral poles, which serve only a few customers and thus disturb service to only a few customers. Fewer than one-third are of major three-phase line poles, with their much more significant impact on customer service. However, for parts of the system built after 1975, the majority of poles and thus the majority of pole failures are of three-phase poles.

This is a difference in character of the reliability impact between these two areas, and of total amount of siSAIDI impact. There are more poles per customer in pre-URD areas. Customers there see, on average, more interruption time due to pole failures: almost 34 minutes SAIDI per year.

By contrast customers in URD areas see fewer interruptions due to poles, period: only about 20 minutes per year. Basically they see the same three-phase pole failure caused interruptions that customers in pre-URD areas see, but experience no failures due to the failure of lateral poles. Therefore, they see about 20 minutes of service interruption per year due to pole failures, or just a bit more than half what customers in pre-URD areas see. Their overall reliability isn't that much better, though: in 2008 they averaged 14.8 minutes of SAIDI due to lateral cable failures, for a total of nearly 35 minutes due to primary voltage failures, so their reliability is only very slightly better overall.

Dealing with Some Details of Impact Differences

When a lateral pole fails, there is most often little that CL&P's field forces can do to restore service prior to repairing the pole. There are limited options to restore service to some of the interrupted customers through switching, although for particularly lengthy repair periods field crews may isolate damaged secondary circuit elements and then switch over local service-voltage level circuits to restore service where possible before actually taking up the repair.¹²

But when a three-phase pole fails, usually a portion – on average about half – of the customers who had their interrupted service can have that service restored by switching of the feeder. Their interruption time will not be the full 422 minutes, but often less than half an hour. While this helps bring down aiSAIDI values, the failure of a three-phase pole affects on average, about 20 times as many customers as the average failure of a lateral pole (311 versus 16:). Half of the affected customers can be restored quickly, but that still leaves about ten times as many who have to wait until the repair is completed, as compared to with the lateral pole.¹³ Thus, the failure of the average three-phase pole still has a much more noticeable impact on reliability than the failure of the average lateral pole.

Thus, even when switching is taken into account, the gradually rising portion of three-phase poles that will be failing in the future means that the overall expected SAIDI impact of distribution pole failures will increase at a rate faster than the failure rate increase – and in fact slightly higher than cost will escalate. CL&P's historical experience shows this trend (Figure 16.18). The calibrated Asset Base Model projects that this trend will continue, with SAIDI impact of distribution pole-related failures increasing by a factor of 2.54 times from 2010 to 2030. It will grow at an annual compounded rate of 4.7%, or over four times as fast as the most optimistic projections of customer growth rate, adding just a bit to CL&P's SAIDI each year. Pole failure SAIDI contribution was projected to reach 45 minutes by 2020, 71 minutes by 2030 and 90 minutes – equal to CL&'s current average SAIDI - by 2038.

Putting a Cost on the Reliability Trend

The level of reliability degradation caused by these trends in pole failures was unacceptable to CL&P. Still, that did not imply that it had to reduce pole failure counts or their escalation rate. Instead, an option would be to just let pole failures happen, and “buy” reliability improvements to make up for the reliability degradation in other elements of this system. For example, CL&P's

¹² If the lateral pole is near the end of a lateral, sometimes the field crew can cut the damaged conductor free and restore service temporarily to customers upstream while they work the repair, but this was not, in 2008, a standard practice at CL&P.

¹³ Readers may be surprised that the SAIDI contribution of a single pole failure is so miniscule, but there are many of them and in aggregate they mount up. 1235 failures in 2008, at 1.3 seconds each, equal 28 minutes.




Figure 16.18 While annual distribution pole failures had increased nearly 50% in the last decade, the SAIDI impact of those failures had increased by over 63%.

Reliability Engineering group had determined that a Smart-Switching program (remote controlled vacuum breakers on feeders communicating to a central DMS system programming to minimize service interruption time) would reduce SAIDI by 8 minutes and aiSAIDI by 11 minutes, at a total initial cost of just under \$25 million. If this program were rolled out over a five year period, its 11 minute aiSAIDI improvement would almost completely counteract the 11 and 2/3 minutes erosion of aiSAIDI that escalating pole failures would create in that period.

The project team was fairly certain the company's executives would not want to take that route. But "How much would it cost to buy the reliability degradation back?" was a natural question in the circumstances, and the project team wanted to have answers before making its progress report. The Smart Switching option indicated the answer to that question was currently around \$5 million per year. Certainly, if the best other option cost *a lot* more than about \$5 million a year, CL&P executives probably would prefer this route. The project team set about refining its knowledge of "other options."

Pareto Curve of Reliability Augmentation Costs.

The "marginal cost of improvement" for the Smart-Switching program was \$2.75 per avoided annual customer minute without power (eleven minutes of improvement, for 825,000 customers, at a cost of \$25 million works out to \$2.75 per customer minute). That figure rises to \$3.46 present worth (PW) when future costs such as maintenance and so forth were taken into account on a time-value-of-money basis.




Figure 16.19 Pareto Curve of reliability improvement options versus costs as determined by CL&P's Reliability and Asset Management groups in 2007, for the 2008 timeframe. Theoretical Pareto curves have only a decreasing slope, but this real-world curve has four "notches" because the program ranking has been non-linearly optimized while taking project dependencies and interactions into account: the four places where slope decreases for a short span represent projects that do little to improve reliability themselves but "enable" other capabilities to be used. An example is the first such project "notch." The broad line shows the Smart-Switching program (11 minutes/\$25 million). The notch right before it includes various operating system and datacom upgrades to the company's DMS and control center system that would be required to make it compatible with Smart Switching. Of themselves, they do nothing to improve reliability and performance. They are simply a cost. However, they are necessary for any of the subsequent Smart-Switching program projects that do improve reliability. The good news depicted here is that CL&P has a host of small programs that are each more effective than that large Smart Switching program: it can buy up to 8 ½ minutes at a total cost of about \$12.5 million (equivalent to an average \$1.78 per avoided customer minute compared to \$2.75 for the Smart-Switching program). The bad news is that the diagram shows only a 32 minute potential improvement among all programs that have been identified. This compares to the 46 minute escalation (from 27 to 75 minutes) that CL&P anticipates over the next 20 years due to the increasing rate of pole failures.




Figure 16.20 Pareto Curve of reliability improvement versus costs for the year 2020, assuming nothing is done about pole failures. The Smart Switching program, among the more effective large programs, has been both “tuned” to maximize its impact and broken into tiers (broad lines) based on effectiveness in different areas of the system.

But the Smart-Switching program was not the most effective possible reliability program available. It was just the most effective *large* program. The Reliability Engineering group had studied a range of smaller projects and programs the company could take to improve reliability, ranked them in order of effectiveness, best to least. The best small project had a cost of only \$1.50 per avoided annual customer minute without power. The worst had a cost of \$10 per avoided annual customer minute without power. Figure 16.19 shows a Pareto Curve of reliability improvement that CL&P’s reliability and asset management groups had developed in 2009. Candidate projects and programs are each a small straight line in this diagram, with slope equal to their marginal cost of reliability improvement (the ratio of their SAIDI improvement over cost). In Figure 16.19 they are plotted end to end, from left to right, the highest improvement/dollar plotted first, then proceeding in order of “bang for the buck” (improvement over cost) through to the least effective program on the far right. The Smart-Switching program is shown as a single broad line, stretching 11 minutes vertically and \$25 million dollars from left to right, and placed in order of its “merit” with respect to the other projects. Smart-Switching was not the most cost-effective program available to CL&P, although it represented a fair

amount of the more cost effective potential. Everything to the left and below it in the diagram was small projects and programs, none over \$5 million, many less than \$250,000 each.

Adjustments to Reliability Analysis

The project team worked with CL&P's Reliability Engineering group to refine and add detail to the analysis of marginal cost of reliability and the study results that lead to Figure 16.20, in two ways.

As a function of pole-failure-caused aiSAIDI. If the number of poles failures each year doubled, there would be roughly twice as many opportunities for a Smart-Switching system to isolate outaged circuit sections and restore service automatically; the effectiveness of that system increases roughly in proportion to problems caused by failing poles.. By 2030, the Smart-Switching program would reduce SAIDI by 19 minutes, rather than 11, and its PW cost would have risen to only \$3.45 (2010 dollars).¹⁴ Similarly, some of the other reliability improvement programs CL&P was considering would also be more effective in a time where pole failures were much more prevalent (guying/reinforcing key poles), but others would not (cable replacement).

Targeting of Smart-Switching. Just as the Smart-Switching program would be more effective on a system with a higher SAIDI, those parts of it installed in areas of the system with a higher pole failure rate would be a bit more effective than those installed where pole failures were less likely. The Reliability Engineering group had intended to propose rolling Smart-Switching out in "problem areas" first. It proposed breaking the program into "tiers" aimed at different areas of the system, starting with those areas that had the highest level of feeder-related (pole and cable mostly) failures. This led to five tiers of Smart Switching: most effective, next most effective, etc.

The overall estimated cost of the Smart Switching program increased slightly when implemented in this tuned and tiered manner (to \$25.8 million). The most effective tier had a marginal cost of \$2.00 PW: it targeted older areas of the system that also had high concentrations of trees near overhead lines. The least effective had a marginal cost of \$5.66 PW, and targeted newer areas of the system.

Figure 16.20 shows the Pareto curve the project team estimated for year 2020, for a scenario where nothing proactive has been done about pole

¹⁴ As the number of operations increases maintenance costs do go up slightly, a fact that CL&P actually estimated and took into account in its study. In addition, there is a risk due to a small but noticeable and very ironic factor: the smart switches are installed on poles, and thus susceptible to damage when the pole they are attached to fails.

Table 16.13 CL&P’s Annual Costs for Pole Failures

Cost Category	2010	2020	2030
Safety and Damage	\$243,500	\$340,500	\$460,500
Cost of Materials and Labor	\$4,162,500	\$6,299,000	\$8,737,500
Reliability	\$5,197,500	\$9,457,000	\$14,550,000
Total	\$9,603,500	\$16,096,500	\$23,748,000

failures. Total possible SAIDI improvement is now 43 minutes. The Smart-Switching program is represented by five shorter segments now; its various tiers. The first three appear to be fairly effective, the final one much less so.

Overall Trends of Pole Failure Cost

Table 16.13 provides a total of the annual costs for pole failures the project team had at this point. Over time, because reliability problems are becoming more numerous, the various means that can be employed to “buy back” reliability become even more effective. But as CL&P implements these various programs, most-effective first as prioritized by its Pareto analysis, it “uses up” the most effective programs first. Thus, the “bang for the buck” of programs remaining to be selected drops over time. The most effective are used first, and so, slowly, inevitably, the cost of buying another minute of SAIDI improvement, to counter further pole aging, rises over time. Coupled with all the other aspects of reliability impact, this increased cost by a factor of 2.8.

Restricting analysis to only the existing pole base and the existing customer count, cost of pole-related failures would rise from \$11.60 to \$28.80 per customer over the next 20 years. CL&P’s accountants estimated it spent about \$380 per year per customer on power distribution. *Pole failures, then, were a very big deal on a corporate scale*, because CL&P was spending about 3% – 7.6% of its gross revenues from delivery of power on the effects of pole failures.

16.9 FINDING THE BAD POLES IN THE ASSET BASE

The Asset Base Model and database described above used a statistical approach applied to the entire CL&P database. Trends that it produced were, within the error limits of the statistical analysis, dependable forecasts. As a result, it provided a useful platform for considering how serious the trends were as well as a starting place for planning what and how to manage the pole asset base as it continued remain in service and its units continue to age.

But the Asset Base Model and its database were not sufficient for execution of any plan once made, because they lacked specificity as to location. The planning database could estimate that in 2010 CL&P had a total of 19,240 poles that were over 65 years old. But if CL&P wanted to inspect, test, or service/replace those poles, it would need something else, too, because that

model and its database would be no help in *finding* those older poles. Common sense would tell a person to look in older areas of the system, and the tacit knowledge of experienced operators in CL&P's field forces could help it find areas where old poles were most likely to be found, but CL&P had no organized, institutionalized knowledge base about where to look for poles of specific ages.

Actual Age Data Since 1989

CL&P actually had accurate and fairly complete data on poles installed since 1989. In the twenty years from 1989 to 2008, it had installed for the first time at a location 54,903 of its 437,226 poles (12.5%). Therefore, in 2008, it knew precisely both the age since first installation and the age of the pole for 12.5% of its poles. And CL&P knew the actual age, but not the age of original installation, for another 21,157 (4.6%) of its poles: those it had replaced during that time, for either NCR or failure reasons. Thus, it knew the age of 16.6% of its poles in detail: for a specific pole, when it was put there. But it did not have that information, or even a hint of it, for the other 84%.

This situation would improve every year, as CL&P kept its comprehensive records and continued to add and replace poles. By 2028 it would have exact data on poles up to 40 years old, etc. But it would be over three decades (2040), according to the Asset Base Model's prediction before it had accurate data on the majority of poles in its system. So, in the very long run, the situation would take care of itself. But until then, CL&P would have to either make do without age data for the vast majority of its pole asset base, or find a way to find to estimate those ages. The project team did not think it could wait that long.

Estimating Age for the 82% of Poles More Than 20 Years Old

CL&P, like almost all North American utilities, was extremely lucky in one respect; a vast majority of neighborhoods in its metro area were "first growth" development. In most areas, the homes, buildings, streets and infrastructure there now were for the most part the first built. CL&P could assume that its electric infrastructure originated from the time those buildings were built: if the oldest homes, stores, and office buildings in a particular neighborhood were constructed between 1954 to 1964, then it was very likely that the electric distribution system was installed around 1954, since basic utility services are installed somewhat "up front" of demand growth but never too early.

That is the basic concept of using geographic neighborhood age data. In CL&P's case, that concept was refined and fine-tuned using statistical analysis backed up by some practical investigation and verification. Figure 16.21 illustrates this concept. Finding and gathering data on the origin dates on a neighborhood basis was not difficult. The city government had road maintenance records going back for decades on a neighborhood by neighborhood basis. The date that a neighborhood first showed on road department records would be a good estimate of when it was first occupied. Tax and property records gave a date of first occupancy addresses. The City Library




Figure 16.21 The concept of developing a geographic map of development in a city is illustrated here with data on a 2 x 2 mile grid. Historical records on building permits issued, road construction, school enrollment, and other clues to past city growth have been used to determine the average “origin date” of each cell in the map providing a rough time for when the infrastructure for each area was first put in place.. The actual CL&P data had much higher spatial resolution (i.e., used a much smaller area size), generally down to neighborhoods of only one or several blocks each. This data provided CL&P an estimated age of origin for the distribution system in every area of the system, since it would have been built in company with the original development in each area.

System even had a long-standing record system mapping readership trends on a neighborhood by neighborhood throughout metro area, which went back to 1900. School enrollment records were a good source of when new homes first appeared in areas, etc. There were many other sources of information buried in various municipal, business, and even some old CL&P record sets.

The city government did not cover all the CL&P service territory, although it represented the area with over 90% of its customers. There were several other smaller municipalities and a small portion of customers in unincorporated areas. Data from those other municipalities, and the county governments, was available, although often in a different data format, etc.

Complexities in Aggregation and Consistency

It was soon obvious that this neighborhood age approach would be very useful to CL&P’s pole aging study (and the study of other equipment ages, too), but that assembling the available data accurately and consistently into one dataset was going to be a challenge. The work was not only intricate, but required a good specialized knowledge of municipal practices as well as an ability to manipulate




Figure 16.22. Origin Data Project used various sources of municipal and public data to estimate the year that the neighborhood along each block face (side of a street) was built out. This information was then made a locational attribute in CL&P’s GIS system.

data in a GIS system. CL&P funded a project led by two professors from a local university, one from its school of Architecture and Urban Planning and one from its Geography (GIS) department. With the aid of several graduate students they assembled records, resolved inconsistencies and data conflicts, and addressed the thornier issue of estimating the age of a sizeable number of “second growth” and even “third growth” neighborhoods – those that had been redeveloped somewhere along the line, etc. The university team applied spatial statistical analysis as well as “urban planning” concepts to help refine the neighborhood age data, resulting in a picture of the metro-area’s growth from 1913 onward. The resulting geographic database, shared with the various governments that contributed data, assigned an “original age” to each block face (side of the street) in the metro area (Figure 16.22).

Summary of GIS Data Use

CL&P used that neighborhood age database as an overlay in its own GIS system. It then determined an “original age” for each pole in the system by assigning that pole the “original age” of the block face nearest to it. This was used as the estimated age of *initial installation* for all poles in that block face area.

Estimating Pole Age from Age Since First Installation

The time since initial installation at a site is not the best possible estimate of the age of a pole that is still there. For example, the original poles in an area where poles were first put into service in 1950 each have a net expectation of about 15% of having been replaced for NCR or failure reasons in the 58 years between 1950 and 2009. There is an 85% probability each is 58 years old, but a slight probability each is only 1 year old (i.e., it was replaced last year). However, CL&P knows of all poles replaced anytime in the previous year, meaning that if this pole is one of those it does not have data for, then it was installed somewhere between 1950, when the neighborhood is estimated to have first been built out, and 1989, when accurate records were first maintained.

Statistical analysis performed by the project team determined that in such a case, the best “guess” at a pole’s age at this location – the value that had a 50/50 chance of being over or under – was 53 years. That is the value CL&P used as “age” in its database for that pole. In that manner it determined an estimated age for every pole in the system that had not been installed prior to 1989.

The trend of expected age versus time since original installation is interesting (Figure 16.23), with a peak around 75 years since installation. When a pole had been installed at a location for 75 years, the probability of that pole being the original one dropped to less than 75% and the best estimated age was around 62 years. At 100 years since original installation, the probability of it being original is only 42%, and its estimated age has *dropped* to 55 because it is much more likely to have been replaced somewhere along the line than to be original.




Figure 16.23 Estimated age of the pole at a site as a function of age since a pole was first installed at that location. See text for details.




Figure 16.24 The distribution of ages of the 437,226 poles CL&P owned as of January 1, 2008, as determined by its statistical Asset Base Model described here (line) and as estimated by the GIS Neighborhood Origins project (grey shaded area). The bulk of poles in the CL&P system were 30 to 50 years old: installed in the 1960s through the 1970s.

Using the GIS Data

CL&P used that neighborhood age database as an overlay in its own GIS system. It then determined an “original age” for each pole in the system by assigning that pole the “original age” of the block face nearest to it. This was used as the estimated age of initial installation for all poles in that block face area.

Using the particular failure rate curve found to best fit the CL&P system – the curve fitted by the Asset Base Model and its database – an existing pole’s expected age peaks at 62 years at around 80 years since the original pole was put there. Where a pole has been there even longer, the expected age will slowly drop and asymptotically approach the sustainable age, 51 years.

CL&P’s project team was aware that these age estimates were likely to have a considerable amount of error in them. However, they were the best the team could do. Standard deviation (uncertainty in) of age estimates for poles, and how that was determined will be discussed in a few pages.

Figure 16.24 compares the age distributions determined independently with the Asset Base Model and that determined solely from this GIS-based block-face analysis of pole ages in the CL&P system. That for the Asset Base Model is deterministic while that from the GIS sources shows a range based on statistical uncertainties in the estimators. The broad agreement of the two distributions gave the project team some confidence that they were correct.

Unambiguous Definitions of Three Terms

To keep all its analysis and planning consistent, it is necessary to define three terms related to the age of a pole very precisely and unambiguously. The project team then always used these in order to assure no ambiguity in communication about pole management:

Time since first installation is the number of years that have passed since a pole was initially installed at this site for this purpose, as determined from the GIS database’s neighborhood block face overlay. This is an estimate of how long a power distribution pole, whether original or the one there now, has been in place at that site.

Age is the number of years the present pole at the site has been in service. Not all poles in the CL&P system are the original poles installed at the site. For poles installed since 1989 CL&P had this exactly. For others it was an estimate based on the statistical method discussed above.

Condition, measured in years, is the amount of deterioration of a pole and is the value CL&P most wants from the standpoint of managing aging effects on its system and performance. Condition of 58 years, for example, is defined as the average amount of deterioration among all poles that were 58 years old.

Condition will vary among poles of any particular age. Some 58 year old poles might be in good condition for their age, equivalent to the average 40 year old pole. Others might have suffered more than the average amount of wear and tear, or just have been more poorly treated or made from below-average materials when new, and have deteriorated to a condition equivalent to 75 years of service. But regardless, while age might be the data item that was easiest to get about a pole, condition was the real item that CL&P wanted.

Thus, CL&P carried all three of these data values forward in its pole database, for every pole: *estimated time since first installation*, *estimated age*, and *estimated condition*. At this point in the narrative, estimated condition was the same as estimated age. Regardless, it was given a separate field in the database in hopes that in the future there might be ways to improve condition estimates so they were better than just estimated age.

Physical Investigation of Pole Condition

The Asset Base Model and the GIS-based block-face database and analysis had led to an “estimated condition” value for each pole in CL&P’s pole database – an estimated age for each pole in its system that was at that point the best estimate of its condition.

There was no doubt in the project team’s mind that average condition, as measured by deterioration and loss of strength, etc., would worsen with increasing age. But the project team also expected that variation about that average would increase, too, as poles got older. That variation in condition-as-a-function-of-estimated-age would shape the effectiveness of any program. If it was narrow, inspection, treatment, and replacement based on estimated age would be effective. But if the variation was wide, it meant CL&P had little clue as to what the condition of a pole was, even if it had a good estimate of age.

To gather information on both average levels and how much poles might be expected vary from the average, early in the project CL&P’s team had initiated a Physical Pole Examination Program (PPEP). Poles were taken for examination from among those that had been removed from service for failure, retirement or NCR reasons. From mid 2007 to 2010 the project team examined slightly more than 1,500 out of roughly 3,800 poles that had been replaced or retired: it examined 32 poles with an estimated age of 30 years, 44 poles that its database estimates had a 50 years of age, and 12 of age 75, etc. In addition, the PPEP sought out, removed, and replaced, then examined, poles thought to be among the oldest in the system: 24 poles that its databases indicated might be up to 80 years or older (very likely these poles needed to be replaced soon at that age). From this physical examination, CL&P determined:

- a. Average deterioration and loss of mechanical strength as a function of estimated age: i.e., what it meant to be a pole of condition = 58. Condition = 58 was defined as the average deterioration of a pole that had been in service 58 years, but what did that mean? This

examination provided that data, examining 27 poles estimated to be 58 years old.

- b. The standard deviation of deterioration and loss of mechanical strength about that mean for every age. How much did deterioration of these thought-to-be 58-year old poles vary?

From a and b above, the standard deviation of condition-measured-in years could be determined as a function of estimated age (measured in years), as could the shape of the distribution of that condition measure about its mean for any age. For example, suppose the average structural strength of poles of age 50 was shown to be 78% or original, and to have value on the low (weaker) side of 46%. Suppose further that 46% was the average deterioration of a 73-year old pole. Then the high-side standard deviation of condition at age 50 when measured in years was equal to 23 years (73-50). This was what was wanted.

Figure 16.25 gives an indication of the standard deviations determined by this analysis and the type of non-symmetrical distribution about the mean seen in the data. Table 16.14 shows standard deviation as a function of age as determined by the analysis. Overall, these results were not surprising to the project team. The standard deviation was quite reasonable up to about 20 years into the past, i.e., during the period for which the company had hard data on pole installation dates. But it increased from there on backwards into history. Between 40 and 70 years, standard deviation doubled, probably due to the compounded effects of the natural variance in deterioration (some poles would deteriorate faster than others) and uncertainty in the data.

Table 16.14 Standard Deviations of Estimated Condition as a Function of Estimated Age (Years)

Estimated Age (Years)	Standard Deviation (Years)	
	Low	High
10	1.6	2.8
20	3	5
30	6	14
40	10	19
50	16	24
60	19	30
70	26	38
80	33	42




Figure 16.25 Average condition (middle line) was a linear function of estimated age because it was *defined* as deterioration equal to the average for that estimated age. Standard deviations high and low varied as shown. Note the increase in standard deviation in both trends between 20 and 25 years. This is the effect of moving from exact knowledge of pole age (up to and including 20 years) to just estimated age (for all poles older than 20 years).

16.10 POLE REPLACEMENT MERIT AND ECONOMICS

This section will present an analysis of the benefits and costs of proactive replacement of poles while neglecting completely the fact that the condition estimates based on age were known to have a good deal of variation (error) that creates a rather wide distribution (large standard deviation) about the mean condition estimate for poles. Essentially, this means working with data that is assumed to be perfect – to have no error (neglecting that the error distribution is equivalent to this). Therefore, the results obtained here will not be of direct practical use as a guideline for replacement, but there are several sound reasons to start the investigation of potential replacement strategies by analyzing in this “perfect data” (no error) manner, as CL&P did:

- 1) The analysis, which is purely deterministic, not probabilistic as required to analyze the data with its error distributions, is much easier to do and particularly to show in a textbook.
- 2) A lot can be learned about the *relative* effects of various factors from this analysis. While the actual numbers – replacement ages, benefit/cost ratios, etc., obtained here will be a bit optimistic, the interactions of various factors and the way the results respond to various actions or changes will be valid in all cases.
- 3) The results provide a useful upper limit on how effective a pole replacement program could be: perfect data gives the best results possible.
- 4) Comparison of any real-world case (with imperfect data) to this “with perfect data” case will determine how much the imperfect data is decreasing the results, thereby putting a value on the additional information. This is a first step toward addressing one of the more important conclusions covered in this chapter: how much should CL&P spend to improve its data, as opposed to actually replacing poles?

An Asset Management Planning Framework

CL&P uses an asset management approach to decide which capital projects to fund and when. Asset management seeks to fund the best set of projects each year to achieve the company’s business objectives, which are defined by targets related to the major goals the company has: meet its obligation to serve all customers in its service territory, comply with all other commitments it has made with governments, maintain reliability (SAIFI, aiSAIDI) within limits, operate and maintain its equipment so it has a maximum service lifetime,¹⁵ keep long-term operating costs low, keep up safety (lost-time injuries within OSHA limits, no public hazards, etc.), and keep within certain financial limits such as

¹⁵ Maximum service lifetime from the standpoint of overall lifetime economics.

debt/equity ratio, etc. It evaluates all candidate projects – projects it might approve – on the basis of how they contribute to these goals: does it provide enough benefits in return for its costs.

This evaluation goes beyond a simple benefit/cost analysis, although that is one aspect of CL&P's planning. For example a new substation project allows the utility to extend service connections to new customers (achieving the obligation to serve), relieves loading concerns for nearby substations (improving expected lifetimes), improves reliability (providing new feed points for load transfers in that part of the system) – all good factors. But the new substation has costs, too. It increases the likelihood of lost time accidents (every hour of labor is an hour of accident risk) and costs money now (for equipment and labor) and in the future (for taxes, maintenance, etc., on the facility and its equipment over the next thirty years) and in operating costs (electrical losses).

CL&P evaluates every project it considers on the basis of how it affects each of these factors. It selects projects each year as a set: "Will this portfolio (group of projects) provide enough new system capability to allow us to meet our obligation to serve, result in loadings and lifetimes within guidelines, contribute enough reliability improvement for us to meet our targets there, keep long-term operating costs where they should be, at the lowest cost and with acceptable levels of risk for the company?"

To see how pole replacement programs would fit into this framework, it will be sufficient to evaluate three aspects of their "performance." First, the capital cost – what will it cost CL&P to do a replacement? Second, is the cost justifiable based on long-term economics? Does replacement drive down future costs enough to justify spending money today? And third, and quite important – what is the impact on reliability (aiSAIDI?). Other factors are largely irrelevant here and will not be discussed.

Table 16.5 presents information on the cost of the average pole failure in the CL&P system for year 2010.¹⁶ "Average pole" as used here means that the data values are the statistical mean in every case. The average pole that failed in 2008 was 58 years old. On average the service to 51 customers was affected by a pole failure. On average, 37% of CL&P customers with service interrupted by pole failures were restored by switching, and so forth.

The majority of the "cost" shown for this failure is due to the cost assigned to the impact the failure has on customer service reliability. Using the 37¢ per avoided customer minute levelized cost (CL&P's marginal cost of long-term reliability as determined for 2008), the 14,325 customer minutes of outage this average pole failure causes will cost \$483. The total cost including the safety and damage value as well as the cost for materials and labor for unscheduled

¹⁶ Unlike analyses in previous pages, which were based on 2008, those here are calculated for 2010, as they were done in late 2008 by the CL&P project team. It anticipated a one-year start-up period for any pole management program it would plan, and thus targeted its analysis at programs beginning in 2010.

Table 16.15 Evaluation of Proactively Replacing the Average Pole that Failed in the CL&P System

Description	
1	pole
58	years condition (equivalent age)
51	customers affected if it fails
36%	restored by switching, automatic or otherwise
30	minutes average restoration time with switching
64%	restored only after repairs are made
422	minutes required on average to “repair” this failed pole
Consequences of failure	
14,325	customer minutes of outage
0.017	equal to this many minutes of system aiSAIDI
\$191	safety and damage assessment
\$3,262	materials and labor to replace on an unscheduled basis
Probability of failure	
0.47%	in the next year, including storms
Costs that could be avoided	
62	customer minutes interruption created by its outage
\$23	value of these avoided customer minutes at .37¢
30-year value of benefits and avoided costs accounting for both time-value of money and escalation of failure rates in the future	
\$483	PW of future avoided aiSAIDI
\$19	PW of future avoided damage and safety
\$297	PW of future avoided unscheduled replacement cost
\$798	Total PW future value
Computation of merit for project	
\$1,940	cost to proactively replace
0.41	benefit cost ratio to replace now
\$31.24	per customer interruption minute avoided

and labor for unscheduled replacement comes to \$798. This is only 41% of the cost of replacing a pole on a proactive (planned and scheduled during normal business hours) basis. Thus, CL&P cannot justify replacing this pole proactively (even assuming, as is done here, it has perfect knowledge of its condition). However, this example highlights two important points.

- 1) *Relatively high present value of future costs.* When CL&P evaluates costs and benefits, a capital project that is going to save \$1 per year every year into the future would be evaluated as having a PW benefit of \$14.29: the multiplier for NPV to annual savings is 14.29.
But in Table 16.14, the value of next year’s expected aiSAIDI reduction (62 customer-minutes) is shown as \$23, but the NPV of future value as \$483, or 21.0 times, not 14.3 times, current annual value. Pole replacement provides a *growing* benefit each year in the future. The old pole was deteriorated, but also deteriorating faster than a new pole would. The expected reduction in aiSAIDI is growing each year farther out. This makes the difference.
- 2) *Relatively short term reliability “bang” for the bucks.* The evaluated cost of avoided customer minutes is \$31.24, which compares to about \$3.34 for Smart Switching. Proactive pole replacement is not nearly as good at quickly driving down reliability problems as other means can be.




Figure 16.26 Benefit-cost ratio for proactive replacement of the average pole that fails: 58 years old and affecting the service to 51 customers when it fails. The determination shown is based on the “perfect data assumption” (see text for details).

Figure 16.26 shows the computed B/C ratio as a function of age, for poles whose failures affect service to 51 customers service (the average for poles failing in 2008 as estimated by the Asset Base model's output. Breakeven age, when B/C first equals zero, is about 77 years (again, given perfect data). Benefit/cost ratio improves to 1.6 by age 90. CL&P will have to wait until a pole of this type is 77 years old before this assessment would say that the B/C ratio is positive as compared to its current costs. Alternatively, if this pole affected service to 175 customers, the reliability improvement would have 175/51 times as much value and it would evaluate was positive ($B/C = 1.02$).

Sensitivity to Age and Customers Affected

Suppose the pole selected to be replaced was again 58 years old, but that it affected the service to only 16 customers – the average number of customers whose service is interrupted for single-phase lateral poles. Fewer customers (16) would see their service interrupted when this pole failed than for the average failing pole (51) but very likely none could be switched back into service prior to replacement of the pole: all would have to wait for the repair to be done to see their service restored. Thus, while fewer than one-third as many customers would be affected, the outage would cause nearly half as many customer minutes of interruption ($16 \times 422 \text{ minutes} = 6,752 = 47\%$ of 14,325). Avoided labor and materials cost for lateral poles is \$2,800, while proactive replacement cost is estimated at the same \$1,940.¹⁷ Benefit/cost ratio at age 58 is only a bit better than one quarter. Breakeven (again assumed perfect data, as used here) is at 88 years. Poles that affect the customer service of fewer customers have much higher breakeven replacement ages.

But if the pole affected service to more customers, the situation would be reversed. If this pole affected service to 180 customers it would have a positive B/C ratio (1.03) at age 58. And if it were at the extreme opposite end of the reliability spectrum, for example, being a feeder riser pole right out of the substation whose failure might interrupt service to 1,500 customers, the replacement age calculated by this approach would be far lower – only 45 years.¹⁸

A Factor Completely Forgotten Until Now: Prioritization of Attention

The computation that the replacement age for a “critical” pole as 45 years struck the CL&P project team as far too low, even allowing for the fact that this determination was using “perfect data.” A small flaw in the preceding analysis

¹⁷ Almost any lateral pole will be of a lesser class and length than any three-phase pole, so the material cost would be lower. But most laterals in the CL&P system run along the rear property line where they cannot be reached with a bucket truck and other heavy equipment. Thus labor is more. A cost of \$1,940 is a reasonable estimate.

¹⁸ This determination took into account that the materials and labor costs for such a pole would be greater than for the average pole, and that very likely about 50%, rather and 37% of the customers for a riser pole could be switched in order to restore service.

is an assumption that the expected repair time for all poles, regardless of criticality, would be the same 422-minute average. Even in the aftermath of a severe storm, the repair of a pole serving 1,500 customers – one that that was so critical to system integrity and switching -- would rise to the very top of the priority list.

The project team had often recognized that the average 422-minute replacement time would vary, but somehow had never really appreciated how important this factor would be to its results. Now it did. Therefore, data from recent pole replacement operations were analyzed to make these determinations by category of pole, leading to a set of estimated average replacement times by pole type. Replacement time for a critical pole, one that could lead to interruption of an entire feeder, fell to only 100 minutes, while those for poles with far lower reliability impacts rose as a consequence. Table 16.16 gives the replacement times the team developed: these are not values determined from only initial historical analysis. The project team did further studies and worked with the company's Reliability Group to determine what priority CL&P *ought* to put on poles in each category in order to minimize customer interruption time (aiSAIDI) and adjusted its results to reflect that. (These times reflect all replacements including during storms, whether "excludable events" or not.)

The revised values changed the estimated avoided cost for replacing a pole of each type, although they did not change that for the average pole, since the average time was still 422 minutes. Because the expected interruption time for customers is now far less for that critical riser pole, the customer aiSAIDI reduction to be derived from early replacement will be far lower, too, and the computed benefit will be in line with reality: breakeven is now at age 51. This might seem low to many people, even given the impacts of the "perfect data" assumption, but it is fully justifiable based on cost due to the substantial impact such a pole's failure would have on customer reliability.¹⁹

The major lessons from this quick analysis of pole-replacement economics using deterministic data is:

- a) Long-term benefit/cost usually looks more appealing than short-term.
- b) The reliability impact of pole failure is the dominating factor in determining replacement value.

¹⁹ More than just replacement at age 42 is justifiable for critical poles. This step in the project coincided with a study that CL&P's Reliability group had independently started a year earlier, to look at upgrading key poles and facilities in the system in a "Hardening" program, along the lines of work pioneered by Dr. Richard E. Brown. That study determined that much heavier poles were justified for key locations in the distribution system. Although steel was considered, the optimum turned out to be a much heavier wooden pole. Interestingly, the new pole specification both raised the cost for key trunk poles, and lowered the estimated failure rate, so that the replacement age predicted by the model for these key trunk poles fell back in line with the earlier, at age 56, identical to that for the average pole.

Impact of Condition Estimate Error on Benefit/Cost of Replacement

To develop an intuitive feel for the consequences of error in condition and what it would do to CL&P's spending on aging effects mitigation, consider three poles, each 58 years old and each serving 180 customers. As discussed earlier, pro-active replacement of these poles evaluates as having a B/C ratio of 1.0119 if the pole has a condition actually equal to 58 years old: total replacement cost is \$1,940 and expected PW of future benefits and savings is \$1,963. But here, while all are thought to be 58 years old and assumed to have condition = equivalent to age 58, only one is actually that condition and the other two are equivalent condition to 44 and 72 years respectively but CL&P does not know that. It replaces all three:

- Three poles are replaced at a total cost of \$1,940 each.
- The pole that is condition = 58 years will see benefits of \$1,963 as expected.
- The pole that has condition = 72 has two elements of change from that expected situation that need to be considered.
 - a) First, at age 72, the future risk of this pole failing in the future is nearly twice as high as is thought based on that (incorrect) condition estimate of 58 years. CL&P will actually see a (time discounted) benefit with a PW of \$3,909.
 - b) This pole should have been replaced about 15 years earlier – one has to say “about” because it might not have had condition equal to 58 years then: no one knows how it became so deteriorated now – it might have been average for its age then and for some reason seen very accelerated deterioration in only the last several years. But assuming the pole was equivalent to 58 year's condition fifteen years ago and that deterioration during the period was linear, gives an (time discounted) estimated PW to the risk that CL&P incurred of \$3,382.²⁰

The benefit of replacing this pole is $\$3,910 - \$3,382 = \$528$

²⁰ This value confuses some people. First, since this pole exists, now (or it could not be part of the three being considered for replacement), it is clear that CL&P gambled and won: the pole should have been replaced 15 years ago but due to data error was not. However, over a large population of poles, one can calculate the risk – how likely was it to fail and what was that worth in today's dollars. Thus, one uses the probability of failure in each of the last 15 years to determine the monetary risk (what would it have cost CL&P had it failed in that year), inversely discounted because a dollar 15 years ago would be worth more than a dollar today.

- Replacing a pole that actually has condition = 44 years means that the actual future risk reduction was about 35% PW less than the \$1,963 expected, or \$1,240.

Thus, CL&P replaces three poles at a total cost of \$5,820 and receives benefits of \$1,963, \$528, and \$1,240 respectively, or a total of \$3,731. Benefit/Cost ratio is 0.64, not the 1.01 expected.

By contrast, the impact of these errors on aiSAIDI reduction the following year is negligible. CL&P expects 218 customer-minutes of reduction per pole (180/51 times the customer minute reduction shown in earlier tables), for a total of 657 customer minutes reduction in the contribution to its aiSAIDI in the next year. It actually receives reductions of 435, 218, and 110 from the poles that are respectively actually more, correctly, and less deteriorated than it thinks, for a total actual benefit of 655. If the distribution of condition error estimates is symmetrical about the mean, the error does not really matter that much as regarding short term reliability gains. The reason for the difference in sensitivity between long-term and short-term sensitivity is due to the effect of the time-value of money.

Testing as a Way to Reduce Uncertainty of Condition Estimates

A range of wood pole testing methods exist, having different costs and limitations and requirements, and providing different types and levels of detail and improved accuracy in the condition estimates they produce. None is precise: all should be viewed as methods that estimate the condition of a wood pole independent of any data or models developed from historical data. In practice, a “second opinion,” even if an only another estimate with some range of error itself, could reduce the uncertainty about condition (the standard deviation of the net condition estimates).

Table 16.16 summarizes test cost and effectiveness data that CL&P used in planning how it would utilize testing, for three representative testing approaches. Cost per pole for testing depends on the type of test done and the number of poles tested (there is a noticeable economy of scale). The reduction in error (standard deviation of the condition estimates) is up to 50% if the most effective testing is used. These rates were determined by “testing the testing” on the poles examined in the Physical Pole Examination Program (to be discussed later in this chapter). CL&P’s tests confirmed what many other utilities had determined, for example, about visual inspection and the “thump” or hammer-test: alone they made no improvement in knowledge of condition.


Table 16.16 shows that even with the best tests available, CL&P could cut the standard deviation of its condition estimates for poles only by one half. No testing method is perfect. All had some range of uncertainty. And this range of uncertainty tended to be wider for older poles, and to correlate well with the condition estimates CL&P had previously prepared. To some extent testing only confirmed CL&P’s data: poles thought to be old generally were. However, the statistics-based condition estimates were a bit more specific, and so overall led

Table 16.16 Revised Practical Replacement Times for Poles of Various Types (Using Imperfect Data for Determination)

Location and Type Of Customers	Estimated Replacement Time	Computed Replacement Age
Lateral – among the least impacts	580	125+
Average lateral pole impact	480	98
Average pole overall	422	86
Trunk – among the least 3 ph impacts	340	71
Trunk - typical impact of 3 ph pole	200	58
Trunk - among the highest impacts	100	52

Table 16.17 Percentage Reduction in Standard Deviation of Condition Estimates for 60-Year-Old Poles Used by CL&P in Planning its Programs

Testing Method	Cost per Pole			Std Dev Reduc. %
	Entire System	Lots of 25,000	Lots of 1,000	
Visual, Sound, Bore	\$32	\$40	\$47	50%
Visual, Sound	\$20	\$27	\$34	27%
Visual/Thump Test	\$5	\$10	\$15	0

**Figure 16.27** Final determination of replacement age benefit cost ratio for poles in the CL&P system, acknowledging the imperfections in the data.

to an improvement. The improvement shown is the *net* when both were combined in any way possible. Each contributed about equally to the quality (or lack of it) in the final condition estimates the company used.

These results had significance to CL&P. Based on them, the company's executive management authorized testing of 15,000 poles in 2009-2010 – the first spending produced by this aging infrastructure pole program. The reason 15,000 were tested, and which poles were targeted, will be discussed later.

Practical Replacement Ages Based on Real-World Data

Even with the improvement of condition knowledge due to pole testing, CL&P's condition estimates would still have a considerable amount of uncertainty. The pole thought to have a condition equal to age = 58 now had a standard deviation of 14 years. One out of six rated as that condition would be in worse condition than the average 72 year old pole. Another two thirds would have conditions between 44 and 72 years, with a tendency to cluster more around 58 than at the ends of that range, and a final 1/6 would be in condition equivalent or better than 44 years of age.

As a result, any replacement program CL&P implements based on that data is likely to make some mistakes, just not nearly as many as before testing was being used. A few poles will be replaced earlier than justified, throwing away useful lifetime. Other poles will be left in service when they should be replaced. Cost will be higher than it would be with perfect data, and failure reduction will not be as great.. Put in economic terms, while all the expected costs of a proactive pole replacement program would be there, all the benefit would not: benefit/cost would be lower than expected.

CL&P's project team performed a more comprehensive and detailed statistical analysis of the impacts of the remaining (i.e., after testing results) error in the company's pole database. This used a probabilistic approach that looked at replacement ages by category of pole, that took non-linearities in deterioration rates into account, that acknowledged the non-symmetrical nature of the condition distributions as a function of age (see Chapter 21), and that considered a number of other factors it thought were relevant, too. This led to revised computations of B/C ratios for replacement of various types of poles as a function of estimated condition with the "practical replacement ages" shown in Table 16.16. That for the average pole is at years, up 11 from the earlier "perfect data" estimate. These are, at least as used initially, the proactive pole replacement ages for CL&P.

16.11 OPTIONS FOR POLE REPLACEMENT PROGRAMS

Plan A: Replace Every Pole When Replacement is Cost-Justified

In mid 2009, when final planning began for the 2010 budget year, CL&P's pole replacement Planning Model estimated that in 2010 there would be 11,218 poles that were at or over 86 years in service (the practical replacement age for the

Table 16.18 Statistics for the Average Pole Justifying Replacement in 2010**Description of the pole and pts role in customer service reliability**

1	pole
69	years equivalent service condition
151	customers affected if it fails
28%	restored by switching
24	minutes average restoration time with switching
72%	restored only after repairs are made
340	minutes required on average to “repair” this failed pole

Consequences of failure

37,980	customer minutes of outage
0.046	equal to this many minutes of system aiSAIDI
\$191	safety and damage assessment
\$3,865	materials and labor to replace on an unscheduled basis

Probability of failure

0.81%	probability of failure

Expected benefits from replacement

295	customer minutes expected to be avoided this year
\$109	value of these avoided customer minutes at 37

30-year value of benefits and avoided costs accounting for both time-value of money and escalation of failure rates in the future

\$2,294.41	PW of future avoided aiSAIDI
\$32	PW of future avoided damage and safety
\$631	PW of future avoided unscheduled replacement
\$2,958	Total PW future value

Computation of merit for project

\$2,510	cost to proactively replace
1.18	benefit cost ratio to replace now
\$8.50	per avoided customer interruption minute following year

average pole that was failing). These were survivors: CL&P had installed an estimated 23,809 poles during the period up to 86 years earlier. Less than half remained. Of those that did not 7,632 (32% of those installed) had failed during the preceding 86 years. The others were replaced for NCR reasons or retired.

But not all those 11,218 poles were ready for retirement. Over half were lateral poles, most of which had much lower than average reliability impacts. Thus, only 7,215 of the 11,218 poles were actually at the practical retirement age for their type. On the other hand, a number of poles in the system, while not yet 86 years in service, were of a type with a higher reliability impact (they serve more customers than the average pole) and thus had reached the practical retirement age for their type. There are an estimated 4,890 in this category, for a total of 12,105 poles at or over their practical retirement age. Nearly all poles that justified replacement were three-phase poles. In total, the replaceable pole list amounted to just over 2.7% percent of all poles in the system, or about one out of every 37 poles. Statistically, almost 1/4 of all failures in the system were expected to be from among these 2.7% of all poles.

Table 16.18 shows the average for these 12,105 poles. They are far from average poles, whether one compares them to the average pole in the system or the average pole that fails. These poles affect, through failure, service to nearly three times as many customers as the average pole that fails: only poles with an above-average reliability impact are going to justify replacement until they are very old. But many of these poles are: average age is 69 years. In 2009, the average age of a CL&P pole was in many cases far beyond their calculated replacement age for its type because CL&P had no proactive replacing poles up to this point, so there were many poles in its system that “should have been replaced a long time ago.”

The bottom line is that these poles are “more important” in terms of reliability impact and cost to replace (whether that be unexpectedly or as scheduled work) than the average pole that is failing. On an economic basis, even accounting fully for the reliability impacts as done here, replacement of the average pole that fails cannot be justified. This is perhaps one of the most important conclusions CL&P reached (and one that is relevant to the entire industry). This will be discussed in more detail in Section 22.6.

Replacing all 12,105 at an estimated average pro-active pole replacement cost of \$2,510 would:

- Cost \$30.4 million
- Provide a calculated long term benefit/cost ratio of 1.18.
- Reduce aiSAIDI in 2010 by about 4.33 minutes compared to what it would have been.

- Improve expected aiSAIDI for 2010 by 3.03 minutes compared to the previous year. While the reduction of 4.33 minutes is only one-sixth of aiSAIDI contribution from wood poles in 2010 (28.2 minutes predicted), it is over three times the 78 second escalation expected in 2010.
- Reduce the cost of materials and labor for unscheduled replacement of poles on the CL&P system by nearly \$700,000.
- Cut escalation of aiSAIDI into the following year will be slightly less, too: 2011's aiSAIDI will be only an expected 79 seconds more than 2010's, rather than 81 seconds more as in the base case.
- Cut cost escalation similarly, too.

In the second year of this program, an estimated 733 additional poles would reach their practical retirement ages and would be replaced under this replace-at-practical-replacement-age guideline, costing about \$1.84 million and reducing aiSAIDI that year by about 16 seconds compared to not replacing this pole. CL&P's aiSAIDI would be B/C ratio this time would be only slightly better than 1.0. These 733 poles have all *only just* reached the point where $B/C > 1.00$.

Every year thereafter, the number of poles reaching their practical replacement age would increase very slightly as compared to the previous year. As a result, the cost to replace all poles meeting that criteria would be a bit greater every year than the year before, but the aiSAIDI reduction obtained would increase slightly, too, as compared to not replacing them.

Figure 16.28 shows the annual aiSAIDI contribution for poles under this program. In 2014, the fifth year of the program, annual replacement cost would have risen to just under \$2.1 million and the cumulative reduction in aiSAIDI contribution of wood pole failures would have grown to 5.2 minutes. At that point CL&P would have spent \$38.34 million on pole replacement (\$36.51 million PW to 2010), and wood-pole aiSAIDI contribution would be back where it had been in 2010. For despite the 5.2 minute reduction the continuing escalation of failures has added 5.5 minutes in this time.

These results were very disappointing to the project team: if CL&P replaced all poles where that were cost justifiable, and continued to do so each year, it could hold off aiSAIDI escalation for only five years. The sobering fact was that:

There was no long-term cost-justifiable replacement plan that would reduce pole-driven aiSAIDI escalation to zero:

at least not with the knowledge and priorities CL&P had at that time.




Figure 16.28. The predicted aiSAIDI reduction of wood poles as reduced by a “replace every pole where replacement is cost justifiable based on long-term economics” plan.

Plan B: Keep Reliability Where It Is As Cost Effectively As Possible Via Replacement of Aged Poles

For a pole that has just made it to its practical retirement age, the B/C ratio for replacement is just barely more than 1.0, as assessed against the spending efficiency that CL&P had in its prior year’s budget. The average B/C ratio estimated by the Planning Model for all 12,105 poles at or over their replacement age in 2009 was 1.18, because quite a few of the poles in that set had survived years past their practical retirement age (many more of their contemporaries had not, but that was a moot point by 2010).

Among these poles that have B/C ratios much better than that average 1.18 and some lower, although all are over 1.00. Figure 16.29 shows the distribution of B/C ratios when all 12,105 are prioritized in order of that ratio, best to worst. To prepare that data, the project team estimated the benefit/cost ratio for each of the 12,105 poles the Planning Model estimated to be at or above the practical replacement age for its type. For each, it:

1. Estimated the reliability benefit, calculating the AiSAIDI impact the pole’s failure would have based on pole type, its position in the system, its estimated condition, the number of customers downstream of it, and the configuration and switching equipment on the feeder. It then estimated the same but assumed the pole was new. The difference was the reliability benefit replacement would provide.
2. Estimated replacement costs. Several factors were considered including pole type, attachments, and access to the pole itself. With these taken into account, replacement costs were estimated more accurately than if just using the averages of \$3,262 unscheduled and \$1,940 scheduled.




Figure 16.29 Benefit/cost ratio of poles thought to have been beyond practical replacement age in the CL&P system, varied from a high of 32 to 1 to slightly less than 1.00. Dotted line shows average 3.88. Nearly 1,500 poles had a B/C ratio of 5.0 or better.

3. Included impact on safety and damage claims, although this was small compared to the other factors involved.

Benefit/cost ratio varied from over 2.1 to 1.0001. The shape of this diagram surprised the CL&P project team. They had expected a plot showing a small number of “glory” poles with B/C ratios six to seven times the average on the left, plunging rapidly to a rather wide and flat portion of the curve with a value below 2.0 over perhaps 80% of the remaining poles. Such results are often seen in asset management. The shape of this curve was a function of when poles had been added over the years, the economics of replacement, and the fact that some poles on the list should have been replaced 20 years or more earlier. This curve indicated that there was a benefit to be had from prioritizing the order of replacement (some poles had nearly twice the average B/C) but that the great mass of poles was relatively close to the average.

The top pole on CL&P’s list was a feeder riser pole seventeen years past the practical retirement age for its type (Table 16.19) that would interrupt service

Table 16.19 Statistics for the Top Rated Pole Justifying Replacement in 2010**Description of the pole and its role in customer service reliability**

1	pole
68	years equivalent service condition
3,683	customers affected if it fails
27%	restored by switching
32	minutes average restoration time with switching
72%	restored only after repairs are made
100	minutes required on average to “repair” this failed pole

Consequences of failure

300,680	customer minutes of outage
0.364	equal to this many minutes of system aiSAIDI
\$191	safety and damage assessment
\$10,600	materials and labor to replace on an unscheduled basis

Probability of failure

0.74%	probability of failure

Expected benefits from replacement

2,127	customer minutes expected to be avoided this year
\$787	value of these avoided customer minutes at 37 cents

30-year value of benefits and avoided costs accounting for both time-value of money and escalation of failure rates in the future

\$11,476	PW of future avoided aiSAIDI
\$21	PW of future avoided damage and safety
\$1,093	PW of future avoided unscheduled replacement
\$11,823	Total PW future value

Computation of merit for project

\$6,000	cost to proactively replace
2.10	benefit cost ratio to replace now
\$2.82	per avoided customer interruption minute following year

to 3,683 customers when it failed.²¹ It would cost a lot more than average to replace, but it was a key to much more than the average customer reliability improvement due to those facts.

Replacing this one pole would reduce expected aiSAIDI the following year by only a sixth of a second. To “buy back” the expected 78 seconds of aiSAIDI escalation expected in 2010 due to increasing failures of wooden poles, CL&P would have to replace 483 poles just as important as this one, at a cost of \$2.9 million. Unfortunately there are not 483 other poles like this one on the list. The second pole on the list has only about 2/3 of the reliability improvement potential as this top pole, the third only half, and so forth: CL&P will need to replace the top 1,750 poles on its list in order to expect 78 seconds of aiSAIDI reduction. The results of doing this would be:

- a) Cost of \$6.0 million. Cost per pole is an average \$3,428: materials and labor are less for this not quite as significant poles.
- b) Expected net improvement in aiSAIDI compared to the base case is 78 seconds, equal to the expected worsening of aiSAIDI.
- c) Reliability will be different: the likelihood that there will be a failure of a critical or important pole will be noticeably less. The probability of interruptions of large numbers of customers at one time is reduced by about 12%.
- d) Benefit/Cost of the lot is 1.52:1, still quite good: CL&P would

²¹ The feeder caused more than a little bit of heartburn and disagreement within the larger Asset Management group at CL&P, many of whom felt it was an anomaly that had slipped through the cracks. This was a *dual feeder* riser pole – a single pole on which two feeders originated out of a substation. Such practice was against company guidelines and had been for over fifteen years precisely because the failure of such a pole would interrupt service to so many customers. This pole has slipped through the system because it had been “grandfathered” before the policy had been implemented. Engineering and executive management were aware of the situation but local conditions precluded replacing this pole with two feeder riser poles and so it had been left in place.

What the CL&P engineers had missed was the details for this pole. Note that from data given in the tables only 27% customers can be restored by switching and that the average time is estimated as 32 minutes. Until this replacement planning caused the company to take a new and hard look at this pole, its reliability studies had been used 66% and 20 minutes: what had been missed in the reliability assessment of both feeders was that each depended on the other for a majority of the restoration switching during its contingencies: if this pole failed that would not be the case, both feeders would fail.

Given everything, the pole was replaced with a steel pole of considerable strength, continuing the deviation from the planning guideline but with an expected impact in reliability including storm withstand capability that was thought to more than compensate, at a cost and with reliability impacts not included here.

be getting 52% more long-term benefit for the dollars spent on this as compared to the least effective dollars it was willing to spend in 2009.

- e) Cost per avoided customer minute of interruption in 2010 is \$5.59.

The following year the situation will have changed in several interacting ways:

- 1) All the poles in the CL&P system will have aged during the year.
 - a. As a result, every pole on the list that was not been replaced in the year 2010 will have a B/C ratio somewhat greater than it had the year before: during the year its probability of failure increased, meaning the benefits of replacing it increased, but its replacement cost remained the same (in constant dollars).
 - b. Another 773 poles will be added to the list of poles that exceed their practical replacement age. All will have B/C ratios just above 1.00 – they have only just reached their replacement age (that age where B/C is first above 1.00).
- 2) The poles proactively replaced in 2010 will not only have very low failure rates compared to what they would have had if not replaced, their failure rates will escalate at a much lower rate than if they were not replaced (they are new and their deterioration escalation will be slower). As a result, the expected aiSAIDI erosion of the entire pole set for 2011 will be 2 seconds less (79 rather than 82) than it would have been if these poles had not been replaced, for many years to come. That is not much but it would take replacement of half a dozen poles from near the top of the list to make that improvement in the following year.

But CL&P picked the “lowest hanging fruit” in 2010. Even with aging, the poles on the top of the list in 2011 don’t offer the same cost effectiveness as the top poles from the year before. CL&P will have to replace 2,600 poles at a cost of \$6.46 million in 2010. Benefit/cost ratio will drop slightly to 1.41, and cost per avoided customer minute rises to \$5.94. This will be the pattern in each succeeding year. Every year CL&P will have to buy more seconds of reliability offset due to the growing escalation of pole failures. Every year the poles remaining on the list will be slightly less cost effective than the year before. Every year it will have to replace more poles and spend slightly more money than the year before, and B/C ratio will be lower and marginal cost of next-year’s reliability improvement higher.

Five years into the program, this plan will hit the same wall that Plan A did. The aiSAIDI contribution from wooden poles will be back at 2010 levels, in fact just slightly higher, and no amount of cost justifiable spending will keep the

reduction to 2010 levels. CL&P will have replaced the same poles over the period, spending exactly the same \$38.24 million but at a rather more constant rate so that the PW will be far better (\$30.68 to year 2010 versus \$36.51 million for Plan A).

From there, both plans would be identical. Either would replace about 1,000 poles in 2015 at a cost of about \$2 million and make a 16 second reduction in aiSAIDI. Each year thereafter, either plan would replace all the poles that aged past their practical replacement age during the past year. That would always be a few more than the year before, at a slightly greater cost, and with a very slightly increased reduction – but not nearly enough to completely stem the escalation of pole-failure caused aiSAIDI.




Figure 16.30 Top, base, Plan A and Plan B aiSAIDI due to wooden distribution pole related failures. Bottom, spending on pole replacements for Plan A and B.

The Cost-Justifiable Sustainable Point

A change in the rules – in this case a decision to proactively replace some poles rather than practice “run to failure” for all, moves the sustainable point (See Appendix B). Both Plans A and B replace a pole when it makes economic sense, using as a guide a marginal cost of \$2.61 per avoided customer minute, with considerations of safety and replacement cost playing into the calculation slightly. Plan B imposes a constraint of only replacing poles up to an amount that counteracts the aiSAIDI escalation rate of the entire set of poles, but that constraint fades to zero within five years and does not affect the long term. Both plans have the same sustainable point.

If CL&P continued either program long enough, annual proactive replacement cost would gradually rise to about \$27 million per year to replace nearly 7,000 poles as the “hump” of poles installed during the growth period in the 1960s and 1970s reached its practical replacement ages. Annual cost would then drop slowly to a sustainable value of about \$6 million per year, replacing around 3,000 poles every year. The aiSAIDI contribution from wood pole failures would continue to peak sometime after 2035 at around 41 minutes (as compared to 52 in the base case).

It would take more than six decades to reach that sustainable point, so it was not too important to CL&P’s project team. What was important was that either Plan A or B holds reliability issues at bay, *cost effectively*, for five years, and “clean ups” the major and trunk poles on the CL&P system by around 2020: within only a decade, escalation of aiSAIDI due to increasing failures of critical and major poles is basically contained. But the cumulative effect of poles that do not individually contribute to reliability problems in a way justifying their replacement “wins”: failures and outages continue to escalate because there is such a large mass of such poles.

And long-term problems are just that: in the long term. The most important point to both the project team and CL&P’s executive management in early 2010 was that, working only with the data, models, and knowledge it had developed over the past two years, CL&P could manage reliability and cost impacts due to pole failures in a *cost-justifiable* manner for five years into the future.

Plan C: Keep aiSAIDI Contribution at 2009 Levels

For completeness, the project team estimated the costs of continuing to replace a sufficient number of poles each year to keep expected aiSAIDI at 2009 levels. They did not expect such a plan to gain much favor with upper management since it required spending money that was not justifiable from within the company’s then current spending rules, but they did expect to be asked what the plan would entail.

Maintaining aiSAIDI escalation at zero for one more year, from either plan A or B, involved simply continuing the overall concept of Plan B: replace enough of the top rated poles to buy the necessary improvement. There would not be

enough poles on the list with $B/C > 1.0$, but by dipping a little lower, enough could be found. CL&P would have to spend just slightly less than \$10 million in 2015, replacing over 4,500 poles, with an average B/C ratio of .92. The following year the expense would be higher (just over \$11.5 million, and the B/C ratio even lower at .87). That trend would continue into the future, peaking somewhere after 2030 at over \$50 million a year with cost-effectiveness less than 50%. At its worst, B/C ratio for this program dips below .50. But it is possible to plan, organize, and implement a program designed to keep pole-related problems at bay forever, in as cost-effective way as possible. That best program is simply not cost effective enough to meet the company's current spending priorities.

Regardless, Plan C moves the sustainable point further and faster than Plans A or B (Figure 16.30). Starting in 2009, it moving the aiSAIDI contribution to 27 minutes, and leaving it there. What varied as a trend with Plan C was spending, which simply escalated, at a much higher rate than any of the trends in the base case, for the foreseeable future.

16.12 CL&P'S PLAN FOR DEALING WITH AGING WOODEN DISTRIBUTION POLES

The Planning Model's trends for overall impacts matched fairly well the past twenty years of CL&P operating experience. Perhaps more important, they made sense intuitively, and they were explainable. The conclusion that there was no "silver bullet" – that there was no cost effective program existed to drive aiSAIDI escalation to zero existed – was disappointing, but had been expected. Therefore, after the project team's presentation to the SVP and a committee of department heads in mid 2009, and following a period of "due consideration" (the executives read the report, considered its results, and talked among themselves), CL&P's Delivery Division "officially" accepted the Planning Model and its results as the basis for its decision making with regard to managing distribution wooden poles. A point here that might be considered a good practice: the SVP arranged to review and accept the model before having it used to produce either of the plans, A or B given earlier, or any other of the consequences known. This was done because the SVP did not want the planning team working on plans with a model not officially endorsed by upper management, and because this definitely looked best from a credibility standpoint: CL&P had accepted the model before knowing whether it would like (or not) the recommendations it made.

Acceptance meant the company acknowledged that it had a noticeable problem with the continued aging of its wooden distribution poles and that the trends of doubling of failures, costs, and reliability impacts discussed earlier were real. There was little doubt they were. Those trends would continue until the company took effective steps to either eliminate the root cause of the problem (aging poles) and/or mitigate the symptoms they caused. CL&P's executives, particularly the new SVP of Power Delivery who had started this

entire evolution, had very specific ideas about how to achieve those goals.

Within only two weeks the project team did the prioritization of the poles that led to Plan B, and had produced Plan A and Plan C. They reported back, going over the plans they had produced in alphabetical order:

A (basic concept of replacement when justified),

B (a plan to minimize PW spending where justified)

C (a plan to minimize spending needed to keep reliability erosion to zero) .

A Strategy Based on Performance and Cost, Not Sustainability

“Our job is not to take care of poles. Our job is to use them in a safe manner to maximize the value we give our customers and stockholders.” That statement by the SVP of Power Delivery, who had set the entire aging pole project in motion originally, defined CL&P’s strategy with regard to poles, reliability, and other aging equipment issues (a similar program on UG cable concluded at the same time as poles). The company would target customer and business performance at levels it thought were satisfactory from a business and regulatory standpoint. Pole replacement would be funded based on if and how it could contribute to that. The SVP acknowledged that some pole replacements were cost effective, particularly the early years of Plan B. But so, too, were a number of other candidate programs and projects CL&P could implement, including the leading elements of the Smart Switching program, certain cable replacement programs, and others. CL&P would consider all of these on a balanced basis, selecting those made the best sense overall: that meant greatest business impact per dollar spent.

Optimized Portfolio Management

CL&P intended to take all the candidate projects it would consider for funding – pole replacements, feeder extensions, transformer upgrades, cable replacements, breaker refurbishments, as well as projects in all other categories, and evaluate all on the basis of their contribution to the company’s business case. It would pick the set of projects that led it to satisfy all its obligations and commitments at the lowest possible cost.

For this it used its EAMPP (Electric Asset Management Project Prioritization) system to optimize the selection of which projects to fund. Through an online form filled out for each candidate project, ESMPP essentially asked “When the company’s obligation to serve and commitments are concerned, and when reliability improvement, cost reduction, customer service, safety, and other goals and considerations were taken into account and valued according to the company priorities, what set of projects provides the most bang for the buck? Give details.” EAMPP would select a *portfolio* -- set of projects from among all those candidates given it as input – that could be funded to assure the company would meet its obligation serve, achieve its SAIDI, SAIFI

and aiSAIDI targets, keep lifetime cost at near-optimum levels, meet the company's public and employee safety goals, and meet all of its other commitments and obligations, at the lowest cost and risk.^{22,23,24} Pole replacement projects would have to stand on their own merit as would all other types of projects. Any particular project would be selected only to the extent that it offered more value toward the company's goals than other candidate project, whatever they may be.

²² ESMPP does not select projects using the purely "merit score" basis approach used to describe many asset management prioritization methods. In that basic approach, which might be termed first-generation asset management prioritization, every candidate project is given a weighted score based on its contribution to reliability improvement, to safety improvement, to cost reduction, etc. – the weighting for each reflecting how important that specific goal is to the company. The various component-weighted scores are added together and then divided by the project's cost to create an overall benefit/cost ratio, a "figure of merit" or merit score for the project. All candidate projects are then ranked from highest score to least. Projects at the top of the list are funded.

ESMPP is a second-generation portfolio optimization method. It computes merit scores for projects but only for display: the scores provided are helpful to planners and executives in evaluating projects and understanding the results. ESMPP *selects* project using a MIP (mixed integer programming) optimization algorithm, its goal being to satisfy a set of corporate targets (SAIDI must be below a certain limit, safety above a certain target, lost time injuries must be kept below the target limit, equipment maintenance costs must be kept within projections, etc.). In selecting projects, ESMPP might initially put a high value on reliability, and "buy" projects that provided reliability, but only up to the point where it had enough reliability to satisfy its target. From that point on, reliability would have no further value and it would select projects for other goals.

It should be noted that ESMPP did not work on targets serially (i.e., buy reliability first, then work on cost, then work on safety, etc.) as the foregoing might imply: it worked on all simultaneously, using the optimization engine to find the best overall set of projects that achieve the targets, meet all the constraints and limits and "must does" it had been given, and that have the lowest *overall* cost to do so. Given the issues CL&P had with reliability, ESMPP was likely to select any project with a very high "bang for the buck" with regard to reliability. But it was more likely still to select a project that had a pretty good bang for the buck for reliability, *and* a good impact on cost reduction, *and* a contribution to safety improvement, and that created no issues for other projects or required scarce resources or commitments, what might be termed an Almost Best Everywhere project (see Chapter 12).

²³ More information on Asset Management and Portfolio Optimization as applied by utilities in such situations can be found in Chapter 19, "Asset Management" by H. Lee Willis and Richard E. Brown, in the *CRC Electric Power Substation Engineering Handbook*, J. D. MacDonald, editor, and *Business Essentials for Utility Engineers*, by Richard E. Brown.

²⁴ The actual CL&P ESMPP does not use a figure of merit evaluation that is directly comparable to B/C ratios given here. However, rather than introduce a new metric, the authors will give all results translated into equivalent bottom line, all causes and costs, Benefit/Cost ratio.

Pole Replacement Projects

To fit into the framework of consideration for candidate projects by ESMPP, the project team could have created and submitted for consideration a candidate project for each pole high on its prioritized list: a project calling for the replacement of the pole in Table 16.18, another for the second on its prioritized list, etc. But that would have created thousands of small projects, and a lot of paperwork.

Instead, the project team created and submitted Pole Replacement Projects (PRPs), each a group of poles taken from the top of its prioritized list. PRPs were designed to have costs between \$250,000 and \$500,000 and to be a set of poles. Forty-four PRPs were submitted, with marginal cost scores from 87¢ to \$2.92, covering replacement of over 7,000 poles for a total of just over \$21 million.²⁵ This was basically the first three years of the plan. Given that the Benefit/Cost ratio cut off limit in the prior year was \$2.61, the project team was thought it possible that ESMPP would select all but a few of these PRPs.

Such submission of projects was a standard way to do things at CL&P. For example, the \$25 million Smart Switching program was broken into 50 sets of smaller projects, each of about \$500,000.

A Heavily Constrained Plan for 2010

The ESMPP spending limited reached for 2010 was \$2.56: at the margin it had to spend about 2% *less* than the \$2.61 limit (in 2010 dollars) it had found the year before, to attain 2010's goals.^{26,27} But in spite of that, ESMPP selected only the top seven PRPs, funding a total of \$2.7 million in 2010. It did not even select all of Plan B's first year, funding a bit less than half of it, and it "left on the table" thirty-six PRPs that had merit scores better than its \$2.64 cut off limit.

²⁵ Recommended approach within CL&P was for departments to submit projects that included the more cost effectiveness (87¢ in this case – close to spectacular) to about 10% above the cut off for the previous year (\$2.87 in this case). By bracketing the previous year's cost cutoff point, the department and company could be assured that spending for that department's area of responsibility had been balanced fairly against spending opportunities in all other departments.

²⁶ Both pole replacement and Smart Switching were new candidate programs for 2010. In past years, neither had been available to ESMPP. While both offered a range of economics (e.g., from 87¢ to \$2.92 for pole replacement) the important point for ESMPP is that each offered some projects that could help attain its goals at much better marginal cost than it had been able to find in the past: both pole replacement and Smart Switching allowed it to reduce spending. Two percent may not seem like much movement, but it is a fairly noteworthy drop in a number that should, if policy and targets and operating methods don't change, stay about the same over time.

²⁷ Readers familiar with Asset Management planning know that many more than one computer run was performed. CL&P planners spent nearly three weeks running and re-running ESMPP to both study variations on the 2010 plan and to work out of the system various characteristics they or management did not like.

Table 16.20 Comparison of Two Projects Selected for Funding in 2010

Factor	Pole Replacement	Smart Switching
2010 spending required	\$448,600	\$479,200
Change in long-term PW operating costs	(\$37,900)	\$32,300
Expected aiSAIDI reduction in 2010 - seconds	6.7	21.1
Expected aiSAIDI reduction in 2030 - seconds	16.8	23.3
\$ per minute of 2010 customer-interruption	\$4.87	\$1.65
Evaluated merit score: \$ per unit benefit	\$1.80	\$2.50
Equivalent B/C score (as used earlier in chapter)	1.42	1.03

The reason for this selection, which surprised the pole-replacement team no end, is a classic lesson in asset management optimization.

Table 16.20 compares the tenth and final (lowest merit score) PRP that was funded in 2010 with the corresponding (lowest merit score) project of the eight Smart Switching programs also selected. The optimization tool had selected eight pole-replacement programs for \$2.75 million that “bought back” nearly half (37 of 78 seconds) of the aiSAIDI escalation for 2010. It funded six Smart Switching programs at a total 2010 spend of \$3.2 million that bought in total 51 seconds of aiSAIDI reduction. This sum of 88 seconds from these two projects exceeded the 78 second aiSAIDI escalation for poles because ESMPP had erosion in other areas (URD cable) to address, too, as well as issues pertaining to risk it addressing, too.²⁸

The pole-replacement program provides the superior long-term bang for the buck. Each unit of long term (PW) multi-attribute benefit it buys costs only \$1.86. By contrast the Smart Switching program is spending nearly 33% more (\$2.52 each) to buy the equivalent. But the Smart Switching program is nearly three times as effective at driving down next year’s aiSAIDI: each avoided customer-minute of service interruption it buys costs only \$1.65, as opposed to \$4.86 for pole replacement.

For 2010, as it had been for several years and would continue, CL&Ps spending was dominated by two goals: reliability targets and budget limits. CL&P wanted to contain aiSAIDI, and SAIFI, and other reliability stats to levels not worse than 2009 – in fact it hoped to slowly drive them down over the next decade. But it also had a budget limit it wanted to meet, a total spend that would keep debt to equity ratios and other financial metrics within limits that executive management thought prudent. ESMPP selected a set of projects that met both these targets and all others.

²⁸ Selecting projects to just meet the mean value of a target, in this case, aiSAIDI, would mean that the company had a 50% likelihood of exceeding the target. A probabilistic approach targeted to achieve 90% compliance was CL&P’s goal.

In doing so, the program determined that it could not afford all the long-term savings it could buy. There were candidate projects input to it that would have permitted it to buy all the reliability improvement it needed from only pole replacement: these would have saved more in the long term, but it would have had to spend more now (about \$4 million more in 2010), and it did not have that much in its budget – not enough to both do that and meet all its other obligations. On the other hand, it could also buy all the reliability improvement it needed from only Smart Switching, and that would result in less spending than in its final plan, but far less long term savings. One way to look at its selected plan is: it achieved 2010 goals while buying as much long-term value as it could.

This was both inevitable and not far different from the results that other utilities would have obtained. An earlier discussion (the “Conundrum”) observed that aged equipment replacement projects often have very good long-term B/C ratios but often do not look as good as other projects from a short-term effectiveness standpoint. The remainder of CL&P’s \$92 million capital budget for 2010 was distributed over two hundred other major projects and several blanket accounts.

Overview of the Long-Term Program

This chapter was written as 2010 came to a close, at a time when 2011’s budget had been almost completely determined. ESMPP funding cutoff for 2012 was \$2.59, higher than 2010 but still lower than it had been in 2009. CL&P and ESMPP found the company had mined pole replacement and Smart Switching for the most cost-effective projects in 2010. What remained was slightly less cost efficient – enough that for 2011 the marginal spending limit rose slightly to \$2.59. The 2011 portfolio was otherwise similar in character to 2010’s with respect to poles and Smart Switching, except spending began what was expected to be a slow, decades-long shift of gradually reduced spending levels while that for pole replacement ramped up – at first slowly and then in a large jump near the end of the decade.

As mentioned above, CL&P had many other issues and concerns than just aging poles to accommodate. Among them aging URD cable and corroding lattice towers were two of the more serious aging-related issues. In 2010’s budget, some of the Smart Switching implemented was directed at making up for aiSAIDI erosion caused by these factors. All that was truly important to CL&P and ESMPP was that the overall targets for aiSAIDI, SAIFI, and other reliability goals be accommodated. About 10 to 15 percent of Smart Switching’s capabilities were used to that effect.

Plans for the Future

In Table 16.21, the authors have done an analysis of the eleven-year period 2010 to 2020 in which all those other considerations have been eliminated and the problem is one of just solving the continuing aiSAIDI erosion from wooden pole

trends using pole replacement and Smart Switching. This was done by holding everything else constant in ESMPP while studying only this aspect of the plan. All cost are in 2010 dollars.

The first column gives the year, the second the spending cutoff for the year. Over time the company had to raise spending by nearly 10% (\$2.81 versus \$2.56). It is costing more to maintain system performance to target each year. There are many reasons but certainly aging effects are a major reason. This is due to the nature of optimization in the face of the long-term trends: ESMPP (or any good optimization tool) buys the most effective projects first, leaving somewhat less effective projects for the future.

Table 16.21 Relevant Spending Statistics for 2010-2020

Year	Spending	Dollars Spent (millions)		aiSAIDI	\$ per avoided	aiSAIDI
Bought	Cutoff	Replace Poles	Install Switches	Bought	Cust-min.	Erosion-sec
2010	\$2.56	\$2.75	\$2.83	78	\$4.29	0
2011	\$2.59	\$2.85	\$2.88	81	\$4.27	0
2012	\$2.62	\$2.95	\$2.93	83	\$4.25	0
2013	\$2.64	\$3.05	\$2.98	86	\$4.23	0
2014	\$2.67	\$3.15	\$3.03	88	\$4.21	0
2015	\$2.70	\$3.25	\$3.08	91	\$4.20	0
2016	\$2.73	\$3.35	\$3.13	93	\$4.18	0
2017	\$2.75	\$3.50	\$3.10	96	\$4.21	0
2018	\$2.77	\$9.40	\$0.76	98	\$6.06	0
2019	\$2.79	\$10.20	-	101	\$6.09	0
2020	<u>\$2.81</u>	<u>\$10.50</u>	<u>-</u>	<u>103</u>	<u>\$6.12</u>	<u>12</u>
Totals	\$2.69	\$54.70	\$24.81	996	\$4.79	12

Columns three and four show annual spending on pole replacements and Smart Switching needed to drive aiSAIDI escalation due to wood pole failures to a net increase of zero. Only \$24.81million (of the \$27.6 million total) of Smart Switching was ultimately devoted to helping accommodate aiSAIDI erosion due to wood pole failures.

While ESMPP always “picked the lowest hanging fruit” available, there is a substantial difference between the two “orchards” it has, pole replacements and Smart Switching. Smart Switching included a total of \$27.6 million in possible implementation, and that varied in both short-term reliability cost-effectiveness and long-term value. Buy the most effective portions of Smart Switching in 2010 and what remain to be bought down the road are forever not as cost effective as those first portions. Its cost effectiveness does not change over time:

what works in 2010 works just as well but no better in 2011, or 2020.²⁹

Like Smart Switching, different PRPs vary in both short- and long-term effectiveness. But in contrast to Smart Switching, pole replacement economics *improve* from one year to the following year. Each year, between 800 (year 2011) to 2000 (in year 2020) poles will move across the $B/C = 1.0$ line into the cost-justifiable set of CL&P poles. Every year, the economics for any pole already on that list improves: assuming it survives it is older and statistically more likely to fail in the future – the payback for replacing will increase but the cost of replacing it (in constant dollars) will not. For a 75 year old pole, if a B/C ratio of 1.05 is not good enough, one has only to wait. Five years later, at age 80, the replacement B/C will exceed 1.25, so much has the probability of failure increased due to expected continuing deterioration.

Over the decade 2010 to 2020, the Smart Switching becomes less and less attractive as the company buys out the best portions of that program. But even at the end, it is still an attractive program from a short term standpoint (last column) and so it buys the entire program by 2017. Meanwhile, because poles that have not yet been replaced improve in replacement cost effectiveness the following year, spending on pole replacements looks just slightly better every year as all poles with $B/C > 1.0$ are not replaced.

Thus, in the early years of this plan, ESMPP spends a bit more on poles each year, but only a bit more. But by 2018 there is not enough Smart Switching left to buy. Pole replacement offers good long-term value – the company is not wasting its money replacing any pole on the list -- but pole replacement is only about a third as effective as Smart Switching was at driving down *next year's* reliability problems. Spending on pole replacements has to almost triple, but a sufficient number of poles can be found to replace that year, and in the next, to buy the required reduction in expected aiSAIDI.

That backlog of poles has been worked through by 2020. CL&P does not have quite enough replaceable poles that year (i.e., not enough poles moved into the cost-justifiable list to give it all it needed to replace that year). Ten million dollars in spending – to replace all the poles remaining on the replaceable list included those that just moved onto the list in 2010, is not quite sufficient to hold aiSAIDI erosion at bay. Thereafter, this plan is similar to Plan A or B after they have exhausted any backlog of replaceable poles: by replacing poles that become cost justifiable each year, the company can reduce aiSAIDI erosion each year by about 20 – 25%.

²⁹ In actual fact there is a rather complicated interaction that makes this statement not altogether accurate. If poles are *not* replaced, aiSAIDI will increase each year. With a higher aiSAIDI, Smart Switching has more to do (i.e., more benefit to provide) each year, and thus its effectiveness will rise, slightly, each year. This interaction with pole replacement is a secondary effective, not as noticeable as effects discussed in the main text, and has been worked out of the plan and the figures given here. The plan shown in Table 16.21 is, within the limits of the constraints and analysis for CL&P done here, the best relative use of Smart Switching and pole replacement in each of the years shown.

But the important point to CL&P's executives, and its Asset Management Department, was that they have an affordable, justifiable plan that would contain "the pole problem" for eleven years.

Looking at the Longer Term

At the time this was written, CL&P was in the midst of finalizing its 2011 -2013 capital spending plan. There is little doubt the company will follow something close to the plan described here, at least for the next two to three years. But there is also no doubt that it will do something a bit different in the long run. The reasons are:

- 1) *More Smart Switching.* CL&P could implement more Smart Switching than its original program laid out. That program had been designed against the \$2.60 marginal cost figure in place at the time. In Table 16.21, marginal cost rose by about 1% a year, a trend expected to continue for some time after 2020. CL&P's reliability team estimated that every 10% rise in marginal cost increases the amount of justifiable Smart Switching by about 12%-15%. This means that there is probably another year or two of aiSAIDI erosion abatement available near the end of the decade depicted in Table 16.21.
- 2) *"Smart Distribution."* Most of CL&P's distribution system was laid out decades earlier to configuration guidelines of feeder routes, branches, and tie points designed for the technologies and reliability expectations of those times. Changes in its basic feeder configuration, adoption of some practices used in Europe, and revision of all those guidelines to maximize the effectiveness of Smart Switching (and vice versa) can make an estimated 30% improvement in cost effectiveness improvement of customer reliability. This will likely buy three or four more years of improvement.
- 3) *Improved testing, pole data, and tracking.* By 2020, CL&P will have accurate data on all poles less than 30 years old in its system, by 2030, on all poles less than 40 years old, since it started recording data completely in 1989. But this alone won't make a noticeable impact on its operations until about 2040, since most failure issues are with poles older than 50 years. Of more importance will be improvements in its pole testing and condition tracking and further improvement of the statistical condition analysis methods it is using. CL&P plans to continue the testing of failed, NCR and retired poles in order to improve its knowledge of how to improve its pole condition estimates. There is no guarantee that CL&P will ever get testing and tracking results equal to the best reported elsewhere (those it obtained in 2009 were only about average), but several vendors of pole testing services are willing to try, and testing technologies are improving all the time.

The three items above, in aggregate, are likely to provide CL&P with a means to cost effectively contain pole-related erosion of customer service and business impacts for at least half of a second decade. From a practical standpoint this is really as far ahead as the utility needs to look: the company is not “kicking the can farther down the road” by taking that approach, as much as taking care of the cans it finds now while acknowledging that there will be others it will find two decades hence.

Beyond 2030, there is reason to expect a solution will be provided by the gradually rising marginal cost limit. Adjusted for other factors, the spending limit cutoff for CL&P’s capital programs had risen from \$2.44 in 2003, to \$2.61 in 2009, and was predicted to rise to about \$2.80 by 2020 (all figures in 2010 dollars) – a compounded growth rate of .8% annually. The reasons for this slow escalation are tied in large part to aging poles, and aging cable, and aging transformers, and aging breakers, etc., as well as to the slowly rising bar of societal and customer demands for reliability and storm response. Regardless, the trend is there. And if it continues, the limit will increase by 10% every twelve years.

As that limit increases, the number of “resources” available to solve problems and make improvements at CL&P grows. The practical replacement ages for all categories of poles drop, bringing many more poles into the cost justifiable range each year, and earlier than otherwise. Smart Switching and other technology innovations that were not quite implementable a decade earlier are now. Operating policies that could not quite be justified now can be, and so forth.

It would take five decades of this trend to move the spending cutoff limit to where CL&P could wholly justify Plan C (replace enough poles each year to keep aiSAIDI at 2009 levels). But CL&P may have that long. The plan in Table 16.21, augmented by improvements (1) through (3) discussed above offers a workable, cost-justifiable solution through the year 2030, at which point the spending cost will be nearly 20% higher – i.e., 40% of the way toward the level needed to sustain Plan C. The 20% improvement will mean another tier of available poles to replace at even earlier ages, justify using even more smart equipment, and using even more innovative and different operating policies to render the system “failure immune.”

In summary, CL&P has a practical, justifiable and affordable plan that will meet its goals for the next three to five years. It has a general indication of what it needs to do to sustain that plan and its successful containment of aiSAIDI for another decade after that, and has the tools and resources in place to gather the information needed and make the decisions required. Beyond that, there is reason to expect that it can manage this infrastructure soundly toward an eventual sustainable point that is acceptable all around. So although there are challenges, not only in implementing the plan but in communicating to shareholders, customers, and regulators, the real long-term need to raise spending, there is no crisis.

17

Guidelines and Recommendations

17.1 INTRODUCTION

This final chapter summarizes and integrates the key points made about aging power delivery infrastructures and the technical and managerial methods covered in the previous sixteen chapters, into a coherent overview of the problem and its solution. It provides guidelines and recommendations as well as lessons-learned from experience and the case study (Chapter 16) which the authors think are generalizable.

Section 17.2 begins by summarizing the five main “causes” of aging infrastructure problems and highlighting how they interact among themselves and with other aspects of a utility’s operations. It then summarizes the four key points of solutions – methods that can mitigate the effects of aging infrastructures and improve the performance of the distribution utility. Section 17.3 then presents a series of recommended approaches, not all of which will be relevant to every utility, but all work well at some utilities. These are organized based on two criteria: how quickly they can be implemented and how much impact they can make. Section 17.4 gives guidelines on setting up and managing an aging infrastructure program at a utility.

17.2 FIVE INTERRELATED FACTORS

An aging power delivery infrastructure is an area of an electric utility system that is mostly old equipment near the end of its lifetime, configured in a layout that is itself quite old and not completely compatible with modern needs. For a variety of reasons, the aged equipment and layout create reliability, customer service, and

Table 17.1 Characteristics of the Aging Infrastructure Areas (same as Table 1.4)

The system layout and design was first put in place more than forty years ago.
The majority of equipment in the area is more than forty years old.
The system is well engineered and fully meets minimum engineering criteria.
The area is seeing steady, if perhaps low, growth in load.
The area is plagued by above average equipment failure rates. Overtime is high due to large amounts of unscheduled repair and restoration.
SAIFI began rising some years ago; SAIDI is beginning to rise now.
Major interruption events always occur due to a bizarre series of outages.
Things go bad very often.

budgeting challenges that if not anticipated and brought under control through proactive management, will eventually overwhelm even the most effective and efficiently run utility company.

The most easily identifiable characteristic of aging delivery infrastructures is that the equipment in the area is old. It will have an *average* service age greater than its design lifetime. Situations vary from one location to another, but most aging infrastructure areas share similar characteristics in other areas, exhibiting a type of “signature” as listed in Table 17.1. This “aging infrastructure signature” was discussed in Section 1.1 and Chapters 7 and 8 and Appendix B.

Aging infrastructures are not just about old equipment. There are actually five interrelated, contributing factors at work in these areas, all of which work against efficient, reliable operation. These are:

1. *Old equipment*, which requires high O&M efforts, and whose higher than normal failure rates degrade customer service quality. Old equipment is discussed in Chapter 7. Methods to measure, manage, and mitigate its effects are discussed Chapters 7 and 8.
2. *Obsolete system layouts*. Many metropolitan utilities have been unable to obtain new substations or ROW into the central core of their metropolitan areas for decades. As a result, they must serve the high levels of both power and reliability that modern society demands, from systems that are far less than ideal with respect to topology and interconnection/configuration strength for contingencies. This exacerbates problems caused by aging by limiting the system’s strength to tolerate the loss of equipment when it fails. This issue is discussed in Chapter 8. Methods that measure,

manage and mitigate problems caused are covered in Chapters 11 – 15.

3. *Old engineering methods.* Many T&D utilities are using engineering methods that worked well in the 1970s, but cannot fully guarantee reliable operation in a world where substation and line loadings are pushing beyond traditional levels and equipment is old. Traditional tools need to be augmented by new ways of searching for problem areas in a system's design and in identifying and prioritizing how and where money is spent to bolster reliability of the system. This problem is discussed in Chapter 8. Methods to fix the situation are covered in Chapters 12 -16.
4. *Uncoordinated and non-optimal use of distribution.* Few utilities coordinate the reliability economics and planning of the sub-transmission – substation level and the primary distribution feeder level optimally or even near optimally. As a result, the distribution system is often an underutilized reliability resource. Although it involves more difficult planning methods,

Table 17.2 Contributing Factors to Aging Infrastructures and Their Impact

Cause	SAIFI	SAIDI	Costs
Aging Equipment			
1a. Equipment failures	X		
1b. Higher repair and restoration efforts		x	X
1c. More inspection and testing			X
Obsolete System Layouts			
2a. Accelerated aging	x		x
2b. Limited contingency capability		X	
2c. Higher delivery operating costs			x
Outdated/Inadequate Engineering Methods			
3a. Inadequate criteria application	X	X	
3b. Weather normalization		X	
3c. Load forecasting		X	
Outdated Engineering Methods			
4a. Traditional, inflexible design rules	X	x	X
4b. Lack of rigorous reliability design	x	X	
4c. Poor plan coordination with other levels		X	
Old Cultural Identities			
5. Inappropriate paradigms			X

coordinated multi-level, reliability – based planning of a power system results in much better reliability and much better use of both capital investment and operating budget.

5. *Old cultures and ideas.* Institutionalized procedures and values – what might be termed corporate culture – at many utilities still tends to operate in terms of the culture and goals appropriate for a purely regulated, vertically integrated power industry in the last quarter of the 20th century. Goals are different today and are continuing to evolve. A fresh perspective must be taken by executive management, and professionals in all areas of a utility, to re-examine the most basic concepts throughout the organization, and to apply innovation and creativity to the methods used to manage the utility, in both aged and new areas of the system. Discussion and methods falling into this category are discussed in Chapter 1, throughout the book, particularly in Chapter 14, and in this chapter.

Table 17.2 shows the impacts these five contributing causes and their sub-areas make on the utility's reliability and cost. The size of the "x: under the three categories of impact shows the degree to which the cause affects that category.

17.3 RESULTS-DRIVEN MANAGEMENT (RDM)

Cross-Functional/Cross-Departmental Results-Driven Management

The authors use RDM here to designate the term "results-driven management," as discussed in Chapter 15, the uniform application of reliability-focused prioritization of spending in a coordinated manner throughout the utility. This is very highly recommended. Results-driven Management means taking that same basic concept, which lies at the heart of reliability-centered maintenance, budget-constrained capital prioritization, and reliability-based optimization, and integrating its use in all aspects of the utility's operation. The basic concept is simple:

Spend money where it will do the most good. Be willing to move funds from one budget to another, from one function to another, and to spend it in non-traditional ways if those changes will buy more reliability for the same money.

Chapter 12 presented a lengthy example of RCM applied to decisions about what equipment to maintain, and how to maintain it – a comprehensive but traditional view of RCM. There, alternative actions that could be taken, from inspection to rebuilding of equipment, were all evaluated on the basis of the ratio of expected improvement in reliability of service to expected cost. The most cost-effective actions were approved. Others were not. The key factors in such analysis are:

- Results are measured against what the utility wants to accomplish, specifically the type of customer service it wants to provide.
- Costs are evaluated in terms of what the utility wants to accomplish. Financial, specifically reduction in revenue requirements or increases in profitability.
- The ratio of the two is the measure of efficiency of any proposed project or action: those with the highest ratio of results to costs are the “winners.”

Similarly, reliability-based planning and prioritization (Chapters 13 and 14) assess capital spending on the same basic basis – evaluate “bang for the buck” and approve the most effective projects. Replacement of equipment, either in kind or with different types or sizes of units, can be evaluated from a similar standpoint (Chapter 15).

The standpoint of overall financial efficiency of the utility, spending, whether for new capital projects, equipment replacement, refurbishment and upgrades of older areas, preventive maintenance, repair and restoration, or for performance-based customer service penalties, is all the same and should be evaluated on the same basis: it absorbs revenues, lowering financial efficiency and potentially reducing the bottom line (profits). Therefore, spending on any of these measures should be coordinated. Money should not be spent on only one at the expense of spending on others, but only if and to the extent that spending on that item generates more improvement per dollar than on the other items.

Implementing RDM

RDM frameworks such as the various reliability-centered approaches have been so effective in improving results that they are standard practice in other industries where reliability is an important consideration (e.g. airlines, nuclear materials processing). Their effectiveness in the power industry is well proven. Therefore RDM, as described here, is highly recommended. It gets results and it minimizes costs.

Implementation can follow the approach discussed in Chapters 12 and 16. Evaluation and prioritization within each function in the company (i.e., preventive maintenance, operations, replacement and repair, capital projects, reliability-augmentation projects) is evaluated on the basis of “bang for the buck” as described in those chapters. Executive management enforces consistency in evaluation rules: all departments use the same definitions of “good results” and consistent definitions of costs, all compatible with upper management’s goals. Results are measured in terms of the type of customer service quality desired. Costs are evaluated in the manner that fits the corporate business model. This is not how many utilities currently evaluate results and costs: different departments use different priorities or rules, and cost bases are different and often incomplete from one department to another.

Utilities cannot take a trivial or “band aid approach” here. This is not an easy evolution. This is because total cost must be comparable across departmental and functional boundaries. For example, certain fixed costs that can be (but shouldn’t be) viewed as continuing are often not included in intra-department cost evaluations. The costs of vehicles and line trucks, tools, fittings, test equipment, and training are often not included in the costs of preventive maintenance projects as calculated at some utilities. These costs are a “given” and evaluation of cost effectiveness of different preventive maintenance actions is evaluated on the basis of only the labor, outside services, and parts required for one versus the other.

While that is a sound approach to deciding among projects in that department, it is not valid for extension and comparison outside the department. For example, how can these expenses truly be compared to contract tree trimming, in which the utility has no fixed costs of comparable type? Maybe the utility should buy fewer line trucks and tools, etc., and transfer funds to more tree-trimming. It can’t tell. Also it’s quite possible that analysis against other departments might show more investment in those resources is needed. The point is that cost-comparisons across departments and functions need to include all costs.

Considerable change in cost-computation and tracking – generally a shift to ABC (Activity-Based Cost evaluation), an accounting method that tracks all costs on a function basis – is a pre-requisite to effective RDM. The authors strongly recommend this since it is worth the effort.

Zero-base all expenses

Thus, the way to implement this includes zero-base budgeting, at least for a period of time until the overall pattern of spending for the “new paradigm” of RDM is identified and appears to be stable. All expenses, for all departments, should be justified on the basis of the same cost effectiveness and nothing should be sacred or excluded from this comparison.¹

Some departments will be very unhappy

Efforts to implement an effective and optimized RDM program *will* result in movement of money from one budget to another and from one function to another, often resulting in an allocation of funds to departments that is considerably different than the traditional pattern. The authors have worked with utilities within which the marginal effectiveness of spending differed by a factor of 8 to 1 from one department to another. The utility could improve the effectiveness of dollars already in its budget by a factor of eight, by cutting one department’s budget and giving the money to another department to fund expenses that could not fit within its original budget.

¹ The authors fully realize that the utility will be obligated to a good deal of spending that does not fit this evaluation directly. For example, equipment and construction are often required purely for safety reasons and to meet code, zoning, and other requirements. The authors do not imply that such expenses should be denied, just that they should not be excluded from consideration.

In the authors' experience something similar is common, although perhaps not to that degree in most utilities. Institutionalized spending patterns still in effect at some utilities were for the most part established under regulation and prior to many of the changes in technology, and need, that characterize today's power industry. It is almost certain that one or more departments will feel very cheated in this process, and that conversely a few will feel vindicated – they always knew their function was more worthwhile. This is a necessary aspect of cultural change which executive management must accept and manage.

17.4 STEPS TO MITIGATE AGING INFRASTRUCTURE EFFECTS

This section presents ideas and concepts – action programs - that will individually mitigate some of the aging infrastructure-related service reliability problems. Implemented selectively as a utility thinks appropriate in its case, they achieve synergies among both their improvements and spending and make the utility's operations and planning a bit more compatible with approaches it will have to take to deal with the aging infrastructure.

Voltage Reduction on the Primary Distribution System

Chapters 7 and 11 pointed out that the application of a voltage, aside of any current that flows, creates a type of “voltage stress” on insulating materials which eventually leads to deterioration, and over enough time, will alone cause failure. Voltage stress contributes directly to aging – other factors of age are multiplied by its effects. Unlike thermal stress caused by current flow, which occurs only during peak conditions and then only when peak loads last for a considerable period of time, voltage stress of roughly the same magnitude occurs all 8,760 hours of the year. In fact, voltage stress is less during peak conditions (because of an increase in voltage sags, voltage is lower).

The relationship between voltage level and deterioration rate (and hence lifetime) is highly exponential. In the case of some types of UG cable, a 5% reduction in applied voltage will decrease voltage-induced degradation by 35%. This means expected failure rate from voltage stress is reduced by 35%. This will not render a 35% decrease in overall failure rate, because voltage alone is not the only factor causing deterioration. However, it will make a substantial difference – perhaps as much as a 15%-20% improvement.

Therefore, one means to reduce the stress, and increase the expected lifetime of UG cables and the insulation in transformers and voltage regulators, etc., is to reduce voltage to the lowest level consistent with operation within required voltage standards. The utility can implement a type of Conservation Voltage Reduction (CVR) as is recommended (and required) by some state regulatory commissions for energy conservation purposes, but aimed at lifetime extension. This is implemented at the primary voltage distribution level, by adjusting the boost on transformers, tap changers, and voltage regulators at the substation.

Methods of analysis and implementation for CVR are applicable for this purpose. The only modification of traditional CVR methods needed for this

application is to analyze the system on a feeder by feeder basis to determine where voltage reduction would in fact have a worthwhile impact on equipment lifetimes. This will tend to be on feeders with substantial UG cable and/or very old, rather highly loaded service transformers (a good TLM program can be used to determine that).

Another factor that needs to be taken into account in assessing when, where, and how to apply voltage reduction for lifetime extension is the impact on revenues. Generally, demand levels are proportional to voltage. A 3% reduction in voltage will render roughly a 3% reduction in demand, although feeders vary greatly in this respect.

Due to the type of loads they serve, some feeders (usually those in industrial areas or residential areas with high saturation of AC and electric heat pump loads) would see no reduction in demand and revenues from a 5% reduction in voltage. A few others (mostly residential areas with high saturation of gas appliances) will see a 10% reduction in demand and revenues. Utility management needs to assess this impact, as it will occur on some feeders. It should be balanced against the expected savings due to the reduction in equipment failure rates and customer interruptions.

Planners and managers should not also use the demand reduction achieved in the pursuit of extended equipment life to justify deferral of required upgrades due to load growth. For example, one can calculate that if the voltage reduction rendered a 5% reduction in load, and load growth is 1.66% per year, then the reduction “bought three years” of deferral on any expansion needs, delaying the time until load exceeds capacity by three years. In truth this does not happen. The current that it carries thermally limits most equipment. Since a reduction in voltage can effect at most a proportional change in current, reducing needed investment in capacity will defeat the desired life extension effects of voltage reduction.

While the above may seem to involve a complicated set of evaluations, it is in fact straightforward. In addition, voltage-reduction can extend the lifetime and reduce the failure rates of equipment that are often both very expensive to replace and virtually certain to cause customer interruptions when it fails. It can be implemented at low or no capital cost, relatively quickly. A very comprehensive assessment can be done using traditional CVR planning methods (Willis, Tram, Powell, 1986), coupled with a detailed spreadsheet analysis of present and future costs and savings on a PW basis. The optimal operating voltage for each feeder is that which produces the most effective impact on the utility’s bottom line: expected costs due to voltage-stress related deterioration.

Adopt Reliability-Based Planning and Engineering Methods

Chapters 8, 9 and 12 discussed why modern power systems, particularly those in aging areas, are particularly prone to customer interruptions due to equipment failures, and how traditional engineering methods cannot always assure that all necessary and sufficient steps have been taken in their design to assure reliable operability at lowest cost.

The authors do not recommend that traditional planning tools (i.e., N-X contingency evaluation loads flows, etc.) be abandoned. Such tools are useful and cost little to use. They provide a good perspective on how the system will deal with each contingency and are often useful for examining what are called post-contingency conditions and determining if recommendations for special operating restrictions should apply after a failure has occurred.

However, utilities must augment their use of those tools with modern planning and engineering tools that work directly with expected reliability of service, tools that can assure measures taken to provide reliability of service are sufficient to reach the levels of customer service quality targeted by the utility. Such tools include a “reliability load flow” analysis of any of several types, and possibly a reliability-based system optimization (see Chapters 12–13).

This use of such tools will have three positive effects on the results that the utility obtains from both its system and all new capital investments:

1. *Improved customer service quality.* Application of these methods does result in better reliability of service. Planning and Engineering produce plans that are more robust with respect to contingency strength and tolerance to equipment failures, and devoid of sensitivities to relatively likely combinations of multiple failures. Very simply, the power delivery system is better designed from a reliability standpoint.
2. *Lower cost.* Reliability-based planning and engineering methods can be used to optimize cost effectiveness of capital spending for reliability. In fact, a considerable improvement in this regard is an almost inevitable result of their use. In a very real and practical sense, money spent on improvements can be optimized to buy the most “bang for the buck” in terms of improving reliability where that is most important to the utility. Money is both used more effectively and targeted more accurately at problem areas.
3. *Improved predictability of results.* The dependability of planning and engineering projections of performance from a reliability standpoint will improve dramatically. Traditional tools deal with reliability implicitly (as capacity margins for contingency backup). Modern reliability-based tools deal directly with reliability of service as a calculated design parameter and thus are much more accurate in determining what to expect in terms of results. Management can be more confident that it will obtain results and can better anticipate where and why problems might develop.

Use Budget-Constrained Planning Methods (Marginal B/C Ratio Prioritization and Pareto Analysis)

Generally, prioritization of capital projects, the selection and approval phase of the planning and engineering process, needs to be done with so-called Budget-Constrained Planning (BCP) methods (see Chapters 5 and 12, and Willis, 1997, Chapter 18). BCP methods should be used in order to fit within both the business model approach recommended here, and to permit valid cross-departmental comparisons of cost that balance capital and operating costs in an effective and completely valid manner.

Budget-Constrained Planning is an approach to prioritization, not engineering or planning method per se, and can be implemented along with planning and engineering tools at any level of analytical rigor and detail. It works best, however, when applied in conjunction with reliability-based methods.

The basic concept behind BCP was described with examples in Section 5.5. Alternatives for each project will be evaluated against one another and against alternatives for other projects on the basis of the same marginal bang/buck target ratio. This is exactly the same approach described for reliability-centered maintenance and used by utilities that practice that in Chapter 15.

Budget-constrained planning has proved effective for both large and medium sized utilities. It has successfully worked within Engineering departments that use only traditional tools and have very limited engineering resources, as well as those that used reliability-based evaluation methods in conjunction with fairly strong labor resources and skill levels in their Planning and Engineering departments. However, it works best in the latter environment.

The authors have seen BCP render capital-spending reductions of up to 30% while still permitting the utility to obtain necessary targeted levels of reliability performance in its system.

In many cases, executive management should expect that when BCP is implemented within the results-driven management structure recommended in Section 15.3, it will result in a noticeable (10% or more) increase in staffing and spending for the engineering and planning functions. BCP dramatically improves the results obtained from engineering effort (the marginal payback at one large IOU in the central United States was over 20:1!). Generally, money spent on more effort and better tools devoted to improving reliability and optimizing (cutting costs while improving performance) is very effective at obtaining results.

Optimize Against PBR Rules

Performance-based rates and penalties or incentives set down by regulatory authorities and/or customer contracts define precisely the costs and value to the utility of customer service performance. If interruptions of service cost \$1/customer-minute, then the distribution utility would be less than maximally prudent from a financial standpoint if it were to spend more than \$1 to reduce customer interruptions by one customer-minute.

Optimizing spending, engineering, and operation policies to minimize

expected pay-out/maximize expected gain under PBR is sometimes referred to as “gaming” the PBR rules, but it is not a game and there is nothing fundamentally unethical or wrong in doing so. Through PBR, the regulatory process identifies the value that customer interruptions have to the utility. If regulators are displeased with the results of a particular PBR formula, they can respond by changing the PBR formula to provide more or less incentive, or to target a different definition of “good.”

Utility Efforts under PBR should be directed at Two Goals

1. *Optimization of results.* First, the utility should implement the type of optimization – minimization of risk – described in Section 13.5. That example was based on an actual investor-owned utility case. There, a hybrid analytical simulation method, admittedly a fairly advanced, highly skilled method for engineering, was used to determine policies and prioritize spending in order to minimize the utility’s total payout under a PBR penalty formula. However, the basic concept will work to good extent with less sophisticated engineering methods.

Proper implementation requires that this be coordinated with the overall RDM approach: the definitions of reliability used throughout all departmental evaluations under RDM must be identical to the definition of reliability or service problems used in the PBR. Ultimately, the PBR sets the utilitywide marginal cost of reliability.

2. *Consistency of regulatory rule.* The utility should make one effort it can to influence regulatory policy, so that rates and PBR rules are consistent with one another. If, through PBR, regulators put a certain cost on reliability of service (or lack of it) they really should approve a base rate structure that gives the utility the required amount of money to implement these changes. Marginal cost of reliability sets the optimum decision-making point for RDM: it is the value used as the target in prioritization and its use does assure cross-functional optimality in decision-making value. That is essentially defined by the PBR rules.

But the budget required to implement reliability improvements up to that level is not necessarily a function of that marginal value. The marginal cost of reliability is the derivative (rate of change) of reliability versus spending. Required budget is the total spending up to that point. Two utilities can reach the same marginal cost with very different budgets, as shown in Figure 17.1.

The point is that total allocated spending, and PBR defined penalties, can be set up so they are consistent. Ideally they should be. The authors have seen cases where that is not the case, and typically quite controversial issues develop as a result. The reason is that in such cases the regulatory rules as a




Figure 17.1 Marginal cost of reliability. Basically defined (very much a function of) the PBR formula, does not necessarily correspond to any specific budget level. Optimized budgets for two utilities are shown here. A theoretical, but realistic, example shows that one needs \$62 million, the other - \$69 million, to work up to the same level of effectiveness in spending for reliability. The reasons for this are numerous. One system may be traditionally allowed (by utility and regulator) to be less reliable than the other and thus need more reinforcement effort. One system could simply be bigger than the other (more customers) and thus need more money. Or, work in one system could be more expensive, for any wide range of reasons, than corresponding work in the other system. The important point is that one utility here needs about 10% more allowable expenses to be covered by rates than the other in order to adhere to the intent of the PBR. That happens only through effort to assure rates and PBR formula are consistent.

whole can be “gamed” in the worst sense of the word. The utility can make decisions which are fully within the rules, yet that benefit some interested parties at the expense of others, in a way not intended by the regulatory process. This might mean “unfairly” favoring industrial and commercial customers over residential, or vice versa, generating extra profits at the expense of customer service levels the regulators expected to be attained, or favoring capital spending (investment that generates future profits) over operating expenses (which are only recovered). In such cases, the financially sound procedure for the utility to follow is to act in its own business interests (that usually means in favor of its stockholders) rather than its customers.

Reexamine and Revise All Standards and Guidelines

De-regulation, competition, changes in customer expectations, new planning and engineering tools, the operational IT tools discussed above, and PBR, all mean that the basis for priorities and decisions has changed dramatically since its standards, guidelines, and policies were last revised. The changes occurring in the period 1995 to 2005 are more dramatic than those that occurred during the five decades from 1945 to 1995. Despite this, many utilities are following basically the same approach they traditionally used. A few have not made anything but incremental changes since the early 1970s.

Make real changes

A rigorous and zero-based approach needs to be taken here. Innovation coupled with a good deal of skepticism both for the need to maintain traditional standards, and for the claims made for new ideas, needs to be applied throughout.

Outside help makes a difference

The authors strongly recommend that this examination be made in part by people who are quite apart from those people within the utility who use the standards and work with the policies. Utilities who perform this re-examination using internal resources only – professional personnel who “grew up” in the utility – tend to make only small changes. It is to be expected that people who have long used the existing standards both have faith in them and know little about competing possibilities. Those utilities who seek the input of experts from outside the company tend to examine more objectively the methods used elsewhere (including outside the United States) and ultimately make much more sweeping, and effective, changes.

Customer Service Quality and Cost are the Two Big Goals

During the 1980s and 1990s, electric distribution utilities began implementing large information systems in the management of their distribution facilities, engineering, and operations. This included AM/FM (Automated Mapping and Facilities Management), ALS (Automated Line Switching), AMR (Automated

Meter Reading), CAD (Computerized Automated Design), OMS (Outage Management Systems), WOM (Work Order Management, RDM systems (which means, in this instance, Results-Driven Maintenance scheduling), and integrated systems called DMS (Distribution Management System) that linked some or all functions together. Some of these systems (OMS) directly affect customer service quality. Others indirectly affect it and are mostly justifiable on the basis of cost reductions (AM-FM), or efficiency improvements (WOM).

Decisions on what systems and how much IT to implement, and how to use it most effectively, should be based upon a uniform customer-service focused basis, using exactly the same bases for “bang for the buck” as applied in all other departments and to all other functions.

This is fairly obvious in the context of previous recommendations in this chapter. What may not be obvious is the priorities assigned to these various software systems will change greatly depending on the importance given to the customer service aspects of their use. How their priority in the eyes of executive management will rank depends greatly on whether the utility’s focus is mostly on reducing SAIFI or SAIDI. Where aging infrastructures are a major concern, this depends on how advanced the problem has become. As mentioned in Chapter 1 and demonstrated with example in Chapter 15, usually degradation in SAIFI leads that of SAIDI by several years. This is due to a number of interactions of failure rates, utilization ratios, and contingency flexibility (or lack thereof) in a traditionally laid out power system.

Therefore, the priority associated with each of the IT systems shown in Table 17.3 will depend on whether the utility’s immediate concern focuses mostly on reducing SAIFI, SAIDI, or some combination of both. For example, OMS systems help reduce duration of outages much more than they can influence frequency of outages (they only come into play once an outage has occurred). By contrast, RCM (reliability-centered maintenance) software tends to work effectively against failure rate (hence SAIFI), as does improved engineering and protection analysis software (CAD), whereas ALS reduces SAIDI, and if quick enough, SAIFI too.

The reason that SAIFI can be reduced by ALS is that reliability definitions often distinguish between momentary interruptions, which are those less than some arbitrary limit of time, and those that are longer. Events less than this limit (the IEEE standard being considered at the moment is five minutes), are classified as momentary interruptions (MAIFI) and do not “count” as events in SAIFI or SAIDI calculations. In such a situation, SAIFI can be reduced noticeably if automated switching within five minutes restores outages that would take more than five minutes to restore. In such cases they do not “count” with respect to SAIFI statistics. Actual SAIDI as experienced by customers will also fall slightly due to the faster restoration capability of the system as a whole.

Table 17.3 Power Delivery System Related IT Systems and Their Interaction with Major Aspects of Customer Service and Utility Cost

IT Systems	SAIFI	SAIDI	Costs
Automated Line Switching (ALS)	x	X	
Automated Mapping/Facilities Management (AM/FM)*	x	x	X
Automated Meter Reading/Power Monitoring (AMR)		X	X
Computerized Planning and Engineering (CAD)	X		X
Dynamic Equipment Loading (DEL)	x	x	x
Feeder automation (FA)		X	x
Online Equipment Monitoring (OEM)	X	x	
Outage Management & Trouble Call (OMS)		X	
Reliability Centered Maintenance (RCM)	X		
Transformer Load Management (TLM)	X	X	X
Work Order Management (WOM)			X

Table 17.3 lists some common distribution-related IT systems and their interaction with SAIDI, SAIFI, and cost control. Size of the X indicates degree of involvement.

Transformer-Load Management Systems

Transformer Load Management (TLM) systems are one of many IT systems which can be applied to power delivery systems. However, they deserve special attention because they can be used quite effectively to improve customer service, they can directly contribute to lowered failure rates among service transformers, and they have the among the highest payback ratios of power delivery IT applications.

TLM programs correlate monthly metered sales to customers with the transformers through which the sales to those customers were made. High levels of sales relative to a transformer's capacity indicate a possible overloaded transformer, and put it on a "jeopardy list" that is sent to preventive maintenance, so the units can be checked and replaced if found to be overloaded. Good TLM programs use formulae based on type and number of customers, kWh sales, seasonal changes in load, and weather to estimate stress on each transformer and whether it is likely in jeopardy.

TLM systems produce very accurate results (about 90%) in the sense that units predicted to be overloaded are in fact overloaded, and few false positives (units predicted to be overloaded that are not). They also miss only a few overloaded units (estimated to be fewer than 5%). As a result, change out of units that will fail and have to be changed anyway can be scheduled at the utility and customers' convenience. This reduces cost and improves reliability.

Effective TLM “programs” are much more than just software. They must include:

- *Good procedures for use*, including monthly inspection of the jeopardy list by field crews and replacement of units found overloaded.
- *Proven feedback from the field*. Information on units that were predicted to be overloaded and prove not to be, and on units that failed due to suspected overloads but were not on the jeopardy list, should be used to adjust the TLM analysis formulae.
- *Adjustments* of the jeopardy list functions to accommodate the corporate-wide marginal cost of reliability.

Extending TLM to include performance management of service

The concept of using monthly operating records in an analysis of possible trouble spots can be extended in two ways into what is called a Distribution Performance Management (DPM) system.

First, the analysis is extended to include not just transformers, but the entire lower delivery system (laterals, service transformers, and service level circuits and drops). Using facilities data, as well as metered data, trouble records, weather and other “monthly” records, a DPM can infer where trouble (failure) is likely to occur, as well as where poor quality problems (excessive voltage drop) is likely to occur in any part of that system. Really “smart” programs even use the results of primary distribution voltage analysis to determine the voltage stress on service transformers. This adds only a slight amount of improvement in the resulting analysis but in an integrated software system, costs next to nothing to add.

Secondly, the reporting of a DPM system is extended beyond what are traditional TLM reports, which include only the jeopardy list and perhaps also a list of severely under-utilized transformers (those where no-load losses are high enough to justify switching out to a much smaller size). DPM typically adds a series of management reports that include:

- *Predictive jeopardy lists* – a “smart” program can look ahead several months. The utility can schedule during the spring change-outs that need to be made before summer.
- *Asset management lists* for the service transformer set, including utilization of units, average monthly loss of life (during peak months the average unit may lose more than one month’s life, during off peak, less, etc.), predicted failures and replacements over the entire set, etc.
- *Predictive trends* about future failures, replacements, loss of life, and low voltage events.




Figure 17.2 The value obtained per unit of investment in a transformer depends tremendously on the way it is used. Here, the ratio of the value/cost for two 32 MVA transformers is shown as a function of the peak load level it will be used to serve. Solid line shows the transformer example from Section 15.4. Value varies by 22% over what could be called the useful range for this transformer (24 – 41 MVA – those loads over which value > cost). Optimum loading is 32 MVA. The dotted line shows the curve for the same type of 32 MVA transformer but at a different site with a different annual load curve shape. Value curve obtained is different as is the optimum peak loading, 36 MW.

TLM and DPM programs often have very high economic paybacks – on the order of less than a year. The major expense is in developing a reliability data link between the customer records and the transformer records (tying each customer to the correct transformer).

17.5 THE AGING INFRASTRUCTURE MANAGEMENT PROGRAM ITSELF

Expect to Need to Use Novel and New Approaches

Despite the fact that it took decades to get into the situation, aging infrastructures are a new and very different problem than those the utility has had to manage in the past. Management and professionals should expect that the ideas, methods, and tools that it ultimately uses will be different, too. An open mind, a willingness to experiment with purpose, and an innovative mindset are needed.

Data are Always Imperfect and Messy

Throughout this book, the authors have discussed the practical realities of data for aging infrastructures: its never enough, its never clean, and its always messy to deal with. Chapter 16's case study make this particularly clear. CL&P's pole asset management team had gathered all the data it could on the longest possible history of the company's pole ownership, some of it coming from innovative places. A considerable amount of manual labor by interns and co-op students was involved in transcribing old records. Ultimately, some data were more useful than others, but a good deal of what proved valuable in the long term was data it gathered "just because it could" up front.

So the lesson is: look wide and far for all the data that *might* be useful, now or in the future. Put measured and reasonable efforts, including labor into obtaining it. Don't obsess on cleaning it up: dirty data (with gaps and missing elements) is dirty data, but still worth something if used with analytical methods that are relatively insensitive to such quality issues.

Long-Term Historical Data is Extremely Useful

One cannot expect to build useful and defensible 60-year models from only 15 to 20 years of data. Some means of finding a few data points from long-ago needs to be found, even if the data's "fuzziness" might be substantial (analytics have ways to deal with that). A key element of success in Chapter 16's case study was that CL&P's project team tried to get the longest historical database they could, even if the older elements of it were somewhat incomplete and less detailed than recent data history.

Don't Assume Old Data are Not Accurate

In that case study, CL&P found data points for some measures of its pole ownership from paper records made in 1933 and 1947. There is a temptation to dismiss old records that are potentially not accurate or important. Although gathered and maintained in a time before computers and precise instrumentation like today's, the data from those past years were managed by people who were, as far as is known, just as dedicated and just as concerned with accuracy as today's work force. Data from long ago have to be kept in perspective, but are very likely as accurate as today's counterparts.

"Bulk" Labor Spent One Time on Data Can Be Worthwhile

A project team should not shy away from just rolling up its sleeves and slogging through data, if the effort looks worthwhile and is one-time only. Here, co-op students and interns can be very useful, as was the case in Chapter 16, where CL&P had such temporary employees transcribe thousands of pages of ancient paper records into a computerized file. Total burdened cost in that example was below \$50,000 for this effort, but the data proved many times more valuable in the pole management and other equipment category effort for years afterward.

Advanced Analytical and Pattern Analysis Methods Are Useful

Most utility engineers and managers have a sound foundation in basic statistical analysis and are confident it can extract what is needed from data. This is not true: basic statistical methods often miss a lot of information hidden in masses of messy data.

But it is the flip side of this assumption that can be fatal to a project team: many utility planners, engineers and managers assume that if standard statistical methods fail to find meaningful information in the data – if regression finds no correlation, if statistical analysis shows no tight distribution, then the data are useless. This is not the case in *most* situations.

Analytics – what the authors often refer to as “statistics on steroids” – can often find surprising amounts and qualities of patterns, correlations, if-then-but-for relationships within large masses of data, such as exist for utility infrastructures. Exactly what “analytics” is and how and why it differs from and goes beyond normal statistical analysis is far beyond the scope of this book. However, in summary analytics involve pattern recognition, logical analysis, and non-traditional statistical methods extended to cover monstrous amounts and dimensions of data. It often finds surprising amounts of information hidden within masses of data as well as surprising facts about non-linearity and special circumstances that are extremely useful.

Advanced Analytical and Pattern Analysis Methods Are Not Enough

Even when using the very best methodology on a considerable body of data, analysis, even advanced analysis, is not going to do a sufficient job of identifying and prioritizing equipment and program focus. The analytical results in the example given here are about as good as they get. Here a very comprehensive (“advanced”) type of analysis was able to build a statistical model that was accurate at predicting trends for the entire population, and as a guide for spatial and location-specific inspection and prioritization.

Chapter 16’s project is a case in point. An information theory based analysis indicates that about 40% of the knowledge ultimately used by the utility to make decisions about CL&P’s wood pole infrastructure management program was generated by the historical data analysis and the models developed from it. Pole inspection and tracking provided roughly the same amount of knowledge. The Physical Pole Examination Program (PPEP) generated about 12%, and the remainder of information used around 10%, came from other sources. The authors believe this is representative of what to expect in most similar studies of large sets of “commodity” level power system equipment.

In general, analysis of historical data is most useful because it determines a fully consistent set of models or information for:

- a) failure rate curves (as in Figure 16.11)
- b) equipment demographics model (Figures 16.12 and 16.14 – 16.16)

c) long-term base case trends (Figure 16.17, etc.)

The last, c, is the cost basis for determination of savings from any program that is developed out of an aging infrastructure initiative, and therefore a cornerstone of the program. A rigorous, traceable method and process from data to conclusions is needed for credibility and defense of the recommendations.

Analysis and Models Cannot Neglect NCR and Retirement

Failures were actually only one of three causes of poles “leaving the asset base.” In Chapter 16’s case study, for poles newer than about 40 years, it is more likely a pole is removed because of NCR (non-condition based) reasons like relocation, or retirement, than failure. NCR and retirements are less an issue for some types of equipment (i.e., MV breakers) than for others. But for poles and other distribution equipment they are often, as was the case here, *larger* effects than failure rate in the early years of modeled lifetime characteristics.

More generally, NCR and retirement are on the order of the failure rate and thus have to be taken into account (see Appendices).

Focus on Developing a Good Basic Model First; Do the Details Later

The authors recommend an approach like that used in Chapter 16. Initially, the project team put aside all the small “ifs, ands, or buts” so it could focus on only the basic framework of its first model. It assumed failure, NCR, and retirement rates would be stationary over time, although it knew that strictly speaking they probably were not. It knew effects were unlikely to be identical in all parts of the system, but neglected that initially. It left out a host of little details that were no doubt real, but would get in the way of quickly determining, testing, and setting up that first basic model. Eventually, the project team went back and re-examined all those details, folding many of them into its model. But this approach permitted it to get a good, basic, workable model done first.

An Asset Base Model Is A Useful and Needed Planning Tool

In Chapter 16, CL&P built failure/survivor Asset Base Model first, a model that could analyze and project trends for its entire pole base, but one that had *no geographic element* to its data: it could project how many poles were at risk, etc., but not tell CL&P where to find them. In Chapter 15’s company discussions, many of the companies had detailed models of the equipment they owned and models, structured around their needs.

Since CL&P wanted to manage its aging equipment to maximum business value, it needed a good Asset Base Model and used the sustainable point survivor model approach (Appendix A) populated with the data developed as described in the first half of Chapter 16. Initially, it was developed as a spreadsheet. Within a year, however, it was converted to a SAS basis in order to use more advanced analytical tools. Functions representing failure rate, NCR rate, and retirement rates as a function of years since installation were found that best fit the historical

record of pole ownership. When these functions were used in this model, the model replicated the CL&P history and ownership statistics with little error.

This is the appropriate first step. Build the model that appraises the problem and puts it in a corporate business context first. The Asset Base Model did, permitting CL&P management to understand the scale and characteristics of the aging pole problem before them, how soon and where and how it would effect their business, and what the scale and characteristics of potential solutions to it would be. At that point, work on obtaining that critical geographic information could be justified.

Geographic Information Is Needed for Implementation

Useful as an Asset Base Model is, eventually a utility will need geographic pole locational data to do anything meaningful. Both Chapter 16's CL&P, and the two electric utilities and university power system owner discussed in Chapter 15, all used GIS technology to target and manage their aging infrastructure programs. In Chapter 16, CL&P use a non-GIS model that treated all 437,226 poles it owned as one big, location-less set. That Asset Base Model was an essential first step. It could tell CL&P that it had 19,240 poles that were over 65 years old, and that these poles would cause an expected amount of trouble and cost over the next ten years. It could answer myriad questions and run dozens of scenarios in supporting management in understanding the problem and the range of options open to them.

But the Asset Base Model could not tell CL&P where to find those poles in order to inspect, service, or replace them as needed. The utility had to go on to develop a geographic and database with three capabilities:

- It could locate every pole in its system: it could find any particular pole in its system database out in its system, and it could find the data for any pole out in its system in its database.
- It could relate each pole record in its geographic pole database to the asset base's analysis of the pole population: this pole is among the 12% thought to be the oldest in the system, etc.
- It had statistics and other measures of the dependability of its analysis. It knew its data were not 100% reliable, but it had mathematics to tell it how reliable and when it was reliable enough.

Geographic Information Can Be Developed Even When Not Directly Available

The authors deliberately picked a case study for Chapter 16 in which the utility had no pre-existing way to tie pole records and age/condition measures of the Asset Base Model together, in order to show how this can be done in such circumstances. The utility there used municipal tax records, street and road department data, etc., to determine when each area of the city was first built out, then correlated that with pole locations and refined the results after a good deal of hard work and statistical analysis. It reasoned that the electric system in an area

would date from about the time the first buildings were built there. From this data, and with a lot of statistical work, it built up an estimated age for each of the 437,226 wooden poles in its distribution systems.

Strictly speaking, that data developed in this way proved only good enough to guide an inspection program. The location age estimations developed in this innovative GIS program were a bit “fuzzy,” but useful enough as a guide to where to inspect in order to tie down condition well enough to be effective in managing pole quality. The inspection could not have been justified or cost effective without this locational data, imperfect as it might have been. This imperfect data would have not led to good results without inspection to help focus mitigation and management programs more effectively. Both were mutually supporting elements of the overall solution to “where are the bad poles?”

Physical Examination of Equipment is Useful, if Focused

Very few aging infrastructure management programs will be effective and near-optimal as to economics without the use of inspection and tracking of condition. This includes testing and research to build up a base of knowledge about the seriousness of various grades of deterioration found by inspection, etc.

In Chapter 16, investigation and analysis of actual poles was carried out on nearly 1,500 poles that had been withdrawn from service due to failure or NCR reasons. These poles were cut apart and the level of deterioration and loss of strength recorded and analysis. Both the average level of deterioration as a function of estimated age, and the standard deviation of deterioration as a function of age were determined and used to refine estimates of condition for every pole in the system. An important point here is that, by the time the utility carried out this work, its project team knew *exactly* what type of data it needed and how it would use that: the work was done specifically too look for information to support the Asset Base Model and Geographic Pole Model.

Inspection and Tracking Are Important

Strictly speaking, inspection and tracking do nothing directly to solve the problems caused by aging equipment. But the information they provide makes programs and tools that do solve the problem *much* more effective. Inspection and tracking are examples of the value of information. They also show the importance of the “new” element. What matters is equipment that needs attention now. Models and databases can come close to providing this, but inspection and tracking focus attention as much as possible. No effective aging infrastructure management system can exist without some element of these activities.

Run to Failure is Often the Only Cost-Justifiable Option

This case study determined that proactive, economically justifiable replacement of most poles in a power system will just never happen. A case just cannot be made for replacement a pole earlier than its end of life in most situations. Any practical, realistic program has to deal with this fact: the “optimum” case for

almost all commodity equipment like poles and service transformers is to simply run it to failure and deal with the failures when they occur. The owner/operator has to find the most effective way possible to deal with the levels of operating issues that that will be created, just as the utility has to find a way to live with the business consequences of run to failure. There are no more cost effective options.

Sustainable Point Is a Useful Management Concept Not Just a Useful Modeling and Analysis Concept

The sustainable point is a simple concept. Eventually, given that current operating policies and inspection, service, and repair and replacement policies are followed, the system will reach a stable point of so many failures, so much cost, such and such a SAIDI, etc. This is a vitally useful concept as an operating framework to use to work on aging infrastructures.

- 1) *Things don't keep getting worse forever.* The system is headed toward a stable situation.
- 2) *If the current situation is better than the sustainable point, then in some way, the system is living on borrowed time.*
- 3) *The sustainable point can be moved* by changing operating policy, inspection, and service and repair replacement policies. In fact, the only way to move it is to change something: the challenge before management is to answer the question, “What do we change, and how?”
- 4) *Things change slowly,* period. Escalating failure rates and repair costs will continue to escalate, but slowly, if inevitable. No catastrophic “cliff’s edge” exists. The system will continue to trend toward the sustainable point, but take years to get there. And all solutions act slowly, too.

The Base Sustainable Point Will Not Be Acceptable

The authors have never seen a base case in which the sustainable point – the stable point of ongoing failure rates and expenses to which the system will eventually deteriorate – that is acceptable from a business standpoint. Medium voltage breakers as maintained at some utilities come as close as anything seen. But the case given in this chapter is a representative example. CL&P was looking at a doubling of age-created problems over the next twenty years, and a thirty to fifty year period during which it gradually would reach the sustainable point in most of its equipment categories, at which point customer reliability, breakdown and failure, and safety issues are all *three to seven times as bad* as today’s.

Driving Straight for the Sustainable Point Is Usually Non-Optimal

By changing operating policies, O&M practices, and service and repair plans, a utility can move the sustainable point to something that it finds acceptable. In fact it has a good deal of choice, depending on what it wants, how it wishes to operate, and, of course, how much it wants to spend, or not. In all cases the authors have seen, there is no doubt that, eventually, spending will have to rise, *a lot*.

But in no case the authors have seen is the optimum business case for the utility drive to the new, long-term sustainable point immediately. In almost all practical cases, the sustainable point, with its catastrophic levels of reliability problems and operating costs, is so far into the future that time-value money economics attaches little value to its business consequences. What is important is the noticeable erosion of business performance over the next few years that current trends will create: a 4% increase in SAIDI contribution next year is more important than a times-five SAIDI contribution increase 42 years from now. “Now” is the time period driving action now.

Leave the Long Term Solution for the Long Term

“Now” dominates spending decisions at utilities and most other businesses, and in almost all cases it should. The focus of any good plan for aging infrastructures will address current and short-term effects cost effectively, while contributing in some way to the long-term eradication of the problem. In the examples covered in Chapters 15 and 16, the only important longer-term aspect of the aging infrastructure programs was that their short-term programs were economically close to optimum, and that they were not only *not* worsening the long-term problem, but actually contributing to their solution, even if they were not entirely solving those problems.

Chapter 16 showed in detail how the wood pole management program developed, there was only the front end of a decades-long delaying action that gradually moves the sustainable point, even as it focuses at any point on “what can we do over the next five years to keep things under control in a cost-effective manner?” But this is good enough – in a practical world, it’s about all the utility should expect.

The Final Product is a Criteria-Driven Program

CL&P’s program in Chapter 16, as well as the three utility examples in Chapter 15 that involved managing aging of an infrastructure – Big State Electric, Mid-State Electric, and Big State University system – all ended up with a fairly straightforward, easy-to-implement program for execution. Pole replacement is driven by criteria based on age, inspection, and other factors. Poles meeting the criteria are replaced. Those that do not, are not.

The intricacy and complexity, and labor and cost, are all in the data gathering and vetting, analysis, model building, and application that determine

the criteria. The criteria can change over time. It is a function not only of the pole population and its characteristics, but of the utility's business situation, needs, and strategy.

This is a recommended approach. Isolate whatever complexity and involved analysis is needed into a process that sets simple and easily executed criteria. Make that process capable of varying the criteria and matching it to business needs. The actual asset management program, then, will be straightforward to execute, and not require complex models and data analysis.

Mitigating the Consequences Often Works Better Than Solving the Problem

An effective tool in the case study "solutions" given in Chapters 15 and 16 was to change the system's *sensitivity* to the problems failing equipment create. In Chapter 16, CL&P's use of a Smart Switching program reduced the consequences of wood pole created outages. Big State Electric made operating system changes to reduce the reliability impacts of aging equipment. These changes were *part* of the most cost-effective solutions to short term business performance issues, even if they did nothing directly to "solve" the problem of aging poles. Mitigating the symptoms is not the full solution in most cases, but it is nearly always an element to be mixed with actual age-control programs, as was the case in those examples.

Continuous Improvement

Annually, the performance of the entire aging infrastructure management program should be reviewed, the Asset Base Model and the Geographic Locational Model and all planning and decision-support tools and processes should be examined as to efficacy and possible improvements that could be made, and improvements should be made where they appear to be cost effective. All the assumptions made and small factors left out initially were, in time, re-examined.

Aging Infrastructures Never Go Away

The examples and discussions throughout this book all make for the same "big picture" conclusion: aging infrastructures will be with the utility forever. No matter that once upon a time every element in every utility system was new: any modern power system has a large portion of very old equipment. When managed well, an aging infrastructure continues to have elements within it that age, and continues to have some elements within it that are quite old. Among the most cost-effective strategies available to "fight" the effects of aging infrastructures is life extension – make the equipment last longer so that eventually it becomes even older.

Therefore, a well-designed and executed aging management program will extend life, and keep old equipment longer. Optimization will result in this management program reaching a sustainable point: the infrastructure will have that portion of old equipment on an ongoing basis.

Therefore, Aging Infrastructure Management Programs Are a Permanent Fixture for Utilities

Today the analysis, planning, and maintenance programs needed to manage aging infrastructures are new to many companies – activities forced on them by the realization that they must manage an aspect of their systems they did not have to address in the past. But these challenges are not going away: the aging infrastructures will be with the utilities forever, and thus the programs needed to manage them will be, too. For this reason aging infrastructure management programs should not be treated as one-time-only projects. Electric utilities need to weave the tools, processes, and resources needed to handle aging into their mainstream institutionalized processes and organization, for they will be a necessary part of their future forever.

BIBLIOGRAPHY

- Brown, R. E., *Electric Power Distribution Reliability – Second Edition*, CRC Press, Boca Raton, FL, 2009
- Brown, R. E., *Business Essentials for Utility Engineers*, Marcel Dekker, Boca Raton, FL, 2010
- Brown, R. E. and J. J. Burke, “Managing the Risk of Performance Based Rates,” IEEE Transactions on Power Systems, Vol. 15, No. 2, May 2000, pp. 893-898.
- Brown, R. E. and Bruce G. Humphrey, “Asset Management for Transmission and Distribution,” IEEE Power and Energy Magazine, Vol. 3, No. 3, May/June 2005, pp. 39-45.
- Brown, R. E. and M. M. Marshall, “Budget Constrained Planning to Optimize Power System Reliability,” IEEE Transactions on Power Systems, Vol. 15, No. 2, May 2000, pp. 887-892.
- Brown, R. E. and J. H. Spare, “Asset Management, Risk, and Distribution System Planning,” IEEE Power Systems Conference and Exhibition, New York, NY, Oct. 2004.
- Brown, R. E. and J. H. Spare, “Asset Management and Financial Risk,” DistribuTECH Conference and Exhibition, San Diego, CA, Jan. 2005.
- Butera, R., “Asset Management For The Distribution Pole Plant - Closing The Performance Gap Between Traditional Maintenance And Asset Management,” IEEE Power Engineering Society Summer Meeting, July 2000, pp. 561–565.
- Center for Petroleum Asset Risk Management, “Section 3: A Brief Background of Real Asst Risk Management,” in Petroleum E&P, 6 pp.
- Humphrey, B., “Asset Management, In Theory and Practice,” *Platts Energy Business & Technology*, March 2003, pp. 50-53.
- Markowitz, H., “Portfolio Selection,” *Journal of Finance*, Vol 7., Number 1, 1952, pp. 77-91. (This paper covered work for which Harry Markowitz was eventually awarded the Nobel Price in Economics).

- McCullough, J. J., "GSU Transformer Health Monitors Combining EMI Detection with Acoustic Emissions Measurement," *Proceedings of the 38th Annual Frontiers of Power Conference*, Oct. 24-25, 2005, Oklahoma State University.
- S. F. McFarlan, "Portfolio Approach to Information Systems," *Harvard Business Review*, Dept-Oct. 1981, pp. 142, 150.
- Morton, K., "Asset management in the electricity supply industry," *Power Engineering Journal*, Vol. 13, No. 5, Oct. 1999.
- J. Ostergaard , J. and Jensen, A.N. Can we delay the replacement of this component?-an asset management approach to the question, CIRED 16th International Conference and Exhibition, June 2001.
- Pareto, V., *Manuale di economia politica*, Centre d'Etudes Interdisciplinaires, Université de Lausanne, 1906.
- Philipson, L. and Willis, H.L., *Understanding Electric Utilities and De-Regulation – Second Edition*, CRC Taylor and Francis, Boca Raton, FL, 2005
- Seevers, O. C., *Management of Transmission and Distribution Systems*, Fairmont Press, Lilburn, GA, 1995.
- Timperley, J. E., "EMI Diagnostics Detects Transformer Defects," *Proceedings of the 2005 Doble Conference*, Boston, MA, April 2005.
- Van Schaik, et al., "Condition based maintenance on MV cable circuits as part of asset management; philosophy, diagnostic methods, experiences, results and the future," CIRED 16th International Conference and Exhibition, June 2001.
- H. L. Willis, *Spatial Electric Load Forecasting*, Marcel Dekker, New York, 1996
- Willis, H. L., *Power Distribution Planning Reference Book – Second Edition*, Marcel Dekker, New York, 2004.
- Willis, H. L., Tram, H. N. and Powell, R. W., "A Computerized, Cluster-Based Method of Building Representative Models of Distribution Systems," *IEEE Trans. PAS*, March, 1986, p. 321b.
- .

Appendix A

Notes on a Spreadsheet-Based Model for Aging Trends

A.1	Introduction	717
A.2	The Basic Concept: Incrementing an Age Distribution	719
A.3	Multi-Sheet Structure of a More Practical Model	723
A.4	Fitting and Calibrating a Model	725
A.5	Real-World Data Situations: Getting a Solution	730
A.6	Comments on Application	735
	Reference	737

A.1 INTRODUCTION

This appendix presents the recommended layout for an electronic spreadsheet template that implements a simulation of electrical equipment set's aging/deterioration over time. This first section gives some basic definitions used throughout this Appendix as well as throughout the book. Section A.2 describes the central concept– the staged use of rows representing years and columns accounting for annual age distributions. Section A.3 outlines how the authors extend that structure to make a basic model suitable for practical application. Section A.4 gives a simple but complete example of how these models are fitted to historical data.. Section A.5 discusses real-world data issues and how to work around them. Finally, section A.6 comments on how these models are best modified or structured for different needs and types of equipment, etc.

This presentation assumes the reader is generally familiar with both the use of electronic spreadsheets, and their “programming.” No software features beyond those available in standard MS Excel® 2007 and no user skills beyond those possessed by the average utility power engineer are assumed.

No Model Included Here

For readers who wish to apply and tailor a model like that described here, there is no substitute for creating it themselves. Therefore, no electronic file or program listing is included here. The discussion and figures given here support readers who wish to do this: a person can build the basic model described in diagrams in Section A.2 by following along with the descriptions given there, checking results against those shown.

For readers who do plan to develop their own spreadsheet models, the authors strongly recommend that they first read this Appendix thoroughly, then read Chapter 16's case study on the use of a model very similar to that described in Section A.3 to gain an understanding of how and why particular features are needed, and only then begin to develop their model.

Model of "It Could Be Anything"

This Appendix describes the concepts and spreadsheet structure built by the authors and used to create all the examples throughout in this book, including the Widgit examples in Appendix B, the various other examples in chapters of the book, and to present the Tables of results in the CL&P pole case study in Chapter 16.¹ As it is described here, it models "age" as a proxy for condition, and is described with respect to modeling wooden distribution poles, as for the example in Chapter 16. However, only slightly more intricacy and structure is needed to model condition distinctly from age, if desired and appropriate (and supported by available data). The method presented here has been applied to many types of equipment. However, considerably more is required to model multi-component equipment like power transformers, etc. Comments near the end will comment on how to structure "multi-component" aging models for such equipment.

Poles were selected as the comprehensive example of the use of this model, given in Chapter 16, because they are both the simplest equipment and yet also the most complicated situation to handle from a modeling and application standpoint. They are simple in terms of *equipment* complexity: poles have but one part (the pole itself – here the authors do not study cross-arms and associated hardware as part of the pole). As a result poles generally can be modeled as having only one failure mode (the pole fails structurally). To see why this is "simple," one merely has to look at the challenge asset managers face when modeling a power transformer well. That requires a model that acknowledges and *explicitly* tracks several component parts of the transformer: the windings and their deterioration as distinct from, say, the rusting of the case or the condition of the bushings, etc. The final section of this Appendix will

¹ The spreadsheet model used by CL&P in the first stages of its pole management study was built by one of the authors and varied slightly from that shown here. The model described here was used, with the CL&P data shown, to re-analyze and produce all the results given in Chapter 16.

address the application of the basic model to such multi-component models.

But while poles are simple with regard to number of components and failure modes, they are a complicated challenge to study for several reasons. First, there is a lot of “noise” in the system. As discussed in Chapter 16, in many cases and in some timeframes, more units of equipment are withdrawn from service for non-condition related reasons than for failures: a successful model must model other ongoing processes and account for how those interact with the utility asset base. Commodity distribution equipment like poles and service transformers are by far the most affected by these other processes: transmission towers, power transformers, etc., may have their complexities, but often they can be analyzed without having to acknowledge those other, non-condition-related processes.

Furthermore, utilities usually have reasonably good data for major equipment like power transformers, but little for poles. The typical situation for poles is to have next to nothing about the history of specific poles or the age distribution of the population as a whole. When dealing with distribution level equipment, like poles, data is sparse, and what there is, is most often not just of poor quality, but of *unknown* quality. Statistical models, like that described here, must be applied in a sometimes painfully complex process of vetting data, as described in Chapter 16.

Finally, with distribution equipment, there are usually *a lot* of units to track and study. Chapter 16’s example concerns a medium size utility, one with fewer than one million connected meters, yet it has nearly half a million wooden distribution poles to analyze and manage. Numerical models that depend on statistics and analytics, and that make up for lack of data by tracking patterns, are the order of the day. The model described here, although ultimately quite simple in concept and structure, fits nicely into that type of environment.

Years versus Age

Throughout this Appendix the words “age” and “year” will have very specific meanings with regard to the spreadsheet model.

“Age” is the period time a pole or group of poles has been in service.

“Year” is, without exception, the chronological year.

Thus, a pole installed in year 1913 has an age of 87 in 2010. A pole installed in 1960 had an age of 1 in 1961, 2 in 1962, and so forth.

A.2 THE CONCEPT: INCREMENTING AN AGE DISTRIBUTION

Figure A.1 shows a spreadsheet row that holds counts of poles by their age: there are 231 new poles (poles less than one year in service), 284 that were new a year ago, 211 that of age 2, etc. The row is considerably longer than shown, going out many columns to accommodate counts poles of much greater ages (to

	A	B	C	D	E	F	G	H	H	I
1										
2										
3										
4	V-Year-V	Total	Age-->	0	1	2	3	4	5	6
5	2010	1613		231	284	211	198	345	344	

Figure A.1. Row 5 holds counts of poles by age for year 2010 starting in column D, and computes the sum (total of all poles) in column B. The row extends much farther than just these first seven years of age (from new to six years), although there are no poles over five years of age in this example. There are a total of 1,613 poles in 2010.

an age of 100 in this case) but in the sample shown, there are only 1,613 poles, of ages zero (less than one year old) up to five years. In Figure A.2, rows holding the “failure rate as a function of age” have been added at the top. In this example, failure rate is just defined as

$$\text{Failure rate in percent} = \text{age} + 1\% \quad [\text{A.1}]$$

But in an actual model, one to be used for aging infrastructure analysis, the cells in row 2 would hold a logarithmic, polynomial, or iterative formula computing failure rate as a function of the age and/or earlier values, or something similar.

Row 6 shows the central concept of the model: similar counts of age for year 2011 are *computed* from those for 2010, and the failure rate, using the formula:

$$\text{count (age, year)} = \text{count (age-1, year-1)} \times (1.00 - \text{Failure rate (age-1)}) \quad [\text{A.2}]$$

In other words: the computation called for in cell E6 is “=D5*(1-D\$2)”. That is copied to the right all the way to the rightmost column in the row. This results in a computation of the values shown for 2011 in row 6. Essentially the values in row 5, representing 2010 counts, are moved down one row and to the right one column, with every value decremented by the amount of failures it would have had in 2010. Only 229 of the 231 poles that were new in 2010 survived to be one year old going into 2011, only 278 of the 284 poles that were one year old are now two years old, etc. And there are now only 1,553 poles, of ages from 1 to 6 years, as shown, in year 2011.

	A	B	C	D	E	F	G	H	H	I
1			Age-->	0	1	2	3	4	5	6
2	Failure Rate (age) -->			1.00%	2.00%	3.00%	4.00%	5.00%	6.00%	7.00%
3										
4	V-Year-V	Total	Age-->	0	1	2	3	4	5	6
5	2010	1613		231	284	211	198	345	344	
6	2011	1553			229	278	205	190	328	323

Figure A.2 Here, rows 1 and 2 hold failure rate as a function of age. Row 6, the counts by age for year 2011, is *calculated* from row 5, the counts for year 2010, as explained in the text.

	A	B	C	D	E	F	G	H	H	I	J	N	N
1			Age-->	0	1	2	3	4	5	6	7	8	9
2			Failure Rate (age) -->	1.00%	2.00%	3.00%	4.00%	5.00%	6.00%	7.00%	8.00%	9.00%	10.00%
3													
4	V-Year -V	Total	Age-->	0	1	2	3	4	5	6	7	8	9
5	2010	1613		231	284	211	198	345	344				
6	2011	1613		60	229	278	205	190	328	323	0	0	0
7	2012	1613		74	60	224	270	196	181	308	301	0	0
8	2013	1613		86	73	58	217	259	187	170	287	277	0

Figure A.3 Here, new poles to make up for failures have been added in Cell D6, which holds the formula “=SUM(E6:T6)-SUM(F7:T7)”:60 units failed in 2010 and so are less than one year old at the start of 2011. This formula has been copied to the right and down, resulting in computations of age distributions for years beyond 2011. Note the spreadsheet keeps the total the same each year, at 1613, maintaining the same number in service. The oldest pole is now 9 year old rather 5 as in 2010.

In Figure A.2, only 1,553 poles remain going into 2011, because 60 failed in 2010. The owner probably needs 1,613 poles in order to operate his system, and so would replace those that fail. Figure A.3 shows this sheet with cell D6 holding the computation “=SUM(D5:Z5)-SUM(D6:Z6).” The number of original poles that failed in 2010 was 60, so an equal amount is added in the “new” column in 2011: they all less than one year old (age zero) at the start of 2011.

Figure A.3 also shows row 6 copied down to compute age distributions for years farther into the future. The result is a spreadsheet that estimates the age distribution of the population for future years, starting with the age distribution in 2010 and using the failure-rate-as-function-of-age values in row 2. One can follow a group of poles of a certain age through time by moving diagonally. For example the 231 poles that were installed new in 2010 (cell D6) become progressively 229, 224, 217, and only 214 in number by 2014.

That is all there is to the basic concept

Thus, the central concept of the model is: columns represent age, rows years, and one computes survivors from year to year and tracks them. The basic spreadsheet shown above needs bolstering to be practically useful, but this is the key concept. Models using exactly this structure were the “first generation” of models of this type, and were successfully applied to analyze transmission lines and structures in several systems (Willis et al., 2004).

Additions for Practicality

Figure A.4 shows the example from A.3 with four columns added. The first two, B and C, concern growth of the asset base. The user of this model may wish to track a set of assets as it has additions made over time. This is also a necessary feature of the model for fitting to historical data, as will be explained later.

	A	B	C	D	E	F	G	H	I	J	K	L	M
1						Age-->	0	1	2	3	4	5	
2						Failure Rate (age)-->	1.00%	2.00%	3.00%	4.00%	5.00%	6.00%	
3													
4	V-Year -V	Growth	Growth	Failures	Avg Age	Total	Age-->	0	1	2	3	4	5
5	2010	1.10%	18	60	2.7	1613		231	284	211	198	345	344
6	2011	1.20%	20	74	3.5	1631		78	229	278	205	190	328
7	2012	1.30%	21	86	4.2	1650		93	77	224	270	196	181
8	2013	1.40%	23	98	4.8	1672		108	92	76	217	259	187
9	2014	1.50%	25	108	5.4	1695		121	107	90	73	209	246
10	2015	1.60%	28	117	5.8	1721		133	120	105	88	70	198

Figure A.4 Three columns (B, C, D, E) have been added to permit the model to apply growth, to report number of failed units directly, and to compute and show average age.

A percentage growth rate is put into column B for the year, and the number of poles that need to be added for “new growth” is computed in C. Column D computes and shows the number of expected failures, calculated in the same way described earlier. Thus, 18 poles represent growth for 2010 (1.1% of 1,613), and 60 poles fail. Their total, 78, is now the total number of poles less than a year old for year 2011 (cell H6), and then total for year 2011 is now 1,631 poles, not 1,613.

Column D computes and shows the average age of all poles for that year. This computation is done in another set of rows (not shown - the authors do this on another worksheet) and the results brought back to this column.

Other Features for Practicality

Often users will add additional columns between E and F to compute and report other statistics they wish to track beyond the three shown in columns C-E in Figure A.3. These may include the asset base aging rate (the increase or decrease in average age from last year to this), the failure rate overall (the number of failures divided by the total number of poles, expressed as a percentage), and perhaps the equivalent age (age at which the failure rate equals that of the entire population in this year).

Dimensions

Moving the age distribution to the right to model each successive year means that an age distribution that begins with poles up to ninety years of age could, when applied to study trends fifty years ahead, need 140 columns. The authors recommend at least two hundred columns, for reasons that will be discussed later. Similarly, the authors recommend at least two hundred rows: one may not need to model two hundred specific years, but copying left and down is easy to do and leaves plenty of room for study of long term trends. (The authors' spreadsheet is 500 by 500.)

Strange results – bizarre errors with a source not immediately obvious - occur with this model if it ever “runs out of columns.” Counts don't add up, predicted results oscillate, etc. The authors sum up the final (500th in their case) column

	A	B	C	D	E	F	G	H	I	J	K	L
1										Age-->	0	1
2									Total Rate (age)-->		0.50%	1.00%
3									Failure Rate (age)-->		1.00%	2.00%
4									NCR Rate (age)-->		0.50%	1.00%
5									Retirement Rate (age)-->		0.25%	0.50%
6												
7	V-Year-V	Growth	Growth	Failures	NCR	Retired	Added	Avg Age	Total	Age-->	0	1
8	2010	1.10%	33	60	30	15	123	2010.0	1613		231	284
9	2011	1.20%	37	72	36	18	145	2011.0	1631		123	227
10	2012	1.30%	42	82	41	20	164	2012.0	1650		145	121
11	2013	1.40%	46	90	45	22	180	2013.0	1672		164	142
12	2014	1.50%	50	96	48	24	194	2014.0	1695		180	161
13	2015	1.60%	53	101	51	25	205	2015.0	1721		194	177

Figure A.5 Final version of the basic model, which now uses multi-worksheets to simplify “accounting” but accounts for failures, NCR, and retirements as separate but simultaneously ongoing processes affecting the equipment. See text for details.

as a test: if its sum is not zero, then, for some reason, the model may not have enough columns. (The authors are fully aware that most readers will question why a 200 by 200 or larger model might need this error diagnostic, but experience shows it is useful).

A.3 MULTI-SHEET STRUCTURE OF A MORE PRACTICAL MODEL

With the basic concept model (Figures A.1 – A.4) established, the authors will now outline a *multi-worksheet structure* that has features needed for more practical study: the same concept is implemented in a slightly different way. Figure A.5 shows the spreadsheet modified so that there are *three* reasons poles might not survive to another year: failures, non-condition-related replacement (NCR), and retirement. This model explicitly represents and tracks all three.

At the top of this “top” worksheet, named Input and Display, where there was previously the “failure rate as a function of age” row, there are now four rows: one each for failures, NCR, and retirements, and one holding their sum. Similarly, columns have been added to track counts of failures, NCR replacements, and retirements by year.

	A	B	C	D	E	F	G	H	I	J	K	L	
1	Failures	Age-->	0	1	2	3	4	5	6	7	8	9	10
2	V-Year-V	Rate-->	1.00%	2.00%	3.00%	4.00%	5.00%	6.00%	7.00%	8.00%	9.00%	10.00%	11.00%
3	60	2010	2.3	5.7	6.3	7.9	17.3	20.6	0.0	0.0	0.0	0.0	0.0
4	72	2011	1.2	4.5	8.2	8.0	9.2	18.9	21.6	0.0	0.0	0.0	0.0
5	82	2012	1.4	2.4	6.6	10.4	9.3	10.1	19.7	21.6	0.0	0.0	0.0

Figure A.6 Nothing pretty here: no extra rows or columns to make display easier. This sheet exists to compute failures by year, by age, and to add them into yearly totals (column A). Its results are used by the top sheet (Figure A.5).

Shown in Figure A.5, beginning with cell K8 and going to the right and down, are the counts by age and year as in the prior model, with the age distribution for each year beginning in column K. But these cells hold a different type of computation than they did in the sheet discussed in Figures A.1 – A.4. Here, *separate worksheets* compute the number of failures, the NCR replacements, and the retirements made each year, by age, and those results are brought back to this sheet. For example, the failure rate function at the top of the sheet in A.5 (row 3 beginning in column K) is copied to the top of a second worksheet named Failure Computation (Figure A6, row 3, beginning in column C there). Failure Computation does just that, computing the number of units failing each year, by year and age, as shown in the figure.

Similarly, NCR replacements are computed on another sheet (NCR Replacement Computation) and retirements are computed on the fourth sheet in the sheet (Retirement Computation). All three have a structure identical to that shown in Figure A.6, copying their respective rate-as-a-function-of-age curve from the top of the Input and Display sheet and computing units affected by their particular effect.

The appropriate values for failures, NCR, and retirements by year and age are taken from each of those respective computation sheets and all three values subtracted from the counts by year and age in the Input and Display sheet. For example, the value in cell L9 on that top Input and Display sheet (the number of one-year-old poles in 2011) is computed as the value in cell (231 number of new poles in year 2010) minus the failures, the NCR, and the retirements computed on those other sheets: 2.3 fail (this is shown in cell C3 in the Failure Computation sheet, Figure A.6), 1.15 are NCR (not shown – that value is in cell C3 on the NCR sheet which is formatted just like A.6 but tracks only NCR), and .58 are retired (again not shown). These result in a total of 4.03 that is subtracted from the 231 to give an expected 226.97 one-year old poles for 2011 (displayed as 227).

The total number of units that fail each year (column A in the Failure Computation sheet) is brought back to column D of the top Input and Display sheet, with NCR and retirement sums taken from those sheets and placed in adjacent columns. The sum of growth, NCR, and failures - the net number of new poles installed – is computed and shown in the top sheet's column G.

Growth covers retirement losses

A subtle distinction should be noted about the poles added for growth in the model shown in Figure A.5. Cell C8 in Figure A.5 (Input and Display) shows thirty-three poles are added for growth reasons in 2010. However, the percent growth for 2010 (cell B8, 1.1%), applied to 1,613 units, gives only eighteen units for growth. The reason is that as applied here, the growth percentage entered by the user in Column C describes the *net* growth of the number of poles in service: 1,613 plus 1.1% equals 1,631 poles for 2011. Since fifteen are retired (cell F8) that year, a total of eighteen plus fifteen new poles must be added to make the total sum up to the desired value of 1,631.

Readers can make their own determination of how any model they construct and use would treat this issue: there are cases where it is appropriate to not “true up” the total with this addition as shown. However, *this approach is necessary for versions of the model used to fit to historical data* (i.e., as used in Chapter 16). This will be discussed in the next section.

A.4 FITTING AND CALIBRATING A MODEL

Given an accurate age distribution for the current year (i.e., data for the row 2010), and an accurate failure-as-a-function-of-age formulae (and if needed, NCR and retirement functions, too), a model as described in sections A.2 and A.3 will produce useful predictions of failure counts and changes in an asset base’s age distribution for future years.² From those projections, planners can develop forecasts of future reliability, cost, and operating impacts.

But to be a practical planning tool this model must have *both* the age distribution for the base year (whatever year is the most recent data - in this example, 2010) *and* the formulae for failures, NCR, and retirements (if used). Any planning approach that expects to dependably project failure and aging trends into the future will require both sets of data. Trending annual failure counts into the future by working with only past failure counts works mathematically, but not practically: without consideration of the age distribution and how failures erode the base of older units that can fail, such projections will be inaccurate over any but the very short term. Furthermore, even with all the recommended data, the fact that the mathematics works – that a solution can be obtained, does not necessarily mean that a practical model has been produced. More than a modicum of common sense is essential: if an experienced asset planner has doubts, the model is probably not ready for use.

In no case, unless it has already solved for it through some means, will a utility have the failure function it needs. In many cases (Chapter 16) the utility will not have the age distribution either, although in cases where major equipment is involved (e.g., large power transformers), it may. Regardless, the utility asset management team will have to solve for what it does not have from data it does have. This section will first demonstrate the mechanics of fitting the model to historical data, showing the basic steps in setting up the spreadsheet to solve for age distribution and failure functions. It will then discuss practical aspects of doing this, including what data is most likely to be available and dependable, and various tricks planners, engineers, and analysts working on the data can use to enhance their ability to reach a solution that is both mathematically sound and practically useful. Section A.5 will follow this with a

² Generally, NCR and retirement functions are not needed when studying equipment that is located on utility sites or ROW, as such equipment is seldom retired or removed from service. NCR and retirement are normally needed only for equipment that is located in easements – distribution equipment.

Table A.1 Units Added For Growth, By Year

Year	Added for Growth
1990	1,000
1991	1,000
1992	1,000
1993	1,000
1994	1,000
1995	1,000
1996	1,000
1997	1,000
1998	1,000
1999	1,000
2000	1,000
2001	1,000
2002	1,000
2003	1,000
2004	1,000
2005	1,000
2006	1,000
2007	1,000
2008	1,000
2009	1,000

Table A.2 Age Distribution for System

Age	Count
New	1126
1	1113
2	1100
3	1087
4	1074
5	1060
6	1047
7	1033
8	1020
9	1006
10	992
11	979
12	966
13	952
14	939
15	926
16	914
17	901
18	888
19	876

with a discussion of issues often encountered in the real world when fitting these models to historical data.

Example of Calibration of a Model

This problem has deliberately been kept simple, with only a 20 year period. The utility added 1000 units each year for growth over the past 20 years (data, Table A.1). Some units failed and the utility replaced those with new as it went along, in addition to the 1000 per year it was installing. It has ended up with the age distribution given in Table A.2. The question is: what is the failure rate as a function of unit age?

Figure A.7 shows the leftmost columns of a spreadsheet like that in Figure A.4 (i.e., with only a failure function, not NCR or retirement, functions) that has been set up to solve for the failure function with the data from Tables A.1 and A.2. Numbers shown in bold are input data or key factors in the solution function:

- Two coefficients for the failure function, A and B, have been created in row 2, columns A and B. These are blank (zero) in Figure A.7.
- A failure function has been set up in row 2 with the formula:

$$\text{Failure rate (age)} = A \times \text{age} + B$$

In other words, the formula in cell G3 is “=G2*\$A3+\$B3” and this has been copied to the right for the entire rest of the row.

Readers will find that other types of formulae, usually logarithmic or exponential functions, are often the best for fitting purposes, but this simple linear function will serve to show the mechanics of solving a spreadsheet model.

- Data from Table A.2 has been entered in column B, Units Added for Growth.
- Data from Table A.3 has been entered in row 26
- Row 27 calculates the residuals (error) between the modeled age distribution for 2010 (row 25), and the actual data (row 26). Row 28 computes the absolute value of the residuals and the average absolute value (AAV) which is the error metric that will be minimized.
- The AAV is copied up to cell C2 for reporting purposes.

The Goal Seeking or Solver function of the spreadsheet can now be used to find A and B. Alternatively, often a savvy group of analysts can simply iterate

manually until they converge on a solution: a “trial at a time” manual process can converge on a result fairly quickly in most single- and dual-variable situations, particularly if the user is familiar with how iterative algorithms are written to iteratively converge on a solution.

Regardless, the correct solution here is:

$$A = .077\%$$

$$B = 0.00$$

These coefficients give an error of zero in this sample case. The solved case is shown in Figure A.8

	A	B	C	D	E	F	G	H	I	J
1	A	B	AAV			Age-->	0	1	2	3
2	0.000%	0.000%	66.7	Failure Rate (age)-->			0.000%	0.000%	0.000%	0.0%
3										
4	V-Year -V	Added	Failures	Avg Age	Total	Age-->	0	1	2	3
5	1990	1,000	0.00		0		0	0	0	0
6	1991	1,000	0.00	0.5	1000		1000.0	0.0	0.0	0.0
7	1992	1,000	0.00	1.0	2000		1000.0	1000.0	0.0	0.0
8	1993	1,000	0.00	1.5	3000		1000.0	1000.0	1000.0	0.0
9	1994	1,000	0.00	2.0	4000		1000.0	1000.0	1000.0	1000.0
10	1995	1,000	0.00	2.5	5000		1000.0	1000.0	1000.0	1000.0
11	1996	1,000	0.00	3.0	6000		1000.0	1000.0	1000.0	1000.0
12	1997	1,000	0.00	3.5	7000		1000.0	1000.0	1000.0	1000.0
13	1998	1,000	0.00	4.0	8000		1000.0	1000.0	1000.0	1000.0
14	1999	1,000	0.00	4.5	9000		1000.0	1000.0	1000.0	1000.0
15	2000	1,000	0.00	5.0	10000		1000.0	1000.0	1000.0	1000.0
16	2001	1,000	0.00	5.5	11000		1000.0	1000.0	1000.0	1000.0
17	2002	1,000	0.00	6.0	12000		1000.0	1000.0	1000.0	1000.0
18	2003	1,000	0.00	6.5	13000		1000.0	1000.0	1000.0	1000.0
19	2004	1,000	0.00	7.0	14000		1000.0	1000.0	1000.0	1000.0
20	2005	1,000	0.00	7.5	15000		1000.0	1000.0	1000.0	1000.0
21	2006	1,000	0.00	8.0	16000		1000.0	1000.0	1000.0	1000.0
22	2007	1,000	0.00	8.5	17000		1000.0	1000.0	1000.0	1000.0
23	2008	1,000	0.00	9.0	18000		1000.0	1000.0	1000.0	1000.0
24	2009	1,000	0.00	9.5	19000		1000.0	1000.0	1000.0	1000.0
25	2010		0.00	10.5	20000		1000.0	1000.0	1000.0	1000.0
26					Target Data -->		1127	1113	1100	1087
27					Error-->		127	113	100	87
28	Average Absolute Value-->			66.7	Error^2-->		127	113	100	87

Figure A.7 Basic failure-function-only spreadsheet model set up with the sample problem. Only the leftmost nine columns are shown: the model runs to the left for at least twenty years of age distribution and row 26 contains all data values shown in Table A.1.

	A	B	C	D	E	F	G	H	I	J
1	A	B	AAV			Age-->	0	1	2	3
2	0.077%	0.000%	0.0			Failure Rate (age)-->	0.077%	0.077%	0.154%	0.2%
3										
4	V-Year -V	Added	Failures	Avg Age	Total	Age-->	0	1	2	3
5	1990	1,000	0.00		0		0	0	0	0
6	1991	1,000	0.77	0.5	1000		1000.0	0.0	0.0	0.0
7	1992	1,000	1.54	1.0	2000		1000.8	999.2	0.0	0.0
8	1993	1,000	3.08	1.5	3000		1001.5	1000.0	998.5	0.0
9	1994	1,000	5.38	2.0	4000		1003.1	1000.8	999.2	996.9
10	1995	1,000	8.45	2.5	5000		1005.4	1002.3	1000.0	997.7
11	1996	1,000	12.28	3.0	6000		1008.5	1004.6	1001.5	998.5
12	1997	1,000	16.86	3.5	7000		1012.3	1007.7	1003.8	1000.0
13	1998	1,000	22.19	4.0	8000		1016.9	1011.5	1006.9	1002.3
14	1999	1,000	28.26	4.4	9000		1022.2	1016.1	1010.7	1005.4
15	2000	1,000	35.05	4.9	10000		1028.3	1021.4	1015.3	1009.2
16	2001	1,000	42.56	5.4	11000		1035.0	1027.5	1020.6	1013.7
17	2002	1,000	50.78	5.9	12000		1042.6	1034.3	1026.7	1019.0
18	2003	1,000	59.69	6.4	13000		1050.8	1041.8	1033.5	1025.1
19	2004	1,000	69.29	6.8	14000		1059.7	1050.0	1041.0	1031.9
20	2005	1,000	79.55	7.3	15000		1069.3	1058.9	1049.2	1039.4
21	2006	1,000	90.47	7.7	16000		1079.6	1068.5	1058.1	1047.5
22	2007	1,000	102.03	8.2	17000		1090.5	1078.7	1067.6	1056.4
23	2008	1,000	114.20	8.6	18000		1102.0	1089.6	1077.9	1066.0
24	2009	1,000	126.99	9.1	19000		1114.2	1101.2	1088.8	1076.2
25	2010		140.35	10.0	20000		1127.0	1113.3	1100.3	1087.1
26						Target Data -->	1127	1113	1100	1087
27						Error-->	0	0	0	0
28	Average Absolute Value-->			0.0		Error^2-->	0	0	0	0

Figure A.8 Basic failure-function-only spreadsheet model showing results with the solved failure function. See text for details.

In many cases, particularly with major equipment like large transformers and high voltage breakers, a utility will have data from which an age distribution can be developed. However, in other cases, when dealing with distribution equipment, it will not be possible to develop an age distribution for any year, recent or past, so the solution approach shown in Figures A.7 and A.8 cannot be used. In such situations, its asset management must find some other stream of *historical* data – counts of something else that is relevant over an appreciable portion of the equipment's history. Examples that can be worked are:

- Data available:
- 1) Units added for growth each past year
 - 2) Units added for failures each past year (or units failing each year)

This is basically the situation in Chapter 16's case study. Both data sets must go back into the equipment history far enough that "the starting solution" – the age distribution prior to the beginning of the data -- is not judged to be significant.

Figure A.9 shows the earlier example set up for a solution with this data. Again, numbers in bold are input data. Here, there is no age distribution data in row 27. Instead, failure count data by year has been added (column F has been used for that). Error is computed as the difference between the numbers in columns C and F (in a column not shown) and the solution is affected in exactly the same manner as before: find A and B to minimize the error. Data available is:

- 1) Units in service each year
- 2) Units added for failures each past year (or units failing each year)

This situation would result in input data for columns E and D rather than B and D as in the situation above. Otherwise it would be the same as that above.

A.5 REAL-WORLD DATA SITUATIONS: GETTING A SOLUTION

The example problem solved in the previous section was only 20 years long, and to compensate used a failure rate curve that was much higher, sooner, than any in actual situations. Finding, vetting, and "fixing" historical data is always a problem in any aging infrastructure investigation and study. Good data may exist for recent years (Chapter 16's situation is fairly typical). But what is really needed is something from long ago. Common sense alone can tell one that if equipment can last 50 years or more, then in order to ferret out failure rates as a function of age, one is going to need data for at least 50 years past. Mathematically, one can fit failure-rate curves to say, only 20 years of data, and extrapolate them out to 50, even 100 years. But no one would expect those curves to be dependable: many functions can be made to fit 20 data points well, but would diverge at 40 and 100 years. Practically, one needs to find some data about equipment, failures, units in operation, or something relevant, from 50 or more years before. Chapter 16 discusses this in detail for one such situations, and includes looking at if and how data that might be suspect can be used even though it is less than perfect.

When Data on Equipment Age Distribution is Available

For major equipment types like large power transformers and high voltage breakers, etc., a utility will have records that include when the unit was purchased, when it was put into service, etc. Further, the utility can usually tie these records to specific units. For example, if not done already, with a bit of clerical work a project team can use past purchase and work order records to determine that the number 2 transformer at Eastside substation, a Westinghouse 22 MVA 138kV to 12.47/kV delta/wye unit, is in fact the unit delivered in March 1954 and installed at Eastside substation the following month. As of

2010, it is 56 years old. Similarly, the team can go through the company records and determine that no one knows where a similar unit delivered a month later, in May 1954 is, but that it is no longer in service – so it does not count in the age distribution. Obtaining an initial date for each unit now in service, in this way, permits construction of an age distribution “row” for the current year.

But again, the age distribution *alone* does very little good– any model for studying aging trends and projecting them into the future *must* have failure function, and if needed (because the analysis is at the distribution circuit level) the NCR and retirement functions. And the age distribution itself provides almost no clue about the failure function. It tells the project team that the company owns four power transformers that are 56 years old, but the company might have bought nine that year, meaning over half have failed. Or it might have bought only four, and there have been no failures.

	A	B	C	D	E	F	G	H	I	J
1	A	B	AAV			Age-->	0	1	2	3
2	0.000%	0.000%	66.7	Failure Rate (age)-->			0.000%	0.000%	0.000%	0.0%
3										
4	V-Year-V	Added	Failures	Avg Age	Total	Age-->	0	1	2	3
5	1990	1,000	0.00		0	0	0	0	0	0
6	1991	1,000	0.00	0.5	1000	1	1000.0	0.0	0.0	0.0
7	1992	1,000	0.00	1.0	2000	2	1000.0	1000.0	0.0	0.0
8	1993	1,000	0.00	1.5	3000	3	1000.0	1000.0	1000.0	0.0
9	1994	1,000	0.00	2.0	4000	5	1000.0	1000.0	1000.0	1000.0
10	1995	1,000	0.00	2.5	5000	8	1000.0	1000.0	1000.0	1000.0
11	1996	1,000	0.00	3.0	6000	12	1000.0	1000.0	1000.0	1000.0
12	1997	1,000	0.00	3.5	7000	17	1000.0	1000.0	1000.0	1000.0
13	1998	1,000	0.00	4.0	8000	22	1000.0	1000.0	1000.0	1000.0
14	1999	1,000	0.00	4.5	9000	28	1000.0	1000.0	1000.0	1000.0
15	2000	1,000	0.00	5.0	10000	35	1000.0	1000.0	1000.0	1000.0
16	2001	1,000	0.00	5.5	11000	43	1000.0	1000.0	1000.0	1000.0
17	2002	1,000	0.00	6.0	12000	51	1000.0	1000.0	1000.0	1000.0
18	2003	1,000	0.00	6.5	13000	60	1000.0	1000.0	1000.0	1000.0
19	2004	1,000	0.00	7.0	14000	69	1000.0	1000.0	1000.0	1000.0
20	2005	1,000	0.00	7.5	15000	80	1000.0	1000.0	1000.0	1000.0
21	2006	1,000	0.00	8.0	16000	90	1000.0	1000.0	1000.0	1000.0
22	2007	1,000	0.00	8.5	17000	102	1000.0	1000.0	1000.0	1000.0
23	2008	1,000	0.00	9.0	18000	114	1000.0	1000.0	1000.0	1000.0
24	2009	1,000	0.00	9.5	19000	127	1000.0	1000.0	1000.0	1000.0
25	2010		0.00	10.5	20000		1000.0	1000.0	1000.0	1000.0

Figure A.9 Basic failure-function-only spreadsheet model set up, before solution, for a situation where the utility has no age distribution data for a recent year, but is using failures by year (column F) instead. See text for details.

Failures happen over time, and the only way to determine how that happens “over time” is to have data over time. The project team must take one of two approaches outlined below.

Age Distributions are Available for Multiple Year

Suppose the project team has age distributions for 2010 (the current or most recent year) and 1980. Units that were new in 1980 have now been in service for 30 years, those 1 year old in 1980 are now 31 years old, and units that were 26 years old in 1980 are now 56 years old, etc. A formula can be fitted to this data to determine a failure function that best fits the data. With age distributions for only two years, it would probably be necessary to restrict the formulae to functions that will have the shape expected (i.e., logarithmic or exponential functions that will be strictly increasing with time). Planners may not have the greatest confidence in that fitted function, but they do have a result, which is a starting point.

In the example given earlier, this would be equivalent to having age distributions for, say, 2000 and 2010. A simple solution equivalent to that covered earlier could be set up as shown in Figure A.10.

	A	B	C	D	E	F	G	H	I	J
1	A	B	AAV			Age-->	0	1	2	3
2	0.000%	0.000%	66.7	Failure Rate (age)-->			0.000%	0.000%	0.000%	0.0%
3										
4	V-Year-V	Added	Failures	Avg Age	Total	Age-->	0	1	2	3
5	2000	1,000	0.00	5.0	10000		1028.3	1021.4	1015.3	1009.2
6	2001	1,000	0.00	5.5	11000		1000.0	1028.3	1021.4	1015.3
7	2002	1,000	0.00	6.0	12000		1000.0	1000.0	1028.3	1021.4
8	2003	1,000	0.00	6.5	13000		1000.0	1000.0	1000.0	1028.3
9	2004	1,000	0.00	7.0	14000		1000.0	1000.0	1000.0	1000.0
10	2005	1,000	0.00	7.5	15000		1000.0	1000.0	1000.0	1000.0
11	2006	1,000	0.00	8.0	16000		1000.0	1000.0	1000.0	1000.0
12	2007	1,000	0.00	8.5	17000		1000.0	1000.0	1000.0	1000.0
13	2008	1,000	0.00	9.0	18000		1000.0	1000.0	1000.0	1000.0
14	2009	1,000	0.00	9.5	19000		1000.0	1000.0	1000.0	1000.0
15	2010		0.00	10.5	20000		1000.0	1000.0	1000.0	1000.0
16	2011	1,000	0.00	0.0	20000		0.0	1000.0	1000.0	1000.0
17	2012	1,000	2126.14	0.0	21000		1000.0	0.0	1000.0	1000.0
18		20,028				Target Data -->	1127	1113	1100	1087
19						Error-->	127	113	100	87
20	Average Absolute Value-->			66.7		Error^2-->	127	113	100	87

Figure A.10 Example set up to solve when age distributions for two years are available.

Age distributions for more than two years will help further identify the failure function; more data always improves the situation. But circumstances where age distribution data is available for even one recent year are rare: the authors have never encountered a situation where dependable age distribution data from several decades in the past is available and dependable. In some cases, *a lot* of clerical work – perhaps several months of slogging through past purchasing, installation, failure reports, and work orders for replacements, can cross-verify records and result in the creation of a timeline for units, so that age distributions for ten, twenty, thirty, and even forty years back can be developed. In such cases, if the existing records were judged to be complete and accurate, the clerical work involved would probably be justified: set of age distributions going back four or five decades is an incredible base from which to work. It is then rather straightforward to fit failure functions to the data.

But most often data is not complete, if even available. Typically, for major equipment like power transformers and transmission lines and equipment, utility records will permit the creation of a reasonable dependable age distribution for a recent year, but nothing more. And typically for distribution equipment, often including medium voltage breakers and voltage regulators but nearly always including poles, service transformers, and line equipment like cap banks and reclosers, creation of a dependable age distribution for even one recent year is impossible from a practical standpoint.

Input-Output Analysis Using Historical Data

Chapter 16's account of CL&P's analysis and management of its wooden distribution poles is a situation where no age distribution data was available for even a recent year. That is much more likely to be the situation with most types of equipment at most utilities. In that chapter, a type of Input-Output analysis was used to *both* find the age distribution for the current year, and solve for the failure, NCR, and retirement functions.

The basic concept of Input-Output analysis was represented in the example shown in Figures A.7 and A.8. More generally, the idea is to go back to the beginning of the utility system (this is never done, but for the moment consider it is feasible) -- to say, 1913 when the company was founded – and build a record of how many poles were bought each year: 472 were bought and installed in 1913, 590 in 1914, etc., through to the current year. And suppose one had a count for how many exist in the most recent year (the current year). Mathematically, this is sufficient to solve for a failure rate function: practically the solution would be of little use, because no one would have a lot of confidence in the failure function (many could be made to fit very well), but one *would* get a solution.

But if one now adds data on the number of poles added just for growth each year (Figure A.9) or the in service each year not just for the most recent year but for prior years (an alternate form of Figure A.9 discussed earlier), the solution might be messier (no formulae will fit actual data as well as it will the situation

in the paragraph above with only the most recent year’s count) but the solution is still straightforward. And the resulting function will typically be much better from a practical standpoint, given more confidence a reasonably useful fit has been achieved. That is the basic concept of “Input-Output” analysis as applied in this book. It will not be described in detail here since Chapter 16 provides a very comprehensive narrative on that application, including how the utility avoided “going all the way back to the beginning,” and provides tables of data and a lengthy discussion of results.

Table A.3 Data most likely to be available for historical analysis

Most likely	Number of units purchased/installed by year
	Number of units that failed by year
	Number of units in service by year
	Number of units serviced or rebuilt by year
	Number of units disposed/parted/surplused by year
	Number of units with certain problems or issues by year
Least likely	Distribution of ages for a past year(s)

An Input-Output approach is most often the approach that looks like it will work best, and often the only approach that is even feasible. As in Chapter 16, often a concerted search will locate historical records that can be used to provide some or all of the data shown in Table A.3. Most typically, except for very large types of equipment (GSUs, etc.) it is not possible to tie specific units in service today to specific purchases/installations long ago. In other words, hard work might determine that the utility bought 28 medium voltage breakers and 4,117 wooden poles in 1978, but exactly where those 28 breakers or 4,117 poles are located now (or even how many are still in the system) is impossible to determine.

Realities of Data Analysis and Modeling

The authors will not try to be more specific about how to “solve” for the formulae and age distributions needed to set up an aging infrastructure model. Every situation they have encountered is a different combination of data sources, data problems and quality issues, modeling needs, and management priorities. Again, common sense, attention to detail, and a certain amount of creativity in how the problem is approached and set up are required.

Readers planning to build aging infrastructure models of any type, whether with models structured as described here, or otherwise, must keep in mind that

these models must be fitted, or calibrated, to historical data to be of any use. That requires either multiple age distributions, or data for many past years - counts of units in service, added, and failing. Ideally, but rarely, a combination of both is available.

But invariably, records that are found will be imperfect, perhaps unreadable in places, or not complete in all cases, or even worse, not provably complete or incomplete. Ancient nomenclature may be, if not undecipherable, at least not completely understood (What does “pretty bad” mean when scribbled in the “As found condition” line on the back of a forty-year-old work order?).

But in the vast majority of cases the authors have seen, the data that can be found and assembled is sufficient if used well. Sufficient is about the best that most project teams should realistically hope to achieve.

“Used well” in the paragraph above means with common sense and attention to detail, and with what many will regard as *very advanced* mathematical methods. Given the “fuzziness” of the data that is common to most situations, care needs to be used in the selection of methods for fitting of formulae. Typically, if a project team looks for and thinks about it, it will find it has useful *information* on data quality that can be used. For example, if a few records are missing from past reports, it is almost certain that counts made from them for long-past years are low rather than high. Numerical methods that take this into account when solving can be used. Pattern recognition and non-numerical methods such as grade logic may be useful in other cases. These and other methods, all part of a new field of *analytics*, use numerical methods including but quite beyond traditional statistical approaches, applied in ways that often require considerable computer analysis (often several hours to obtain a good fitting). But they typically yield superior results.

Often a probabilistic model is desirable, not just because the processes being modeled are probabilities (of failure, etc.), but for other reasons, too. One major factor behind many failures, weather, is a process dealt with best on a probabilistic basis. Also, probabilistic methods can be fitted using techniques that sometimes work better in “sparse data” situations: they can use the variations from year to year in the data, and what often looks like random patterns, to their advantage in ferreting out trends, etc. Finally, many asset management approaches want to analyze the likelihood of outcomes, to do risk analysis for planning, and probabilistic models are essential for that. The decision to use a probabilistic model, along with the desire to access more powerful analytic methodology, was what motivated CL&P, in Chapter 16’s example, to move on from matured and developed forms of the deterministic spreadsheet model described here, to a probabilistic of the same basic concepts and structure, but built within SAS software.

A.6 COMMENTS ON APPLICATION

A model like that described in Sections A.2 through A.4 is appropriate to study commodity distribution equipment like poles and service transformers:

equipment which exists in very large quantities and for which little or no specific data on each individual unit exists. Chapter 16 gives a comprehensive storyline of how such a model is fitted and used, so a lengthy discussion of the “how” was not given here. As applied in Chapter 16, and as needed for application as a practical planning tool in any situation, the model will invariably grow much larger than that shown in the figures here. The spreadsheet used in Chapter 16 had more than a dozen columns similar to D through G in Figure A.5, reporting a wide assortment of statistics on a year by year basis. Further, it had 300 rows for years and 300 columns for age, giving it the dimensions to study of up to 300 years of aging trends. While nothing like that size is usually required, it is easy to extend the dimension of a spreadsheet to this size, once it is set up, the increase in computation time is miniscule, and the large size avoids some potential problems. Readers should realize that the historical analysis period used in Chapter 16 was not atypical. There data from 1933 to 2010 (77 years) were used, and another 20 years of history was synthesized (“assumed”) so that a total of nearly 100 rows (years) were used just to set up the historical basis for solving the current (2010) age distribution and the failure, NCR, and Retirement functions of age. Add fifty years of modeling long term trends, and one has used half the 300-year capacity.

Beyond that need for dimension, some steps that analysts may wish to perform at stages in the search for and solving of the best formulae, etc., often require running the model far into the future a hundred years or more, to diagnose how and how fast various trends asymptotically approach the sustainable point (if it oscillates, etc.). Close inspection of very long-term results can help solve problems with the fitting or reveal clues to trends and patterns buried in the data. The authors have used models with as many as 500 columns and rows for some investigations.

Application to Major Equipment

The model discussed here is useful, but generally not the only model used, when dealing with major equipment. There, the typical situation is all of the following:

- *More data.* As discussed earlier here and in Appendix B and Chapter 16, typically a utility will have a lot of data for its major equipment like large power transformers.
- *But fewer units.* There are far fewer major units than units like service transformers or capacitor banks. For example, Chapter 16’s CL&P had over 435,000 poles but only 503 large transformers.
- *More complexity of condition modeling.* Typically utilities need models that represent state and condition in a more comprehensive manner than just “age is a proxy for condition.”

- *Multi components.* Most utilities will want to track and model the state and condition of key components in the equipment: in a transformer, the windings, case, bushings, LTC, etc.

In such situations, models along the lines of those discussed in Appendix B and presented in Chapter 16 are used as the major tool to track, model, and project how the equipment will behave. However, sustainable point models as outlined here are useful to solve for the failure rates and in some cases the age distributions of equipment. A utility may own 1,000 power transformers, and intend to represent them with multi-component, multi-state condition models as required to model the various parts of a large and complex asset such as a power transformer or a high voltage breaker. However, that means that it owns 1,000 *transformer cores*: with a little work analysts can put together data for a model like those discussed here, that represents only the history of the operating cores. That can then be used to solve for failure rates and age/condition distributions, and to project future trends, that will be useful in the larger multi-component modeling effort.

REFERENCE

H. L. Willis, M. V. Engel and R. E. Brown, "Equipment Demographics-Failure Analysis of Aging T&D Infrastructures," *Proceedings of the Canadian Power Conference*, 2004, Toronto.

Appendix B

Sustainable Point Analysis

B.1	Introduction	739
B.2	Elements of Aging Infrastructure Analysis	742
B.3	Quantitative Examples	745
B.4	Metrics for Measuring Aging Infrastructures	762
B.5	Effect of Different Failure Rate Curve Shapes	769
B.6	Reading Distributions	773
B.7	Applying Economic Analysis	780
B.8	Summary of Key Points	782
	References and Bibliography	785

B.1 INTRODUCTION

Sustainable point analysis is a statistically based method for studying the aging, condition deterioration, and condition improvement options for a large set of T&D equipment (or for that matter, any other type of equipment). It is based upon a fundamental concept that some people find surprising or even counterintuitive:

Any large set of equipment in which units that fail are replaced by new units will eventually reach a sustainable point after which average equipment age, annual failure rate, and service and replacement costs will be roughly the same from year to year.

In other words, *any* set of poles, structures, transformers, cutouts, vaults, breakers –any collection of equipment where failed units are replaced with new or rebuilt units – will deteriorate only to a point. From there onward, its overall situation will remain stable, with the same average condition, the same expected failure rate, the same reliability impacts, and the same expected repair and maintenance costs every year, unless the equipment owners change their operating, repair and




Figure B.1 The concept of the sustainable point. In this example, a group of transformers – perhaps 1,000 – are installed in the same year, and average age that year is zero. Over time as they spend time in service, the average age increases. Some fail and are replaced, so the average age of units remaining in service does not increase by exactly one year per year. Here, average age gradually increases until a sustainable point average age of 50.0 years is reached, but this takes far longer than 50 years, as can be seen. Similarly, the average failure rate for the entire set, and all other statistics for it, will reach a stable value at the sustainable point, too.

replacement policies, maintenance practices, or the type of equipment used.

It can to reach this sustainable point. Figure B.1 illustrates the key characteristics of overall aging trend exhibited by nearly any set of equipment. When a group of, say 1,000 poles or 1,000 transformers is installed, its average age is zero (all new). Neglecting initial flaws (infant mortality) the average failure rate is as low as it will ever be: it will slowly increase over time. The age of these units increases as they remain in service. But the wear and tear of service slowly eats away at their condition. A few units experience more stress or are weaker than others and fail early after only a few years. They are replaced with new units, which are of “zero” age, and those new units bring down the average age of the group just a bit from where it would otherwise have been. Thus, ten years after installation, the average age of the entire set is not exactly 10.00 years, because there have been a few replacements that are newer – average age is perhaps only 9.98 years. The average failure rate will be just a bit less than if all were exactly ten years old, because it is a combination of the failure rate for the ten-year-old units and the newer replacements.

The newer, replacement units begin to age, too. They are “behind” the

original units in their race to their end of service life, but are aging just as fast – every year they stay in service they all get a year older, etc. The whole group, old and new, continues to age, with the average age increasing and the average condition deteriorating every year. Units fail and are replaced, etc., always with new units that bring down the average age of the group compared to if none had failed, and improve the average failure rate compared to that of the oldest units.

But at some point the original units – and perhaps even some of the early replacements – age and fail to the point where there aren't very many of them left: so many years have passed that, even at a very low failure rate, many have already failed and been replaced. All the original units that are left are so deteriorated that they are failing at a high rate – so high that *a lot* of replacements are being made – enough that the number of new (zero age) units added each year counterbalances the continued aging of the remaining old units. From thereon, the equipment group is essentially at a stable point. (For the reader who cannot picture this, or wishes to see this, the bulk of this chapter will look at quantitative examples, which will make this and other points much clearer.)

Theoretically there are *no* exceptions. Regardless, for most operating electric and gas utilities the sustainable point is not *directly* of great practical value: despite the utility's "aged infrastructure" and the problems it creates, the sustainable point is still far in the utility's future and represents a condition-reliability-cost point that is much worse than today's, and completely unacceptable to it. And that is, surprisingly, part of the reason the sustainable point is so important:

Things will get worse. If the utility's T&D infrastructure is not at the sustainable point, then it is still trending toward that point: average age, reliability problems, and repair and replacement costs will continue to rise every year, slowly but inexorably. Analysis of the sustainable point and the trends toward it, based on historical and equipment data, provides useful information, on "how bad it will get and how soon we'll get there."

The sustainable point can be moved. But the sustainable point is not a function of the equipment alone. It is also affected by the utility's operating, maintenance, repair, and replacement policies and procedures. Changes in the utility's focus and procedures can both move the sustainable point and change the rate of increase of those reliability and cost problems – and sometimes even reverse them.

An Effective Tool for Management of Aging Infrastructures

Thus, the sustainable point does have some practical value to an operating electric or gas utility, mostly as an analysis and decision-making tool: it provides a powerful, effective foundation for analysis and optimization of strategies for dealing with aging infrastructures and the costs and reliability problems they create.

More accurate analysis. Statistical and pattern inference methods that assess utility asset, operating, and historical data are more accurate and effective if built around “sustainable point” equations. The knowledge that the system being studied will eventually asymptotically approach some sustainable point – even if that point is not known in advance – provides an “analytical fulcrum” with significant advantages: statistical fitting error is often cut in half. Patterns that otherwise might not be found are. Overall, the analysis of condition, failure and deterioration trends, and the expected improvement due to various mitigation methods, are all greatly improved.

Better decisions. Planning methods based on *moving* the sustainable point by changing inspection, service, maintenance, refurbishment, replacement, operating, or other policies, turn out to be particularly effective for studying various strategies and approaches to handling an aging equipment inventory. A utility can evaluate the effectiveness, in both the short-term and long-term, of various approaches to “fixing” its problems. In many cases, optimization can be applied to provide practical solutions to maximize results.

Damage from external causes. Most T&D equipment is so robust that in the presence of only deterioration from normal service, equipment lifetime will be many decades. But a utility’s planners must account for and anticipate the effects of “external damage” due to storms, earthquakes, floods, and similarly intense if infrequent natural events. While the cumulative effect of wear and tear of service is, ultimately, the bigger effect, external damage is a significant issue that can affect decisions and policies on equipment service, repair, and replacement. Furthermore, external damage complicates analysis and identification of condition deterioration trends from within historical and equipment data. Sustainable point analysis’ “mathematical fulcrum” helps here, too. The sustainable point method leads to planning and decision methods that better distinguish and fine tune policies aimed at each of these “lifetime shortening” effects.

B.2 ELEMENTS OF AGING INFRASTRUCTURE ANALYSIS

Two Types of Data Analysis

In almost all cases where a utility, industrial equipment owner, or government agency is reviewing and analyzing an aging infrastructure it owns, the equipment has been installed over a period of many years and is therefore of many different ages. Figure B.2 shows the number of overhead service transformers installed, and the number failing for a large power delivery system in the United States, for the years 1940 to 2010. This is an example of a *year by year* (YY) analysis: the units are analyzed with respect to specific times and with data representing each year at that time.




Figure B.2 Number of OH transformers installed (solid line) and failing (dotted line) in Big State Electric’s system from 1940 to 2010. Such *year by year* study is one type of analysis done on aging equipment sets.




Figure B.3 Failure rate curve for the set of one million OH service transformers from Figure B.1. Usually, failure rate curves are determined by *synchronized installation* analysis, in which equipment is analyzed on the basis of time since installation regardless of its age now – as if it were all installed in a single year.




Figure B.4 Same data displayed in the year-by-year data plot in Figure B.1, represented as simultaneous installation (bottom) manner. See text for details.

However, in many cases statistical analysis needs and data review are best carried out by looking at equipment purely as a function of age, regardless of when it was installed or what age it is now, in what is called a *simultaneous installation* (SI) analysis. Basically, in such situations, data on each equipment unit is used to “slide” it backwards or forwards in time, until the analysis data set has been adjusted so that all units appear to have been installed in a single year. Here, trends over time are much more apparent, etc. A failure rate curve such as shown in Figure B.3 is a product of SI analysis.

Invariably, thorough analysis of an aging equipment set requires both types of analysis. Figure B.4 shows the data from Figure B.1 displayed in both manners: the bottom data set was formed by simply using time since installation to shift all data points (for both installation and failure, if it happened) for each unit. Simple manipulation of the data on failures in the bottom plot then leads directly to the failure rate curve. *But regardless of the analysis approach used – YY or SI – for any set of equipment and operating conditions the sustainable age is the same.*

A Year Is a Year is a Year

In almost all practical aging infrastructure studies, and in all data plots and analysis given in this book, no distinction is made of time *within* a year, and ages are given, and analyzed, on only a year basis. Thus, 1,000 transformers might have been installed at various times in 2004, but all were treated as of identical age. Their age is zero years (new) anytime in 2004, 1.00 year old anytime in

2005 (i.e., they are zero years of age up through December 31, 2004, and one year of age from January 1 to December 31, 2010. Beyond the reduced data and analysis needs that this approximation brings, it is most often forced on planners and engineers by the available data, which often gives only the year of purchase or installation, with no further identification as to when within a year.

This simplification in data and analysis occasionally adds a small but noticeable deviation from the results one would get using a finer, perhaps daily, reckoning of time. However, there is no loss of generality of conclusions reached, and from a practical standpoint, *no loss of value* in any results from a decision-making and planning standpoint. Planning decisions that are specific to times within a year can be planned using such “yearly only” analysis, for example when a utility determines it will replace 840 units in year 2012 and then schedules them during periods when forces are expected to be available and maintenance outages can most conveniently be taken.

B.3 QUANTITATIVE EXAMPLES

Examples with Equipment Installed All in One year

The examples given here all involve looking at a large body of equipment, one million units, that is all installed in a single year and thus starts out as identical in age and condition. Here, year-by-year and synchronized installation analysis are identical. Chapter 4 will look at actual utility cases where equipment was installed over many years.

These “install one million units and watch as they age and fail over the next seventy-plus years” cases simplify introducing some important characteristics of sustainable point trends, and permit comparison that provides several important lessons. And they are not completely unrealistic. An equipment manufacturer might study in detail the failure and operating history of 1,000,000 units it made and sold in, say, 1984, to try to determine the lifetime trends of its products, etc.

Example 1: Wigits with a Constant Failure Rate

In this example, 1,000,000 wigits are installed in year zero. Wigits have a 1.955% annual failure rate, regardless of age -- in other words a flat failure rate “curve” -- one that does not vary with time in service.¹ This is very unrealistic -- almost all utility equipment has a gradually increasing failure rate over time. But a constant failure rate provides a simple starting point, creates a set of results that will be useful for comparison in subsequent examples, and actually to a very interesting *practical* result that will be explored before the chapter is over.

One million wigits are installed over the course of year zero. During their first year of service, 1.955% of these million units fail and are replaced with units that are, on average, a year newer. This means that at the end of one year of

¹ This particular value was picked so that a certain result works out to a certain value as will be described later.

service, 980,450 remain in service and 19,550 units are replacements. Thus, at the end of that first year of service, the average age of the units is:

$$\text{Average age after one year} = 980,450 \times 1 \text{ yr} + 19,550 \times 0 \text{ yr} = .9805 \text{ yr} \quad [\text{B-1}]$$

A very useful metric to track in aging infrastructure studies is the *average aging rate* – if and how fast the entire set of equipment grew older in a year. It becomes zero at the sustainable point. But in this first year it is rather high: this entire equipment set of 1,000,000 widgets aged by an average of .9805 year in this first year of service.

During the second year of operation, 1.955% of the 980,450 original widgets that survived through year one fail: they are replaced by 19,168 new units, leaving only 961,282 original units, each now two-years old, to go into the next year. In addition, 1.955% of the 19,550 units added the year before (382) fail, leaving 19,168 year-old units. So, there are a total of 19,550 (19,168+382) new replacements. Average age of the units in service increases again, but by not quite by as much as in year 1.

$$\text{Average age} = 961,282 \times 2 \text{ yr} + 19,168 \times 1 \text{ yr} + 19,550 \times 0 \text{ yr} = 1.9417 \text{ years} \quad [\text{B-2}]$$

$$\text{Increase in average age during this year} = .9613 \text{ year}$$




Figure B.5 Average age of one million widgets as a function of time after they are put into operation in year zero, with a 1.955% annual failure rate regardless of age, and where failed units are replaced with identical new units. While numbers differ, any real equipment set where failed units are replaced with new units will exhibit this general trend, with average age increasing a bit less each year until it stops when the sustainable point is reached.

After a third year of service, average age will be 2.8842 years; an increase of only .9425 years. Each year thereafter the original set of units is depleted by more failures, but by a bit less than the year before because there are fewer and fewer units each year that have survived so they can fail at this age. Each year, a larger portion of the units in service are replacements made at some time in the past. But regardless, each year, a total of 1.955% of one million units fail. What varies is that over time fewer are original units, and more and more are replacements, and even replacements of prior replacements.

Figure B.5 shows the average age of the units in service as a function of time after the initial installation. The equipment set starts out with an average age of zero (all one million units are new). Over time, in the manner described above, the average age increases each year, but each year by just a bit less than the year before. It is asymptotically trending toward an average age of 50.00 years – the average age associated with the sustainable point. *This behavior over time is a general characteristic of any large equipment set: while the numbers might vary from one case to another, qualitatively all large groups of equipment exhibit this characteristic, with average age asymptotically approaching a “sustainable age.”*

Distribution of Ages

Average age of equipment is a metric nearly always computed, tracked, and reported within utilities, even though everyone involved knows that it does not communicate all the detail needed; there are many distributions of age that could give the same average. Nonetheless, average age is a useful index and an important element in fitting statistical models to real equipment datasets.

Greater detail is provided by looking at the *distribution of ages*, or estimating it when actual data is not available or dependable (statistical fitting in such cases is one area where sustainable point analysis provides measureable benefits). Figures B.6 and B.7 show how the distribution of unit ages in this example changes over time. When operation begins, in year zero, the age distribution is a spike: one million units of age zero and none of any other age. After one year it consists of 19,550 units of age zero and 980,450 units that are one year old. Figure B.6 shows the distribution at year 2, with 19,550 units of age zero (given the constant failure rate this will never vary), 19,168 units of age 1.000 year, and 961,282 units of age two.

Figure B.7 compares the ten year age distribution with those at 50, 100, and 150 years. At the end of ten years the average age of the million units in service is 8.99 years: over the first decade of operation, the average age of the one million units has increased by nine years. Over those ten years of operation, 195,500 replacements were made, but of those, only 179,168 remain in service, because 16,332 of the replacements have already failed, too. This tremendous number of “replaced replacements” is unrealistic, a function of this example’s flat failure rate curve and its very high early-year failure rates.

Over time there is a complicated dynamic working within this set of one million units, with original units failing, and some replacements failing, and even




Figure B.6 Distribution of ages of the one million units in service at the end of two years.




Figure B.7 Distribution of ages for various other times after initial start up.




Figure B.8 Distribution of age at the sustainable point, which is not reached for centuries (an unrealistic characteristic of the failure rate curve that does not increase with age in this example). Horizontal axis here is 300 years compared to 150 for the previous two plots.




Figure B.9 Same data as in Figure B.8 plotted in a normal rather than log scale. This is the characteristic “age profile” associated with constant failure rates. It may be blurred in actuality because all equipment was not installed in a single year, but proper analysis and data manipulation can account for that.

a few replacements of replacements of replacements failing each year. But the net effect of the shape of the distribution of ages (Figure B.7) is quite easy to explain: every year the “spike” of original units shrinks by 1.955% of its value and moves one year farther out – what was 380,056 49 year old original units in year 49, becomes 372,626 50 year old units a year later, and leaves all other values stay the same (this is because of the constant failure rate as a function of age, and not normally seen in the real world).

This trend of the spike moving out and shrinking continues forever (Figures B.8 and B.9), until the original spike of one million units eventually disappears. This does not happen in this example even after three centuries, due to the unrealistic nature of the failure rate curve used here (the distribution at 300 years looks like Figure B.7’s plot of the sustainable age distribution, but the “spike” is only 2,425 units left at year 300, there being that many units out of one million expected to last three centuries

What is Common and Generalizeable About this Example

Over time for any set of equipment, with any type of failure rate curve, there will be:

- Fewer original units left in service.
- A higher proportion of replacements are in service.
- A higher proportion of replacements of replacements in service.

And thus, all equipment sets exhibit these characteristics over time

- Average age trends a to an average “sustainable age.”
- Aging rate begins at a high rate and decreases over time.
- Even if the equipment is all installed at one time, eventually there is a wide distribution of ages of equipment: *age distribution is not primarily a function of when equipment was installed.*

What Is Not Common or Generalizeable

First, age distributions almost never behave as in Figure B.7, with a spike just moving to the right each year and leaving behind it the same shape of distribution year after year. Instead, most real distributions vary in shape slightly everywhere along their curve, “evolving” in shape (the next example will show this).

And finally, there is a “coincidence” in this example that is *not* general and should *never* be assumed, but which many people mistakenly see in the data presented here and which many expect, and few even insist upon, as having to always be the case. In this example, the sustainable age, 50 years, and the

average failure rate, 1.955% are reciprocals:²

$$1/(\text{failure rate}) \approx \text{sustainable age} \quad [\text{B-3}]$$

$$1/(2\%) \approx 50 \text{ years}$$

The authors realize that, intuitively, it might appear that this should always be how things work out. *But this is almost never the case.* Usually $1/(\text{failure rate})$ is *much* greater than the sustainable age, and many engineers and managers new to aging infrastructure analysis find puzzling (this will be explained in the next example). A few almost refuse to accept it. But the fact is that the two values, average age and failure rate, are reciprocal *only* when failure rate is not a function of time in service, which should almost never be the case.

Example 2: A More Realistic Failure Rate Curve

In Example 1, the expected failure rate did not increase with the age of the units, which is quite unrealistic because that would mean that either deterioration is not taking place over time (never the case) or that it is but has no effect on unit reliability (again, just not ever the case). In this example failure rate does increase with age – quite dramatically in fact. The failure rate curve shown in Figure B.10 is based on actual utility data, with its values only slightly modified by the authors so that the sustainable age works out to exactly the same 50.00 years as in Example 1, which will permit several revealing comparisons later in this section.

Figure B.10 shows the failure rate for Ever-There Corporation's Dura-Widgits, which are functionally identical to the common widgit of Example 1, but made of a breakthrough synthetic material – dried and spun fluff, reinforced with carbon fiber and then infused with a long-polymer cross-linked polyethylene epoxy. As a result they are much more robust than standard widgits, so that in their first full year of service Dura-Widgits have an expected failure rate of .0250% -- or one quarter of one tenth of one percent. Out of 1,000,000 to be installed, only 250 are expected to fail in their first year of service, as compared to 19,550 widgits in Example 1. Clearly these Dura-Widgits are a real boon for a utility – they cut replacement rate in the first year after initial installation by more than 75.

But as Dura-Widgits age their expected annual failure rate increases, at first very slightly each year but then by greater increments from year to year. The actual formula fitted and adjusted here is rather simple: the annual expected failure rate shown increases by 5.339% of its prior year's value each year. In year two it will be 105.339% of .0250%, or .0263%. In year three it will be

² Note: the failure rate *should* be 2% in this example: with a constant failure rate the sustainable age should be the reciprocal of the failure rate. But 1.955% is needed in this example because the analysis is done on a yearly basis as explained earlier. Not all years installed in year X were installed on the same day, etc, but they are treated that way here.




Figure B.10 Failure rate curves as a function of time in service for Dura-Widgits (solid line) as compared to Example 1's normal widgits (dashed line). See text for details.

105.339% of .0263% or .0277%. And so forth. Looking at Figure B.10, one can see that Dura-Widgits provide a rather dependable 70-80 years of durable service, after which failure rate begins to increase significantly.

Aging Characteristics Caused By Typical Utility Equipment Failure Rate Curve Shapes

The failure rate curve shown in Figure B.10 is quite similar in values, escalation rate, and shape to many the authors have encountered in work with actual utility data. Despite its slowly escalating trend, its value is quite low for decades, because it starts out at a *very low* value.³ After 50 years of steady escalation, it has reached only one-third of a percent – still just a sixth of the value in Example 1. It reaches half a percent 8 years later, a full percent at 71 years, and finally passes Example 1's 2.00% value in year 85. At age 100 years it reaches 4.54%, and exceeds 50.00% at 146 years. It reaches 100% at age 160, so there is no expectation at all that Dura-Widgits will last more than 160 years.

The effect that the difference in failure rate curve is dramatic, as illustrated in Figure B.11, which compares Examples 1 and 2 with respect to the number of "survivors" – original units still in service – as a function of operating time. In

³ This example deliberately neglects infant mortality issues with equipment.

Example 1, the number of original units still in service drops below 900,000 in just six years – meaning more than one in ten has already failed and been replaced. By contrast in Example 2 the count of original units stays above 90% until year 61.

But as the Dura-Widgits age, their failure rate escalates. The failure *rate* for this example “catches up to” Example 1’s at age 85: after 85 years of service the average Dura-Widgit has an expected failure rate of 2%. Despite this, the total number of original units remaining in year 85 is still far higher than in Example 1, because failure rate was so much lower for over eight decades prior to that time. At age 85, Example 2 has 679,889 survivors, Example 1 had only 186,708.

Every year after 85 years in service, the failure rate in Example 2 increases a bit more. But the number of original units that have failed in Example 2 does not surpass that of Example 1 for another 35 years, until year 120 (both have about 95,000 original units, or just less than 10%, left in year 119).

Also worth noting is that the number of replacements in service versus the number of replacements that have been made as a function of time. At year 119, both examples have nearly the same number of original units in service (95,416 for Example 1, 95,331 for Example 2). This means that both examples have roughly 904,500 non-original units – replacements – in service. Yet at that time, the utility has replaced a total of 2,326,450 units in Example 1, compared with




Figure B.11 Number of original units remaining as a function of years since initial installation. Note that unlike previous unit count plots this is not a log vertical scale.

only 942,999 units in Example 2. After 119 years, over 96% of all replacements ever made are in service in Example 2, versus fewer than 39% in Example 1: its failure rate means that are many more replacements of replacements.

Significant technically, but of little practical importance, is the fact that normal widgits actually have a much higher likelihood of lasting a long time. While the failure rate for Dura-Widgits reaches 100% at 160 years, so there is no possibility of them surviving past that age, it is so high in the two decades prior to that that only 1 out of a million is expected last to 150 years. By contrast, with its straight 2% annual failure rate, Example 1 still has 51,739 of its original one million widgits left in service at 150 years.⁴ At the end of 300 years, it will still have nearly 2,500 original units in service.

Aging Rate in Example 2

Figure B.12 shows the average age of installed units over time for this example, comparing it to the same statistic from the previous example. Readers will naturally be curious about the oscillations, which will be discussed in a few paragraphs. For the moment, this examination will focus on what happens during the first fifty years of operation.

From the beginning of their service lives, the average age of the 1,000,000 installed Dura-Widgits increases much faster than did the average age of the normal widgits in Example 1. Far fewer Dura-Widgits fail and have to be replaced with new units each year, so there are far fewer newer units to bring down the average age. In their first full year of service, only 250 Dura-Widgits are expected to fail as compared to 19,555 in Example 1. These failed units are replaced with 250 new units of the same type. At the end of the first full year of service, there are 250 new Dura-Widgits and 999,749 that are a year old. The average age is .99975 years. Note that, with a substantially lower failure rate at this age, they aged faster than the units in Example 1: average age increased by .99975 as compared to .9805 year.

At the end of ten years, the average age of the Dura-Widgits in operation is 9.98 years, compared to 8.99 years in Example 1: after only a decade Dura-Widgits are providing a full year of additional service life. At fifty years of operation the average age of an installed Dura-Widgit is 47.24 years compared to only 31.46 years for normal widgits. Dura-Widgits provide an average of more than 15 years of additional service life over their first half century of application. No doubt Ever-there Corporation would advertise this heavily.

⁴ The authors say this is of little practical importance for two reasons. First, it is so far into the future, that the net present value of continued service in that timeframe is nil. More important, this analysis considers only failure from normal deterioration. Over a period of 150 years, the likelihood of a storm, forest fire, earthquake, flood, or other “external” event causing the end of life of a unit is greater than this probability of it lasting to 150 years.

Oscillations

In some cases the trends oscillate in time before settling down to the sustainable point. Oscillations are for reasons that are not, at this point, obvious. Readers and students in the authors' seminars initially are intrigued by them to the point of being distracted from more important basic lessons of sustainable point analysis, but almost invariably dismiss them as interesting but unimportant. In fact, oscillations are very important from a practical standpoint, as they sometimes occur with cable.

The oscillations shown in Figure B.12 are exacerbated by the unrealistic nature of this example – all the equipment was installed in a single year – but oscillations do occur in real utility systems and are much more common than is often recognized. Here, while the sustainable age is 50.00 years, initially the average age initially shoots past 50 years and continues up to over 64 years (at year 81) before plunging to only slightly more than 35 years (year 125), etc. Over the next decades the average age of installed Dura-Widgit will oscillate in ever gentler increases and decreases toward a final value of 50.00 year average age.

Oscillations occur only when there is a *combination* of two characteristics in the equipment set. *Both* must be present for oscillations to occur:

- a) The failure rate curve is not constant or even linear: at some point the failure rate increases rather steeply as a function of time in service.




Figure B.12 Average age of Dura-widgits (solid line) compared to standard widgits from the previous example. Both equipment sets have a sustainable point average age of 50.00 years.

- b) All or a substantial portion of the units were installed over a short period of time. This is seldom the case when equipment and aging is analyzed over an entire utility system. But it is often true of the equipment in a specific sub-system or area of a system, such as all cable installed in an office park built out in the late 1960s.

Table B.1 lists typical characteristics of oscillations. The most extreme oscillation possible will exist if both of conditions listed above are in their extremes: if the failure rate “curve” is a step function and all equipment was installed in a single year. For example, suppose the failure rate “curve” is zero until past age 40, at which point it becomes 100% (all of the units fail in year 41). If one million units are installed in year 0, the average age curve (the equivalent of Figure B.8 for that example) would show a saw-tooth pattern. It would start at zero in year zero and increase at a rate of exactly 1.000 year per year up through an average age of 40 in year 40. In year 41, all units would fail and be replaced with units of zero age and the trend would start to repeat itself endlessly. Important points about oscillations are given in Table B.1

Table B.1 Key Points About Aging Trend Oscillations

1) Oscillations occur in real utility systems.
2) Whenever they occur they affect all aging statistics: average age, failure rates, overall operating costs, and all other age-related statistics oscillate in a related manner.
3) They are periodic about the sustainable point of the equipment set and asymptotically die down, but often only gradually.
4) The waves are not necessarily symmetrical, but can have a somewhat sharper rise than fall, or in rare cases, vice versa.
5) Some of the “asset wall” shapes of age distributions seen in actual utility systems are an oscillation (usually the first).

Oscillations That Occur in Real Utility Systems

Figure B.13 shows data from a portion of a metropolitan utility system in the eastern United States. Several decades earlier the city government, in partnership with private investment companies, “redeveloped” a 34-block tract in an area of town that had deteriorated into largely abandoned neighborhoods of old warehouses and piecework factories. The master plan included high rise offices and condominiums, large institutional buildings like hospitals, a new police

headquarters and school administration buildings and several similar government offices, a sports complex, and a convention center and three hotels.

Although total build out for the redevelopment plan took more than fifteen years, as is often the case the utility installed the bulk of its primary cable for this area in a period of slightly less than six years -- thousands of feet of new ducted UG primary cable, splices, buses and vaults, installed, if not in one year, over a relatively short span of time compared to the cable useful lifetime. As a result, the primary distribution “sub-system” in this area – comprising miles of cable and its associated equipment, and serving over 110 MW of peak demand – exhibits very definite oscillations.

Other portions of this utility’s UG cable system exhibit similar oscillations peaking in different years because they were built out at different times, and often with different periodicities because they include different types of cable and equipment. Oscillations are not seen in overall statistics for the whole utility system because there is a type of “diversity of peaks” working within the whole – different areas of the system, of different ages of initial installation, are cycling up and down at different times. Overall, the oscillations tend to cancel one another and smooth out.

Norm for small areas of most utility systems. In almost all utility systems, different areas of the system were “built out” during specific times – there are areas that developed in the 60s, in the 70s, the 80s, the 90s, and in areas where construction activity is intense now. Usually, the utility’s electrical equipment in each area is dated to that last development period: it was installed at the same time.




Figure B.13 Oscillation of failure rates and operating costs for a region of a utility system in the eastern United States served by an all UG distribution network. Trends shown are fitted and predicted data, black dots indicate actual utility operating cost data.

Oscillating failure/replacement cost behavior as shown in Figure B.13 *is the* and to serve that development. Thus, every such region will exhibit a certain amount of historical and predictable future oscillations in failures and operating costs. The ability to track, anticipate, plan, and manage equipment, failures and replacement on a local basis, for all localities in the system, is a key part of good aging infrastructure management.

Age Distribution in Example 2

Figure B.14 shows the distribution of ages of installed Dura-Widgits at various times during the first century of operation. At ten years, there are of 996,801 original Dura-Widgits left, and fewer than 4,000 replacements have been made to date, all of them, of course, having ages less than ten years in service. (Readers should keep in mind this plot has a vertical log scale. The slight drop in the height of the “spike” at year ten as compared to year zero is equal to the area under the curve for all nine years leading up to it.)

Nothing that is really qualitatively different happens until about year 100. Each year until then, the spike moves to the right by one year and decreases in value only slightly, but by a bit more each year. Cumulatively, over a full century, the escalation in the failure rate begins to show its effects: only 399 replacements are needed in year 10, growing to 866 in year 25, 3,062 in year 50, 9,647 in year 75, and 19,924 in year 100.

In example 1, the number of replacements needed in any year was always 19,555 per year, due to its constant 1.955% failure rate regardless of age. In this example, annual needed replacements first exceeds 19,555 in year 99 (19,690 replacements needed), even though the escalating failure rate for Dura-Widgits first exceeded 1.955% nearly 15 years earlier, at 85 years in service. At 99 years, failure rate is actually more than 4% and climbing at nearly a quarter percent per year. But as discussed earlier, it is taking time for Example 2 to “catch up” with Example 1.

Part of the reason is the *distribution of ages* in the two examples as they both approach 100 years in operation and the impact that has on the number of “replacements of replacements” needed. In example 2, fewer than 4,000 replacement units are needed in the first decade of operation compared to 199,550 in example 1. Thus, ninety-nine years later, example 2 has only 1/50th as many units that are between 90 and 100 years of age: so there are far, far fewer old replacements that fail and need to be replaced. And while there are still many original units (over 425,000 remain at 100 years), there are simply not quite enough units, that are old enough that their failure rates all add up to greater than 2% until year 99.




Figure B.14 Distribution of ages at various points during the first 100 years of operation. As time passes the spike representing the original units moves to the right and slowly decreases in value, dragging a distribution of various ages of replacements behind it.

The Failure Curve Tipping Point: A Change in the Character

Beginning around year 100, a number of very interesting changes take place in the character of equipment aging in Example 2. By this time (year 100) the failure rate of the original units has been above 2% for fifteen years, and is now above 4%: more than one out of every 25 original units is failing each year. Failure rate will climb to over 5% by year 104 and continue to grow from there. As a result, in the following 25 years the spike representing the original units in Figure B.14 will “melt” as the number of original units drops by over an order of magnitude, from more than 425,000 in year 100 to less than 39,000 in year 125. But despite this, *the total number of replacements required each year begins to drop*; it peaks in year 104, at 20,424, and then decreases significantly each year, falling below 10,000 by year 125.

What the utility is seeing is the first of the oscillations: the number of failure replacements required each year is actually dropping even though there are many old units still in operation. As stated earlier, Chapter 5 will deal in detail with the oscillations and their causes and characteristics in actual utility systems. But for the moment what is happening during this very brief period after a century of operation began provides a useful lesson.

In year 104 the failure rate curve reaches a “tipping point”⁵ *Its value exceeds*

⁵ The authors apologize for using a term they don’t like themselves, but there is no better mathematical example of a tipping point than this.

its annual growth rate. As described earlier, here annual failure rate grows at a compounded rate of 5.339%. In year 104, the failure rate, which started out at only .25%, finally grows to over 5.399% (it is 5.30% in year 103 and 5.59% in year 104). The fact that this happens in year 104, and at a value of 5.339%, is entirely a function of the escalation rate and curve shape selected by the authors. *But in real systems this always happens at some point.* And from there on the total number of annual failures for the original equipment of this age will go down.

There is always a “tipping point” or *age of maximum failures* after which the number of failures for any group of equipment continuing to age will drop. This tipping point is entirely a function of the failure rate curve shape and occurs at the point where failure rate first exceeds the annual growth rate of the failure rate curve.

Going into year 104, there are 347,687 of the original one million units still in operation. Their failure rate is 5.59%, meaning that 19,427 of them are expected to fail, leaving 328,260 to go into year 105. In year 105 the failure rate escalates to 5.89%, but it will be applied to only 328,260 units, meaning slightly fewer failures are expected to occur – just 19,320. Failures are always the product of number of units times failure rate. Prior to year 104, the latter grew more each year than the former shrank, so the net product increased each year. From year 104 on, the opposite occurs, and the total number of failures in the original set of units begins to shrink, at first slowly and then more quickly from year to year.

While the number of original units that fail is shrinking rapidly in year 104, there are still enough of them that failures among them dominate the overall failure statistics for the entire set of one million operating units, even though there are more replacements in service at this point than there are original units. This is due to the distribution of ages of the units. Even at 104 years into operation, there are still few *early* replacements – ones that have reached an advanced age where their failure rate, too, is above 2% and climbing rapidly. Overall, there are just not enough in this category to cause enough failures to make up for the drop in failures among the original units. As a result, from year 104 to year 125, the number of annual replacements needed for the entire population *drops* by a factor of more than two.

While this is happening, and because of this, the shape of the age distribution diagram shows a considerable change in character. Up to 104 years, if one neglects the spike of original units, the largest age group of equipment is always the latest replacement units – they outnumber units of any other age, except for that spike of original units. But beginning in year 105, a “hump” begins to occur in the distribution of (Figures B.15 and B.16). Prior to year 105, excepting the spike of original units, the most populous age was always zero- the new replacements. Now this is not the case. The system is on the downside of an oscillation.




Figure B.15 Starting at around 100 years of operation the character of the age distribution begins to change: the spike of original units starts to shrink rapidly and a “hump” begins to develop in the distribution of relatively newer units.




Figure B.16 Data plotted on a linear rather than log scale shows the large difference in distribution by age.

Both widgeit examples started out with many units installed in a single year – of a single age. Both end up, decades later, with a wide distribution of units of all ages, nonetheless. This distribution is the result of the slow, decades-long deterioration of the original equipment set and its gradual replacement over time with newer units. Eventually, replacements fail and are replaced, etc., until the age distribution reaches that for the sustainable point, after which it does not vary.

It is worth noting that this same distribution would be reached if the units were not installed in the same year. To see this, suppose they were installed half in one year, and half ten years later. Each half would change and its distribution of ages evolve as the other (and as in the original, all-million-installed-in-one-year example), but one lags the other by ten years. One would reach its sustainable point first after which its values would not change, the second lot would reach that point ten years later, and after that point they would have an identical age distribution as compared to if they had all been installed at one time.

The point here is that the basic sustainable point behavior of an infrastructure – how it ages toward a statistical distribution – does not depend on when the equipment was installed, or the particular age distribution of that equipment in a particular base study year (e.g., 1990). It is a function *only* of the equipment, the way it is used, and the way it is cared for, repaired, and replaced.

As mentioned in earlier, often a utility or industrial power system owner does not care about the sustainable point – it is far into the future, what matters most is the exact characteristics, aging, and failure characteristics now, or the next few years. But such situations are never generalizeable for “big picture” planning and comparison of major differences in policy or operating procedure.

The sustainable point behavior of an infrastructure is the most generally useable characteristic of its overall potential for aging and deterioration-related problems caused by the aging.

B.4 METRICS FOR MEASURING AGING INFRASTRUCTURES

Average Age and Average Condition Are Not Highly Useful Metrics

This section will examine various measurements – metrics – that can be applied to an aging equipment set to evaluate and track its characteristics and help model and predict its future behavior and performance. As stated in the introduction, average age is not a particularly useful metric. The authors do not recommend that readers and their companies abandon use of it. Average age is so widely reported, and such an intuitive and expected measurement, that an owner pretty much has to report and track it. But not much importance should be given to it.

Table B.2 shows various numerical statistics that have been computed for Example 1 (normal Widgeits) and also for Example 2 (Dura-Widgeits). This section will discuss each in turn and its use in aging infrastructure analysis and modeling.

Table B.2 Statistics for Example 1 and 2 at 25, 50, and Infinite Years

Measurement	At 25 Years		At 50 Years		At Sust Pt.	
	1	2	1	2	1	2
Average equipment age	19	25	31	48	50	50
Average remaining lifetime	50	8	50	48	50	50
Average failure rate	2.0%	.03%	2.0%	.30%	2.0%	1.1%
SAIDI contribution ⁶ , minutes	3.3	.16	3.3	0.6	3.3	2.0
Equivalent Failure Age, years	50	24	50	48	50	72

Average Age

Although Examples 1 and 2 have the same average sustainable age, they have wildly different expected service lifetimes for equipment. Widges (example 1) have an expected lifetime of 50.00 years. In any case where failure rate is constant as a function of service life then:

- The sustainable average age and expected service age are the same.
- Expected service lifetime and failure rate are reciprocals: $50 = 1/(2\%)$.

The expected service life of Example 2's Dura-Widges is 92 years, or just over 80% more than the expected lifetime of a normal Widge. This is actually a remarkably low value considering the very different failure rates the two systems exhibit in service.

During the first year of service, 78 times as many Widges will fail.

During the first decade of use, 61 times as many Widges will fail.

During the first half century, 17 times as many Widges fail

During the first full century, 2.5 times as many Widges fail.

Clearly, the trend over time is toward parity in failure rates, but the Dura-Widges will never cumulatively catch up. After 320 years – exactly twice the theoretically possible maximum lifetime of a Dura-Widge – the ratio of total replacements needed is 2.1:1. Given long enough (1,000 years), the ratio settles down to just over 1.8 – the ratio of the expected service lifetimes of the two devices. Due to the oscillations that in Example 2, there periods (between 100-125 years, etc.) where a few more Dura-Widges fail than Widges, but overall and on average about 80% more Widges will fail as compared to Dura-Widges.

Thus, two systems that have identical average ages in the long (sustainable point) run, and have expected equipment service lifetimes and failure rates that

⁶ Computed as: average of 180 minute outage caused for 3 customers per transformer for every failure, divided by 1,000,000 transformers times 3 customers each.

differ by a significant margin.⁷ This alone demonstrates only that average age is not a particularly dependable statistic for study of the long run situation.

The objection the authors have to the use of average age as an infrastructure metric is that it can be very *misleading*. Figure B.17 repeats Figure B.12 so it is convenient to consult, and shows average age as a function of time for the two examples. The reader should consider:

1. *When tracking any one set of equipment, increasing average age generally means the situation is getting worse.* In Example 2 as average age increases it indicates the situation is, in some sense, getting a bit worse. As operating time passes and the Dura-Widgit infrastructure's average age increases, operating costs and failure rate go up in company with average age. Those metrics might not increase proportionally to average age, but they increase steadily as it increases.

Given the choice, any utility would prefer to own Example 2's equipment set when it has an average age of, say, 15 years, than when it has an average age of 45 years, etc. If a utility were considering the purchase of this equipment set, it would be willing to pay more for it if its average age were 15 years, and its failure rates and operating costs lower, than if its average age were 45 years. Here, and in just about any real-world situation, as a measure of change over time, an increase in average age indicates a deterioration in overall quality.

2. *But when comparing two different types of equipment, or two different systems, average age tells one little.* Example 1's Widgits reach an average age of 25 years after 35 years of service, whereas Example 2's Dura-Widgits reach an average age of 25 years at essentially 25 years of service. The two examples are, as evaluated by this metric, equal. Yet Widgits are failing at 22 times the rate of Dura-widgits at this point.

This is literally a textbook example, and as such rather dramatic in the contrast in draws, but in actual utility study, a change in proposed inspection, repair, or replacement policies, or a shift to a different type of equipment, might replicate in small measure the situation demonstrated here: average age will be misleading as a metric in such studies.

Average age, then, can mean one thing in some situations and another in others. It is best to avoid using such a potentially confusing metric.

⁷ The authors define "significant" here and elsewhere as a margin that would make or break a utility company. An increase of 80% in operating repair and replacement costs for its T&D equipment would ruin the finances of most electric utilities.




Figure B.17 Average age of Dura-Widgits (solid line) compared to standard Widgits from the previous example. Both equipment sets have a sustainable point average age of 50.00 years.




Figure B.18 Expected remaining service life as a function of years already in service.

Expected Lifetime and Expected Remaining Lifetime

Widgits have an expected service lifetime of 50 years. Dura-Widgits have an expected lifetime of 92 years. As a reasonable comparison of total value (regardless of the cost of the items) this is a useful comparison: it identifies which last longer and provides a rough idea of the ratio to the benefits one can expect.

Figure B.18 shows the change in expected remaining lifetime of an installed unit as a function of time since its installation. When new, Widgits (Example 1) have a 50-year expected lifetime, and since failure rate is not a function of time already in service, that is what a unit of any age has, too. By contrast age has a great deal to do with both the failure rate and the expected remaining lifetime of a Dura-Widgit. When new expected remaining lifetime is 92 years. After 50 years in service it has dropped to 46 years remaining (which counting time already in service, equals 96, or a four year increase in expected total lifetime given that the unit is a survivor of 50 years). If a Dura-Widgit survives 100 years of operation, it has only 12.5 years of expected lifetime left (an increase of 20 years in expected total service life given it has already survived a century).

for Widgits and Dura-Widgits.

Figure B.19 compares a slightly different metric, the average *remaining service lifetime for all members of the infrastructure*. Again, Example 1 is a constant 50 years regardless of time: whether when new, or after ten years of operation, or 100, or at its sustainable point, the average remaining lifetime of all the units in service in Example 1 is 50 years. By contrast, the remaining lifetime of the Dura-Widgits changes over time, individually as in Figure B.18, and here, as a group. When operation begins and all one million units are new, average expected remaining lifetime is 92 years. That gradually falls as the bulk of units (those that don't fail) age. This metric oscillates as all this examples statistics do, and reaches 50 years, the same as example 2's at the sustainable point.

Expected remaining lifetime is a very useful value to compute and use in the management of individual or small groups of units, as will be shown elsewhere in this book. But *average remaining lifetime* of an entire equipment set is, again, not a terribly revealing metric: here, at the sustainable point, and at many other points in between, it gives a distorted perspective on which example is better, or rates them the same.

Average Failure Rate

Average failure rate (AFR) of an infrastructure is just the total number of failures in a year divided by the total number of units in that year. In example 1, with its unrealistic constant failure rate as a function of time in service, AFR is equal to the constant failure rate of 1.955%. In example 2, and in almost all real-world situations, average failure rate changes with time, but it is not equal to the failure rate curve value. Figure B.20 shows that for Example 2. In the first year after installation, the one million Dura-Widgits' average failure rate is identical to the failure rate curve value for year 1: all of them are new and share the same year 0




Figure B.19 Expected average remaining service lifetime of units in operation, as a function of time since operation began.




Figure B.20. For Example 1's Widgits with their unrealistic constant failure rate over time, the Average Failure Rate is the same (dashed line). Example 2's Dura-Widgits with their realistic failure rate curve exhibit a characteristic more like most utility equipment. After installation the average failure rate (broad solid line) is very nearly the failure rate curve's value (thin solid line). See text for details.

failure rate. As time passes however, most stay in service and their failure rate increases with their age. But a few fail, and are replaced with newer units with lower failure rates than would otherwise be the case. As a result, the average failure rate for the entire set of one million units falls below that given by the failure rate curve for total time since operation began. At fifty years after operation, a good many units are replacements, with lower failure rates, and average failure rate for the entire set is .31%, as opposed to the failure rate curves value of .34% at 50 years. Over the long term, average failure rate, like all other statistics, oscillates in waves of decreasing height.

Equivalent Failure Rate Age (EFRA)

After 80 years of operation, a bit more than 25% of Example 2's one million Dura-Widgits are original units, the rest being replacements installed sometime during those 80 years. Most replacements, in fact, are rather new: the failure rate curve was so low early on that few replacements were made in the first two decades. There are few "old replacements." In fact, the average non-original unit is a bit less than five years old: most replacements have been made only during the most recent decade.

As a result, mix of Dura-Widgits in operation is about 75% original, 80-year-old units, with the 1.6% probability of failure of units that age, and newer units – most much newer – with failure rates as much as $1/50^{\text{th}}$ as low. Average age is 65 years in operation. Average failure rate is 1.174% -- 1,174 units of various ages, mostly original units but some new, too, will fail in year 80.

This average failure rate of 1.174% is equivalent to the failure rate curve's value at 71 years. The units in operation might be an average of 64 years old but they are failing as if they are an average of 71 years old: a unit 71 years old has a failure rate of 1.174%. The entire population is failing at a rate equivalent to an age roughly 10% greater than its average age because of the non-linear shape of the failure rate curve. Its slope increases with age, and that means that as units age their failure rate increases more than proportionally.

This type of behavior is a common situation in actual utility analysis cases, because the failure rate curve for Example 2 is quite like most utility failure rate curves, which increase exponentially.

Note the complexities here, in what is, essentially, a real-world case with a very realistic failure rate curve. First, average age gives no dependable indication of failure rate trends. After 80 years of operation average equipment age is 64 years but failure rate is equivalent to units that are, on average, 71 years old. Second, and much worse, over the next 24 years, average age will improve, from 64 years to just 50 years average age, *but equivalent failure rate age will continue to increase*. At 104 years of operation, average age is only 50 years, but average failure rate is 2.04% -- units are failing with a rate equivalent to that of an 85-year-old unit.

This behavior is caused by the non-linearity of the failure rate curve, and because most real-world failure rate curves have a shape something like Example

2's, this behavior is common, occurring in more than 90% of utility cases the authors have seen, in the time periods 35 to 80 years after installation depending on specifics of the case. This reason, more than any other, is why the authors do not rely on average age as an indicator in most of their aging infrastructure work. It not only tells one little useful information, it can be misleading.

SAIDI Contribution

In most utility applications, failures of the equipment being studied lead to service outages which contribute to SAIDI (System Average Interruption Duration Index) – the total time without power experienced by the average utility customer in the year. Most utilities compute SAIDI contributions for various equipment sets and parts of their infrastructure.

The one million units in these examples is a number roughly equal to the number of service transformers in use in the largest US utility systems. Calculating SAIDI values here as if they were service transformers assumes: three customers per transformer and an average outage time (that to report the outage, respond, replace the unit and restore service) of 180 minutes. Example 1's Wigits give a constant-over-time SAIDI contribution of 3.3 minutes, or a bit less than 5% of the 75 minute SAIDI target that the authors encounter most often in their work with utilities. By contrast, new Dura-Wigits provide a negligible 10 seconds (.16 minute) rising to 36 seconds at 50 years, and giving 2 minutes at the sustainable point. Not shown in the table: at the height of the first oscillation Dura-Wigits would briefly exceed Wigits in SAIDI contribution (from year 99 to 110), peaking at 3.5 minutes contribution before plunging again at the bottom of the first oscillation (year 139) to 44 seconds.

This and similar reliability statistics are all affected by aging infrastructure failure trends, and most utility aging infrastructure management programs have to pay attention to service reliability impacts and plan to incorporate SAIDI and SAIFI targets. However, SAIDI, SAIFI, and other reliability statistic impacts are outcomes of an aging infrastructure. They often need to be analyzed and forecast for planning purposes, but are not part and parcel of the infrastructure analysis itself. More basic metrics, linked to characteristics of the equipment itself, are most useful in understanding the infrastructures' aging and what the utility can do to manage that in its best interests.

B.5 EFFECT OF DIFFERENT FAILURE RATE CURVE SHAPES

Multiplying the Failure Rate Curve by a Constant

When failure rate is constant (Example 1) the impact of doubling or halving its values is quite simple and intuitive: double the failure rate and the sustainable age is halved and the failure rate at the sustainable point is doubled. Triple it and the sustainable age is cut to one third and the sustainable failure rate is tripled, etc. Cut it in half, and sustainable age doubles and sustainable failure rate drops by half. It's that simple.




Figure B.21 Scaling an exponential failure rate curve has surprisingly small effect on long-term failure behavior of an infrastructure. Multiplying the entire failure rate curve by two cuts sustainable age by only about 20% and increases sustainable failure rate by less than 5%. It takes a multiplier of 16 to cut sustainable average age in half. See text for details.

The impact of scaling an exponential failure rate curve is quite different and much more complicated, as illustrated by Figure B.21. Double all values in Example 2's failure rate curve, and sustainable age drops from 50 to just over 41 years, while failure rate at the sustainable point increases from 1 1/4 % to a bit less than 1 1/3 %.

Triple the values and those values change to 39 years and a bit more than 1 1/3 % respectively. Increase them by an order of magnitude, and they change to only 29 years and 2% respectively -- despite a tenfold increase in the curve's values, the impact on both values (and all others) is less than twofold. With Example 2's failure rate curve, it takes a scaling factor of *sixteen* to cut the sustainable point average age in half -- which makes the sustainable point failure rate an average 2.3%, still not even twice the base value.

Scaling an exponential failure rate curve will affect the oscillations, too, but again, not in a linear fashion. Scaling the failure rate curve upward makes the oscillations a bit shorter in "wavelength" and a bit more dramatic in relative magnitude: doubling the curve cuts oscillation duration by about 30%. The first

oscillation peaks at 71 years in this case. Its peak value is about 33% higher than the sustainable point, not 20% higher as in the base situation. Tripling the failure rate curve results in even shorter duration oscillations with higher relative peaks as compared to the sustainable values, and so forth.

There are two reasons for this very non-linear behavior, both important for aging infrastructure planners to bear in mind. First, one cannot increase failure rate beyond 100%. This fact alone assures non-linearity of results. But most important is the period leading up to reaching 100%. In this period any exponential curve has very high values: over 50%, then over 75%, and then over 90%, all in a fairly short period of time just prior to its reaching 100%. Very little equipment survives this period to make it to that final year where failure is guaranteed. Failure behavior over time and the sustainable point values are dominated by what happens in this time period.

Thus, when one scales an exponential failure rate curve as defined as Example 2, the 100% period moves forward in time and perhaps sharpens its rise somewhat (i.e., it shortens the time that it takes to go through that high-value period). The overall impact can be summarized in this way: The very low failure rates early on are increased, but they were so low that these increases make little overall impact, so the biggest impact is that the time to reach really high values near 100% moves forward a few years, resulting in a noticeable but not proportionate reduction in sustainable age and increase in overall failure rate.

Summary: The less exponential a failure rate curve, the more it is affected by scaling (multiplying by a constant). If it is a constant, its failure statistics will be linear or have an inversely linear affect. But if it is exponential, the impact will be less than proportional: the higher the order (exponential rate of increase) the less sensitive it will be to multiplying it by any particular factor. To see this point, the reader can imagine the ultimate is sharp or high-order failure rate curves, a step function. Consider a failure rate curve with a value of 0% out to 50 years and then a value of 100% afterwards. Its sustainable point and all other failure statistics it generates will not be affected at all by any multiplication.

Generally, Failure Rate Curve *Impacts* Don't Add

If one takes the failure rate curves from Examples 1 and 2 and adds them together, the resulting failure rate curve begins with a value of 1.980% and increases exponentially as did Example 2's curve. This curve gives a sustainable point failure rate value of only 2.36%, far below the 3.2% sum of the failure rates for the two examples. Adding 1.955% to Example 2's curve increases its sustainable failure rate by only about 60% of what that 1.955% value gives alone.

As was the case with scaling, adding a constant to a highly exponential failure rate curve does not proportionally affect it, and the more exponential the curve is, the less sensitive its results will be to the addition of a scalar. The reason, again, is what this does to values at and near 100. With 1.955% added to every value, the curve still reaches a value of 100% in year 160. As values for Example 2's curve tend toward 100%, the rate of increase is about 5% per year so 2% does not

make a big increase here - the curve reaches 100% about 5 months earlier and still in year 160. Adding nearly 2% to all makes very little relative to increase the high-value parts of the curve: failure rate in year 159 is already very close to 100%, this just pushes it nearly to the limit.

Rate of Change of Average Failure Rate Age

Since average age is not a particularly useful measure, tracking its rate of change usually does not provide a good deal of meaningful information either. (The annual increase in average age was tracked in earlier figures mainly to demonstrate convergence toward the sustainable point.) Both examples started out with an annual change in average age each year just shy of 1.00 and ended up at 0.0 at their sustainable points.

Rate of change of equivalent average failure age (EFRA) as shown in Figure B.22 is much more meaningful as an indicator and predictor of it and how badly the situation is changing for the worst (or for the better). This example shows complicated behavior, but this can be forecast (it is just the result of the failure rate curve's shape).




Figure B.22 Annual change in average failure rate equivalent age for Example 2 exhibits a complicated behavior including oscillations, all a function of the *shape* of the failure rate curve, and perhaps best shows if and how fast the situation is getting better or worse. Complicated as this may seem, trends like this are *predictable*.

B.6 READING AGE DISTRIBUTIONS

Invariably, everyone involved in aging infrastructure analysis, tries to find meaning in the shape of the infrastructure's age distribution, "reading" its profiles for clues about what has happened, what is happening, and what will happen in the future. This is not always possible— sometimes the age distribution's profile is ambiguous or even misleading without other information. Therefore it is not recommended. But people do it, so this section will review how the shape of an age distribution develops, comment on significant aspects of profiles, and draw some conclusions and recommendations.

Figure B.23 shows a plot of number service transformers by age – an age distribution plot – for a large utility delivery system is the United States. Average age is 32.5 years. Many utilities have gathered data and produced similar plots for poles, service transformers, cable, breakers, power transformers and other facilities in their system.

The two features labeled on the figure are often seen in utility infrastructure studies. These are the Age Hump - a high construction of equipment in some age bracket, and the Asset Wall, a very sharp drop in number of units at a particular age. The particular age distribution shown has an age hump in the 25-45 year range and an asset wall between 45 and 55 years. A majority of utility age distributions have one or both of these features but the timing differs: in another system or for another set of equipment types in this system, the age hump might be at 20-35 years, and the asset wall at 35 years, etc.




Figure B.23 Ages of poles in service for a utility in the United States. Average pole age is 32 years. Age humps and asset walls are features seen in the infrastructures of many utilities.

Failure Rate Curve Shape Dominates Long-Term Behavior

The shape of an age distribution – its profile – is a function of two factors:

1. The failure rate curve
2. The timing of unit installation

But the longer the period since units were installed, the more the age distribution is dominated by the failure rate curve alone: old systems do not have a memory of when the initial units were added. The longer the time since the infrastructure began (since initial installation and start up), the more the age distribution of equipment in place is dictated *only* by the shape of the failure rate curve: when the original units were installed, if and when there was growth, if there were storms that necessitated many replacements – none of that matters given that enough time has passed.

Given there is no growth, or that it has been a constant growth rate, *eventually* the failure rate curve shape is the only factor that determines the profile of the age distribution curve.

The key point here is “eventually.” This truism applies with mathematical certainty at the sustainable point, but that is usually decades beyond not just the present, but the future period of most interest to planners and management (the next 10-20 years). But those characteristics the infrastructure will have at the sustainable point are often seen, in less than complete measure, decades earlier. From a practical standpoint, it is more proper to say that the failure rate curve shape has more influence over age distribution profile than is generally recognized.

Exponential Failure Rate Curves can Create Asset Walls

Figure B.24 shows five very different failure rate curves that vary from constant (not a function of time) to a very highly exponential function of time. These five curves vary in the sustainable point average age they give, in how long the infrastructure will take to get to the sustainable point, and in a number of other factors, but all give an identical 1.24% average failure rate at the sustainable point. They are the authors’ attempt to normalize failure rate curves. Though they are very different shapes they all have, ultimately, the same cumulative effect on the equipment owners – eventually the owners will have to replace 1.24% of their equipment each year.

Figure B.25 shows the five sustainable point age distributions these curves produce. All five are normalized to 1,000,000 units – the area under the curve is the same for all five. They are completely different distributions, and the major lesson the reader should take away from these figures is that failures are due to the *combination* of age distribution and failure rate, and many very different combinations of each can give what are essentially identical results.




Figure B.24 Five failure rate curves that lead to the same sustainable point failure rate. Over the long run, these produce equivalent reliability results in spite of their different shapes. See text for details




Figure B.25 Distribution of ages at the sustainable point for the 1,000,000 units in service resulted from the five failure rate curve shapes shown in Figure B.6. When the units were installed does not affect these at all: each is a function of its failure rate curve only.

Understanding Why the Asset Wall Exists

The term “wall” does not refer to the rightmost vertical edge of the profile being like a wall, but to the fact that, immediately to the right of that, there is an “invisible wall” of perhaps five to twenty years duration that will not let equipment survive through to its other side. All it takes to create an asset wall is a failure rate curve that stays low for several decades (letting a lot of units survive to that point) and then climbs rapidly to levels that fail many units each year, i.e., exponential behavior.

A highly exponential curve never has a *lengthy* period during which failure rate increases slowly. Instead, after being quite low for years or even many decades, it starts to increase at an increasing rate, skyrocketing in only a decade or a bit more to levels that deplete a lot of units annually.

Failure rate during this “depletion period” does not have to have reach anywhere near 100%. An average failure rate of 5% for a decade depletes about half the units – make failure rate grow during this period, too, and fewer than 1 in 5 units survive. The asset wall is just a short period during which failure rate is high enough that most equipment won’t survive, that follows a long period over which most of them do. Highly exponential curves create this situation, usually right before or around their tipping point.⁸ It’s as simple as that.

In the long run, an exponential failure rate
curve shape, alone, can create an asset wall.

Recent High Rates of Growth Will Dominate an Age Distribution

Figure B.26 compares three age distributions, all three normalized to 1,000,000 units in current operation, and all three based on the same failure rate curve (that from Chapter 5’s utility case study). These are labeled in the figure and represent:

- a) 1,000,000 units were installed 70 years earlier and replacements have been made since, so that there are a million in operation now (a case similar to Example 2).
- b) Far fewer units were installed 70 years ago, and units have been added to a steadily growing system over the past 70 years, with replacements made along the way as needed. Number in operation only reached 1,000,000 this year.
- c) Whatever happened in the past happened long, long ago – there are 1,000,000 units in operation and at their sustainable point.

⁸The “tipping point” is the point where the failure rate curve’s value exceeds its derivative.




Figure B.26 Three age distributions, each of one million units operating now, with far different past histories. See text for details.

Age distribution (b) shows the characteristic shape of an age distribution dominated by recent growth. Note that this distribution looks a good deal like the sustainable point distribution in Figure B.25 for a linear failure rate curve. It is often impossible to tell the difference between a linear failure rate and no growth and an exponential failure rate curve and recent growth based only on the age distribution's profile shape. However, truly linear failure rate curves are nearly non-existent, so past growth is much more likely.

Age Distributions Can Be Shaped by “First Oscillation” Effects

The age distribution shown in Figure B.27 is dominated in part by an oscillation created by significant growth over a period that ended about 30 years earlier. This has created an age hump. Something similar is seen in the infrastructures of many utilities. An age hump is caused by

- Some time in the past, there was a sustained period during which a healthy amount of load growth occurred: many new units were added in this time period.
- The time since the end of that period is less than the time to the failure rate curve's tipping point.




Figure B.27 Profile of an age distribution that is almost surely in an oscillation caused by a sustained period of customer and load growth that ended about 30 years ago (point A). When it began cannot be determined. The asset wall’s deletion period (B) means that few units survive past this period – perhaps growth began 60 years ago but early units added then have been replaced by now.




Figure B.28 Age distributions do not have a “memory” for the original units’ installation dates, etc., that goes past the asset wall’s depletion period, so the uptick in recently added units (Point C) is ambiguous. See text for detailed discussion.

Figure B.27's age distribution shows an oscillation caused by a sustained growth period in the past.

Periods of high growth in the past can
also create an asset wall.

The shape of Figure B.27 shows that the high growth period *ended* about 30 years earlier (point A), but one cannot determine when it started: it could have started a decade or more before point B, but that would not affect the age distribution now enough to leave an unambiguous sign, because the start of that period would be in or past the asset wall. So if there was high growth 60 years ago, and if many units were added that far back, they already have or are failing rapidly now, and thus are not in the age distribution.

The age distribution of an
infrastructure loses its memory
beyond the time of the asset wall.

Figure B.28 shows why an age distribution often cannot be fully “read” as to what happened in the past. The uptick in units near the Y-axis – an increase in units added over the last decade – *might* be because:

- The growth period that ended 30 years ago started 60 years ago. Many units added between 45 and 60 years ago have failed recently as they reach the asset wall and the uptick in recently added units is due to their replacement.
- The growth period that ended 30 years ago began 45 years ago: there was no mass of earlier units that has died as it entered the asset wall during the last ten years. The recent uptick in units added is due to recent growth.

From these age distributions alone, one cannot tell which of these situations is the case: age distribution alone does not provide unambiguous information. Similarly, taken alone, the distributions do not reveal if a large mass of units is about to fail in the next ten years – maybe there aren't any units over about 45 years of age simple because no one added to the system 45+ years ago. Or maybe tremendous numbers of units were added back then and they all failed in the last decade – that means the huge mass of units now between 35 and 45 years of age will likely all disappear in the next decade. Without more information, one cannot tell.

To fully understand what has
happened, what is happening, and
what probably will happen, one needs
to study both the age distribution and
the failure rate curve.

B.7 APPLYING ECONOMIC ANALYSIS

Present Worth Of the Difference in Value

Clearly Example 2 has a very large technical advantage in expected service lifetimes and average failure rates over Example 1. Initially, Dura-Widgits fail only 1/75 as often as Widgits, and although their failure rate increases over time, it eventually settles down to a level of only 80% as much as its sustainable point. Dura-Widgits are simply more robust than normal Widgits.

How much is this advantage worth to a utility? The answer will depend on a host of factors unique to the specific company, and examples covering that are given in some of the final pages of this book. However, a general idea of the value of Dura-Widgits versus normal Widgits can be developed by looking at the differences in failure rates over time from a time-value of money basis. This type of analysis discounts long-term future values for very legitimate and proven reasons. Costs and savings ten or twenty years in the future are simply not as much of a concern to a business as costs and savings this year or next year. And beyond the financial issues at work, there is always uncertainty about the future. Equipment that lasts for decades may become obsolete even if it still works dependably, or may not be needed as technology or societal needs change. Again, very long-term advantages are not as important as near term.

A reasonable way to estimate the relative advantage of Dura-Widgits versus Widgits is to compute the present worth cost of replacements into the future. Using a 10% discount factor (a dollar next year is worth 90 cents today, one two years in the future is worth 81 cents today, etc.) and for the moment assuming that Widgits and Dura-Widgits cost the same, Dura-Widgits give a 1/39 total “present worth of future replacement value.” While in their first year of service they require only 1/78 as many failures as Widgits, by year 75 they require 1/2 as many, and by year 104 they require 1.1 times as many. But when adjusted for how important cost savings are in the future, the average is 1/39 – Dura-Widgits have a total future replacement need that is only 2.55% that of Widgits.

Is that really worth a great deal to a utility? In the real world, Dura-Widgits would very likely not cost the same – even if they were cheaper to build and ship, the manufacturer would charge more for them given their longer lifetime. Some idea of the total value can be provided in this manner.

- 1) Widgits have an equipment cost of \$1.00 and an installation cost of \$1.00.
- 2) Replacements have the same equipment cost but labor cost is \$1.50.
- 3) Dura-Widgits cost \$1,513 per 1,000 rather than \$1,000 a piece.

These costs roughly mirror actual utility costs for service transformers, for example: installation and equipment costs are roughly equal and replacement labor is about 50% more and very often at time and a half or more, for roughly twice total labor cost.

Table B.3 Cost and Present Worth - Wigits versus Dura-Wigits

Cost Category	Wigits	Dura-Wigit
Equipment cost, unit	\$1.00	\$1.513
Initial installation labor	\$1.00	\$1.00
Replacement labor	\$2.00	\$2.00
Total installed (incl. replcmnts)	2,995,000	1,588,444
Total cost (undiscounted \$)	\$7,865,000	\$2,527,839
Total cost (discounted PW\$)	\$2,527,836	\$2,527,836
Replacement cost only (\$)	\$5,865,000	\$2,067,140
Replacement cost only (PW\$)	\$527,836	\$14,949

Table B.4 Analysis of Present Worth for Wigits and Dura-Wigits

Year	PWF	Wigits				Dura-Wigits			
		Equip.	Labor	Total \$	PW Total \$	Eq Cost	Labor	PW \$	PW Total \$
0	100%	\$1,000,000	\$1,000,000	\$2,000,000	\$2,000,000	\$1,427,405	\$1,000,000	\$2,427,405	\$2,427,405
1	90.0%	\$19,550	\$29,325	\$48,875	\$43,988	\$357	\$375	\$732	\$659
2	81.0%	\$19,550	\$29,325	\$48,875	\$39,589	\$376	\$395	\$771	\$624
3	72.9%	\$19,550	\$29,325	\$48,875	\$35,630	\$396	\$416	\$812	\$592
4	65.6%	\$19,550	\$29,325	\$48,875	\$32,067	\$417	\$438	\$855	\$561
5	59.0%	\$19,550	\$29,325	\$48,875	\$28,860	\$439	\$462	\$901	\$532
6	53.1%	\$19,550	\$29,325	\$48,875	\$25,974	\$463	\$486	\$949	\$504
7	47.8%	\$19,550	\$29,325	\$48,875	\$23,377	\$487	\$512	\$1,000	\$478
8	43.0%	\$19,550	\$29,325	\$48,875	\$21,039	\$513	\$539	\$1,053	\$453
9	38.7%	\$19,550	\$29,325	\$48,875	\$18,935	\$541	\$568	\$1,109	\$430
10	34.9%	\$19,550	\$29,325	\$48,875	\$17,042	\$570	\$598	\$1,168	\$407
11	31.4%	\$19,550	\$29,325	\$48,875	\$15,337	\$600	\$630	\$1,230	\$386
12	28.2%	\$19,550	\$29,325	\$48,875	\$13,804	\$632	\$664	\$1,296	\$366
13	25.4%	\$19,550	\$29,325	\$48,875	\$12,423	\$665	\$699	\$1,364	\$347
14	22.9%	\$19,550	\$29,325	\$48,875	\$11,181	\$701	\$736	\$1,437	\$329
15	20.6%	\$19,550	\$29,325	\$48,875	\$10,063	\$738	\$775	\$1,513	\$312
16	18.5%	\$19,550	\$29,325	\$48,875	\$9,057	\$777	\$817	\$1,594	\$295
17	16.7%	\$19,550	\$29,325	\$48,875	\$8,151	\$818	\$860	\$1,678	\$280

Table B.3 provides an overview of the comparison on this basis. Table B.4 shows the first seventeen years of an analysis spreadsheet by year. Only the first 100 years of operation were analyzed: discounted only about the first 30 years really have any impact. The resulting break-even cost for Dura-Widgits is just over \$1.51 per unit: they may be 80% as durable in the long run and have as little as 1/78 the failure rate in the short run, but those advantages are only worth a bit more than 50% per unit.

A time value of money evaluation tends to focus analysis and attention on the present and the short run, which is proper and necessary for financial planning and budgeting purposes. Such time-discounted financial analysis assumes the technical evaluation of actual events and costs for future years has been done correctly. While financial evaluation and “business-based” cases built upon it are the correct way to ultimately use the results of aging infrastructure analysis, it must be based on sound technical evaluation of the infrastructure. For now, present, short-term, and long-term behaviors will be looked at all within a purely technical analysis. Many utilities will need this for a simple business reason: *looking at their aged infrastructure, as opposed to aging, they are already at the long run.*

Equipment that differs in expected service lifetime by 85% (92 versus 50 years) has the same sustainable average age in service (50 versus 50 years) and will eventually require about 1/185% fewer replacements. The reason has to do with the shape of the failure rate curve. Over the very long term the difference in the two failure rate curves results in 85% fewer failures, but the *distribution* of when failures are likely as a function of useful service life means that the real practical impact of is much, much less for Dura-Widgits: failures when far into the future are simply not worth as much as failures sooner.

B.8 SUMMARY OF KEY POINTS

This appendix has reviewed a key concept of aging infrastructures: an equipment infrastructure kept in service by replacement or repair when it fails trends, slowly but surely over time, toward a stable *sustainable point*, at which the average age, the distribution of ages, the average condition and the distribution of conditions, the average failure rate and the distribution of units failing at various ages, and all operating costs remain relatively constant from year to year thereafter (see Table B.5).

Sustainable point analysis provides an unbeatable foundation for planning and decision making with regard to managing an aging infrastructure. The sustainable point is a function of both the equipment and its owner’s operating policies. Change the loading standards or operating voltages, alter the inspection, tracking, service and repair processes, implement a different refurbishment plan, or change replacement policy, and the sustainable point *will* move. Sustainable point analysis can thus show how revisions in operating policies and criteria can lead to changes in both the eventual sustainable point, and the annual rate of decline or improvement the utility can expect. It is thus a

Table B.5 Summary of Key Points

Average Age Provides Little Useful Information and Can Be Misleading. Average age of a group or set of equipment tells planners and management next to nothing useful about the equipment's trends or their importance to key performance indicators.

Number of failures or average failure rate provides the most relevant measure. Most problems, be they service quality, reliability, or cost-related, stem from failures. Tracking that and predicting that provides the most useful data.

Equivalent Failure Rate Age is the age in the failure rate curve that matches an infrastructure's average failure rate; regardless of the mix of ages in operation, the whole is failing like a group of units of this age. As a single age-related characteristic, this is the most meaningful and useful value.

Scaling and Adding to Failure Rate Curves does not proportionally change their impacts. Multiplying or adding to highly exponential failure rate curves changes the resulting failure rates and sustainable points they produce, but not in proportion to the multiplier or value added. Higher order curves are less sensitive than lower order curves.

Sustainable point values are the best long-term characterization generally available of an infrastructure's aging situation, potential, and propensity to react to changes in equipment of operating policies.

Age Distribution provides more information than averages or other statistics.

Failure rate curve shape alone determines long term age distribution shape.

Age distributions can be read for indications of past, present, and future infrastructure behavior.

Some age distributions cannot be read – some are ambiguous.

Economic analysis tends to emphasize short-term aspects of aging infrastructure performance, and to focus management attention on changes in equipment or operating policies that impact the short term.

Sustainable Point: In *any* group or set of equipment, be they transformers, breakers, tires in a fleet of trucks, or tableware used by a restaurant chain, if operating conditions are kept the same from year to year and failures and unserviceable units are always replaced by new, eventually a *sustainable point* will be reached, a point beyond which the average age, failure-replacement rate, and operating and replacement costs do not change, but instead stay the same from year to year thereafter.

powerful operating and business planning tool for a utility or industrial system owner trying to optimize its bottom line. In addition, mathematically, the fact that a sustainable point *has* to exist is a kind of “statistical lever” that can help fit numerical models to a utility's historical operating and equipment asset databases. Fitting error for failure rate analysis and for models of impact on utility performance is often halved using such techniques.

And finally, and perhaps most important, the concept of a sustainable point is important to managers, planners, and operating engineers at a utility. The system

Any equipment infrastructure asymptotically tends toward its sustainable point over time.

Theoretically a system never makes it to the sustainable point, but practically, it does. Mathematically, a system will always be asymptotically moving toward that point with some tiny increment left. Practically, many real systems could be “close enough it does not matter” within a time period less than the power industry has been in existence.

Rate of change toward the sustainable point is steady and continuous: as a system ages its statistics will move a bit closer toward the sustainable point each year. Nothing can stop this other than a change in fundamentals that determine the sustainable point: the type and design of the equipment, the way it is used and loaded, or the way it is inspected and serviced.

Rate of change toward the sustainable point is always slow: infrastructures do not “fall off the cliff” with respect to rising reliability and cost issues: increases occur *gradually* from year to year.

Rate of change slows as a system ages: as a system ages, the rate of change toward the sustainable point grows a bit less each year.

Failure rate curves have a “tipping point” where their value first exceeds their derivative (percent change in value from year to year). This will be the point of maximum total failures – the age where the product of [amount of equipment that has made it to this age] times [failure rate for this age] is at a maximum. Generally, a utility will find that equipment cannot be made to not stay in service long after this point.

Oscillations in average age, failure rate, and cost can occur when a) the failure rate increases over time (which is almost always the case) and b) the equipment was installed in a relatively short time compared to the expected lifetime of the equipment. In such situations average age, failure rate, and cost will go through a series of gradually less dramatic “swings” with a periodicity of a decade or more.

Oscillations almost never occur on a system-wide basis. Oscillations are just not seen in the overall failure, age, and operating costs statistics for an entire system because the equipment was installed over many decades.

Oscillations are common in isolated areas of the system where local development was built out over a relatively short period of time.

will not “fall off a cliff” with regard to reliability and failures in only a few years; deterioration is a slow-acting poison, if an inevitable one. The good news is that the decline will be slow and the trends predictable. The bad news is that the trends are inevitable, and that improvement methods, unless quite expensive, act no more quickly than the trends themselves. But good or bad, the sustainable point is where the system is heading, and the bottom line is that the utility can move it to a more acceptable combination of failures and costs by changing operating policies and methods.

A Practical Warning Sign of Systemic Problems

When average age increases at a log-linear rate, the shape of the distribution ages will tend to remain constant over time. *This is a sign that failure is not due to deterioration.* It is highly likely there is some systemic problem with the equipment, such as a poor design or installation practices, or grounds that fail very early, etc., or that there are very high rates of random external damage.

REFERENCES AND BIBLIOGRAPHY

- P. F. Albrecht, and H. E. Campbell, "Reliability Analysis of Distribution Equipment Failure Data," EEI T&D Committee Meeting, January 20, 1972.
- J. B. Bowles, "Commentary-caution: constant failure rate models may be hazardous to your design," *IEEE Transactions on Reliability*, Vol. 51, pp. 375-377, Sept. 2002.
- R E Brown, *Electric Power Distribution Reliability*, Marcel Dekker, New York, 2002.
- R. E. Brown and J. R. Ochoa, "Distribution System Reliability: Default Data and Model Validation," *IEEE Transactions on Power Systems*, Vol. 13, pp. 704-709, May 1998
- R. E. Brown, G. Frimpong, H. Willis, "Failure Rate Modeling Using Equipment Inspection Data," *IEEE Transactions on Power Systems*, May 2004, p. 782, IEEE New York.
- R. M. Bucci, R. V. Rebapragada, A. J. McElroy, E. A Cheblis and S. Drller, "Failure Prediction of Underground Distribution Feeder Cables," *IEEE Transactions on Power Delivery*, Vol. 9, pp. 1943-1955, Oct. 1994.
- EEI Transmission and Distribution Committee, "Guide for Reliability Measurement and Data Collection," October 1971, Edison Electric Institute, New York.
- P. Gil, *Electrical Power Equipment Maintenance and Testing*, Marcel Dekker, New York, 1998.
- Institute of Electrical and Electronics Engineers, *Recommended Practice for Design of Reliable Industrial and Commercial Power Systems*, The Institute of Electrical and Electronics Engineers, Inc., New York, 1990.
- A. Gupta, A. Pahwa, R. E. Brown, and S. Das, "A Fuzzy Model for Over Distribution Feeders Failure Rates," in Proc. 34th Annual North American Power Symposium, Tempe, AZ, Oct. 2002.
- A. D. Patton, "Determination and Analysis of Data for Reliability Studies," *IEEE Transactions on Power Apparatus and Systems*, PAS-87, January 1968.
- D. T. Radmer, P. A. Kuntz, R. D. Christie, S.S. Venkata, and R. H. Fletcher, "Predicting Vegetation-Related Failure Rates For Overhead Distribution Feeders," *IEEE Transactions on Power Delivery*, Vol. 17, pp. 1170-1175, Oct. 2002.
- H. L. Willis, *Power Distribution Planning Reference Book, Second Edition*, Marcel Dekker Publishers, New York, 2004.
- H. L. Willis and R. E. Brown, "Substation Asset Management," Chapter 19 in *Electric Power Substations Engineering Handbook – Second Edition*, CRC Press, New York, 2007.

AGING POWER DELIVERY INFRASTRUCTURES

Second Edition

Good aging infrastructure management consists of optimizing the choice of equipment and its refurbishment while also making compatible changes in all those operating and ownership policies, the whole combination aimed at optimizing the business results the power system owner desires. Both a reference and tutorial guide, this second edition of **Aging Power Delivery Infrastructures** provides updated coverage of aging power delivery systems, the problems they cause, and the technical and managerial approaches that power systems owners can take to manage them.

See What's New in the Second Edition

- All chapters have been updated or are completely new
- Comprehensive discussions of all issues related to equipment aging
- Business impact analysis and models and engineering business studies of actual utility cases
- Strategy and policy issues and how to frame and customize them for specific situations

This book looks at the basics of equipment aging and its system and business impacts on utilities. It covers various maintenance, service, and retrofit methods available to mitigate age-related deterioration of equipment. It also presents numerous configuration and automation upgrades at the system level that can deal with higher portions of aging equipment in the system and still provide good service at a reasonable cost.



CRC Press

Taylor & Francis Group
an informa business

www.taylorandfrancisgroup.com

6000 Broken Sound Parkway, NW
Suite 300, Boca Raton, FL 33487
711 Third Avenue
New York, NY 10017
2 Park Square, Milton Park
Abingdon, Oxon OX14 4RN, UK

K10222

ISBN: 978-1-4398-0403-2

90000



9 781439 804032

www.crcpress.com