


Practical Lighting Design WITH LEDs



RON LENK • CAROL LENK

 WILEY

 IEEE
IEEE PRESS


Mohamed E. El-Hawary, Series Editor

Practical Lighting Design With LEDs

IEEE Press
445 Hoes Lane
Piscataway, NJ 08854

IEEE Press Editorial Board

Lajos Hanzo, *Editor in Chief*

R. Abari
J. Anderson
F. Canavero
T. G. Croda

M. El-Hawary
B. M. Hammerli
M. Lanzerotti
O. Malik

S. Nahavandi
W. Reeve
T. Samad
G. Zobrist

Kenneth Moore, *Director of IEEE Book and Information Services (BIS)*

Practical Lighting Design With LEDs

Ron Lenk
Carol Lenk



Mohamed E. El-Hawary, *Series Editor*



A John Wiley & Sons, Inc., Publication

Copyright © 2011 the Institute of Electrical and Electronics Engineers, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey
Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Lenk, Ron, 1958–

Practical lighting design with LEDs / by Ron Lenk, Carol Lenk.

p. cm. – (IEEE Press series on power engineering ; 67)

ISBN 978-0-470-61279-8

1. Light emitting diodes. 2. Electric lamps—Design and construction. 3. Electric lighting—Design. I. Lenk, Carol. II. Title.

TK7871.89.L53L46 2011

621.32—dc22

2010048267

Printed in Singapore

oBook ISBN: 978-1-118-00821-8

ePDF ISBN: 978-1-118-00820-1

ePub ISBN: 978-1-118-01173-7

10 9 8 7 6 5 4 3 2 1

To our children, for being so patient

Contents

Figures **xiii**

Preface **xvii**

1 Practical Introduction to LEDs **1**

What Is an LED?	1
Small LEDs versus Power Devices	2
Phosphors versus RGB	3
Inside an LED	4
Is an LED Right for My Application?	6
Haitz’s Law(s)	8
The Wild West	10
LEDs and OLEDs and . . . ?	11

2 Light Bulbs and Lighting Systems **13**

Light Sources	13
Incandescent	13
Halogen	14
Fluorescent	14
Induction Lighting	15
High Intensity Discharge (HID) Lamps	15
Characteristics of Light Sources	16
Light Quality	16
Efficacy	17
Timing	17
Dimming	18
Aging	19
Types of Bulbs	19
Bulb Shapes	19
Bulb Bases	21
Specialty Bulbs	21

viii Contents

History of Lighting 22
Governments 23
Lighting Systems 23

3 Practical Introduction to Light 27

The Power of Light 27
 Background: Light as Radiation 27
Radiometric versus Photometric 28

Luminous Intensity, Illuminance, and Luminance (or Candela, Lux,
and Nits) 32
 Luminous Intensity 32
 Illuminance 34
 Luminance 36
 Summary of Amount of Light 37

What Color White? 37
 MacAdam Ellipses 39
 A Note about Color Space Standards 41

Color Rendition: How the Light Looks versus How Objects Look 41
 The Human Factor 43

4 Practical Characteristics of LEDs 47

Current, Not Voltage 47
Forward Voltage 48
Reverse Breakdown 49
Not Efficiency—Efficacy! 51
LED Optical Spectra 54
Overdriving LEDs 57
Key Datasheet Parameters 58
Binning 58
The Tolerance Game 60

5 Practical Thermal Performance of LEDs 61

Mechanisms Behind Thermal Shifts 61
Electrical Behavior of LEDs With Temperature 62
Optical Behavior of LEDs With Temperature 63
Other Performance Shifts With Temperature 64
LED Lifetime: Lumen Degradation 66
LED Lifetime: Catastrophic Failure 67
Paralleling LEDs 67

6 Practical Thermal Management of LEDs **71**

Introduction to Thermal Analysis	71
Calculation of Thermal Resistance	72
The Ambient	74
Practical Estimation of Temperature	75
Heat Sinks	76
Fans	78
Radiation Enhancement	79
Removing Heat from the Drive Circuitry	79

7 Practical DC Drive Circuitry for LEDs **81**

Basic Ideas	81
Battery Basics	82
Overview of SMPS	85
Buck	86
Boost	88
Buck-Boost	89
Input Voltage Limit	91
Dimming	92
Ballast Lifetime	94
Arrays	95

8 Practical AC Drive Circuitry for LEDs **99**

Safety	99
Which AC?	101
Rectification	103
Topology Selection	106
Nonisolated Circuitry	109
Isolation	110
Component Selection	112
EMI	114
Power Factor Correction	117
Lightning	119
Dimmers	120
Ripple Current Effects on LEDs	122
Lifetime	125
UL, Energy Star, and All That	126

9 Practical System Design With LEDs **131**

PCB Design	131
Getting the Light Out	136

LEDs in Harsh Environments	139
Designing With the Next Generation LED in Mind	141
Lighting Control	142
DALI Protocol	142
DMX512 Protocol	143
802.15.4 and ZigBee Open-Standard Technology	143
Powerline Communication	144

10 Practical Examples **145**

Example: Designing an LED Flashlight	145
Initial Marketing Input	145
Initial Analysis	146
Specification	148
Power Conversion	150
Thermal Model	153
PCBs	155
Final Design	156
Example: Designing a USB Light	157
Initial Marketing Input	157
Initial Analysis	157
Specification	158
Power Conversion	159
Thermal Model	161
PCB	163
Final Design	163
Example: Designing an Automobile Taillight	165
Initial Marketing Input	165
Initial Analysis	165
Specification	167
Power Conversion	167
Thermal Model	172
PCB	172
Final Design	174
Example: Designing an LED Light Bulb	175
Initial Marketing Input	175
Initial Analysis	175
Specification	178
Power Conversion	181
PCBs	184

11 Practical Measurement of LEDs and Lighting 187

Measuring Light Output	187
Lux Meter	187
Integrating Sphere	189
Goniophotometer	192
Special Considerations in Measuring Light Output of LEDs	192
LED Measurement Standards	192
Luminaire Light Output (LM-79)	192
LED Lifetime (LM-80)	193
ASSIST	193
Measuring LED Temperature	194
Measuring Thermal Resistance	195
Measuring Power, Power Factor, and Efficiency	197
Accuracy versus Precision	197
Measuring DC Power	197
Measuring AC Power	200
Measuring Power Factor	201
Measuring Ballast Efficiency	201
EMI and Lightning	202
Accelerated Life Tests	203

12 Practical Modeling of LEDs 207

Preliminaries	208
Practical Overview of Spice Modeling	209
What Not to Do	212
What to Do	213
Modeling Forward Voltage	214
Reverse Breakdown	219
Modeling Optical Output	221
Modeling Temperature Effects	224
Modeling the Thermal Environment	226
A Thermal Transient	228
Some Comments on Modeling	229

References 231**Index** 233

Figures

Figure 1.1 T1¾ (5 mm) LEDs	2
Figure 1.2 Fluorescent Tube’s Spectral Power Distribution	4
Figure 1.3 LEDs Can Be Used Everywhere	6
Figure 1.4 Haitz’s Law	8
Figure 2.1 Currents in a Fluorescent Tube	15
Figure 2.2 Various Bulb Shapes	20
Figure 3.1 The Electromagnetic Spectrum	28
Figure 3.2 Scotopic Vision is much more Sensitive than Photopic Vision	29
Figure 3.3 Emission Spectra of Four Common Light Sources	30
Figure 3.4 Solar Radiation Spectrum	31
Figure 3.5 One Steradian Intersects 1 m ² of Area of a 1 m Radius Ball	32
Figure 3.6 Solid Angle in Steradians versus Half Beam Angle in Degrees	33
Figure 3.7 Definition of Beam Angle	34
Figure 3.8 Typical Lambertian Radiation Pattern	36
Figure 3.9 Dimensions for a USB Keyboard Light Design	36
Figure 3.10 Spectra of Neutral-White and Warm-White LEDs	38
Figure 3.11 CIE 1931 (x, y) Chromaticity Space, Showing the Planck Line and Lines of Constant CCT	38
Figure 3.12 (x, y) Chromaticity Diagram Showing CCT and 7-Step MacAdam Ellipses	40
Figure 3.13 Cool White Fluorescent 4100 K, CRI 60; Incandescent, 2800 K, CRI 100; Reveal® Incandescent 2800 K, CRI 78	41
Figure 3.14 Approximate Munsell Test Color Samples	42
Figure 3.15 Circadian Rhythm Sensitivity	44
Figure 3.16 Identical Gray Boxes Look Different Depending on Their Background	45
Figure 4.1 Reverse Bias Protection	50
Figure 4.2 LEDs with Reverse Bias Protection	50
Figure 4.3 Light Output as a Function of Current	52
Figure 4.4 Forward Voltage as a Function of Current	53
Figure 4.5 Efficacy versus Drive Current	53
Figure 4.6 Light Output as a Function of Wavelength	54
Figure 4.7 Many LEDs Have Poor R9	55
Figure 4.8 (x, y) as a Function of Current	56
Figure 4.9 Different Output Light Distributions Are Available	56
Figure 4.10 Neutral-White Bin Structure	59

Figure 5.1 Brightness as a Function of Temperature	63
Figure 5.2 LED Temperature Profile for Parameters Given in Text	65
Figure 5.3 Forward Voltage as a Function of Current	68
Figure 6.1 Thermal Model for LED Example	73
Figure 6.2 Thermal Model of Two Parallel Thermal Paths	73
Figure 6.3 LED Temperature as a Function of Time	73
Figure 6.4 There Are Many Thermal Paths to Ambient	74
Figure 6.5 Estimating Temperature Rise from Power Density	76
Figure 6.6 An LED Heat Sink	77
Figure 7.1 12V Battery I-V Curve	83
Figure 7.2 Alkaline-Cell Battery Voltage as a Function of Time with a Resistive Load	84
Figure 7.3 Operating a Transistor in Linear Mode Is Inefficient	85
Figure 7.4 When the Transistor Is On, Current in the Inductor Increases; When the Transistor Is Off, Current in the Inductor Decreases	86
Figure 7.5 LM3405 Schematic for Buck	87
Figure 7.6 FAN5333A Schematic for Boost	88
Figure 7.7 HV9910 Schematic for Buck-Boost	90
Figure 7.8 Pulse Width Modulation Turns the Current Rapidly On and Off to Get an Average Current	92
Figure 7.9 Dimming Circuit	93
Figure 7.10 The Effect of the Current Sense Resistor Is Compensated by Putting One in Series with Each String	96
Figure 7.11 LED Forward Voltage Variation Can Be Compensated at the Cost of Additional Power	96
Figure 7.12 Ballasting LED Strings with Total Current Sensing	96
Figure 8.1 Block Diagram of AC SMPS for LED Lighting	103
Figure 8.2 A Bridge Rectifier	103
Figure 8.3 Half-wave Rectification	104
Figure 8.4 Reducing the Ripple from a Bridge Rectifier with a Capacitor	105
Figure 8.5 Running LEDs Directly Off-Line	107
Figure 8.6 How the Off-Line Buck Works	108
Figure 8.7 A Nonisolated Off-Line LED Driver	109
Figure 8.8 Adding a Transformer Makes the Converter into a Forward	111
Figure 8.9 Adding a Transformer Makes the Converter into a Flyback	111
Figure 8.10 Protecting the HV9910 from High Voltages	114
Figure 8.11 Resistors Balance Voltages for Series Capacitors	114
Figure 8.12 Normal Mode EMI Filtering for a Two-Wire Input	115
Figure 8.13 Common Mode EMI Filtering Added for a Three-Wire Input	115
Figure 8.14 Current Loops May Cause EMI Problems; Reducing Loop Area Helps	116
Figure 8.15 A Big Capacitor Maintains Constant Voltage During the Line Cycle, Generating Large Peak Currents and Bad Power Factor	118
Figure 8.16 A Smaller and Cheaper PFC	118
Figure 8.17 Simple Power Factor Correction Circuit	119

Figure 8.18 Adding an MOV to the Design Protects It Moderately Well from Lightning	120
Figure 8.19 Output Waveform of a Triac Dimmer	121
Figure 8.20 Keeping an IC's Power Alive During the Off-Time of a Dimmer	121
Figure 8.21 As Ripple Current Increases, Power Loss in the LED Also Increases	123
Figure 8.22 Forward Voltage Increases with Increasing Current	123
Figure 8.23 Increasing the Current Does Not Proportionally Increase the Light	124
Figure 8.24 200mApp Ripple Current on a 350mA DC Drive	125
Figure 9.1 A Schematic which Could Be Improved	132
Figure 9.2 A Good Schematic	133
Figure 9.3 Poor Grounding Layout and Improved Layout	135
Figure 9.4 Soda-Lime Glass Optical Transmission	138
Figure 9.5 DALI Topologies	143
Figure 10.1 FAN5333A Schematic for Flashlight	151
Figure 10.2 FAN5333A Final Schematic for Flashlight	153
Figure 10.3 Thermal Model	155
Figure 10.4 PCB for LED Flashlight Ballast	156
Figure 10.5 First-Cut USB Light Schematic	160
Figure 10.6 LM3405 Schematic for USB Light	161
Figure 10.7 USB Light PCB, Top View and X-Ray Bottom View	164
Figure 10.8 Final USB Light Schematic	164
Figure 10.9 HV9910 Schematic for Taillight	168
Figure 10.10 LS E63F I-V Curve	170
Figure 10.11 Final HV9910 LED Automobile Taillight Schematic	171
Figure 10.12 LED Taillight: Whole Board, Topside Zoom of Ballast Area, Bottom Side X-Ray Zoom of Ballast Area	173
Figure 10.13 BR Bulbs	176
Figure 10.14 Initial Thermal Model for BR40	177
Figure 10.15 First Attempt at a BR Design	182
Figure 10.16 Final HV9910 Schematic for LED BR40	184
Figure 10.17 BR40 PCB: Topside and Bottom Side with X-Ray Vision	186
Figure 11.1 Typical Lux Meter	188
Figure 11.2 Sketch of Operation of an Integrating Sphere	189
Figure 11.3 Walk-in Integrating Sphere	190
Figure 11.4 Gooch & Housego 6-Inch Integrating Sphere with OL770-LED Spectroradiometer	191
Figure 11.5 Location of Thermocouple to Measure LED Temperature	195
Figure 11.6 High Current Drawn at AC Line Peak Gives a High Crest Factor	201
Figure 11.7 Phase Shift between Current and Voltage Also Gives Low Power Factor	202
Figure 12.1 Luxeon Rebel I-V Curve	214
Figure 12.2 First Model of Luxeon Rebel	217

xvi Figures

Figure 12.3 Setting the GMIN Option to 1.0E-20 Fixes the Current at 0V	220
Figure 12.4 LED Model Including Reverse Breakdown	220
Figure 12.5 With the Breakdown Modeled, the Complete I-V Curve Is Shown	221
Figure 12.6 Normalized Light Output versus Current	222
Figure 12.7 Data versus Equation for Optical Output versus Drive Current	222
Figure 12.8 LED Model with Light Output, Temperature Not Yet Included	224
Figure 12.9 The Effect of Temperature on Light Output	225
Figure 12.10 LED Model with Temperature Effect on Both Forward Voltage and Optical Output	226
Figure 12.11 Complete LED Model	228
Figure 12.12 The Thermal Response of Two LEDs to Current Pulse	229
Figure 12.13 The Optical Response of Two LEDs to Current Pulse	229

Preface

THE LIGHTING REVOLUTION

LEDs are bringing in a new era in lighting. Similar to the evolution of computing power that computers went through, from vacuum tubes to the silicon-based semiconductor brains of modern-day computers, lighting is now poised to ride an exponential growth wave in efficacy. From oil lamps to the invention of the Edison light bulb 100 years ago to the fluorescent lights of 50 years ago to the LEDs of today, lighting technology is finally joining the modern world of solid-state technology.

In the near term, LED-based lighting will increasingly become the efficient light source of choice, replacing both incandescent and fluorescents. The hurdles that keep consumers from adopting present-day energy efficient lighting, such as shape, color quality, the presence of toxic mercury, and limited lifetime, are all better addressed by LEDs.

In the long term, LED-based lighting will be better and cheaper than every other light source. It will become the de facto light of choice. LED lighting will be cheap, efficient, and used in ways that haven't been imagined yet. It will transform the \$100 billion lighting industry, and with transformation comes opportunity.

THE LAST VACUUM TUBE

Lighting is the last field that still uses vacuum tubes. All electronics today use integrated circuits because of the enormous benefits in performance and cost. But an incandescent bulb is a type of vacuum tube, and so is a fluorescent bulb. LEDs are solid-state devices, the same as the rest of electronics. The amount of light that an LED can convert from 1 W of power is already on par with the best fluorescent tubes. The future is even brighter as LEDs are anticipated to double that performance in the next decade, and then go on to reach the physical limits of electricity to light conversion. We look forward to seeing the last ceiling-mounted vacuum tube in the not-too-distant future.

GREEN LIGHTING

The benefits of using LEDs for lighting are many. The most obvious is their efficiency. Lighting accounts for 20% of total electricity use throughout the world today. Using LEDs could cut this down to 4% or less. As LEDs become the dominant light source over the next decade, the reduction of energy used and greenhouse gases

emitted will benefit everyone. Consumers will save hundreds of dollars every year from reduced energy use. Building owners will save even more. Utilities will be better equipped to manage growth. And the earth will experience the accumulation of fewer greenhouse gases, as well as a reduction in the emission of toxic mercury found in fluorescent lighting.

A LIFETIME OF LIGHTING

As solid-state devices, LEDs have extremely long lifetimes. They have no filaments to break. They can't leak air into their vacuum because they don't use a vacuum. In fact, they don't really break at all; they just very gradually get dimmer. Imagine changing your light bulb only once or twice in your entire lifetime!

LIGHTING THE WORLD WITH LEDS

Just as microprocessors got cheaper and more powerful, LEDs will also benefit from the cost-reduction techniques developed in the semiconductor industry. LED light prices will eventually decrease to be on par with incandescent bulbs. Taken together with LEDs' reduced energy usage, this will enable the universal availability of lighting. Imagine every child in the poorest village having a light to read by.

The design of LED-based lighting systems is an exciting field, but they are fairly technical devices. With this book, we hope to enable the reader to do great things with lighting, both for him- or herself and for the world.

RON LENK
CAROL LENK

*Woodstock, Georgia
February 2011*

Chapter 1

Practical Introduction to LEDs

Light bulbs are everywhere. There are over 20 *billion* light bulbs in use around the world today. That's three for each person on the planet! We expect that within the next 10 years, the majority of these bulbs will be light emitting diodes (LEDs). This is because LEDs can provide efficiency dozens of times higher than incandescent light bulbs. They can be as efficient as the theoretical limit for electricity to light conversion set by physics. This book is all about the practical aspects of LEDs and how you can make practical lighting designs using them.

WHAT IS AN LED?

The purpose of this book is to tell you practical things about LEDs. So in this section, we're not going to regale you with jargon about "direct bandgap GaInP/GaP strained quantum wells" or such. Let's directly address the question, What is an LED?

The name "light emitting diode" tells you a lot already. In the first place, the noun tells you that it's a diode. A diode conducts current in one direction and not the other. And that's what an LED does. While we'll explore the details of its electrical behavior in Chapter 4, the only thing to note for the moment is that it has a much higher forward voltage than the diodes usually used in electronics. While a 1N4148 has a drop of about 700 mV, an LED may drop 3.6 V. This is because LEDs are not made from silicon, but from other semiconductors. But other than that, an LED's electrical characteristics are very much like those of other diodes.

The words "light-emitting" tell you a lot more. Now all diodes emit at least a little bit of light. You can open up an integrated circuit (IC) and use a scanner to see which parts of the circuit are emitting light. This tells you which parts are conducting current. IC designers use this to help debug their ICs. However, the amount of light emitted by ICs is very small. Since the purpose of LEDs is to emit light, they have been carefully designed to optimize this performance. That's why, for example, they have a much higher forward voltage than normal, rectifier diodes. Rectifiers have

been optimized to minimize their forward voltage while maximizing reverse breakdown voltage. LEDs are optimized to produce the most light of the right color at the lowest power, and things such as forward voltage (by itself) don't matter. Of course, forward voltage does enter into how much power the LED dissipates, and we'll see in Chapter 5 how to characterize the light emitted versus the power dissipated.

SMALL LEDs VERSUS POWER DEVICES

Present-day thinking divides LEDs into two classes: small devices and power devices. Small LEDs became widely used in the 1970s. They come in all different colors, such as red, orange, green, yellow, and blue. They are the small T1 $\frac{3}{4}$ (5 mm) devices shown in Figure 1.1. Nowadays, there are literally tens of billions of them sold each year. They go into cell phone backlights, elevator pushbuttons, flashlights, incandescent bulb replacements, fluorescent tube replacements, road signage, truck taillights, traffic lights, automobile dashboards, and so on.

What characterizes these small devices is their power level, or as the industry thinks of it, their drive current. The typical red small LED, for example, has a drive current of 20 mA. At a forward voltage of 2.2 V, this is only 44 mW of power. (The efficacy is so low that this is just about equal to the heat dissipation as well.) Small white LEDs have a higher forward voltage (3.6 V, corresponding to 72 mW), and some small LEDs can be run as high as 100 mA. But fundamentally, this type of LED is used as an indicator, not a real light source. It takes fourteen of them to make a somewhat reasonable 1 W flashlight, and hundreds of them to make a (dim) fluorescent tube replacement.

While the information in this book is applicable to these small LEDs, the main focus is on power devices. Power devices are typically 1–3 W devices that are usually



Figure 1.1 T1 $\frac{3}{4}$ (5 mm) LEDs.

run at 350 mA. Their dice (the actual semiconductor, as opposed to its package) are substantially larger than those of small LEDs, although their footprint need not be. These devices are typically used in places requiring lighting, rather than as indicators. Applications include flashlights, incandescent bulb replacements, large-screen TVs, projector lights, automotive headlights, airstrip runway lighting, and just about everywhere lighting is used. Of course, not all of these applications have yet seen widespread adoption of power LEDs, but they will soon.

PHOSPHORS VERSUS RGB

Most lighting designs are going to be made with white light (which includes incandescent “yellow” light). For this reason, this book concentrates primarily on white LEDs. However, what is described here for white LEDs can be straightforwardly applied to color LEDs. Color LEDs are very similar to white, albeit with differing forward voltage. The reason for the varying forward voltages is that the colored light (red, yellow, blue, etc.) is generated directly by the semiconductor material. The material is varied to get differing colors and the differences in material in turn cause differences in forward voltages.

However, white light cannot be directly generated by a single material (we are ignoring special types of engineered materials that are not yet in production). White light consists of a mixture of all of the colors. You already know this because white light can be separated into its constituent colors with a prism. White light thus has to be created. There are currently two main methods of generating white light with LEDs. In one method, an LED that emits blue light is used, and the blue light is converted to white by a phosphor. In the other method, a combination of different color LEDs is used.

The first method is the most common. A typical wavelength for the blue light generated by the LED is 435 nm. Why use blue light? This has to do with the physics of the way the white light is generated. The blue light is absorbed by a phosphor, and re-emitted as a broad spectrum of light approximating white. For the phosphor to be able to absorb and re-emit the light, the light coming out has to be lower in energy than the light going in. That’s just like any electronic component. Energy goes in, some is dissipated as heat, and the rest comes out again, transformed. So to get all of the colors in the spectrum that humans can see, the phosphor needs to have input at a higher energy (shorter wavelength) than the shortest color’s energy. For humans, this is about 450 nm, and so a 435 nm blue LED is the most energy-efficient way of generating white light using a phosphor.

Before turning to the second method of generating white light, we should say a few more words about the phosphor. There are various types of phosphors. Phosphors are designed to absorb one specific wavelength of light, and re-emit it at either one or more different wavelengths or in a band of wavelengths. LED phosphors are typically designed to do the latter. But there are limits to how broad a band of colors a phosphor can emit. So many LEDs use bi-band or tri-band phosphors to better cover the spectrum of light needed to approximate white. These phosphors




Figure 1.2 Fluorescent Tube's Spectral Power Distribution.

Source: http://www.gelighting.com/na/business_lighting/education_resources/learn_about_light/pop_curves.htm?1. See color insert.

are mixtures of two or three primary phosphors. These more complicated phosphors are typically used when better color rendition is needed (see the discussion of color rendering index [CRI] in Chapter 3).

As a side note, we can comment briefly on fluorescent lights. In some ways, a fluorescent light is quite similar to an LED, but its fundamental mechanism of light emission is different. It generates a high-temperature plasma inside a tube, which emits light in the ultraviolet (UV) range (254 nm) rather than in the blue range. But after that, it too uses a phosphor to absorb the light and re-emit it in the visible range. Note that since the wavelength of the light is considerably farther away from the visible spectrum than the 435 nm generated electrically by the LED die, the efficiency ultimately possible for a fluorescent is intrinsically lower than that possible for an LED. (At the moment, fluorescent lights and LEDs have roughly the same efficiency.)

But also interesting is the type of phosphor the typical fluorescent light uses. These phosphors are of the type that re-emits in just one or two narrow wavelengths, not in a band of colors. The specific wavelengths emitted have been very carefully chosen to make the light emitted give a good specification for the CRI. But the spiky nature of the emission spectrum (see Fig. 1.2) means that colors at wavelengths other than these are poorly reproduced by the fluorescent lamp. Of course, there is no reason (we know of) that fluorescents can't have the same spectra that LEDs do. But for the present moment at least, LEDs have the potential to give much better color rendition than do fluorescent lamps.

INSIDE AN LED

This book is about designing lighting with LEDs, not about how to make them. Nonetheless, some aspects of their construction are worth knowing. It helps to understand some of the design aspects of different manufacturers' products. It also

helps to understand some of their claimed improvements in lifetime. We'll be talking about white LEDs made with phosphors, although much of the information is the same for other types.

The first thing to realize is that while almost all of the devices currently used by engineers—diodes, transistors, logic gates, microprocessors—are made of silicon, LEDs are not made of silicon. (There used to be some germanium devices around, but they don't work very well when they get hot, and so were abandoned.) However, it has proven difficult to get silicon to emit light. Thus a number of different semiconductors have been put to use. While it's not important to know the details, you should realize that there are a variety of different materials being tried. Not all of the physics is understood yet, and the aging processes are unclear as well. Different types are in use for different devices from different manufacturers. What this means practically is that you should expect changes ahead. The device you buy today will probably be different from what is available tomorrow.

The fundamental semiconductor device in an LED is relatively large, a few square millimeters. This device emits blue light (for white LEDs), and two things must be done to it: the blue light has to be converted to white light with high efficiency, and the white light has to come out without being blocked. So the normal ceramic package that ICs come in won't work, because it (intentionally) doesn't let any light through.

What most manufacturers do is to add some transparent silicone (a rubbery polymer) on top of the die. This lets the light come out without much absorption or color change, bending the light as needed, and providing a degree of mechanical protection for the die. At least one manufacturer then adds a piece of glass on top of the silicone, although it's not clear to us that this offers much advantage.

To accomplish the color conversion, a phosphor is used, which is a complex molecule that absorbs the blue light that the LED is emitting and radiates it out over a band of other colors. It takes two or three different phosphors to make a reasonable white color; you should expect to see phosphor blends with even more components in the future.

Some manufacturers put the phosphor directly on top of the die, with the silicone going on top of that. Others stir it in to the silicone before putting the mixture on top of the die. Putting it directly on the die increases the amount of blue light that is absorbed, but makes the phosphors sit at the same temperature as the die. Phosphors tend to degrade with high temperature. Indeed, phosphor degradation is one of the major reasons why LED light output decreases with age. Putting the phosphor in the silicone reduces the temperature the phosphors have to survive, but decreases the amount of blue light that is absorbed and converted. You could add more phosphor to compensate for this, except that phosphors are relatively expensive.

The die, phosphor, and silicone are all in a package. (And every manufacturer has its own package and footprint.) The package includes bond wires that connect the die to the leads so you can put current through the LED. Even though it's just a single device, multiple bond wires are used in parallel to accommodate the relatively high currents.

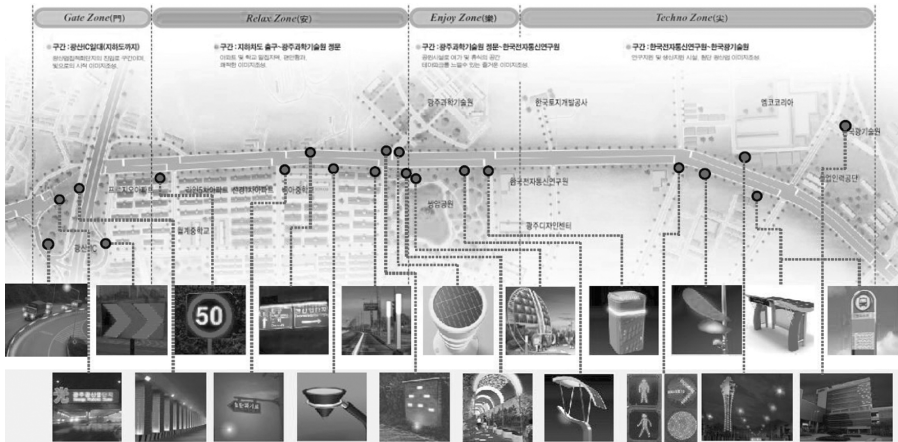


Figure 1.3 LEDs Can Be Used Everywhere. Source: Kaist, KAPID. See color insert.

Now the package has an unwanted side-effect. Since the LED emits light over a broad angle, some of the light is intercepted by the package. This affects efficacy somewhat, but also some of the intercepted light is reflected and emitted. That's okay, except that as the package ages (it's sitting at 85°C for 50,000 hours), it yellows. As the package yellows, the absorption of light by the package increases, which decreases the efficacy. And the reflected light is also yellowed, causing the correlated color temperature (CCT) and CRI of the emitted light to shift. In some devices, this package aging is one of the major reasons why the LED time to 70% light output is 50,000 hours and not longer.

Some LEDs also include some optics in their package. This may take the form of a lens and/or a mirror. The optics may be used to increase light extraction or to shape the emission direction of the light. If you don't care about the emission direction of the light (e.g., if you're building an omni-directional light bulb) you should try to avoid using devices with extra optics. (Why pay for the extra cost?)

Thus LEDs are complicated devices. It's well worth your while to ask detailed questions of your vendor about how the devices are made and how they will stand up to high temperature aging. You may even need to speak to people at the factory to get sufficient information.

IS AN LED RIGHT FOR MY APPLICATION?

To listen to enthusiastic marketing, it seems that LEDs can be used everywhere. But even though this book is about LEDs, we have to acknowledge that not every application will be best served by them. As LEDs continue to increase in efficacy

Table 1.1 Checklist of Considerations on Whether to Use LEDs for an Application

Question	LED	Fluorescent	Incandescent
Is energy efficiency top priority?	LEDs are probably best.		
Is cost an important factor?		Fluorescents should be considered.	
Is cost the <i>only</i> thing that matters?			Best to use an incandescent.
Does the application need long life?	LEDs, properly designed, are the best choice.	Fluorescents may be good enough.	
Are there lots of on/off cycles?	LEDs should definitely be used.		
Are there temperature extremes?	LEDs are better than fluorescents, and usually good enough.		For really extreme conditions, incandescent bulbs are even better.
Is the heat generated used for other purposes?	LEDs may not dissipate enough heat, e.g., to melt snow off a traffic light.	Fluorescents also may not dissipate enough heat, e.g., to melt snow off a traffic light.	Incandescent bulbs may remain a good choice.
Is good color rendition needed?	LEDs are sometimes good enough.	Fluorescents almost never are.	Incandescent bulbs remain the best.
Do colors need to be changed in operation?	LEDs are the only choice.		
Is a new form factor needed?	LEDs are the only choice.		

and drop in price, more and more applications will benefit from them. We expect that ultimately fluorescent tubes will become obsolete. But we also expect that incandescent bulbs will be around for a long, long time. Here's a checklist of things to think about in deciding whether an LED solution is right for your application.

HAITZ'S LAW(S)

You've probably heard of Moore's law. This was the prediction by Moore in 1965 that the performance of microprocessors would double every two years. It was based on observations, but proved to be remarkably accurate for the next 40 years. It is only now that it has finally slowed, as ICs reach some fundamental physical limits.

A similar prediction for LEDs was made by Roland Haitz (2006). This is backed by much more historical data (see Fig. 1.4). As currently stated, it predicts that the luminous output of individual LED devices is increasing at a compound rate of 35% per year and that the cost per lumen is decreasing at 20% per year. To the extent that current manufacturers seem to have settled on 3W as the maximum practical power in a small device, we can read this as also meaning an increase in efficacy of 35% per year.

This predicted rate of performance increase would be utterly unbelievable, except that it appears to be true. The authors began tracking the prediction a number of years ago, calculating where efficacy would be each month. Year after year, we have verified the numbers, and efficacy indeed continues to increase.

We talked to Haitz a couple of years ago about his law. His opinion was that it still had a long run ahead of it. And while he may be right that the lumens per device will continue to increase, in the next few years the efficacy will certainly start deviating, of course due to fundamental physical limitations.

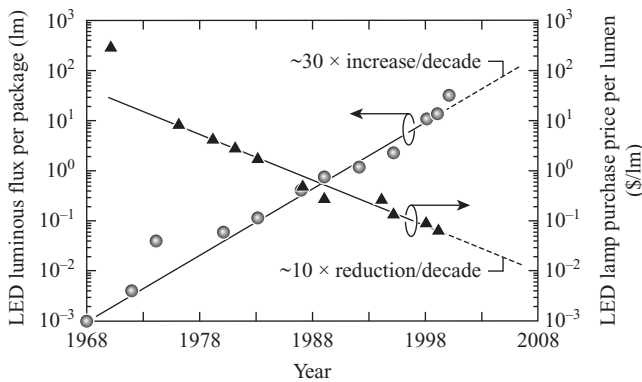


Figure 1.4 Haitz's Law.

Source: <http://i.cmpnet.com/planetanalog/2007/07/C0206-Figure3.gif>. Reprinted with permission from Planet Analog/EE Times, copyright United Business Media, all rights reserved.

To understand Haitz' law, we need to consider the meaning of "lumens" (see Chapter 3 for more detailed information). Lumens is not exactly a measure of light, but is rather a measure of how much light humans see with their eyes. As such, it very much depends on how eyes work. In particular, human eyes are most sensitive to green light. Thus, if you produce 1 W of light at 555 nm, you have 683 lumens. There's no possible way to increase this number; it is really almost a definition. The same is of course true for LEDs. If an LED gets 1 W of power, and converts it entirely to light at 555 nm, it will have no heat power dissipation at all. (Obviously, this is not really possible because of the Second Law of Thermodynamics.) All of that light then is equal to 683 lumens. So efficacy is limited to 683 Lm/W no matter what.

Now the reality is that we don't normally want intense green light. We want white light. And since white light consists of many different colors, the lumens and efficacy must be less than 683 Lm/W. What then is the real limit on efficacy?

There are two different limits, depending on how the white light is generated. Recall that white light can be made either by directly combining lights of different color or by emitting low wavelength light (such as blue or UV) and converting it with phosphors into white. The phosphor method is limited by the physical efficiency of phosphors. Since they absorb low wavelength (=high energy) light, and emit higher wavelength (=lower energy) light, the difference in energy is lost as heat. This is described by Stokes' law. While the exact limit is subject to details (such as what CRI light is acceptable), phosphor conversion of white light from LEDs is limited to about 238 Lm/W. Note that since fluorescent tubes are also phosphor-converted, but starting from 235 nm rather than 435 nm, their ultimate efficacy is considerably lower than that of LEDs. While they too have room for improvement currently, ultimately LEDs will be more efficient than fluorescent lighting.

Direct emission of various color lights can be more efficient, because there is no absorption and re-emission involved. But since colors other than green are needed, the human eye response means that 683 Lm/W isn't achievable this way either. A seminal paper by Ohno (2004) shows that, to get acceptable CRI, white light cannot be made at higher efficacy than about 350 Lm/W.

Haitz's law as extrapolated to efficacy thus has several more years to run. As the 200 Lm/W limit is reached, blue converted by phosphors will plateau in efficacy. To continue increasing in efficacy, red-green-blue (RGB) systems will need to be implemented. But if 35% per year is continued, only 2 years will remain before the ultimate limit in efficacy is achieved. After that, Haitz's law may still apply to the cost per lumen. Indeed, the figure shows that it is not until after 2015 that the cost per lumen of LEDs will approach that of 60 W incandescent bulbs.

Ultimately, then, LEDs can be expected to reach the theoretical limit of efficacy, and their cost can ultimately drop below that of incandescent bulbs. And what happens after that? Since efficacy can't be increased because of physics, it might be reasonable to suppose that LEDs are here to stay for the long term. Nothing can be better than LEDs, only cheaper.

THE WILD WEST

The LED lighting industry, and the LED industry in particular, is currently like the “Wild West”: there aren’t many rules, and most people aren’t paying attention to them anyway. All sorts of claims are being made that are obviously wrong, and plenty more that you need special equipment to detect.

Looking first at LED device production, we should start out by saying that there are *some* reputable manufacturers. These tend to be the largest ones, although you can’t assume that that’s true, either. They produce what they say they do, and their datasheet contains information from measurements they’ve taken. The problems start rather with their marketing departments.

The biggest players are presently in a contest to demonstrate that they have higher luminous efficacy white LEDs than their competitors do. As a result, they routinely release press announcements proclaiming their progress. Now everyone in the industry measures efficacy at a temperature of 25°C. That’s just a given. But actual operation is always at elevated temperatures, since LEDs heat up in operation. And the press announcements never mention how much that wonderful efficacy rolls off at higher temperatures. Different manufacturers’ processes have different roll-offs, so you don’t know what you would get from this new device. What’s more, it’s routine to announce results from a single lab device. It’s not in production, and very possibly not producible without major changes. So it’s all a bit of a cheat.

Moving on, even the big manufacturers tend to have problems with efficacy roll-off with aging, which is to say, lifetime. The truth is that the various manufacturing processes appear to create LEDs that age differently. And the aging varies greatly depending on the drive current, the die temperature, and even the package temperature. The fundamental problem is that 50,000 hours is 8 continuous years. There’s a new LED process a couple of times a year (recall the 34% increase in efficacy per year). So there isn’t time to collect data before the part is obsolete. You would think that you could extrapolate data from, say, the first 1000 hours. But the truth is that this works so poorly that the committee writing the specification for LED aging gave up on it. LED lifetime? It’s anybody’s guess.

Further, many of the LED manufacturers have problems with data. We’ve seen datasheets for products that have been in production for a year that still have forward voltage copied from a competitor’s datasheet. We see efficacy numbers that came from hand-held meters. In some cases, the parts don’t match the datasheets either in color or in efficacy. The sad story goes on and on. Thus, “Caveat emptor!” The only way to be sure of what you get is to measure it yourself. Read Chapter 11 to find out how.

Moving on now to LED bulb manufacturers, the situation is even worse, if possible. We tested a couple dozen different bulbs. Only 5% of them generated the lumens they claimed, with a majority of them being wildly off! In some cases, it was apparent that no measurement had been made at all. They calculated that each LED is rated at 60 lumens, and they put three of them in the bulb, and so the package says it is 180 lumens! No thought had been given to the drive current, the optics,

the packaging, not to mention the temperature effects. The U.S. Department of Energy is making efforts to clean this up. We hope for progress in this area.

We feel that all of these problems are characteristic of an infant industry. Doubtless all of this will improve. We just hope that consumers aren't so disappointed early on that the industry never gets to maturity.

LEDS AND OLEDS AND ...?

Incandescent bulbs replaced candles and kerosene lamps. Fluorescent tubes replaced incandescent bulbs for many purposes. It seems likely that LEDs will replace both fluorescent tubes and incandescent bulbs. What's next after LEDs?

There's been a lot of talk about OLEDs being the next big thing in lighting. The "O" in the front of the acronym OLED stands for "organic." But it's really still an LED. The difference is that this particular type uses organic rather than inorganic material. The OLEDs' claim to fame is that they are more mechanically flexible than inorganic LEDs. Perhaps they could be made directly into light bulb shapes or printed onto mechanical forms of light bulb shape.

As we indicated in the section on Haitz's law, LEDs are probably going to reach the maximum theoretical limit for efficacy of any light source. So if OLEDs are going to supplant LEDs, it can't be on the basis of efficacy, because it's impossible to be better. The same is true for any other new light source. Once the theoretical limit is reached, nothing can be better.

The way that OLEDs *could* supplant LEDs is if they were cheaper. Once there are a variety of possible ways of achieving the maximum efficacy, the market will ensure that the cheapest one is the one that dominates. In our view, OLEDs are really just another type of LED, and their progress is part of Haitz's law. So we don't know if OLEDs or LEDs will prove the eventual cost winner. But our opinion is that there probably won't be any newer technologies for lighting that end up completely replacing LEDs. LEDs will end up being so inexpensive that cheaper won't matter to consumers. We think LEDs are here to stay.

Chapter 2

Light Bulbs and Lighting Systems

This book is about lighting design with LEDs. While the rest of the book is about the LED part, in this chapter we present some background on the lighting part. The reason for this is that light bulbs have been around for more than 100 years. In that time, there have been many people working on them, and much technology has been developed. While we can't claim that this is a comprehensive survey, there's probably information in this chapter that you'll be happy you have.

A few words about terminology are in order. Wikipedia¹ says that "A lamp is a replaceable component such as an incandescent light bulb, which is designed to produce light from electricity." As you can see, there is a general confusion about what to call light-producing devices. Most consumers call the device a light bulb, and the unit that holds it a lamp. Manufacturers usually call the device a lamp, and the unit holding it a fixture. In this book, we will usually try to follow consumer usage. But the reader should be aware of the difference when reading publications.

LIGHT SOURCES

Incandescent

Light-emitting diodes are merely the newest in a long list of different types of lighting devices. Ignoring truly ancient devices such as candles, all of them use electricity. The first and still most common light source is incandescent. An incandescent bulb works by heating a piece of metal, the filament, until it glows. By adjusting the power level, it can be made to glow different colors. The typical incandescent filament runs at about 2850 K, resulting in the familiar yellow color.² When you

¹[http://en.wikipedia.org/wiki/Lamp_\(electrical_component\)](http://en.wikipedia.org/wiki/Lamp_(electrical_component)), under license. Accessed December 2010. <http://creativecommons.org/licenses/by-sa/3.0/>

²Could a consumer device that runs at half the temperature of the surface of the sun be introduced to the market today? And by the way, fluorescent tube plasma runs at 1100°K, much hotter than your oven.

dim the bulb it receives less power. This not only produces less light, it also reduces the temperature of the filament. This is why dimmed incandescent bulbs look reddish.

Note that the glass shell in an incandescent bulb is used to maintain a partial vacuum, preventing the filament from oxidizing and failing. There has been some research into altering the mixture of remaining gas in the shell to enhance bulb life.

The incandescent bulb runs very hot. The surface temperature of a 40 W runs about 120°C. That's why you have to wait a bit after turning it off to touch it. The common failure mode for an incandescent is for its filament to break. This typically happens after about 1000 hours of operation. Switching incandescent bulbs on and off a lot can also cause the filament to fail, but in typical operation this is not the dominant failure mechanism.

Before leaving the topic of incandescent light sources, a comment on safety is in order. If you unscrew an incandescent bulb from its socket without turning off the light switch, sticking your finger in the socket will connect you with 120VAC. This is life-threatening. If you try to hold an incandescent bulb when it's on, it will burn you. It's hard to imagine a device with these sorts of extreme problems being introduced today. Conversations with engineers at UL suggest that incandescent bulbs are "grandfathered in." They were there before regulations existed, and so they can't be easily eliminated. But it certainly seems like the time has come for engineers to come up with something better.

Halogen

Halogen bulbs are also incandescent. The difference between halogen and normal incandescent bulbs is that halogens contain a small amount of halogen. The halogen makes the filament burn hotter, which slightly increases the efficacy of the bulb. It also makes the CCT higher than in a normal incandescent. An additional benefit is that the halogen helps the filament to survive longer (by redepositing the filament material).

Fluorescent

Fluorescent bulbs work entirely differently from incandescent bulbs. They too have a partial vacuum inside a glass tube. In this case, though, the tube intentionally has some mercury vapor in it. When the filament inside the bulb is heated, it emits electrons. These ionize the mercury, forming a plasma arc at about 1100 K. The mercury emits UV light to go back to its normal state. The UV light hits a phosphor coated on the glass tube. This is the white coating on fluorescents. The phosphor absorbs the UV light, and emits visible light, which is the output of the bulb. The phosphor is carefully designed to produce just the color light that is desired, and is usually a mix of different phosphors.

To run this complicated device requires a special circuit called a ballast. The ballast is connected to the AC line as input, normally either 120 or 277VAC in the




Figure 2.1 Currents in a Fluorescent Tube.

United States. At its output the typical one bulb ballast has two pairs of wires. Each pair is heating one of the two bulb filaments. Additionally, current flows from one pair to the other. This latter is the current that produces the plasma arc. Figure 2.1 shows the currents in this lighting system.

Fluorescent tubes run much cooler than incandescent bulbs. The typical surface temperature is about 40°C. You can easily touch them and pull them out of their fixture while they're running. For this reason, they typically have an electrical interlock system. If the tube is not present, the ballast typically is designed to turn off to avoid shocking you if you stick your finger into the socket.

Fluorescent tubes have a variety of failure modes. The most common failure mode is for the filaments to break. This typically happens after about 10,000 hours, 10 times as long as incandescent bulbs. Because the metal of the filament runs so hot, the metal gradually burns off, weakening it. Additionally, every time the fluorescent is turned on, the sudden heating blows off some of the metal. This material lands on the glass, causing the end blackening seen in old tubes. Fluorescent ballasts also fail, but this is typically on the order of 10 to 25 years.

Induction Lighting

Induction lighting is a type of fluorescent lighting that was designed to overcome the lifetime limits of the filaments in normal fluorescent lighting. Induction lighting doesn't use filaments. The energy is introduced into the plasma through a transformer. In this case, the ballast is the primary of the transformer, and the plasma arc is the load on the secondary. Coupling between the primary and secondary is through the air, so the ballast needs to be close to the bulb.

Induction lamps have a rated life of 100,000 hours. Since there are no filaments to fail, the lifetime is determined by the vacuum seal on the bulb and by the time it takes the ballast to wear out. This sounds like a really good light, so why is it uncommon? We don't really know, but one has to wonder about the safety of being exposed to 13.6MHz radiation from these systems.

High Intensity Discharge (HID) Lamps

High intensity discharge (HID) lamps are fundamentally similar to fluorescent lamps. The major difference is that instead of generating UV light and converting it to visible, these gas discharge lamps emit visible light directly. For example, sodium vapor lamps use sodium instead of mercury. Sodium emits yellow light,

which is often seen in lights in parking lots. The term “HID” covers a variety of lamps that differ in the material used to generate the light, such as metal halide, sodium vapor, and xenon.

Because there is one less conversion step in these bulbs, HID bulbs tend to be more efficient than standard fluorescent tubes. They can be 100Lm/W versus 60Lm/W for fluorescent tubes. As with all lighting, there is a trade-off of light quality versus cost. The higher efficacy of these bulbs is offset by higher cost. And achieving higher CRI also reduces the lifetime of the bulbs, adding to cost.

CHARACTERISTICS OF LIGHT SOURCES

All of these light sources have various pros and cons. That’s why there are many different types of light source available. No one type has proven suitable for all applications. This section reviews some of the good and bad characteristics of these various light sources. This will help to give some perspective on the prospects for LED lighting.

Light Quality

The first characteristic to consider for a light source is the quality of the light. We’ll be getting into detail about light quality in Chapter 3. For the moment, we’ll address two simple measures: correlated color temperature (CCT) and color rendering index (CRI). While imperfect in a variety of ways, these two give a broad overview of a light source. The CCT describes the color temperature—for example, yellow is hotter than red. CRI describes how well a variety of different colors are reproduced—for example, a CRI of 82 is better than 60, and a CRI of 100 is perfect.

Noon sunlight has a CCT of about 6500K. A typical incandescent bulb has a CCT of 2850K and a CRI of 100. This latter is the basic yellow light bulb, and pretty much the standard against which other light bulbs are compared. A variation of this is the “daylight” incandescent bulb. In this bulb, the glass is tinted with neodymium, which absorbs some of the red light from the filament. The result is a higher CCT and a lower CRI. But for some lighting applications, this less yellowish light may be preferred.

Fluorescent lights are available in a variety of CCTs and CRIs. The most common type, used in offices, has a CCT of about 4100K and a CRI of about 82. This gives the cool white color supposedly best for working. It should be noted that the CRI is a misleading number when applied to fluorescent lights. (It’s misleading for LEDs, too, but that’s a subject for later in the book.) The reason is that CRI was designed for a spectrum of light that is smooth as a function of wavelength. But fluorescents’ spectrum is full of spikes; it emits light at a number of very specific colors. These colors have been carefully picked to give a good CRI number, but that doesn’t at all reflect the quality of the light as perceived by humans.

Efficacy

The other big characteristic of light sources is their efficiency, or rather efficacy: How much light they produce for how much energy you put in. Incandescent lights are the bottom of the heap here. They are basically just big resistors. For example a 60 W bulb produces 830 lumens, which is only 14 Lm/W. Higher power bulbs have slightly higher efficacy, but not by much. A “daylight” incandescent is even worse, only three-quarters of this efficacy, because one-quarter of the light is intentionally absorbed in the glass to change the color.

Fluorescents have considerably better efficacy. A 4-ft T8 tube produces 2700 lumens for 32 W, an efficacy of 84 Lm/W. But don’t be misled. You can plug an incandescent directly in to 120 VAC, but the fluorescent tube requires a ballast to convert the power. With a ballast running about 89% efficient, input power is actually 36 W, so efficacy for this fluorescent system is 75 Lm/W. Compact fluorescent lamps (CFLs) are even lower, around 60 Lm/W.

Timing

Timing covers both turn-on time and flicker. Incandescent bulbs are simple. When you apply power to them, they turn on within half a line cycle—at any rate, faster than your eye can see. And right away, they are at full brightness. Fluorescents are more complex. To ensure good lifetime of the tube, the filament is preheated before the plasma arc is created. This preheat time is typically 700 msec, which is quite noticeable, and sometimes stretches into seconds. Once the tube is on, it can then take minutes for it to come up to full brightness. This delay time is one of the major objections to compact fluorescent light bulbs (CFLs).

Some types of lamps have even longer start times.³ Sodium streetlights can take minutes to turn on. Since they are turned on at dusk, which is only roughly defined by a photo-sensor, this turn-on time is not perceived to be a problem. HID lamps in general don’t turn on again right after being turned off; you have to wait 10–15 minutes before you can turn it back on. This presents problems when there are power outages. Getting around this requires extra money, and so is part of system cost.

Flicker refers to what happens when a light turns off every time the AC line goes through 0 volts. Incandescent lights of course are subject to this, but you don’t notice it because the filament takes so long to cool down that the change in light isn’t noticeable. (Fundamentally, the filament has a long thermal time constant. You can see this when you turn off an incandescent bulb. Some light continues to be emitted for a noticeable fraction of a second afterward.)

Fluorescent tubes, including CFLs, however, extinguish their plasma arc within about 100 μ sec. (Incidentally, this is why operating a fluorescent tube above about

³http://ecmweb.com/lighting/electric_voltage_variations_arc/. Used by permission of Power Electronics Technology, a Penton Media publication.

10KHz gives a 10% efficacy advantage over a 60-Hz operation.) They thus can be seen to turn on and off 60 (or 50) times per second. This produces an annoying flicker in the light. The same problem potentially affects LED lights, since they turn off even faster than fluorescent plasmas do.

The problem here is conflicting engineering requirements. Of course, you could keep the lamp on by supplying some internal power storage, such as a capacitor. But putting a capacitor on the AC line results in a bad power factor. At least for LED lamps, the U.S. government is requiring a good power factor, and so this option is not available, at least not cheaply. You could also add a capacitor at the output. Some fluorescent ballasts do include a capacitor, but since it adds cost it is uncommon.

Dimming

Many lights are on dimmers. The way most dimmers work is simple. They disconnect the AC line from the load during part of the line cycle. This produces less power and therefore less light, albeit nonlinearly.

All light sources have some types of problems with dimming. Incandescent bulbs drop in CCT as they are dimmed, making them look progressively redder. Fluorescent tubes and CFLs generically just turn off if put on a dimmer, since they perceive the missing voltage as a decrease in average line voltage. Further, reducing the line voltage applied to a standard fluorescent ballast means that not only the arc current but also the filament power is reduced. This can enormously shorten tube life. In extreme cases, fluorescent ballasts put on dimmers have been known to catch on fire.

LEDs potentially have the same problems. If they are designed to produce constant light output, they won't respond at all to a dimmer. As less of the line voltage is present, the current drawn during those times will go higher. At some point, the current being drawn during the remaining part of the line cycle will be so high that the driver has to shut down to protect itself. LED lamps could also be designed to dim as the dimmer cuts out part of the line cycle, but then they need some way of knowing what percentage of time the line is missing. Furthermore, the driver circuit needs to stay on during the time when the line is zero, and so needs some hold-up capacitance. This is potentially another power factor problem.

A more fundamental problem for both CFLs and LED bulbs is their energy saving characteristics. Most dimmers have some minimum load specified, such as 30W. Since most incandescent bulbs are at least 40W (in the United States), they work fine with dimmers. But CFLs and LED bulbs may be considerably less than 30W. The dimmers don't work properly with these very light loads. They may turn on and off erratically, causing dimming not to work properly. In extreme cases, the dimmer can burn up. A solution that has been tried is to add in some extra load to maintain the dimmer power—but this then defeats the goal of energy savings. We'll talk more about dimming in Chapter 8.

Aging

If you replace only one of multiple incandescent bulbs in a fixture, it is immediately apparent that the older incandescent bulbs have grown dimmer over time. The same is true for fluorescents and LEDs. The differences between them are in how long this aging takes, and what happens at end of life of the bulb.

When the manufacturer of an incandescent bulb states its lifetime to be 1000 hours, this is the average time until a bulb stops working. As it turns out, this is also approximately the time until the bulb is about 70% as bright as when it started out. In the lighting industry, 70% as bright is considered to be the point at which a bulb is noticeably less bright. So this is good design; the incandescent bulb fails about the time that it starts getting noticeably dim.

Fluorescents are more complicated than incandescent bulbs, and thus their lifetime is more complex as well. Since fluorescent bulbs' lifetime depends on both operating hours and on/off cycles, their average lifetime depends on their usage pattern. So it may be that the recent reduction in the claimed lifetime of fluorescent tubes (from 10,000 hours to 8000 hours) wasn't related to design changes for the purpose of reducing cost, but rather to a re-evaluation of the typical usage pattern. If you put a fluorescent tube on a motion detector circuit, its lifetime may be reduced compared with that of an incandescent due to the number of times it has to turn on.

This variation in usage produces more of a spread in lifetime for fluorescents than for incandescent bulbs. But it's even more pronounced for LEDs. Since they are semiconductors, LEDs have extremely long lifetimes, probably hundreds of thousands of hours. But they dim long before this happens. A good part run in a good design might have 50,000 hours till it reaches 70% brightness. But unlike the incandescent bulb, the LED bulb keeps right on running after it reaches 70% brightness. So the U.S. government has decided that an LED bulb is "dead" after it reaches 70% brightness. That's the rating that is required on the box. But the reality for consumers is that it's going to keep on working almost forever. As a caution, we might note that no one yet has any real idea what the distribution of lifetimes for LED bulbs is going to be; not enough time has passed yet.

TYPES OF BULBS

Bulb Shapes

There are a surprising number of different bulb shapes available, as shown in Figure 2.2. The most common type in the world is the A shape, the standard pear shape. And in the A shape, the A19 is by far the most common. The number 19 after A signifies the diameter of the bulb in eighths of an inch (for the United States). Thus an A19 has a diameter of $19/8$ in = $2\frac{3}{8}$ in. In the rest of the world, the number is the diameter in millimeters, so that a common size is A55. The figure doesn't show it,

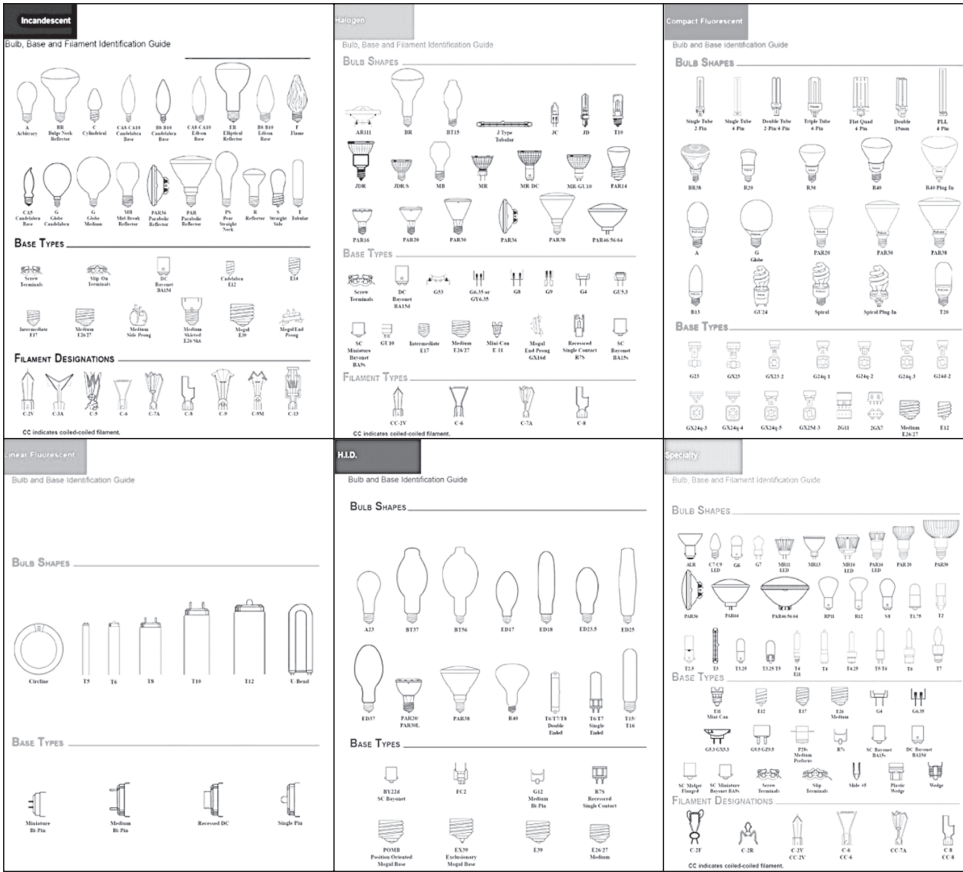


Figure 2.2 Various Bulb Shapes. Illustration courtesy of Halco Lighting Technologies.

but many of the bulb shapes come in various sizes. The A shape is available in A15, A19, and A21. The larger size is typically used for the higher wattage bulbs.

Other common shapes for incandescent bulbs include the BR and the PAR, which are used for floodlights and spotlights. As indicated by the name, spotlights have a narrow beam and are used to light a specific area. Floodlights have a broader beam. The G bulb is a sphere, and is commonly seen in residential bathrooms. Candelabra lights come in a variety of styles.

Halogen bulbs have a variety of shapes, some of which resemble those of incandescent bulbs. Popular types include the BR and PAR, used as substitutes for incandescent of the same shape, and MRs, which are used as spotlights. They are also used for track lighting. Track lighting comes in two types, one that uses 120VAC directly, and the other with a ballast that converts 120 VAC to 12 V. The bulbs for this type then are designed to work on 12V rather than 120V.

Fluorescents come in tubes and in compact format. The most common is still the T12, though in recent years these are being replaced by T8s. The number after

the “T” again refers to the diameter of the tube in eighths of an inch. A T12 has a diameter of $12/8$ in = 1.5 in. The normal length of a fluorescent tube is 4 ft, but there are also 8-ft types used in very high ceilings, such as the high output (HO) and very high output (VHO) tubes. There are also circlines and U-tubes, seen in ceiling panels that are too small for the normal 4-ft length of tube. CFLs are now most commonly seen in the spiral shape, but there are many multiple tube types available. (These are better referred to as “biax” and “triax.”) CFLs are also being put inside a variety of incandescent bulb shapes, sometimes with their own letter designers.

Bulb Bases

In addition to the various types of bulbs, there are also a variety of base types, but most bulbs come in only one or two base types. The most popular bulb shape is the A19 (or A55). In the United States and European Union, this comes in a bulb base that is called the medium screw base, the E26 (or E27 in the E.U.). This is thus probably the most common bulb base. However, in the United Kingdom and Ireland, the A55 bulb comes in a bi-pin base. And of course, fluorescent tubes also have two pins at each end.

In California, a pin base is used for government policy reasons. Under Title 24, new and remodeled houses are required to use energy-efficient lighting. Energy-efficient in practice means CFLs. But to prevent people from swapping out the CFL for an incandescent bulb, the fixtures are required to be a special type needing a pin base rather than a medium screw base. Anecdote suggests the requirement is well-founded. We have repeatedly heard of people swapping out the entire light fixture after the inspector has finished, to get rid of the hated CFL.

Specialty Bulbs

There are a huge variety of specialty bulbs. We call them specialty because they are not sold in billions, but they may still be sold in tens of millions. LEDs fall in this category for the moment. For example, cars use a variety of different lights: headlights, taillights, panel lights (for the dashboard), dome lights (for overhead lighting). All of these types have been incandescent in the recent past, and presumably will be LED in the near future. Taillights on trucks and buses now are almost all LED-based. Buses and airplanes use a small fluorescent tube for interior lighting. And at least in the United States, 80% of all traffic lights have converted from incandescent to LED. There’s actually been a problem reported with LED traffic lights. They generate so much less heat than the old incandescent bulbs that they don’t melt the snow off in the winter. We’ve also heard complaints from people in the Northwest Territories of Canada that without incandescent lighting, the house was not as warm in the winter.

Emergency exit signs are now nearly 100% LEDs. Here the reason is purely economic. An emergency sign is required to run for 90 minutes in the event of a power outage. This determines the size of the battery, the most expensive

component. So in this case a change to a light source that uses less energy saves money for the manufacturer.

Similar energy-saving considerations are behind a change from incandescent lighting in refrigerator cases to LEDs. The heat generated by the incandescent is not only wasted energy; it also has to be removed by the cooling unit! There is thus a double hit on the cost.

The desire for energy saving in computers and cell phones has also led to the use of LEDs for backlighting, even though at first it wasn't entirely economical. The reason in both cases was the desire to save the battery energy. For the same reason, flashlights now almost universally use LEDs. Here the reason is not only to save the battery, but also to increase the light output without overheating. Televisions are also headed in the same direction, although for general energy-saving reasons.

Almost everything named in this section has been about conversion to LEDs. One area that may not change is oven lights. Since ovens run up to 550°F (300°C), incandescent bulbs seem like a natural choice. They are not affected by heat, but semiconductors are. They may never be changed; there are still vacuum tubes used in some specialty areas.

HISTORY OF LIGHTING

Edison made the first practical incandescent light bulb in the 1880s. Fluorescent tubes became popular in the 1950s. LEDs are just starting to become popular in the 2000s. What will the future bring?

Jeff Tsao, a Principal Member of the Technical Staff at Sandia National Labs, researched this question, which deserves to be more widely known (Tsao, 2009). He looked at worldwide consumption of artificial light over the last several hundred years, covering candles, kerosene lamps, gas lighting, incandescent bulbs and fluorescent tubes. Obviously, much of the older data can't be that accurate. But when plotted on a log scale, these inaccuracies aren't that significant. What was clear is that over the entire time period, the world has spent an approximately constant percentage of its GDP on lighting, 0.72%.

Assuming this continues into the future, this has a surprising consequence. One of these is that as lighting becomes more efficient, total light used will increase. The introduction of LED lighting at, say, twice the efficacy of fluorescent lighting will mean that twice as much light is used, not that half the energy is saved. Energy is saved according to the model only if GDP decreases or the cost of electricity decreases.

This conclusion seems to fit with what is actually happening in the world today. Governments reasoned about energy-efficient lighting as follows. "If you replace all of the 60 W incandescent bulbs with 15 W CFLs, the energy saved will be 45 W times the number of bulbs." But these energy savings have not been entirely realized. Anecdotally, people report that they want brighter lights, not lower energy consumption. The reasoning apparently goes like this: "My closet was lit by a 40 W incan-

descent. With a CFL I get the same light for only 10W. But my closet was always dim. I can upgrade to a 15W CFL and still save energy!”

Our expectation is this. As Haitz’s law continues, the exponential increase in efficacy and decrease in cost will drive almost all lighting to convert to LED. This will be translated into a huge increase in the total amount of light used in the world. As the ultimate physical limits on efficacy are reached, the amount of light used in the world will stabilize. Increases will thereafter be small, driven by ordinary increases in GDP and population.

GOVERNMENTS

Governments worldwide have become interested in regulating lighting. A simple calculation shows why. Twenty billion light bulbs times 60W each times 4 hours a day equals over 1 *trillion* KWH a year! This is about 20% of all electricity used in the world. This huge energy consumption has serious consequences for the economy, for national security, and for the environment.

Such considerations have led to numerous government regulations. We’ve already mentioned Title 24 in California. Even more draconian measures are being implemented worldwide. Many governments have, or shortly will, ban the sale of incandescent bulbs entirely.

Why not just let the economics speak for itself, and let people switch to energy-efficient lighting to save money? The problem is that individually the change doesn’t make a lot of sense. Consider changing a 60W incandescent bulb to a 15W CFL (or a 10W LED bulb). That saves 45W. The bulb is on an average of 4 hours a day. That’s 180Wh a day. With electricity costing 12 cents/kWh, that switch to energy-efficient lighting has saved 2 cents! While this adds up to \$7.88 for the whole year, in the developed countries such an amount may not be considered worth the time spent—and many people feel it is not worthwhile going to a bad-looking light source for this amount of savings.

LIGHTING SYSTEMS

So far we’ve discussed light bulbs and their usage. But bulbs aren’t used by themselves; they need infrastructure to work, such as electricity and ballasts. This comes under the heading of lighting systems. Details about the electricity supply will be discussed in Chapter 8. As for ballasts, corresponding to the large number of different bulbs there are a large number of ways to run them. In this section, we’ll consider ballasts for fluorescents and then for LEDs.

As we’ve indicated, fluorescent tubes have special voltage and current requirements that mean they can’t be powered directly by the AC line. Instead, the AC line powers a special electronic circuit called a ballast, which in turn powers the fluorescent lamp. Fluorescent ballasts come in two varieties, instant start and rapid start. The difference lies in how the filaments are treated. You’ll recall that a filament is used to provide electrons to the plasma arc. In a rapid start ballast, the filaments are

preheated, that is, they have power applied to them and are warmed up to about 1100 K before they have to start emitting electrons. This has the benefit of making the filaments last a long time, but the downside of requiring extra power. With an instant start ballast, the filaments are not preheated. This saves power, but makes filament life, and thus tube life, much shorter.

In principle, you could get the best of both worlds by preheating the filaments, and then turning them off after arc ignition. But this isn't common, probably because of the cost of the extra circuitry to do this. Something similar is done, however, with dimming ballasts. (Normal fluorescent ballasts don't work with dimmers; you have to buy a special type that does.) As the tube is dimmed, the plasma arc heating the filament heats it less and less. To maintain filament temperature, which is what keeps them running for a long life, filament power is increased in a dimming ballast inversely proportionally to the degree of dimming.

One other differentiation among fluorescent ballasts is that some are magnetic, while others are electronic. Magnetic ballasts are basically a large piece of iron in the form of a transformer, with a big capacitor on the output to limit current into the fluorescent tube. They of course run at line frequency, which is why they are big and heavy. Electronic ballasts are switch-mode power supplies. They are much lighter, but have the same mechanical size as magnetic ballasts for backward compatibility. They tend to run above 20 KHz, for reasons of size and to avoid audio noise, as well as to take advantage of the small gain in efficacy of fluorescent tubes above 10 KHz.

Most LEDs are similar to fluorescent tubes in not being directly compatible with the AC line. (Some are, and this may well be the future of LED systems.) They also require a ballast to convert AC power to the constant current that is best for LEDs. Chapter 8 in this book covers design of these ballasts. It is to be noted that some people reserve the word "ballast" for fluorescent tube electronics, and prefer the word "driver" for LEDs.

We have the AC power line, the ballast, the lamp; surely that's everything in the lighting system? No, there's also the physical environment the lamp operates in. It too influences light production and efficacy. Since incandescent bulbs operate at thousands of degrees, they are relatively insensitive to environmental temperature; not so fluorescent and LED bulbs. Fluorescent systems require a longer time to start when it's cold. The fluorescent tube in your garage has trouble starting in winter for this reason. There are special fluorescent ballasts designed for cold weather. Conversely, CFLs can have problems with high ambient temperature. Placed into a ceiling can, the temperature can rise enough to cause the CFL to fail, sometimes catastrophically. Since the actual temperature in the can depends on many factors (how much power is being dissipated, whether the can is insulated or enclosed, what the ambient temperature is), in general CFLs are not recommended for can usage.

LEDs lights have similar problems. Although they are not as sensitive to cold, their ballasts may be. And if they get hot, their light output decreases and their lifetime also goes down. In an insulated enclosed can, they can get hot enough to fail. This can be particularly true if you design an LED system to have constant light output. As the temperature increases, their efficacy decreases. Thus you put more

power into the LEDs to compensate for their decrease in efficacy. But this causes the temperature to rise even more.

Another environmental factor is the fixture into which the bulb is placed. This can dramatically influence the light output of the lighting system. As an extreme example, consider a fluorescent tube mounted into a black fixture. Since half of the light from the tube goes upward, this half will be absorbed. The actual light output from the system will be only half that generated by the tube, thus also reducing the effective efficacy by half. While tube fixtures are not usually black, they can lose as much as 30% of the light output. LEDs can potentially avoid this loss because they are not omnidirectional.

One more environmental factor almost never discussed is dust. There is some residual voltage present on fluorescent tubes that tends to attract dust. (The same is more noticeable on your laptop computer screen.) Of course, no one dusts their lamps. Over years of operation, this dust can cut the actual light output from the lighting system by 30%. Whether LED bulbs also suffer this problem, or can claim additional efficacy over fluorescent lighting systems, remains to be seen.

Chapter 3

Practical Introduction to Light

Light is simply electromagnetic waves whose wavelength is in the visible range. The energy associated with light waves is measured in watts. The quantities used to characterize light physically are *radiometric*. The way light looks, for example, what color or how bright is it, depends on the human eye and brain. The human eye responds to roughly 380nm to 780nm wavelengths. But it gets quite complicated since the eye's response varies depending on relative conditions and differs from individual to individual. The quantities used to characterize how light is perceived by humans are *photometric*.

A light source is characterized by its emission spectrum. Lighting engineers care about the amount of light and the color of the light, not just from the emitter but also from absorbing and reflecting surfaces illuminated by the light. Most importantly, we must understand how the end user perceives the light.

To architects, often only lumens and color temperature are specified. To cockpit instrument designers, light intensity is the most important factor. To LCD screen engineers, brightness is the most important measure. To lighting designers, irradiance is the important quantity. But to the consumer, color rendition may be of greatest importance. We shall cover all these concepts at a practical level for the LED lighting engineer.

THE POWER OF LIGHT

Background: Light as Radiation

The energy produced by light (e.g., of a certain wavelength) is measured in joules; the rate of joule usage is given in watts. This is the radiometric quantity, the radiated energy per second. The energy is proportional to the frequency of light via the Planck-Einstein equation, $E = hf$, where h is Planck's constant, and f is frequency.




Figure 3.1 The Electromagnetic Spectrum. See color insert.

Most of the time, light is referred to by its wavelength. For example 475 nm is blue light and 650 nm is red light. Since frequency f , wavelength λ , and speed of light c are related by $f\lambda = c$, we have $E = hc/\lambda$. Therefore the shorter wavelength blue light has higher energy than the longer wavelength red light. If we extend the spectrum farther out, even shorter wavelength than blue light is ultraviolet (UV) light at 230–400 nm. UV has much higher energy than visible light. On the other end of the spectrum is infrared light (IR), with a wavelength longer than that of red light, at 700 nm–1 μ m.

Light of a single frequency is *monochromatic* light, such as that from a laser or a color filter. But most sources of light have a mix of colors. The radiant power of such a light is the sum of all of the power at each frequency.

When referring to light sources, the term “radiant power” is often replaced by the term *radiant flux*. Flux refers to total power radiating out in all directions, regardless of how much area it intercepts. Imagine a candle inside a balloon. The inside surface of that balloon sees all the flux of the candle. If the balloon is blown larger, it would still see the same flux.

RADIOMETRIC VERSUS PHOTOMETRIC

So far we’ve only talked about the light source, which is measured by radiant flux. Since we are only concerned with how that light is perceived by the human eye, we use the radiant flux as weighted by the spectral response of the eye to derive the *luminous flux*, measured in lumens (lm). All quantities weighted by the eye response are photometric quantities. For every radiometric measure, there is an equivalent photometric measure. Photometric quantities are measured in “lumens.” The word “photometric” refers to the visible range of light to which the human eye is sensitive.

Figure 3.2 shows the eye response curve as standardized by the Commission Internationale de l’Eclairage (CIE). Photometric quantities are weighed by *photopic vision*, or bright light vision, as opposed to night vision, or *scotopic vision*. You’ll

Table 3.1 Radiometric versus Photometric Description of Light Power

Quantity		Description	Unit	Symbol
Radiometric	Real Power	Radiant Flux, Radiant Power	Watts (W)	
Photometric	As seen by Humans	Luminous Flux	Lumen (lm)	Φ

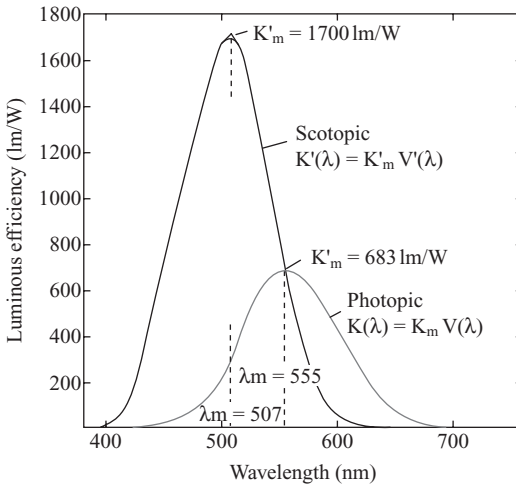


Figure 3.2 Scotopic Vision is much more Sensitive than Photopic Vision.
Source: Kalloniatis and Luu (2007).

notice that the photopic vision peaks at 555 nm yellow green. Human eyes are most sensitive to this color green. This is why a green laser pointer looks much brighter compared with blue or red laser pointers of the same power.

Scotopic vision has a similar shape sensitivity curve, but is shifted toward the blue, shorter wavelength side by 45 nm, to peak at 510 nm, or pure green. Scotopic vision is also much more sensitive to light than photopic vision. Its peak efficacy is around 1700 Lm/W, as opposed to 683 Lm/W for photopic vision. A 15 lm light won't look as bright as it looks in a dark room where the scotopic vision takes over.

Intermediate between scotopic and photopic vision is a level of light where both are active simultaneously. This is called *mesoscopic* vision. An example of mesoscopic light levels is the light outdoors at night, where there are a lot of background sources at a distance. If you are designing for mesoscopic lighting conditions, remember that the usual lighting levels are specified for photopic vision. This may result in producing more light than is actually necessary for the mesoscopic application.

The conversion between radiant and luminous flux depends on the wavelength spectrum of the light. Returning to photopic vision, if we take a 1 W green light at the peak of the human eye response, 555 nm, then the luminous flux is 683 lm.

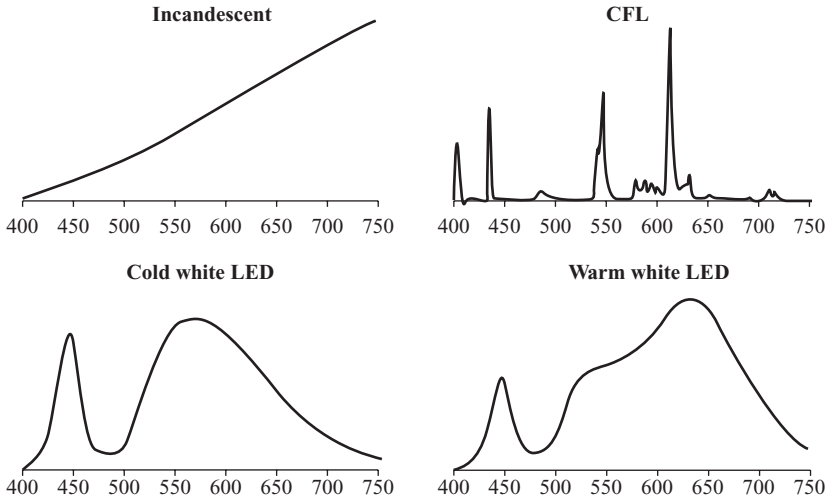


Figure 3.3 Emission Spectra of Four Common Light Sources.

At this peak, there is 100% conversion from radiated optical watts to lumens. This is the maximum efficacy in the sense that 1 W of any other wavelength or color of light will convert to fewer lumens.

Let's take a look at how the spectrum affects the conversion from radiant flux to luminous flux. Figure 3.3 shows the spectra of four light sources. Incandescent light bulbs are black body radiators and weigh heavily toward the long wavelength (red) end and into the infrared range. Compact fluorescent light bulbs have a spike at the edge with ultraviolet and then spikes at each of the phosphor wavelengths. The warm white LED has a phosphor emission that is more toward the red end when compared with the cold white LED.

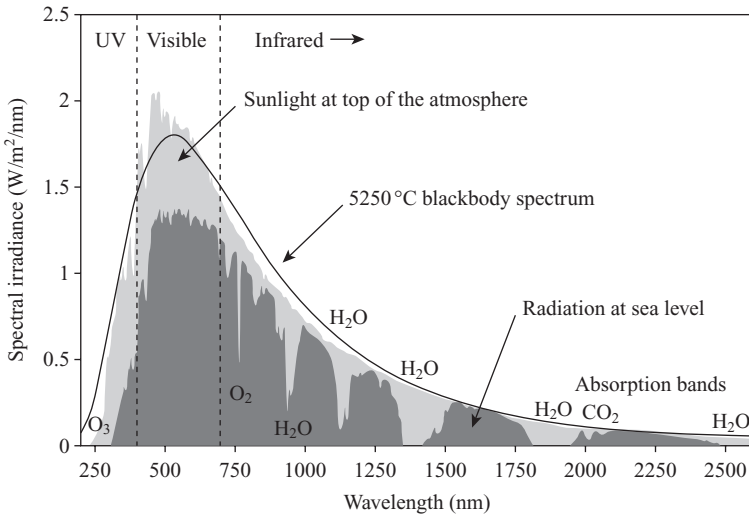
Let's compare radiated power of each light to the number of lumens. Table 3.2 lists some typical values. In the last column, we see that the incandescent bulb has the lowest conversion factor from radiant flux to luminous flux. This is as we expected since its spectrum was the worst match for the eye response spectrum. Another interesting thing to notice is that CFLs and LEDs are twice as efficient as incandescent bulbs, but more than four times as efficacious due to the closer matching between their spectrum and the eye response. Obviously, CFLs have a narrow spectrum and consist of only a few frequencies. LEDs on the other hand are much wider spectrum. This difference has implication in the color rendition, as we shall see later in this chapter.

For comparison, the solar radiation spectrum is shown in Figure 3.4. It is fairly evenly distributed save for the violet colors near the UV.

An incandescent bulb radiates light approximately evenly in all directions (it is *isotropic*). This qualifies it as a perfect light source to be specified in lumens. In fact some boxes of bulbs do list the lumens along with the usual wattage. In Table 3.3, you'll notice the efficacy is higher in the higher wattage incandescent bulbs. This is

Table 3.2 Radiant and Luminous Flux

Light Source	Input Power	Color Temperature	Radiant Flux	Efficiency	Luminous Flux	Efficacy	Luminous Flux/ Radiant Flux
Symbol	P	CCT	Φ_{rad}	Φ_{rad}/P	Φ_{lum}	Φ_{lum}/P	$\Phi_{\text{lum}}/\Phi_{\text{rad}}$
Units	W	K	W		Lm	Lm/W	Lm/W
Incandescent	58	2700	5.38	9.3%	749	12.9	139
CFL	9	2600	1.50	16.6%	511	56.8	342
CW LED	7.5	3900	1.77	23.7%	619	82.6	349
WW LED	8	3050	1.83	22.6%	518	64.1	284
555 nm	—	—	1	—	683	—	683
Green Light							

**Figure 3.4** Solar Radiation Spectrum.

Source: http://en.wikipedia.org/wiki/File:Solar_Spectrum.png under license <http://creativecommons.org/licenses/by-sa/3.0/>. Accessed January 2011. See color insert.

Table 3.3 Incandescent Efficacy Increases with Wattage

Wattage	Initial Lumens	Initial Efficacy (Lm/W)
40	505	12.6
60	865	14.4
75	1190	15.9
100	1710	17.1

due to the higher temperatures of the filament, and thus the higher black body radiation efficiency.

LUMINOUS INTENSITY, ILLUMINANCE, AND LUMINANCE (OR CANDELA, LUX, AND NITS)

Unless you're an expert in lighting, chances are very good that you do not know the difference between luminous flux, luminous intensity, illuminance, and luminance. We're going to go into detail in this section to remove the confusion, but here is a summary. *Luminous flux* is how much light is perceived by humans as coming out of a source, and is measured in lumens. *Luminous intensity* is how much light is coming out into a given solid angle, and is measured in candela = lumen/steradian. *Illuminance* is how much light is coming out onto a given area, and is measured in Lux = lumen/m². And finally, *luminance* is how much light is coming out onto a given area in a given solid angle, and is measured in candela/m² = nits.

Luminous Intensity

The quantity of luminous intensity (I_{lum}) adds a directional component to luminous flux. It is the amount of light radiated into a certain two dimensional angle. This two dimensional angle is called the solid angle and is measured in *steradians* (sr). The solid angle traces out a cone radiating out from the source. Using SI units, a unit sphere of 1 m radius has a surface area of $4\pi m^2$. So a sphere is 4π steradians, half a sphere is 2π sr, and so forth. 1 sr intersects the unit sphere in a circular bowl shape. When that bowl has a surface area of $1 m^2$, then the solid angle is defined as 1 sr. Or more generically, a sphere of radius r is defined to be intersected by a solid angle of 1 steradian when the surface area described by that solid angle is r^2 (see Fig. 3.5).

The SI unit of luminous intensity is the *candela* (cd). One candela is equivalent to one lumen/steradian. 1 candela is roughly the luminous intensity of a candle. Since 1 cd is 1 lm/sr, to find the lumens in one complete sphere, we multiply by 4π . A candle thus emits roughly 12.6lm.




Figure 3.5 One Steradian Intersects $1 m^2$ of Area of a 1-m Radius Ball.

Source: <http://commons.wikimedia.org/wiki/File:Steradian.png>
under license

<http://creativecommons.org/licenses/by-sa/3.0/>.

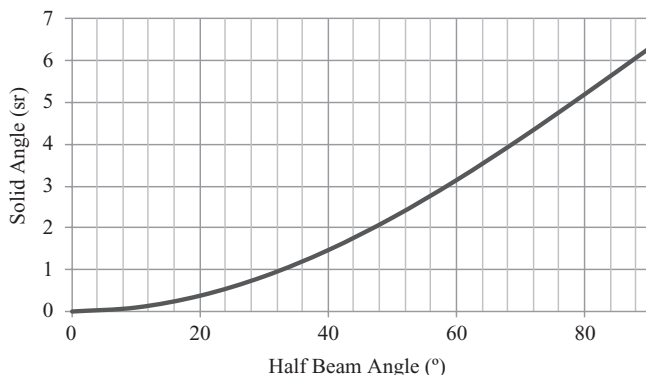


Figure 3.6 Solid Angle in Steradians (sr) versus Half Beam Angle in Degrees ($^{\circ}$).

A typical 100W incandescent bulb emits 1700lm at start of life. These lumens are radiated over the entire sphere of solid angle, or 4π sr. To calculate the luminous intensity for this bulb, we divide 1700lm by 4π sr to get 135 candelas.

Let's take the same bulb and put a perfect reflector behind it such that all the light is emitted into a single hemisphere, a solid angle of 180° or 2π sr. Then the luminous intensity is $1700\text{lm}/2\pi$ sr = 271 cd. It is double the intensity because the same amount of light is going into half the area.

Light Source	Luminous Flux (lm) Φ_{lum}	Angle 2θ	Half Angle θ	Solid Angle (sr) Ω	Luminous Intensity (lm/sr = cd) $I_{\text{lum}} = \Phi_{\text{lum}}/\Omega$
Candle	12.6	360°	180°	4π	1
100 W Incandescent	1700	360°	180°	4π	135
100 W Incandescent with 180° reflector	1700	180°	90°	2π	271
100 W Incandescent with 36° reflector	1700	36°	18°	0.3075	5528

For an arbitrary angle θ , the solid angle Ω can be calculated by the formula $\Omega = 2\pi(1 - \cos(\theta))$. Be careful to define this angle as half of the total angle. For example, in the hemisphere, the beam angle $2\theta = 180^{\circ}$, so that the half angle $\theta = 90^{\circ}$, and thus $\Omega = 2\pi(1 - \cos(90^{\circ})) = 2\pi$ steradians. At a beam angle of 60° , $\theta = 30^{\circ}$, and we find $\Omega = 2\pi(1 - \cos(30^{\circ})) = 0.84$ sr. See Figure 3.6 for a graphical representation of this function.

Now let's take that same 100W bulb and extend the reflector such that all the light is emitted into a 36° flood light. Setting $\theta = 18^{\circ}$, $\Omega = 0.3075$ sr, and $I_{\text{lum}} = 5528$ cd.

Lower power indicator style LEDs that emit light mostly in a forward pattern will be specified in candelas rather than lumens. For example, an LED might have a **view angle** of 15° , intensity of 13 cd, and flux of 0.7lm. All of the lumens




Figure 3.7 Definition of Beam Angle.

Table 3.4 Common Types of MR-16

MR-16	Beam Angle	I_{lum} (candela or CBCP)	Solid Angle Ω (sr) = $2\pi(1-\cos(\text{beam angle}/2))$	Luminous Flux (lumens) $\Phi = I_{lum} \times \Omega$
20 W Spot	12°	3350	0.0344	115
20 W Narrow Flood	24°	880	0.1373	121
20 W Flood	36°	450	0.3075	138
35 W Narrow Flood	24°	1700	0.1373	233
35 W Flood	36°	950	0.3075	292
35 W Wide Flood	60°	370	0.8418	311
50 W Spot	12°	8600	0.0344	296
50 W Narrow Flood	24°	2700	0.1373	371
50 W Flood	36°	1500	0.3075	461
50 W Wide Flood	60°	620	0.8418	522

are emitted within the view angle. So with $\theta = 7.5^\circ$, we get $\Omega = 0.05$ sr, and $I_{lum} = 0.7 \text{ lm}/0.05 \text{ sr} = 14 \text{ cd}$. The 1 cd difference between our number and theirs is probably due to rounding, since there’s only 1 significant figure.

Let’s say that you want to build an MR-16 equivalent light bulb with LEDs. Table 3.4 shows common types of MR-16 bulbs and their lumen output. *Beam angle* is defined by the point where the intensity drops to 50% at maximum intensity (see Figure 3.7). A 20 W, 12° spot light has 3350 CBCP (center beam candle power). “Candlepower” is an old term for candela. So let’s first calculate how many lumens are emitted. We set $\theta = 6^\circ$ to get $\Omega = 0.0344$ sr. Then we multiply $I_{lum} \times \Omega = 3350 \text{ cd} * 0.0344 \text{ sr} = 115 \text{ lm}$. Remember that candela = lumens/sr. So for an equivalent system, we need to have the LEDs emit 115 lm into a 12° beam angle.

We summarize the various measures of light so far in Table 3.5.

Illuminance

So now let’s take a look at how this light illuminates a surface, such as a book, a picture on the wall, or a keyboard. The amount of light Φ landing on a surface of

Table 3.5 Measures of Light

Quantity	Symbol	Radiometric	Unit	Photometric	Unit
Amount of Light	Φ	Radiant Flux, Radiant Power	Watt (W)	Luminous Flux	lumen (lm)
Intensity of Light	$I = \Phi/\Omega$	Radiant Intensity	Watts/steradian (W/sr)	Luminous Intensity	lumens/steradian = candela (cd), or candlepower
Light on a Surface	$E = \Phi/A$	Irradiance	W/m ²	Illuminance	Lux

area A is defined to be *Illuminance* E_{lum} and is measured in lm/m^2 or *Lux*. In the English unit system, illuminance is measured in *foot-candle* (fc) which is lm/ft^2 . The conversion factor is that one fc = 10.76 Lux.

Taking our MR16 example above with 115 lm, let's calculate the illuminance. The solid angle was 0.0344 sr, which is $0.0344/4\pi = 0.0027$ of the total solid angle of a sphere. Half a meter away, a sphere of radius one half meter has a surface area of $\pi \text{ m}^2$, and so the area illuminated by the MR16 is $0.0027 * \pi \text{ m}^2 = 0.0086 \text{ m}^2$. The illuminance is thus $115 \text{ lm}/0.0086 \text{ m}^2 = 13,372 \text{ Lux}$. You should note that a Lux meter is less than \$100 and is the easiest way to measure light. However the meter is quite sensitive to the exact distance to the source via the inverse square law. So if the Lux meter moves 1 cm farther away, then its reading will change to $13,372 \text{ Lux} * [(0.50)^2/(0.50 + 0.01)^2] = 12,853 \text{ Lux}$.

As another example, consider the following. Our company wants to build a USB light for lighting up a computer keyboard and for reading printed material. The USB port can run an LED at 37 lm (see Chapter 10 for this number and its usage in a real design example). Is this bright enough?

For this calculation, we need three numbers. We will take the distance from the LED source perpendicularly down to the keyboard to be $d = 0.33 \text{ m}$ (see Figure 3.9). We will take the larger edge of the keyboard to be $x = 0.3 \text{ m}$. The smaller edge we will take to be $y = 0.12 \text{ m}$.

Now the key to the calculation is to know how the LED light is dropping off as you get further away. In general, the *Lambertian* source (Fig. 3.8) can be adequately modeled as a cosine: At 0° it is 1.00, and at 90° it is approximately zero. Without going into all the math, the answer turns out to be that the part of the light that is intersected by the keyboard will be proportional to $x^2/(x^2 + 4d^2)$. We thus find that the luminous flux on the keyboard is $37 \text{ Lm} * (0.3 \text{ m})^2/[(0.3 \text{ m})^2 + 4 * (0.33 \text{ m})^2] = 6.34 \text{ Lm}$. In other words, the size of the keyboard is small compared with the height of the LED above it, so that only 17% of the light emitted by the LED is shining on the keyboard. This is then fairly poor efficiency usage of the light. Since we know the total luminous flux on the keyboard, the illuminance is the flux divided by the area. And thus we have illuminance = $6.34 \text{ Lm}/(0.3 \text{ m} * 0.12 \text{ m}) = 176 \text{ Lux}$. Looking at Table 3.6, we see that this is enough for casual reading.

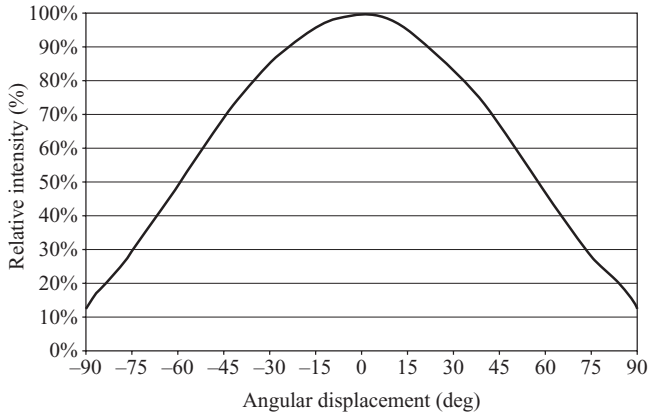


Figure 3.8 Typical Lambertian Radiation Pattern.

Source: Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Co., 2007.




Figure 3.9 Dimensions for a USB Keyboard Light Design.

It is worth noting that if the LED were brought in closer, the efficiency would be much better. For example, if the distance were halved, $d = 0.17\text{ m}$ instead, and we would have $37\text{ Lm} * (0.3\text{ m})^2 / [(0.3\text{ m})^2 + 4 * (0.17\text{ m})^2] = 16.2\text{ Lm}$. The keyboard then uses 44% of the light instead of 17%, more than 2½ times better.

Luminance

In this book we won't actually be using luminance for any of our examples, as its usage is slightly specialized. But a brief overview of the concept is necessary to round out your knowledge of lighting. Luminance is sort of a combination of luminous intensity and illuminance. It is how much light is striking (or coming out of)

Table 3.6 Examples of Illuminance in Lux

Illuminance (Lux)	Example
0.00001	Light from Sirius, the brightest star in the sky
0.0001	Total starlight, overcast sky
0.002	Moonless clear night with airglow
0.01	Quarter Moon
0.27	Full moon on a clear night
1	Full moon overhead in tropical latitudes
3.4	Dark limit of civil twilight under a clear sky
50	Family living room
80	Hallway/toilet
100	Very dark overcast day
320–500	Office lighting
1000	Overcast day; typical TV studio lighting

Source: <http://en.wikipedia.org/wiki/Lux> under license <http://creativecommons.org/licenses/by-sa/3.0/>

a surface per unit area (like illuminance) per solid angle (like luminous intensity). Its units are thus candela/m² (also called a *nit*).

Fundamentally, luminance is a measure of how bright a surface appears to a human eye at a specified angle to the surface. Its main present-day use is in characterizing the brightness of computer monitors.

Summary of Amount of Light

To reiterate from the opening of this section, there are four common measures of light. Luminous flux is how much light humans perceive as coming out of a source, and is measured in lumens. Luminous intensity is how much light is coming out into a given solid angle, and is measured in candela. Illuminance is how much light is coming out onto a given area, and is measured in lux. And luminance is how much light is coming out onto a given area in a given solid angle, and is measured in candela/m².

WHAT COLOR WHITE?

With monochromatic light sources such as lasers or nearly monochromatic ones such as color LEDs, color is a straightforward quantity. With a white LED, a continuous spectrum of different colors is mixed, as is shown in Figure 3.10. Notice how dramatically different spectra can all be considered to be white light.

The wavelength spectrum can be converted to a point on any color space to represent the resultant color as in Figure 3.11. The (*x*, *y*) values are referred to as *chromaticity coordinates*. In this 1931 RGB color space, a color is represented as a




Figure 3.10 Spectra of Neutral-White (Left) and Warm-White (Right) LEDs.
 Source: Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Co., 2007.




Figure 3.11 CIE 1931 (x, y) Chromaticity Space, Showing the Planck Line and Lines of Constant CCT.
 Source: http://en.wikipedia.org/wiki/Color_temperature under license <http://creativecommons.org/licenses/by-sa/3.0/>. See color insert.

mix of varying amounts of red, green, and blue colors. The colors on this chart are highly saturated to show the colors better.

In the center of the color space, there is an area of generally white light where the red, green and blue are added about equally. A useful reference for white light is a *black body* or *Planckian* radiator. A Planckian radiator will radiate different color light at different temperatures (think of “red hot”). The colors of these temperatures between 2700 K and 20,000 K are all perceived as white light by the human eye. This forms the Planckian locus and is also plotted in Figure 3.11 as the curving line in the center. The temperatures in Kelvin are referred to as *color temperature*. For example, the sun is a black body radiator. The sun at noon on the equator has a color temperature of 6500 K. At sunset, the color temperature will be much lower and have a reddish color. Sunlight color also depends on time of day, latitude, and season.

The incandescent light bulb is also a black body radiator and has a color temperature of about 2700 K. The halogen light bulb is also an incandescent black body. The difference between them is a small amount of halogen the halogen light bulb contains, which improves its efficacy. The halogen allows it to burn hotter and therefore has a color temperature of 3000 K.

Light sources whose spectral distribution is different from a Planckian radiator will not land on the Planckian locus. They can, however, be color matched to the most similar color temperature of the Planckian locus. For these sources the term *correlated color temperature (CCT)* is used. CCT follows the *isotemperature* (equal color temperature) lines, also shown in Figure 3.11.

Fluorescent tubes for use in offices are generally 4100 K, although both lower and higher color temperature tubes exist. The term “warm white” was defined by ANSI to help consumers get a better handle on fluorescent light colors and refers to a CCT of 2700 K and 3000 K. “White” is 3500 K. “Cool white” is 4100 K. “Daylight” is 5000 K and 6500 K. Compact fluorescents have improved the phosphors for warm white versions so that their color temperature can be as low as 2700 K.

LED manufacturers are following the same line of thought and getting creative with terms such as “neutral white” and redefining what cool white refers to. The ANSI C78.377-2008 standard for LEDs simply lists a “nominal CCT” of 2700 K, 3000 K, 3500 K, 4000 K, 4500 K, 5000 K, 5700 K, and 6500 K. The names of colors are mentioned as a side comment. The lesson here is that just looking at the name of the color may not be that informative. You need to check what “CCT warm white” from a manufacturer really means.

MacAdam Ellipses

It is important to realize that a point plotted on the color space is simply an equivalent color as perceived by the human observer. In actual experiments, a human observer is presented a source light on one side of a screen. On the other side, a light is formed by varying the amounts of red, green, and blue light until the two lights look like the same color. The limited sensitivity of the eye gives rise to an ellipsoidal area in which color differences cannot be perceived. These are called *MacAdam ellipses*

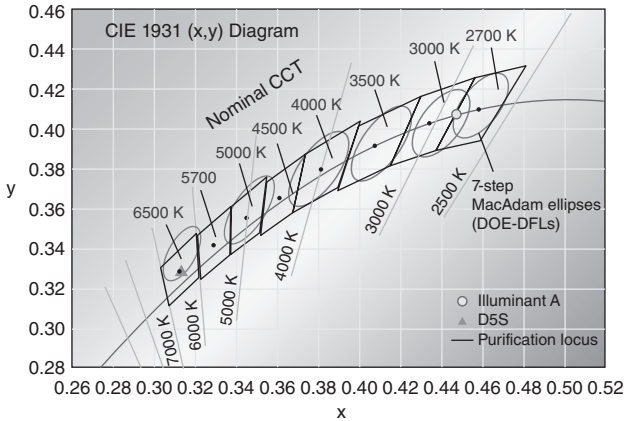


Figure 3.12 (x, y) Chromaticity Diagram Showing CCT and 7-Step MacAdam Ellipses.
Source: <http://www.photonics.com/Article.aspx?AID=34311>. See color insert.

(MacAdam 1942) (see Figure 3.12) and are used as the foundation for color bins for LEDs and fluorescent lights.

MacAdam ellipses are said to have “steps” which are standard deviations in the general human population.¹ For example, 1-step ellipse means the color difference between the center of the ellipse and the edge of the ellipse can only be noticed by one standard deviation, or 68% of the population. From one edge of the ellipse to the point directly opposite is two standard deviations, or 95%. Later work, however, showed that a 3-step ellipse is a more representative assessment of color discrimination. Accordingly, major manufacturers use a 3-step ellipse for linear fluorescent lights.²

For LED color bins, ANSI C78.377, indoor lighting chromaticity chart uses 7-step MacAdam ellipses to define CCT bins. 3-step ellipses are about a quarter of the size of 7-step ellipses. LED manufacturers will sometimes split each CCT bin into further color bins. If color uniformity is important to you, make sure to choose from a single color bin.

The authors maintain that the ellipses in use today for LEDs are still too large. There are noticeable variations in color even within a single LED color bin. ASSIST (see Chapter 11) cites evidence that a four-step MacAdam ellipse is perceptually noticeable (ASSIST, 2005). This issue needs to be resolved before LED lights will be ready for widespread adoption. No one wants to buy two identical LED light bulbs and have them appear different colors. This has been an issue for CFLs. For example, some appear pinkish and some appear greenish, the difference apparently being due to small deviations above and below the Planck line. The system still isn't perfect and research is continuing. If you're building consumer lights, consider getting the tightest binning available that you can afford.

¹<http://assets.sylvania.com/assets/documents/faq0026-0999.f4172b60-cde8-4cb2-8a0a-37d3490d9e49.pdf>

²<http://www.citizen.co.jp/english/release/10/100302ecl.html>

A Note about Color Space Standards

The color space used above is based on the 1931 CIE standard. This was based on an observer with a 2° field of view. There were a variety of problems with this original research. In particular, a very small sample population was used, leading to spuriously large standard deviations (or “steps”). Furthermore, not a single observer was found to give reproducible results.

It was discovered that at a 10° field of view more of the rods were used to detect color. In 1964 a 10° field of view standard was defined. Both standards use the (x, y) coordinates, so it is important to note which color space is being used. The more common standard is still that of 1931.

Another problem was also realized about the (x, y) space. MacAdam ellipses in the green region are much larger than ellipses in the blue area. So a CIE 1976 uniform chromaticity diagram was defined using (u', v') coordinates. The result is that comparisons in chromaticity in the (u', v') space are less skewed from the bluer colors. This is called the CIELUV color space. L refers to the luminance of the colors used. A related color space is the CIELAB. The exact differences between these spaces are beyond the scope of this book. It is sufficient to know that one should be careful to know which color space an LED manufacturer is using. The software that comes with spectroradiometers will crunch all the numbers and inform you of all the coordinates.

COLOR RENDITION: HOW THE LIGHT LOOKS VERSUS HOW OBJECTS LOOK

While light sources of different color temperatures all look white, the objects they illuminate can look very different (see Figure 3.13). The light sources all look white to the eye because the brain interprets it that way. Even the limited spectrum of the fluorescent light looks white because the eye is basically a three vector sensor.



Figure 3.13 Left: Cool White Fluorescent 4100 K, CRI 60; Middle: Incandescent, 2800 K, CRI 100; Right: Reveal® Incandescent 2800 K, CRI 78.

Source: http://www.gelighting.com/eu/resources/learn_about_light/pop_color_booth.html. See color insert.

In other words, three monochromatic colors, red, green, and blue, can be mixed to arrive at any color. However, the perceived color of the object being illuminated depends on the spectral distribution of the light source. The object can reflect only the colors that are present in the light source. If a ruby of 650nm wavelength is illuminated by a fluorescent light with a spike at 615nm, then the reflected light will not be the same deep red as it will appear in daylight or incandescent light with significant power at 650nm. Light from an LED source depends on the particular phosphor used, and varies by manufacturer. But in general LED sources have a broader spectrum than fluorescents and will have better color rendition. So a richly colored brown sofa will look better under incandescent and LED lights than under even a warm white fluorescent light.

The CIE measure for color rendition is the *Color Rendition Index (CRI)*. This is the industry standard for now but, again, it is far from perfect. The reference remains the black body radiator, as colors in daylight look natural. To find the CRI, a calculated comparison between the test light source and the reference is done with eight Munsell pastel colors (see Figure 3.14, which lists the approximate Munsell). The difference on the color space is scaled by a factor to give 100 for the black body radiator and 50 for fluorescent lights. The General CRI, R_a , is the average of these eight numbers, R1-R8. This limited number of color choices limits the usefulness of CRI as a measure of color rendition. Also used are six other moderately saturated colors that represent colors common to the human experience such as flesh tone. These constitute the *Special CRI* that give additional information and are not used in the R_a calculation.

There are obvious problems with this system. It can give high CRI to a source that does not render saturated colors. For example, some LEDs have a negative R_9 ,

Name	Appr. Munsell	Appearance under daylight	Swatch
TCS01	7,5 R 6/4	Light grayish red	
TCS02	5 Y 6/4	Dark grayish yellow	
TCS03	5 GY 6/8	Strong yellow green	
TCS04	2,5 G 6/6	Moderate yellowish green	
TCS05	10 BG 6/4	Light bluish green	
TCS06	5 PB 6/8	Light blue	
TCS07	2,5 P 6/8	Light violet	
TCS08	10 P 6/8	Light reddish purple	
TCS09	4,5 R 4/13	Strong red	
TCS10	5 Y 8/10	Strong yellow	
TCS11	4,5 G 5/8	Strong green	
TCS12	3 PB 3/11	Strong blue	
TCS13	5 YR 8/4	Light yellowish pink (skin)	
TCS14	5 GY 4/4	Moderate olive green (leaf)	

Figure 3.14 Approximate Munsell Test Color Samples. Source: http://en.wikipedia.org/wiki/Color_rendering_index under license <http://creativecommons.org/licenses/by-sa/3.0/>. See color insert.

value, yet have an R_a of 80. The reason for the negative value is that the delta on the color chart is so far off that the scaling factor doesn't really work. On the other hand, CRI can give a low value to a light source that enhances contrast, which gives preferred color rendering. For example, some RGB LED systems render the pastel colors more saturated than the reference incandescent and will yield a CRI in the 60s even though it really ought to be greater than 100.

Typical fluorescent tubes for office use have CRI of 60. CFLs can be as high as 80. The better LEDs tend to be just above 80, as this seems to be the level where the color rendition starts looking natural. At 80, there will be slight color differences compared with incandescent sources, but the number is not so low as to render objects as a different color. There also exist CRI 90 LEDs where the red phosphor is increased to a longer wavelength to improve the R_9 value.

CRI and CCT do not give adequate guidance about how objects will look when illuminated. Although work is ongoing, it seems to us unlikely that any small set of numbers will ever be found that can do the job. As an example, look back at Figure 3.13, which shows the same scene with various CCTs and CRIs. It's not just that one looks worse than the other. Rather, the reds seem subdued on one; one seems brighter than the others; one "pops" more than the other, and so on.

The Human Factor

There are many other factors involved in human perception of light and color. Color perception in the eye is the job of *cone receptors*. The cones are heavily concentrated in the center 1.5° of the retina called the *fovea*. This is one of the reasons why color perception changes when viewing a small sample of color versus viewing, for example, the color of the wall. On the other hand, *rods* sense only light intensity. Rods are everywhere except the fovea. In the center 1° of the eye, there are no rods. This is why when observing the night sky, where scotopic vision kicks in, there is a blind spot right in the middle of the retina and one must look slightly off center to see a faint star.

There is some variation in the human population in color perception. The physiological cause for this difference is a yellow spot over the center 4° of the retina. This spot can vary slightly in individuals and affect the color of the light reaching the retina. The cornea and the layers of cells in front of the retina can also vary in individuals and as they age. There may also be genetic differences in color perception between men and women (Jameson et al. 2001).

Lighting designers need to realize that there is a human factor in perceived lighting as well. The amount of light perceived by humans is *not* the same as the amount of light measured by a meter. The difference between the two is caused by eye dilation. As it gets darker, your pupil dilates, admitting more light. This causes the perceived light to be higher than the measured light. A typical conversion between the two is that the perceived light is the square root of the measured light. As an example, suppose that the measured light is 10% of full brightness. The perceived brightness is then approximately $\sqrt{0.1} = 0.32 = 32\%$. This factor is important

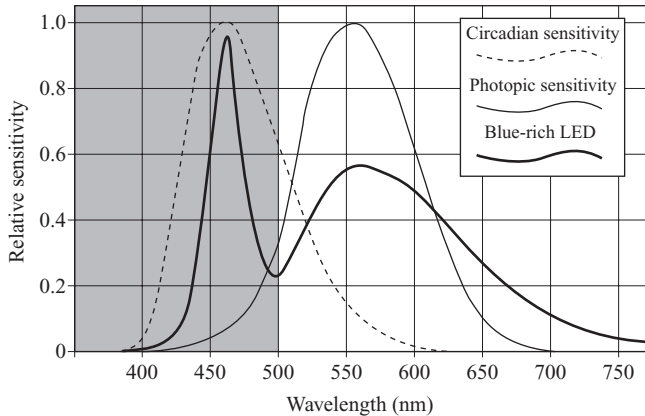


Figure 3.15 Circadian Rhythm Sensitivity. Source: “Visibility, Environmental and Astronomical Issues Associated with Blue-Rich White Outdoor Lighting,” May 2010, IDA. Image copyright of IDA.

in dimming lights. You need to have a wider dimming range than is suggested just by measuring lumens.

You should be aware of the possibility that the type of light produced by current LEDs may have some adverse effects on humans. The International Dark-Sky Association released a statement³ that “Unfortunately, bluish light produces high levels of light pollution with significant environmental impact. These lights are known to increase glare and compromise human vision, especially in the aging eye. Short wavelength light also increases sky glow disproportionately. In addition, blue light has a greater tendency to affect living organisms through disruption of their biological processes that rely upon natural cycles of daylight and darkness, such as the circadian rhythm. ... Circadian rhythms are controlled by light emitted within the dashed curve [see Figure 3.15]. The color of light emitted by a typical bluish-white 5500 K LED is depicted by the bold line. A large portion of light emitted by this light source falls outside of the human photopic vision range, and falls within the circadian rhythm curve.”

On the other hand there may also be benefits to the use of blue light as present in LEDs. Some studies have suggested that low levels of blue light can be effective at treating mood swings or low energy levels in some people. Red LED lights may be effective against skin aging. And the fact that scotopic vision has greatly higher sensitivity than photopic may mean that the light intensity required for street lighting may be lower with bluer light than might otherwise be supposed.

Color perception is also a psychological phenomenon (see Figure 3.16). A gray color will look darker on a light gray background and look lighter on a dark gray background. The brain tries to compensate for the lighting condition. The light gray background signals to the brain that the illumination is very bright. Therefore

³“Blue light threatens animals and people”, <http://www.ledsmagazine.com/press/20192>. But also see http://docs.darksky.org/Reports/IDA_Blue-Rich_Light_White_Paper051710.pdf and <http://docs.darksky.org/SB/LED-SB-v3i1.pdf>.



Figure 3.16 Identical Gray Boxes Look Different Depending on Their Background.

the gray color is actually darker than it would appear and so the brain interprets the gray color to be darker. In Figure 3.16, if you stare at the gray box on the dark background, you will train your brain to think that the lighting is dark. Over a period of a couple of seconds, you will notice the center gray square becoming brighter and seem to pop out more. This effect also happens with different colors. It is almost impossible to identify the real color of a sample when it is surrounded by different colors.

Chapter 4

Practical Characteristics of LEDs

CURRENT, NOT VOLTAGE

The first thing to know about LEDs—and all diodes—is that they are current devices, not voltage devices. What does this mean? With a resistor, voltage and current are proportional. If you place a voltage across a resistor, a certain current will flow. By Ohm's law ($V = IR$, or volts equal ohms multiplied by amps), this current will be the voltage divided by the resistance, $I = V/R$. If you double the voltage, the current will double. Conversely, if you push a current through a resistor, it will develop a voltage across it. Again, by Ohm's law, this voltage will be the current times the resistance, $V = I * R$.

But with diodes, Ohm's law doesn't apply. Voltage and current are not proportional. In fact, they are more like exponentially related. Specifically, the current through a diode (to first approximation) can be modeled by

$$I(V_f) = I_0 e^{k * V_f}$$

while its inverse models the voltage through it

$$V_f(I) = \frac{1}{k} * \ln\left(\frac{I}{I_0}\right)$$

In these equations, V_f is the voltage from anode to cathode of the diode, I is the current through the diode, k sets the scale for the voltage, and I_0 sets the scale for the current. k and I_0 are constants. As typical values for white LEDs (at room temperature), we will take, $I_0 = 3.2 \mu A$ and $k = 3.64/V$. These values are those for actual commercially available devices we will estimate looking at computer modeling of LEDs in Chapter 11.

Let's look at the equation for current for a moment. The thing to note is that the current through the diode is a very strong function of the voltage. Using the typical values for white LEDs as above, in order to put 2.80V across the diode we

Practical Lighting Design With LEDs, First Edition. Ron Lenk, Carol Lenk.

© 2011 the Institute of Electrical and Electronics Engineers, Inc. Published 2011 by John Wiley & Sons, Inc.

have to put 85 mA of current through it. Adding just a quarter volt, to 3.05 V, more than doubles the current, to 212 mA. And another quarter volt, to 3.30 V, again more than doubles the current, to 527 mA.

What this means is that it isn't reasonable to try to control the diode's operation by controlling its voltage. Let's look at it the other way, with the second equation. When you put 100 mA of current through the diode, the forward voltage is 2.843 V. Doubling the current by adding 100 mA more increases the voltage by only 191 mV, less than 7%.

It is in this sense that we say that a diode, and an LED in particular, is a current device and not a voltage device. To a practical engineering first approximation, the forward voltage is always the same, regardless of how much current you put through it, so the performance of the device is determined by how much current you put through it.

This also means that, to the same approximation, the power into the diode is determined by the current. Since power is voltage times current, and voltage is constant, power is just proportional to current. As we'll see in the chapters on power supplies for LEDs, this is the way they are actually driven. The optical output of an LED is specified in lumens/watt. After deciding how many lumens are needed, this tells you how many watts are needed, and then this tells you how much current is needed. Power supplies for LEDs are typically designed to drive them with a constant current.

FORWARD VOLTAGE

For easy estimates, the forward voltage of a diode is a constant. But the forward voltage does vary somewhat with current. And each diode is different, so where do the numbers come from? How can different diodes be compared?

Let's talk first about ordinary silicon diodes used as rectifiers. They have typical V_f s ranging from 500 mV for small-signal diodes to 1.2 V for large rectifiers (Schottky diodes are a different type of device). These numbers are a sort of practical guide to the performance of the diode. For a given amount of current, the diode with a lower forward voltage will dissipate less power. (Why wouldn't you always use the diode with the lowest forward voltage? Lower forward voltage usually means a physically larger device.)

Suppose you want to have a 500 mA current through your diode. If power dissipation during conduction is the most important thing, you could go to a database and assemble a list of all the 1 A diodes and see which has the lowest forward voltage at 500 mA. Easier at this level of approximation is just to look at forward voltage at the rated current of 1 A—a number that is always right on the front page of the datasheet. And indeed, this is how manufacturers and distributors have their online databases set up: Forward voltage at rated current.

Now of course, nothing analog is really that simple. Forward voltage also depends on the temperature of the die, and this depends on how big the package is. The same diode in a bigger package will stay cooler, and thus have a higher forward voltage. Conversely, you can get a somewhat lower forward voltage by moving to

a diode rated at a higher current. For example, a 3A diode rated at the same forward voltage as a 1A diode will have a lower forward voltage when operated at 1A. But for a first cut, a simple spreadsheet lineup of V_f will do.

The same principles apply to LEDs. For white and blue LEDs, typical forward voltages range from about 3.1V to 3.8V. Yellow LEDs are somewhat higher, and red LEDs are down around 2.2V. It's actually a little easier comparing forward voltages of LEDs than ordinary diodes. It's an industry standard to report the V_f at a current of 350mA, even for devices capable of carrying 1A.

As an aside, note that the underlying reason that LEDs have much higher forward voltages than silicon diodes is that they aren't made of silicon. Their bandgap is different than silicon, that's how they generate light. And the reason there's such a range of voltages for white LEDs is that there's a number of different semiconductors being used in the industry now. Each has its own bandgap. One of the major areas of research for all LED die manufacturers is how to reduce the forward voltage of the device. Reducing this would increase the light output of the LED per watt.

REVERSE BREAKDOWN

All diodes will conduct current when a voltage is applied from anode to cathode. All will also conduct if enough voltage is applied from cathode to anode—whether intentionally or not. Zener and avalanche diodes fall into the intentional category; this type of conduction is their main operation mode. You select one of these diodes based on what voltage you want conduction to occur at. The manufacturers hold the tolerance on this voltage to 5%, or even better.

Rectifier diodes and LEDs fall into the unintentional category. If they conduct in the reverse direction, there's an excellent chance that they've broken. Now with rectifier diodes, there's an easy solution. The reverse breakdown voltage is one of the parameters you select for when choosing the diode. 200V, 400V and 600V are all commonly and inexpensively available, usually with very similar other characteristics. So if your diode breaks down, you just replace it with a higher voltage part.

LEDs are not so convenient, unfortunately. Most of them have a reverse breakdown voltage of only 5V. This very low voltage can present a serious problem in a practical circuit. Even with several LEDs in series, the breakdown is still only tens of volts. There isn't any room for error. Any glitch or noise in the control loop of the power supply, or a hiccup in the AC line voltage, may be enough to momentarily generate breakdown. It doesn't take long—diodes can be broken in microseconds under the right conditions.

One obvious solution to try is to put a regular rectifier diode in series with the LEDs, as in Figure 4.1. This prevents any reverse current from flowing through the LEDs. Unfortunately, it also dissipates power when the LEDs are operating normally. If the LEDs are being driven at 400mA, the forward voltage of the blocking diode of about 1V will dissipate about 400mW of power all the time. This may be unacceptably high, either for thermal or efficiency reasons. Thermally, it may be




Figure 4.1 Reverse Bias Protection.




Figure 4.2 LEDs with Reverse Bias Protection.

difficult to get rid of the excess heat of 400mW that isn't doing anything for you most of the time. For efficiency on small light sources, 400mW may be a significant fraction of the total power into the device.

The other choice is to put a diode in anti-parallel with the LEDs, as in Figure 4.2. Now if there should happen to be a voltage applied from cathode to anode of the LEDs, it is clamped to the forward voltage of the diode, about one volt. This is entirely safe for the LEDs. It also doesn't dissipate any power in normal operation, as the diode is reverse blocking when the LEDs are conducting normally.

As long as we're on the topic of LED current, it's worth a quick look at ratings. Being semiconductors, all LEDs have an absolute maximum current rating from the manufacturer. These are typically 700mA or 1A, but there are many others. With a normal diode, you know that the current rating is really for steady-state operation. For a pulse, you can usually have much more current. For example, the 1N4973 is a 1A diode, but it can handle a 30A surge for half a line cycle, 8.3msec. So what about LEDs? The answer is different for each manufacturer and for each type of LED.

Some white LEDs can be driven above their rated current by a large amount. Their phosphors saturate, and the light starts to turn blue, but if this doesn't bother you, they can take quite a bit more current. In general, current generation LEDs seem to have plenty of bond wires to carry a lot of current. So the real limitation is whether the die can take more current without burning up or developing hot spots. This is primarily a thermal question until you get to very high currents. And as we'll see in the next section, very high currents are bad for efficacy, and so you probably don't want them anyway. But a little bit of overdrive is probably okay.

NOT EFFICIENCY—EFFICACY!

Efficiency usually refers to electrical conversion. You put a certain amount of energy in to a power supply; you get a certain less amount out. The electrical efficiency is defined to be the output power divided by the input power. It is usually expressed as a percentage. Since there are always some losses, it is a number always less than 100 and not less than zero.

Things are not so straightforward for LEDs. You put power in, but what you want out is not power, but light. Now light, just like electricity, *can* be measured in watts. And indeed, deep blue LEDs' light output is measured in watts. Efficiency is then straightforward. It is the output power in watts divided by the input power in watts. As an example, the LXML-PR01-0275 from Lumileds is a royal-blue LED (455 nm). It has an optical output power of 275 mW, and an input power of $350 \text{ mA} * 3.15 \text{ V} = 1.103 \text{ W}$. Its efficiency is thus $0.275 \text{ W}/1.103 \text{ W} = 24.9\%$.

As discussed in Chapter 3, humans have very little vision at 455 nm. Nobody would try to use such a light to see by. It therefore makes sense to specify its output in lumens. Other colors, including white, are different. For example, for traffic lights we will want to know how bright red, yellow, and green appear to people. Similarly, for light bulbs using white LEDs, we want to know how bright people perceive them to be, not how much optical power is coming out. So for these nonblue LEDs, watts aren't used, lumens are. And so efficiency isn't the right term either. Instead, we use efficacy. (Although the authors don't know this for a fact, it is worth observing that the lumens output of a 455 nm LED is very low, because humans don't perceive that color well. As a result, the lumens output of such a device would be very low. Perhaps this would look bad on a datasheet compared with devices emitting light that people see?)

White and other LEDs' efficacy is measured in lumens per watt (Lm/W). Lumens are used because this reflects the amount of light that humans perceive, as described in Chapter 3. Now this has some strange effects. To start, the human eye response peaks in the green part of the visible spectrum. This means, for example, that for a red and a green LED with the same efficacy, the red must be emitting a lot more light than the green. You thus can't directly compare two different LEDs based on efficacy unless they are the same color. Even two quite similar colors—say greenish yellow and yellowish green—aren't directly comparable.

Additionally, LEDs don't emit exactly one wavelength. They're not lasers. For example, Lumileds's LXML-PM01-0040, a green LED, has the dominant wavelength specified to be between 520 nm and 550 nm. Additionally, the spectral half-width is 30 nm. So the light is somewhere between 490 nm and 580 nm. Again, human eye response peaks at 555 nm, and so the efficacy of these devices will vary widely, even when in terms of optical power they are similar.

When you get to mixtures of colors, the confusion is even worse. And white light is by definition a mixture of all the visible colors. As discussed in Chapter 1, white LEDs are made by adding two or three different colors together. So the efficacy of such a device depends very strongly on the exact mixture of colors, on the exact

brightness of each color in the mix, and on exactly where in the spectrum each color is emitting.

Finally, white LED vendors consider efficacy to be the parameter on which they are most significantly competing. So, they tend to play games with the spectral composition of white LEDs. Shifting closer to green will increase the efficacy, so shifting the CCT from 2750 K to 2900 K produces a better number.

Efficacy as specified in datasheets is measured at nominal current, typically 350 mA. As you change the current, the efficacy also changes. Now when you decrease the current, there isn't a huge change, maybe a 10% increase if you go to 100 mA. So, practically, you can use the datasheet number. But if you increase the current, the efficacy starts to decrease substantially. So if you're concerned with efficacy, beware. Running an LED at more than its nominal current does produce more than nominal light, but the efficacy suffers.

Let's take a look at how this works out practically. Lumileds' LXML-PWW1-0060 has a guaranteed minimum output of 60 lumens at 350 mA (at 25°C—we'll delay talking about temperature effects till Chapter 5.) To know the efficacy, however, you need to also know the power level, not just the current. We find in the datasheet that the nominal forward voltage of the device at 350 mA (and—again—at 25°C) is 3.15 V. Thus its nominal efficacy is 60 lumens / (350 mA * 3.15 V) = 54.4 Lm/W. Note that the efficacy is *not* 60 Lm/W—some sales people have been known to confuse lumens and lumens per watt.

Next let's figure out how the efficacy changes with drive current. The temptation is to look at the curve (Figure 4.3) "Relative Luminous Flux" and say, "The light output increases this amount, and therefore so does the efficacy." But that's incorrect on two scores. The efficacy depends not just on the light output, but also on the drive current and the forward voltage. And while the flux is increasing, the drive

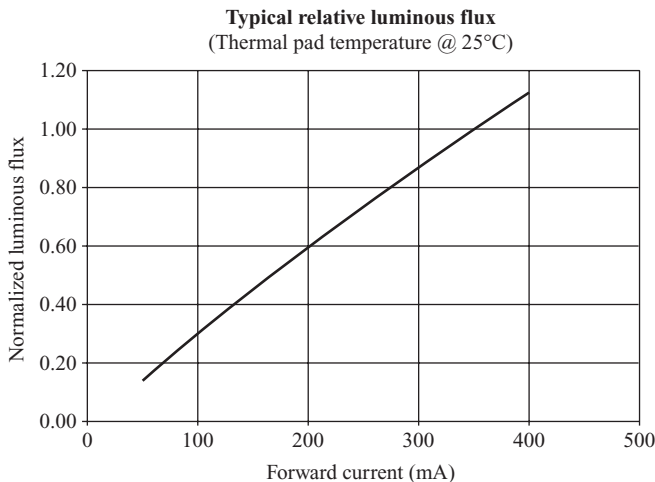


Figure 4.3 Light Output as a Function of Current.

Source: Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Company 2007.

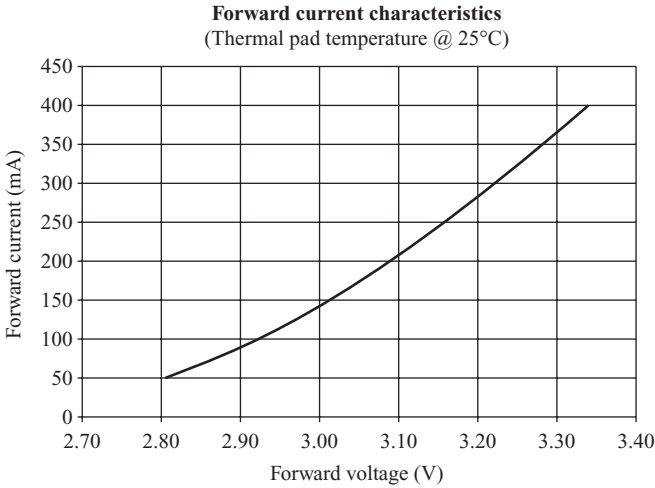


Figure 4.4 Forward Voltage as a Function of Current.

Source: Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Company 2007.




Figure 4.5 Efficacy versus Drive Current.

current is also increasing, and so is the forward voltage (Figure 4.4). The efficacy as a function of drive current is not shown in the datasheet!

We use the data in the two figures provided by the datasheet to construct the curve of efficacy versus drive current in Figure 4.5. You should note that above the nominal drive current, the efficacy of the LED starts to decrease quite a bit. To get a feel for the effect of this, suppose you want to get 100 lumens. At 350 mA, you need $100\text{Lm}/(54.4\text{Lm/W}) = 1.84\text{W}$. You could do this by running two LEDs both at something close to 350 mA. Suppose instead you do this with a single device. Since you need something more like 700 mA, the efficacy is down to 42Lm/W —at which you need $100\text{Lm}/(42\text{Lm/W}) = 2.38\text{W}$. You're spending considerably more power (29% more) to get the same light output, and perhaps falling below specification on required efficacy of the light source. Of course on the other hand,

you're paying less money for one LED instead of two. It's one of those trade-offs engineers always face. We'll see more of this in Chapter 9.

LED OPTICAL SPECTRA

Since the goal of an LED is to produce light, its optical spectrum is a key performance parameter. Let's take a look at a typical spectral curve for a warm-white LED, shown in Figure 4.6. Along the x-axis is wavelength in nm and along the y-axis is relative spectral power. Note that the manufacturer doesn't label units on the y-axis: All you can say are statements such as "the output at 675 nm is half that at 600 nm."

A few things stand out. While the bulk of the curve resembles the human eye response curve, one part is quite different. There's a big spike in the blue region. The origin of the spike is obvious. This is the native emission of the die, and the spike represents the portion of that emission that is not being converted by the phosphors. This doesn't really hurt the color of the LED, since humans see very poorly in blue. But it does represent one of the areas in which LED performance can be improved. Converting all of the blue light to the more useful part of the spectrum would increase efficacy.

The other noticeable thing about the curve in Figure 4.6 is how quickly it tails off in the red. Getting a little more red up there is the difference between 2800K and 3500K CCT. Even more importantly, it's also the difference between reddish objects looking brown versus looking red. So this part of the curve gives some idea of how reddish red objects will look. Unfortunately, making deep red phosphors is currently the most challenging task for LED manufacturers. This is why many present-day white LEDs have poor CRI. They drop out in the reds (particularly R9; see Figure 4.7).




Figure 4.6 Light Output as a Function of Wavelength.

Source: Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Co., 2007.




Figure 4.7 Many LEDs Have Poor R9.

Source: http://www.yegopto.co.uk/LightingLEDs/CRI_Seoul_Semi.

The optical spectrum of the LED tells you everything about its light. But there's no practical way to look at the spectrum and decide if the light is a reasonable color for a given application. So as has been discussed in Chapter 3, what people have done is to make up a set of numbers that characterize important features of the spectrum. These are numbers such as CCT, CRI and (x, y) . As discussed, these have pluses and minuses. The CCT gives some idea of how "cold" the light is, at least for white light. The CRI gives some idea of how well colors are reproduced by the light—to the extent that the spectrum is blackbody. The various characterizing numbers are useful because they give *some* idea about what the light looks like. But the real reason they're useful is because they are what LED manufacturers can characterize. Ultimately, there's no substitute for building your design with the actual LEDs and trying them out.

Despite the limitations of the characterizing numbers, we can discuss how they vary. Variation should not be a surprise. If you drive too much current into a white LED, it turns blue. The phosphor has saturated, and the blue light of the LED comes through without being converted. In general, phosphors are complicated molecules, and how much and how efficiently they capture and convert blue light must depend to some extent on things such as drive current. RGB LEDs are even more susceptible to variation with current, since their efficacy varies with current.

The color coordinates of white LEDs vary with drive current. Look at Figure 4.8, from an Intematix datasheet. As drive current increases, both x and y increase (400mA is maximum rated drive current for this device). This variation, while not huge, is enough to change the appearance of the LEDs. And even with this information, we keep noting that the x and y coordinates alone aren't enough to specify what

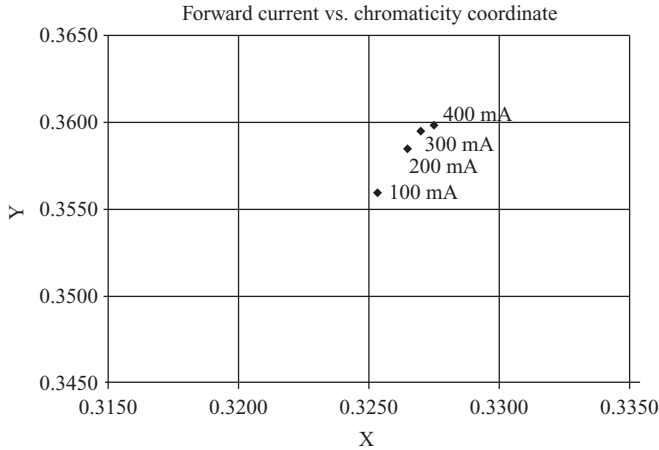


Figure 4.8 (x, y) as a Function of Current.

Source: C6060-16014-CW/NW datasheet, Intematix Technology Center Corp., 3/2008.

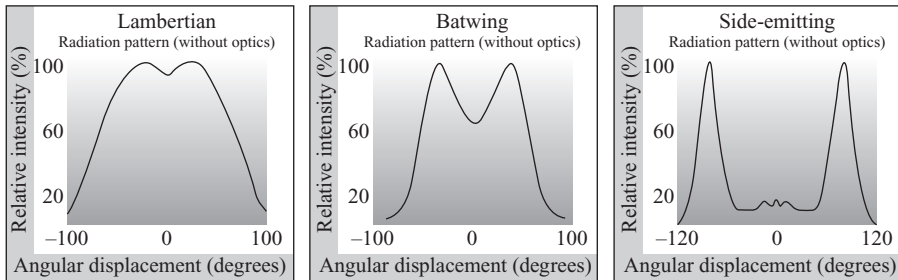


Figure 4.9 Different Output Light Distributions Are Available.

Source: <http://www.philipslumileds.com/technology/radiationpatterns.cfm>.

it’s going to look like. Most manufacturers don’t specify even this amount of information, so if you plan to run the LEDs at anything other than their nominal current, you have to build and measure the system to find out what the light will look like.

One more optical characteristic to consider is angle of emission. The light doesn’t come out the back of the LED, because it’s mounted on a package there. In some LEDs, the light is emitted as you would expect, from the front of the die. But other lights are emitted perpendicular to the die. Further, the light emitted from the LED die usually goes through some optics in the package, which further modifies the light distribution. The result is that the light output is not the same intensity at all angles; there is more light in some directions than others.

There are three common light emission distributions for LEDs: Lambertian, batwing, and side-emitting. Lambertian is the most common. Looking at Figure 4.9, the fundamental difference is clear. Lambertian is primarily forward-directed light. Batwing is also forward-directed, but has an annulus of light, with a dip in the middle (remember that the diagrams need to be rotated out of the plane of the paper). Finally,

as the name indicates, side-emitting LEDs emit most of the light out of the sides, and very little is forward-directed.

Most manufacturers specify a radiation pattern, without being too specific about what this exactly means. But there are at least two important measures of angular light emission. Lumileds specifies both a viewing angle and an included angle. The included angle is probably the more important of the two, and is probably what other manufacturers mean when they specify an emission angle. The included angle is the (circular) angle in which 90% of the light is emitted. What this means practically is that if the angle from the LED doesn't run into anything up to this angle, then 90% of the light emitted by the LED will actually escape from the device. Lumileds defines this to be "total angle," so that their 160° specification corresponds to $\pm 80^\circ$. For other manufacturers 120° is more common, and this presumably is also total angle.

Viewing angle is here the angle at which the brightness is half what it is at the maximum. (This is what was called "beam angle" in Chapter 3.) For most applications, this number is less important. Note the fundamental difference between the two angles. Viewing angle gives you the same information as the emission curve, the angle at which relative intensity is half. Included angle is how much of the total curve is included—the integral of the area under the curve. If the angle from which you extract light is anything smaller than the included angle, you have to look at the emission curve to figure out how much light you're extracting. If it's bigger, then you know you're getting more than 90%.

OVERDRIVING LEDs

We just said that driving LEDs at currents higher than their specification results in color shifts. And of course it also means the LEDs get hotter from the additional current, plus a small additional factor from the forward voltage rising with additional current. A hotter LED has a shorter lifetime. Given all these negatives, why would you want to overdrive an LED?

The main driver is usually cost (though size can be a factor, too). Five LEDs cost 25% more than four LEDs. If instead I can drive each of the four LEDs 25% harder, I save the cost of that additional LED. So what should you do? Color shifts are usually not that important for today's market, so it still comes down to LED temperature. If you can hold the appropriate die or case temperature down to a number that gives you adequate lifetime, go ahead and do it. We think this is the important factor as long as the absolute maximum current rating of the device isn't exceeded.

And if that extra 25% would exceed the absolute maximum, then you have to consult with the vendor. Of course they will tell you that you can't do it, that there is no guaranteed performance beyond that rating, and that your system will fail horribly. This is certainly possible, but some digging may reveal which part of the LED is the first to go. The bond wires? Die hot spots? Phosphor saturation? There's failing such that the part no longer works, but there is also failing where the part no longer

exactly meets all of its specifications. Finding out which one the absolute maximum current rating is based on may allow you some leeway to exceed even this rating, at least by a bit.

KEY DATASHEET PARAMETERS

The previous sections have talked about quite a number of parameters for LEDs, as well as their temperature variations. It would be nice if LEDs had just a single parameter to characterize their suitability for an application. That way you could just say, “This LED is a #8, and I need a #9.” Unfortunately, LEDs are deeply embedded in the analog world. There are half a dozen parameters that are important, and no one device is going to exactly meet all of your criteria.

Let’s try to summarize the most important criteria in a practical sense. The first thing to do is to select the CCT (or color, if you’re not building a white light source). This usually eliminates 2/3 of the choices. CRI is less important, because most devices are about the same unless you want an unusually high CRI.

After CCT, the next criterion is usually the efficacy at operating temperature. We have to re-emphasize that this should not be at the 25°C manufacturers specify, because different devices’ optical outputs change differently with temperature. Take the light output at 25°C, multiply this by the decrease in light at 85°C (usually a good starting point), and divide by the current specified and by the forward voltage, again compensated for temperature:

$$\eta(85^{\circ}\text{C}) = \frac{\text{Lumens}(25^{\circ}\text{C}) * \text{RelativeLight}(85^{\circ}\text{C}/25^{\circ}\text{C})}{\text{Current} * V_f * \text{Relative}V_f(85^{\circ}\text{C}/25^{\circ}\text{C})}$$

You want to select a device with a minimum efficacy given by how much light you want to get out and how much heat you can dissipate. Usually, but not always, going to higher efficacy devices beyond this increases cost.

BINNING

Realizing how much the variation in some of these parameters influences performance, manufacturers offer binning. Binning means that the manufacturer offers the same part with differences in the key parameters. The buyer selects which bin the parameters should be in. This allows some degree of control over the parameters. The parameters are otherwise usually fairly broadly defined. Part of the reason for this is that it’s hard to control the parameters over different production runs; the vendors want to make sure they sell everything they make.

As an example, all manufacturers offer CCT binning. If you look at “neutral-white” on Lumileds’ datasheet, it means the CCT is between 3500 K and 4500 K. But this is such a broad range that the parts actually available are binned. Figure 4.10 shows how this neutral-white space is binned. There are limits on both x and y , resulting in 12 color bins.




Figure 4.10 Neutral-White Bin Structure.

Source: “Technical Datasheet DS56, Power Light Source Luxeon Rebel,” Philips Lumileds Lighting Company, 2007.

There are more bins offered above the Planck line than below. Four of the bins straddle the Planck line; half are above and half below. This means that LEDs from the same bin may produce some greenish and some reddish white light. It should be noted again that this is an ANSI standard (ANSI C78.377), and so the problem is not with Lumileds: All the vendors are specifying bins this way.

Note that even within a single bin there can be variation in the color of the bin, because of course a single (x, y) coordinate cannot be produced. The key question is whether this unit to unit variation is perceptible by customers. The generic answer you get from manufacturers is “no,” within a single bin you can’t tell the difference in color. However, the authors believe that variation is noticeable, at least to some portions of the population. Not everyone sees color the same, and some are more sensitive to variation than others (see Chapter 3). This viewpoint is backed up by the fact that some manufacturers are going to substantially smaller bin sizes. In any case, this is something for the marketing department to work on. Since there are differences in color perception among humans, be sure you are not the sole judge of whether the binning is good enough!

We note that there’s some trickery involved with this color binning. Certain manufacturers specify all of these bins, but then you can’t buy just one of them. For example, they may require you to accept any one of four or even six adjacent color bins. Every time you get a shipment, its bin will be labeled. But you may get stuck with bins you can’t use. Make sure you talk with your vendor about what bins they actually ship, versus what the datasheet says.

Other binned parameters may include forward voltage and brightness. Brightness is always binned, because this is what drives the price. Brighter lights cost more. But it may turn out to be more cost-effective to buy a lower brightness bin, while paying for a lower forward voltage bin. If what matters is efficacy, both the light output and the power input matter. Again, you have to talk with the vendor to understand how the pricing structure is set up.

THE TOLERANCE GAME

You've selected a part based on consideration of its luminous output and color (and of course cost). Its forward voltage is whatever it is. Now you look at the datasheet to find out minimum efficacy at the current you're going to drive it, and figure out how many of these LEDs you need. (We'll have detailed examples of how you actually do this in Chapter 10.) Everything's ready to go, right? Wrong! Hidden inside the datasheet is a spec that can kill your calculation. It's hidden away as a footnote and is easy to miss.

This hidden menace is the measurement tolerance on the light output. Of course all parts have a tolerance on their light output, or on their efficacy. You've already accounted for this by calculating light output from your device at minimum guaranteed LED light output. The problem is rather that the manufacturers don't really guarantee a minimum light output. They guarantee a minimum light output, *given their measurement accuracy*. This accuracy is usually specified as $\pm 7\%$, or $\pm 10\%$. So a datasheet that says the LED has an output of 90–100 lumens could have an output anywhere from 81 to 110 lumens.

Now by an interesting coincidence, whenever the authors have measured lumens from various manufacturers' LEDs using their own calibrated equipment, the parts are always at the bottom of their bin and at the bottom of their measurement tolerance! So the part that you might think was supposed to be a typical 95 lumens and a minimum of 90 lumens is almost always 81 lumens. Now if this were really a tolerance problem in their optical measurement, you would expect a distribution of brightness. Since number of lumens is an important factor in marketing, you have to wonder about truth in advertising. In any case, if you need to guarantee the brightness of your device, you need to account for the measurement accuracy as well.

Chapter 5

Practical Thermal Performance of LEDs

In the previous chapter, we looked at a variety of important parameters of LEDs, both electrical and optical. All of these parameters are affected by temperature and by temperature-related aging. And the dependencies need some detailed explanations in order to be useful. For that reason, we've separated out discussion of thermal performance of LEDs into this chapter. Measuring the temperature of the LED is yet another concern. While there is some information on that in this chapter, you should look at Chapter 11 for a detailed discussion of experimental technique for this tricky measurement.

MECHANISMS BEHIND THERMAL SHIFTS

You'll recall that a white LED consists of a number of different components. There's the die itself emitting blue light. There's a phosphor that converts part of the blue light into white light. There's the silicone encapsulant to protect the die and the phosphor mechanically. And then there's the package into which the whole thing is mounted. All of these components potentially contribute to thermal performance of the LED. And they each contribute differently again to thermal aging.

Let's start with the die. Since there are many different materials used for dice, some proprietary, manufacturers haven't revealed information about their thermal performance. (Thermal performance specified in datasheets is for the LED as a whole, not the individual components of the LED.) But we can take some educated guesses based on what we know about silicon. There is certainly an absolute maximum temperature beyond which the die simply fails. There is probably also a thermal run-away temperature, at which certain parts of the die become hotter than other parts (hot spots) and fail, causing the whole die to fail. Finally, it seems likely that the wavelength of the emitted blue light shifts with thermal aging. The shift is probably fairly small, but then the phosphors absorb light in a very tight band of wavelengths. A small shift in the emission wavelength may be enough to

significantly affect the ability of the phosphor to absorb the light, resulting in a decrease of efficacy with thermal aging.

The phosphors are an obvious source of temperature effects. They are tuned to absorb at a certain wavelength, and this obviously must vary with thermal excitation of the molecules. They also degrade with temperature, red-emitting phosphors being notorious in this regard. Fixing this is a major research effort for phosphor companies.

The encapsulant is supposed to be mechanically strong, and at the same time optically clear. The latter is the problem. At first, epoxy was used. It has a number of desirable properties, including being inexpensive. But it turns yellow with age (especially fast-cure types). It was then replaced with silicone, now universally used for high-brightness LEDs. Silicone is better than epoxy. But it is still a polymer, and it too eventually turns yellowish with heat and time. This affects the color of the light being emitted, and reduces the efficacy of the LED (yellowish appearance means that other colors are being absorbed). Finding better formulations is a major effort for encapsulant companies.

In silicon, the package isn't usually a source of problems. But for LEDs, it can be. If the package is not perfectly reflective (and nothing is), then it is partially absorbing the light, reducing efficacy. And the package also has to survive high temperature for tens of thousands of hours. In some cases, the package gradually yellows, again affecting absorption and color emission. Cree (2010) says "...the primary mode of degradation for HB LEDs is the package itself." Here again is an engineering tradeoff, cost of the package versus thermal aging performance.

ELECTRICAL BEHAVIOR OF LEDs WITH TEMPERATURE

The most immediately noticeable performance change of LEDs with temperature is that their forward voltage drops. This is a die phenomenon, and so varies from manufacturer to manufacturer. Datasheets specify the typical decrease of forward voltage, which runs between $-2\text{ mV}/^\circ\text{C}$ and $-4\text{ mV}/^\circ\text{C}$, but they generally don't show a curve of V_f versus temperature.

What does this mean practically? If you run the LED at constant current, as the device warms up the V_f goes down. Thus the power into the LED also goes down with temperature. And this in turn means that the light output also goes down. A back-of-the-envelope calculation gives a feel for what this means. Suppose the 25°C forward voltage is 3.6 V , and suppose to make the calculation easy that the dV_f/dT is $-3.6\text{ mV}/^\circ\text{C}$. If the temperature goes up 100°C to 125°C , then the forward voltage drops by $-3.6\text{ mV}/^\circ\text{C} * 100^\circ\text{C} = -360\text{ mV}$, to 3.24 V . Since this drop is 10% of the room temperature forward voltage, the power into the device is also reduced by 10%. Even if the efficacy of the LED were unaffected by power—which isn't true—the light output would be down 10% at temperature just because the input power is down.

Now, a way to ensure that the light output is not affected by this V_f drop is to run the LED at constant power. We could measure the voltage across the LED, and increase the current as the voltage dropped to maintain constant power. Suppose we were putting 700 mA into the LED at room temperature, a power of $3.6\text{ V} * 0.7\text{ A} = 2.52\text{ W}$. At 125°C , the voltage has dropped to 3.24 V, so we need to increase the current to $2.52\text{ W}/3.24\text{ V} = 778\text{ mA}$, an 11% increase.

That the datasheet shows just a typical dV_f/dT might make us suspicious. Is there actually a nonlinear curve of V_f versus temperature? We've done a few measurements that suggest that the curve is more or less linear. Indeed, measuring the forward voltage is one way of measuring the temperature of the die. But experience with silicon suggests that this may be a good approximation over some range of common operating temperatures, and may vary beyond those limits. Testing the actual performance of the specific LED you intend to use at its actual operating temperature is the only way to be sure of the temperature's effect on the LED.

OPTICAL BEHAVIOR OF LEDs WITH TEMPERATURE

Thermal effects on electrical performance of LEDs depend only on the die. Optical effects involve all the components of the LED, and so are more numerous. To start with the most commonly discussed effect, the brightness and efficacy of an LED decrease with increasing temperature. Now if you look at a typical datasheet, what is shown is a curve of (relative) brightness versus temperature at a constant current (see Figure 5.1). Since constant current is the common way to drive LEDs, this is a

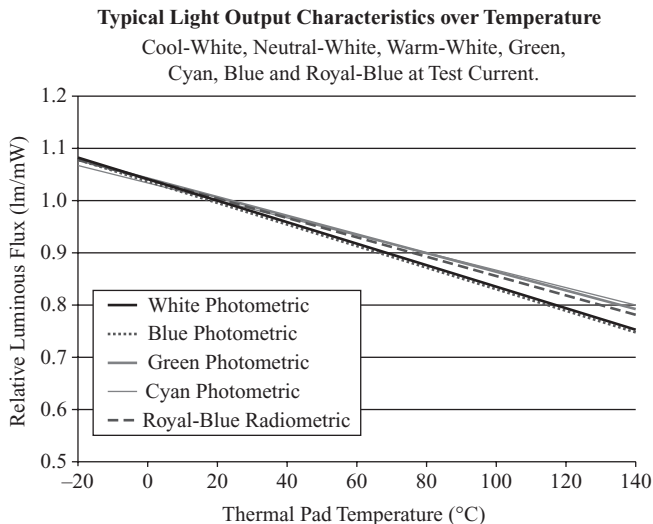


Figure 5.1 Brightness as a Function of Temperature.

Source: "Technical Datasheet DS56, Power Light Source Luxeon Rebel," Philips Lumileds Lighting Company, 2007.

useful curve. But the reality is that this mixes together two effects. At temperature, the efficacy decreases, but the constant current means that the input power is also decreasing. So the datasheet shows you the brightness decrease, *not* the efficacy decrease.

Let's find the efficacy as a function of temperature. The graph shows that from 25°C to 125°C at constant current the light drops from 1.00 to 0.78. (Note that this is flux, not efficacy. The "lm/mW" on the y-axis means "lumens or mW," and is because "royal blue" is measured in watts, and all the other colors are measured in lumens.) Now this is a 100°C delta. The datasheet also shows that the $dV_f/dT = -3 \text{ mV}/^\circ\text{C}$, so the forward voltage drops from a typical value of 3.15 V to 2.85 V. This means the input power drops from $3.15 \text{ V} * 350 \text{ mA} = 1.103 \text{ W}$ to $2.85 \text{ V} * 350 \text{ mA} = 0.998 \text{ W}$. This is a drop to 90% of input power. Efficacy thus has dropped from 1.00 to $0.78/0.90 = 0.86$, or, in other words, efficacy drops 14%. The other 8% of the light drop is, again, due to the constant current drive circuit producing less power for the LED at high temperature.

There are other effects of temperature on the light beyond the efficacy drop. The color of the light changes to some extent with changes in temperature. This is usually not specified on the datasheet, although once in a while you will see a graph of (x, y) versus drive current. This seems to be due to small shifts in the phosphor emission with temperature. Since we know that red is the most susceptible to these changes, we may expect there to be shifts in both the CCT and the CRI. Since there is rarely data on this, you have to go into the lab and measure it, and then decide whether it is a significant factor for your application.

OTHER PERFORMANCE SHIFTS WITH TEMPERATURE

We haven't said anything yet about absolute maximum temperature ratings. Some LEDs are rated at 150°C, others at 85°C. This is probably due to a phosphor degradation effect. Our guess is that lifetime is the driver of this specification. For a variety of reasons, datasheets must specify that the LED drops to 70% of initial brightness after 50,000 hours of operation (see the next section for further discussion of this). The only variable that can be changed is the temperature. By specifying a lower temperature, phosphor degradation is decreased, and so the time to 70% light is lengthened. We suspect that the effect is not all that dramatic. If the 70% point is after 40,000 instead of 50,000 hours, are you going to reject the part? So it's not absolutely necessary to keep the LED below 85°C. Indeed, it's hard to run any decent power through an LED and keep it this cool. You need to talk directly with the manufacturer and see what data they have on lifetime at higher temperatures.

There's also some subtlety in reading this specification. Exactly what temperature is it that is supposed to be kept below absolute maximum? Some datasheets say it is the die, others the pads. There's a big difference between the two. If you're putting 3 W into an LED with a thermal resistance from junction to pad of 10 K/W,




Figure 5.2 LED Temperature Profile for Parameters Given in Text.

that's a temperature difference of 30°C, so when your pad is at 85°C, your die is at 115°C! Be sure you know at which point the temperature is measured for the lifetime rating.

One more performance consideration is the effect of pulsing current into the LED. Pulse-width modulation (PWM) of the current into the LED is a preferred method of dimming. This is because the current into the LED does not vary, but the amount of time the light is on does. If you dim the LED by reducing its current, there are some color shift issues, as described above. Now if you pulse the LED at high frequency, the device reaches a temperature dependent on the average power, as you would expect. But if the pulses are slow, there may be time between pulses for the die to cool. If even part of the time the temperature is lower, this may have a significantly positive impact on the 70% life of the LED.

What sets fast against slow degradation is the thermal time constant. As we discuss in Chapter 6, thermals are exactly analogous to RC networks. And so in addition to thermal resistance there is also thermal capacitance. The two together form a thermal time constant, which is how long the device takes to heat up when power is applied, and how long it takes to cool down when power is removed.

The thermal time constant of a typical LED is around 10 msec.¹ This means that after power is applied, it takes 10 msec to reach 63% of the temperature difference between start and end. Suppose you are putting pulses of current into an LED at 120Hz (say from a dimmer circuit). Suppose also that its on-state duty cycle is 80%. We'll assume that when the LED is on, the power level is such that in steady-state the die would reach 125°C. Look at Figure 5.2 to see a simulation of the LED temperature profile with these parameters. The peak temperature the LED reaches is 111°C. So the lifetime of the LED is considerably longer than it would be at 125°C. In fact, since 111°C is the peak temperature, the lifetime is longer than if the LED were steady-state at 111°C. But since it is complicated to figure out the actual

¹Inferred from Figure 4 of X. Poppe, "When Designing with Power LEDs, Consider Their Real Thermal Resistance", *LED Professional Review*, Nov/Dec 2009, p. 41. This number is order-of-magnitude confirmed by the standard pulse test for determining 25°C efficacy, which is 25 msec long.

temperature(s) at which the LED is operating, it is practical to use the lifetime at 111°C.

LED LIFETIME: LUMEN DEGRADATION

The previous sections have talked about temperature effects that can be quickly measured in a lab. Now we turn to very-long time thermal effects. As currently defined, an LED lamp is said to have a lifetime that is equal to the time required for half of the lamps to get to 70% of their initial light output (L70/B50). For better or worse, LED lifetime is defined the same way.

The first thing to observe is that this is difficult data to collect. Fifty thousand hours is about eight years of continuous operation. With new generations of LEDs coming out every six months, there isn't time to measure life before the part is obsolete. So lifetime measurement has turned to extrapolation from short-time data.

In a very general way, we expect everything to age on a logarithmic time scale. If a parameter decreases by 5% in 1000 hours, we expect that in the next 9000 hours it will decrease a further 5%. Unfortunately, this Arrhenius law applies to a single aging mechanism. In an LED, we've identified at least four independent aging mechanisms: the die, the phosphor, the encapsulant, and the package. For a given device, there are potentially at least four different logarithmic aging time constants.

As an example of this, consider the odd case of the brightening light bulb. Some LEDs have had light curves showing that there is an initial *increase* in light output at constant current and temperature over the first 1000 hours of operation. This appears not to be just a fluke measurement. Apparently some of the light bulbs the DOE tested have also gotten brighter in the first 1000 hours. After that, they started their decline as expected.

In another case, the manufacturer found that there were two different log slopes. During the first 1000 hours, the light decreased relatively rapidly. After that, the log slope decreased, and the light decreased very, very slowly. The conclusion here is that those various different aging mechanisms are interacting. That interaction is really what makes this problem so hard.

This problem is so trying that the Illuminating Engineering Society (IES), on which the government relies, gave up on it. LM-80, "Measuring Lumen Maintenance of LED Light Sources," was intended, as the name says, to define how well LED lights maintain their light with age. There was a great deal of committee work on the subject, from many of the manufacturers of LEDs. The conclusion was that there was no general way of estimating the 70% lifetime. The only way to know the lifetime is to actually measure it.

There's an additional complication in measuring the L70/B50 time. The definition says that this is the time at which half of the LEDs measured have reached 70% of their initial light output. So that's the mean time. But there's no specification of the standard deviation. Suppose that 48% of the sample fails after the first

2000 hours, and then there aren't any more failures for 48,000 hours. The L70/B50 time is 50,000 hours, but that seems incorrect. Since no one knows how to extrapolate the data yet, this isn't a current problem, but someday it will have to be addressed.

LED LIFETIME: CATASTROPHIC FAILURE

The lifetime of LEDs is treated differently from that of other bulbs. When the lifetime of an incandescent bulb is specified to be 1000 hours, this means that after 1000 hours half of a sample of such bulbs is burned out. Their filaments are broken and they produce no light whatsoever. Similarly for a fluorescent bulb, when their lifetime is specified to be 8000 hours (even though this number appears to be inflated), it means that after 8000 hours half of a sample of such bulbs is burned out. Their filaments are broken or their vacuum is lost.

LED lifetime is different. With industry agreement, the government has mandated that the "lifetime" of an LED light is the time after which half of a sample of such bulbs has lost 30% of its initial light. The LED light bulbs are not burned out, they are just a little bit dimmer. Their lifetime to 50% failure is large, probably mostly set by the driver circuitry in the bulb. The LED itself has a lifetime that is probably enormous, perhaps hundreds of thousands of hours. This lifetime is set by catastrophic failures, such as bond wires breaking or lightning strikes. For example, one manufacturer's test showed that after 50,000 hours, a sample of 68 devices had only one real failure.

How did such an inequality of treatment arise? Fluorescent bulbs have been very carefully engineered. It turns out that at roughly the time when they fail, their light output is about 70% of initial. And 70% is significant because it's the level at which people can start to tell that a bulb is dimmer. It thus appears that the problem is not with the definition, but with the LEDs' failure mode. Here's a proposal. Can manufacturers find one of the construction parameters to adjust such that $L70 = B50$? That would bring LEDs into parity with the other light sources. And it would be even better if that adjustment would result in a cost reduction.

PARALLELING LEDs

One thing that is rarely done is to put diodes in parallel. As LEDs are diodes, we should expect the same to be true for them. Let's take a quick look at why. If LEDs were absolutely identical, paralleling them would work. For example, in a Spice simulation, each LED has exactly the same forward voltage versus current. If you attach two of these simulated LEDs to a current source, each will take exactly half of the current. The reason this works is the negative feedback of the forward voltage. If the current in one were higher than in the other, it would have a higher forward voltage. But they are in parallel, so that can't happen.

In a real circuit, however, this negative feedback effect is overbalanced by a positive feedback effect. The diode with the higher current gets hotter, reducing its

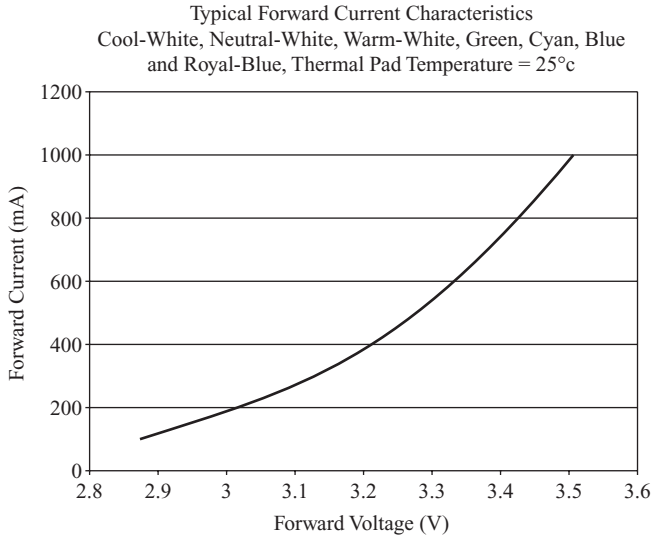


Figure 5.3 Forward Voltage as a Function of Current.

Source: “Technical Datasheet DS56, Power Light Source Luxeon Rebel”, Philips Lumileds Lighting Company, 2007.

forward voltage. This reduced forward voltage forces the lower current diode to also have a lower forward voltage, thus further reducing its current.

Let’s see what happens in a real case. We’ll do successive approximation. Assume that the two LED packages are in perfect thermal contact. We could assume that they are both hard-mounted to a 1-in thick piece of copper. (Don’t assume that your 10mil thick aluminum foil has infinite thermal conductivity; it doesn’t.) Let’s take two LEDs from the same forward voltage bin, which are a maximum of 100 mV different. For ease of reading Figure 5.3, we can assume that one has a forward voltage at 700 mA on the nominal curve of 3.38 V, and the other will then be 3.48 V. (Note the inconsistency in the datasheet; the parameter section says typical forward voltage at 700 mA is 3.60 V.) Now we have 1.4 A to distribute between the two, and they have to have the same forward voltage since they’re in parallel. If the nominal device takes 800 mA, it would have a forward voltage of 3.42 V, and the other device gets 600 mA, giving it a forward voltage of 3.32 V + 100 mV = 3.42 V, the same. (We got a little lucky here, guessing the 800:600 ratio right out of the box.)

The nominal device is dissipating 800 mA * 3.42 V = 2.74 W, while the other device is only dissipating 600 mA * 3.42 V = 2.05 W. At 10°C/W thermal resistance from junction to case, the die in the one device is up 27° from the case, and the other is up only 21°. Now consider that the forward voltage drops –3 mV/°C. The first device is hotter by 6° than the other, so it has a forward voltage 18 mV lower than the other. This causes a further small increase in the current that the first one is taking, but the change is small enough not to require us to go through a second iteration.

The conclusion is that, at least for this particular LED in this configuration, the positive feedback is not overwhelming but worrisome. One LED gets 200 mA more than the other. It probably won't be damaged, but we're exceeding its 700 mA rating by 14%. But if the metal plate connecting the two were thinner, or they were further away, the temperature difference might start becoming serious. In the worst case where the two LEDs are not connected at all, one could imagine one diode hogging most of the current, getting much hotter than the other, and eventually failing. In the chapter on DC drive circuitry, we'll discuss the feasibility of improving paralleling performance by adding a resistor in series with each LED.

Chapter 6

Practical Thermal Management of LEDs

In the previous chapter we discussed the numerous effects that temperature has on LEDs. This leads us to think about how to manage these thermal effects. We start off by giving an easily understood analytical method for calculating thermals. We then turn to an in-depth look at the environment in which LEDs operate, and a variety of methods for keeping them cool.

INTRODUCTION TO THERMAL ANALYSIS

Think about what happens when you put hot coffee into a ceramic cup. Within ten minutes, the coffee is noticeably cooler. If instead you put the coffee into a thermos, it's still hot ten hours later. The unit that measures the difference between the ceramic cup and the thermos is called *thermal resistance*. You can think of it as being analogous to electrical resistance. Instead of resisting the flow of electricity, it stops heat from flowing. The ceramic cup has low thermal resistance, so heat flows through it quickly, and the coffee's heat is lost quickly. The thermos has high thermal resistance, so heat flows through it very slowly and the coffee retains its heat for a long time.

Next think about boiling the water to make the coffee. If you have just a little water, it reaches boiling temperature quickly. If you are making coffee for 10 people on the same burner, it takes much longer to reach boiling temperature. The unit that measures the difference between a little water and a lot of water is called *thermal capacitance*. You can think of it as being analogous to electrical capacitance. Instead of storing up charge and increasing the voltage, it stores up heat and increases the temperature. A little bit of water has a small thermal capacitance, and so the temperature rises quickly. A lot of water has a great thermal capacitance, so its temperature rises slowly.

These examples are perfectly analogous to the principles discussed in this chapter. The underlying reason is that resistors and capacitors are linear devices, and

Table 6.1 Analogy between Electrical and Thermal Components

Electrical Parameter	Electrical Unit	Thermal Parameter	Thermal Unit
Resistance	Ohm	Thermal Resistance	°C/W
Capacitance	Farad	Thermal Capacitance	J/°C
Current	Amp	Thermal Power	W
Voltage	Volt	Temperature	°C

so are thermal resistors and thermal capacitors. The same equations apply to both, only the units are different. We thus have the direct analogy shown in Table 6.1.

You can readily check that this works properly. Ohm's law for electricity says that $V = IR$, which is that volts equal ohms times amps. Similarly, the unit for thermal resistance is °C/W, and the unit for thermal power is W, so their product is temperature with the units of °C.

Furthermore, there is an analog to an RC time constant. Ohms times farads equals seconds. Similarly, thermal resistance times thermal capacitance has units of J/W; since a joule is a watt-second, the product has units of seconds. As an analogy, think of heating up the coffee by heating the end of the spoon you use to stir it. How long it takes for the coffee to heat up to the proper temperature for drinking depends both on the thermal resistance of the spoon (metal has much lower thermal resistance than plastic), and the thermal capacitance of the water (more water has greater thermal capacitance).

CALCULATION OF THERMAL RESISTANCE

Using this analogy enables us to compute even complex thermal situations. Suppose you have a heat source that can flow through two different paths to the environment. Each path is a thermal resistor, and so they are in parallel. The total heat transfer is the parallel combination of the two thermal resistors, $R = R_1 * R_2 / (R_1 + R_2)$, just like in electricity. The ambient is at a fixed temperature, so it's analogous to a fixed voltage source. No matter how much current (thermal power) you put in to it, it's going to stay at the same voltage (temperature). You can thus figure out the temperature at which the heat source will sit. It sits at $V = I * R + V_{DC}$, and thus $T = \text{thermal power} * \text{thermal resistance} + T_{\text{ambient}}$.

Let's take some specific examples. I have an LED that is dissipating 3W. Its packaging has a thermal resistance of 10°C/W from junction to its case. The case has a thermal resistance of 12°C/W to the ambient temperature, which is at 30°C. At what temperature is the LED die?

We draw the schematic in Figure 6.1. The ambient temperature is at 30°C. There's 3W flowing through the thermal resistance of 12°C/W, and so the temperature at the case is $30^\circ\text{C} + 3\text{W} * 12^\circ\text{C}/\text{W} = 66^\circ\text{C}$. The same 3W is flowing through the thermal resistance of 10°C/W, and so the die temperature is at $66^\circ\text{C} + 3\text{W} * 10^\circ\text{C}/\text{W} = 96^\circ\text{C}$.




Figure 6.1 Thermal Model for LED Example.




Figure 6.2 Thermal Model of Two Parallel Thermal Paths.




Figure 6.3 LED Temperature as a Function of Time.

Let’s try an example with two parallel thermal paths. I have an array of LEDs dissipating 10 W mounted to a heat sink consisting of two flanges. One flange goes to the 40°C ambient, is 3 cm long, and has a thermal resistivity of 100°C/W-m. Its thermal resistance is thus 100°C/W-m * 0.03 m = 3°C/W. The other flange is made of the same material but is 4.5 cm long. Its thermal resistance is thus 4.5°C/W. What is the LED temperature?

We draw the schematic in Figure 6.2. The two thermal resistances are in parallel from the LED to the ambient, since heat flows through both of them. Both ends of both are at the same temperature. Thus the net thermal resistance = $3 * 4.5 / (3 + 4.5) = 1.8°C/W$. We have 10 W going through this resistance, and so the temperature rise is 18°C. The ambient is at 40°C, and so the LEDs are at 58°C.

Note how the heat flow splits between the two paths. We have a temperature drop of 18°C. The one path has a thermal resistance of 3.0°C/W, and so it carries 18°C/3.0°C/W = 6 W. The other path has a thermal resistance of 4.5°C/W, and so it carries 18°C/4.5°C/W = 4 W. The total heat power is (of course) 10 W, and the lower thermal resistance path carries more of the total heat power.

As a final example, suppose we have an LED in a small light bulb sitting at 25°C. The LED is dissipating 3 W. We measure its temperature as a function of time using a thermocouple, and get the curve shown in Figure 6.3. What are the thermal resistance and capacitance of the bulb?

The final, steady-state temperature of the LED is 55°C . This is a 30°C rise from ambient. Since this is caused by 3W of power, the thermal resistance is $30^{\circ}\text{C}/3\text{W} = 10^{\circ}\text{C}/\text{W}$. To compute the thermal capacitance, note that 63% of that 30°C rise equals 19°C . The figure shows that it reaches $25^{\circ}\text{C} + 19^{\circ}\text{C} = 44^{\circ}\text{C}$ in 140 seconds. So the LED rises from ambient to 63% of its final temperature at a time of 140 seconds, which is its thermal time constant. The thermal capacitance is equal to the thermal time constant divided by the thermal resistance (because $t = R * C$), $140\text{sec}/10^{\circ}\text{C}/\text{W} = 14\text{J}/^{\circ}\text{C}$. From the thermal time constant, we can also estimate how long we should continue measuring the light bulb to get to a steady-state temperature. In electronics, five time constants is considered a general rule, and so $5 * 140\text{sec} = 700\text{seconds} = 12\text{ minutes}$ is a reasonable amount of time to measure this system.

THE AMBIENT

We've been a bit cavalier about the "ambient" temperature. In many cases it's clear enough. Think of setting up a 1W resistor in your lab to dissipate $1/2\text{W}$. Thermally, the ambient temperature is the air in your lab, perhaps 25°C . If your resistor temperature rises by 50°C , its temperature will be 75°C .

But what's really going on here? The resistor has two wires going to your power supply, and those conduct some heat, although not very much. They are also dissipating some heat from their own I^2R losses, although again not very much in a properly designed setup. If the resistor is sitting on a lab bench, then some heat is being conducted by the bench. If the bench is metal, this may be the dominant heat removal mechanism for your resistor. And finally, the air itself carries away heat, but not at an infinite rate. The air too has an equivalent thermal resistance.

We end up with the thermal schematic shown in Figure 6.4. While we don't have values shown for the thermal resistances and capacitances, it's clear that we could measure them. This is quite a complex system, and the "ambient" temperature doesn't have a very clear meaning. So let's simplify this model further. Let's assume that the resistor has a battery built in to it, so it doesn't have any connecting wires

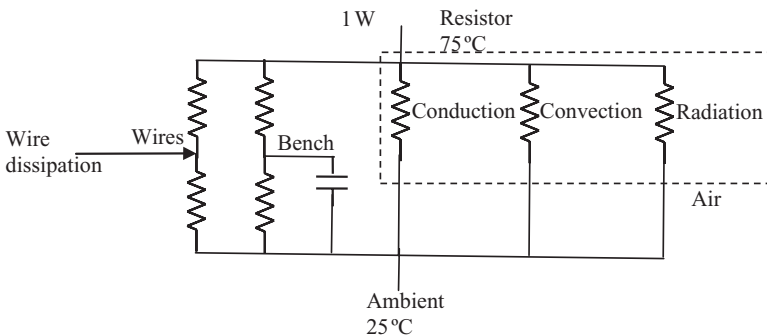


Figure 6.4 There Are Many Thermal Paths to Ambient.

to carry or generate heat, and that the resistor is suspended in air by very high thermal resistance strings, so that the bench isn't in contact with it. Now what happens?

Clearly, the temperature rise of the resistor in this case is going to be higher than before, because several of the parallel thermal resistances have been removed. The only one remaining is the air. We know from experience that the resistor surface is not going to be sitting at 25°C, the air temperature in the room. It gets hotter. How much hotter depends on a number of details.

The ambient air actually has three methods of removing heat from the resistor: conduction, convection, and radiation. The first two depend on there being air; the radiation doesn't. In a practical sense, conduction in air is unimportant for LEDs. The thermal conductivity of air is miniscule, 0.02 W/m-K. This is 1/10 the conductivity of plastics, which aren't very good thermal conductors. In short, air is a pretty good insulator, and we can safely ignore conduction in it.

That leaves convection and radiation. Convection is very complex. It depends on the size and shape of the object in the air, the characteristics of the air, and particularly the size, position and shape of objects surrounding the object under consideration. Simulations we've seen suggest that the convection flow is usually laminar, not turbulent, which simplifies things a little. But it is too hard to get a practical estimate of convective cooling in air without a specially built simulation model.

Finally, there is radiation. Radiation is governed by the Stefan-Boltzmann law. It says that a black-body in free space radiates at a rate set by the fourth power of the temperature. This is somewhat more complicated than it seems at first, because you have to subtract out the effect of the ambient temperature. A body at 298°K (=25°C) radiates 288 mW/in², but it also absorbs radiation at the same rate, so that it is in thermal equilibrium. A body at 25°C sitting in a 25°C environment doesn't radiate any net power.

PRACTICAL ESTIMATION OF TEMPERATURE

We can still get an upper bound on the temperature by supposing that there is no convection, and only radiation is removing heat. Once we calculate the temperature based only on radiation, we know that the actual temperature must be lower than this when we add the convection back in. Better yet, there are some empirical reasons to suppose that in the range of interest for LED systems, radiation is about 1/3 the total power removed. The other 2/3 is convection of the air.

Rather than dealing with the math, let's look at Figure 6.5. It gives you an estimate of the temperature of an object dissipating a certain number of W/in² in air. To use it, you calculate the power the object is dissipating and divide it by the radiating object's surface area. Reading the curve for the appropriate ambient temperature gives you an estimate of the temperature rise of the surface.

Let's take a 5 W resistor as an example. Its surface area is approximately the area of its cylindrical body plus its two end caps. We have $\text{Area} = (8 \text{ mm} * \pi * 24 \text{ mm}) + 2 * (\pi * (8 \text{ mm}/2)^2) = 703 \text{ mm}^2 = 1.09 \text{ in}^2$. At 5 W, the power density is




Figure 6.5 Estimating Temperature Rise from Power Density.

$5\text{ W}/1.09\text{ in}^2 = 4.6\text{ W}/\text{in}^2$. Reading the curve for a 70°C ambient suggests the resistor body should rise 142°C above ambient, to 212°C .

The datasheet shows that the resistor is rated to dissipate 5 W at 70°C and 0 W at 235°C (the curve is cut off at 220°C for unrelated reasons). This means that the 5 W causes a temperature rise of 165°C . This is reasonably close to our chart's answer.

Let's also try out our original resistor experiment. We have $\text{Area} = (2.8\text{ mm} * \pi * 6.35\text{ mm}) + 2 * (\pi * (2.8\text{ mm}/2)^2) = 68\text{ mm}^2 = 0.11\text{ in}^2$. At $\frac{1}{2}\text{ W}$, the power density is $0.5\text{ W}/0.11\text{ in}^2 = 4.7\text{ W}/\text{in}^2$. We get the same power density, suggesting the resistor body should rise 142°C . But of course it doesn't. The resistor is rated for 1 W , so $\frac{1}{2}\text{ W}$ shouldn't heat it up that much. The problem can be traced to the measurement conditions. For the test, the body of the resistor is elevated above the board. For such a tiny part, the leads contribute significantly to the surface area, and are a significant part of the cooling. The effect wasn't noticeable when we calculated for the 5 W resistor because it is so much bigger.

The conclusion is that for reasonably sized objects dissipating reasonable amounts of power, Figure 6.5 gives a reasonable first guess at surface temperature rise. But there is no substitute for a real measurement.

HEAT SINKS

Back in Figure 6.4, we showed all of the various thermal paths to ambient. A good way to reduce the temperature of the device is to dramatically lower just one of the resistances. Since they're all in parallel, this one will then dominate the others, just like a $1\text{ K}\Omega$ resistor in parallel with a bunch of $100\text{ K}\Omega$ resistors is still about $1\text{ K}\Omega$. Since in this case there will be only one significant thermal resistance, this also makes calculation a lot easier.

In this section we will look at reducing the resistance of the thermal conduction path. This is typically done with a heat sink, a piece of metal attached either directly or indirectly to the LEDs. As an example of what can be done, look at the spec for



Figure 6.6 An LED Heat Sink.

Source: <http://www.aavidthermalloy.com/cgi-bin/stdisp.pl?Pnum=569000b00000g>. Courtesy Aavid Thermalloy.

Aavid Thermalloy's part number 569000B00000G (see Fig. 6.6). The big basket in the middle is to allow multiple LEDs to be mounted, the fins are to increase the amount of convective cooling, and the black anodization is to increase the thermal radiation. Together they achieve a thermal impedance of $5.5^{\circ}\text{C}/\text{W}$. This will doubtless be the lowest thermal impedance in the system, and so design is easy. If the LEDs are dissipating 8 W, the temperature rise will be 44°C .

Now of course there's a catch. In order for the heat sink to work as advertised, the LEDs have to make proper contact with it. And in this case, proper contact means proper thermal contact. If the LEDs are loosely suspended above the heat sink, of course it won't work properly. You need to attach the LEDs with something that will both hold them in place and have low thermal resistance to the heat sink.

While there are a number of such methods available, a common one is a thermal epoxy. A typical one has a thermal conductivity of $1\text{ W}/\text{m}\cdot\text{K}$. To figure out the thermal resistance, *don't* multiply this by the thickness, even though you get the right units! You need to multiply by the area covered by the epoxy, and *divide* by the epoxy's thickness. This is thermal conductivity. Then you invert the number to get thermal resistance.

To take an example, suppose we are bonding a 4 cm^2 area, and spread the epoxy to be 1 mm thick. Then the thermal conductivity will be $1\text{ W}/\text{m}\cdot\text{K} * 0.0004\text{ m}^2/0.001\text{ m} = 0.4\text{ W}/^{\circ}\text{K}$. The inverse of this is $2.5\text{ K}/\text{W}$. This is in series with the thermal resistance of the heat sink, giving a total thermal resistance of $8.0\text{ K}/\text{W}$ from LED case to ambient. With our 8 W example above, the temperature rise is actually 64°C , rather than 44°C .

Note that increasing the area of the thermal contact increases the conductivity, which decreases thermal resistance, just as you would expect. And similarly, decreasing the thickness of the epoxy also decreases thermal resistance. So the goal should be to get the maximum surface contact area possible, and then spread the epoxy as thin as possible.

While we're on the subject of heat sinks, we should mention metal-core PCBs (MCPCBs). These are made by specialist manufacturers, and consist of a regular PCB bonded to a piece of metal, usually aluminum. The bonding is usually isolated. MCPCBs have found some favor for mounting LEDs. The attraction of the method is that thermal performance is better than that sometimes accomplished by

“do-it-yourself” methods. The downside, of course, is that the cost is higher than that of regular PCB. The authors have considered using them in a number of projects, but usually end up not doing so because of the additional cost.

FANS

Another thermal resistance that can be minimized is the convection. Convection in air is moderately effective at cooling, but it can be dramatically enhanced by a fan (for our practical purposes, a blower is the same as a fan). Fans and blowers work by forcing air to move across the hot surface. But how do you select a fan?

Fans are typically rated by how much air they move, with units of liters/minute. The reality is that there are a lot of complex motions of the air that determine how much cooling you get. The size and shape of the fan, the size and shape of the object being cooled and its orientation to the fan, as well as other objects in the path of the air flow, all contribute to the actual cooling. So the best we can offer is an estimate of how much fan you need.

From an article by Mike Turner (Turner 1996), we can estimate the flow rate from the equation $G = P/(\rho C_p \Delta T)$, with G the volumetric flow rate in m^3/sec , ρ and C_p characteristics of air, and ΔT the temperature rise. Plugging in air values at 25°C , and converting to liters and minutes, we end up with the estimate

$$G(\text{l/min}) = \frac{12P(\text{W})}{\Delta T(^{\circ}\text{K})}$$

Note that this is independent of the surface area of the object being cooled, and so we are assuming that all of the air blows past it. As an example, consider the 5W resistor from before. It rose 165°C in 70°C still air. Suppose we want the temperature rise to be a much milder 30°C . Then we need airflow of $12 * 5\text{W}/30\text{K} = 2\text{l/min}$. Looking at actual devices, we see the Sunon¹ UF3A3-500 has a rated flow of 3.43l/min, and so will probably do just fine. Of course, this formula is only an estimate; you need to actually build the system to verify that temperatures are what you calculate.

Now there are a number of issues with fans. In the first place, the air has to have some place to go. If you put a fan into a totally enclosed system, it will help to make sure the whole inside is at the same temperature, but it won't help much with getting the heat out to the ambient. To work best, you need to have an inlet and an outlet for the air. This may or may not be an aesthetic system design problem.

For one thing, you have a high-speed blade whirling around (17,000rpm for the fan we just picked). What happens if a piece of dust or paper gets sucked in? Aside from affecting your cooling, it could also cause the fan to stall. Is the fan protected so that it doesn't burn up when it stalls indefinitely? PTCs are often used for this purpose. And what about safety issues? What if a child sticks his finger into the fan? This must be considered when using a fan as a coolant.

¹http://www.sunon.com/mmfan/mmfan_e.htm

The fan also requires power. AC can be convenient for powering fans, but then you have to make sure customers can't come into contact with the live AC. It's much more common to use DC brushless fans, which run on 5 V or 12 V and also have reduced EMI compared with other types. But the 5 V has to come from somewhere, this somewhere presumably being the LEDs' ballast. Do you have 5 V or 12 V already in the ballast, and if not, how are you going to generate it? Fans can draw substantial current; the one we just picked needs 100 mA at 3 V, 300 mW. This is more than you want to generate from a zener. And now you've added an extra 300 mW to the power dissipation of your light. That means additional cooling is needed, and that there's an additional drop in efficacy of the light.

A final concern is acoustic noise from the fan. Everyone is familiar with noisy fans in computers. Less noisy ones cost more. How will your customers react to a light that has a fan running all the time?

RADIATION ENHANCEMENT

One more area that you should consider for reduction of thermal resistance is the thermal radiation. You can't affect the laws of physics, but notice that the estimated temperature rise in Figure 6.5 depends both on the amount of power and the surface area of the object. If you could increase the surface area, that would decrease the temperature.

To get a feel for how much benefit this could have, consider that 5 W resistor again. Its nominal size is 8 mm diameter by 24 mm length. But the maximum size allowed by the tolerance is 9 mm diameter by 25.5 mm length. The nominal values have a surface area of 703 mm², while the maximum is 848 mm². This is an increase of 21%. It drops us from 4.6 W/in² to 3.8 W/in², dropping our expected temperature rise from 165°C to 124°C.

The conclusion here is that even small increases in size can have dramatic effects on temperature. The manufacturer of the resistor could increase the power rating of his device without affecting its size by holding the mechanical tolerances tighter, and specifying the nominal dimensions more towards the top of the current range. It's worth asking whether your device could be 10% bigger without customer rejection.

REMOVING HEAT FROM THE DRIVE CIRCUITRY

While the focus on this chapter has been on thermal management of the LEDs, you may also need to do some thermal management of the ballast. Consider a ballast delivering 30 W to a set of LEDs. Getting rid of the heat from the LEDs is probably the most important concern you have, but it shouldn't be the only one. If the ballast is 85% efficient, the power it dissipates is $(30 \text{ W}/0.85) - 30 \text{ W} = 5.3 \text{ W}$, 18% of the power the LEDs are getting.

What are the alternatives? You might be told to increase the efficiency. If the efficiency were 100%, then it wouldn't dissipate any power at all! Of course, 100%

efficiency is physically impossible. And as we'll see in the next chapters, it is hard to increase efficiency very much. 85% is a good compromise. You can get to 90%, or maybe even 92–93%, but it costs increasingly more money and design time to do so. At some point, it doesn't make sense to pursue efficiency improvement further.

There are specific actions you can take to remove heat from the ballast. The first is simply to realize that only a few components in the ballast generate significant heat. These are the power components, typically the transistor and diode, though the IC may in some cases also dissipate significant power. Putting a good amount of copper on the PCB, including a ground plane, can help to distribute the heat more evenly throughout the ballast, helping the temperature of these components.

Another important thing to realize is that ballasts are fairly immune to high temperatures. Most components used in a ballast are rated to 125°C. The important exception is electrolytic capacitors. These can have lifetimes very seriously degraded by elevated temperature, and can become the lifetime-determining component in the entire light. Methods of dealing with this, as well as ways of not using electrolytic capacitors, are detailed in the following chapters. Given this, there may not be a grave disadvantage to letting the ballast run hot, as long as customers can't burn themselves.

The same things that can be done to alleviate temperature in an LED system can also be applied to the ballast. The power components can have heat sinks attached to them. The ballast can have a heat sink attached to it on the outside case. In extreme cases, a fan can be used to blow air over the ballast case. But again, the easiest thing to do is to make the ballast mechanically larger. This has two positive effects. It allows the heat-generating power components to be placed further away from the heat sensitive components, and it provides increased surface area for the ambient to cool the ballast.

Chapter 7

Practical DC Drive Circuitry for LEDs

LEDs need to be electrically driven in order to emit light. In this and the next chapter, we're going to discuss how to design drive circuitry for LEDs. In this chapter we're specifically interested in DC drive circuitry. A typical DC source is a battery, as, for example, used in a flashlight. It could also be the output of a switching converter, for example 12VDC. In any case, what distinguishes DC from AC for practical purposes is that DC typically has a much lower voltage than AC. This means that the regulations governing usage are much easier to comply with. There's no EMI to worry about, the voltages are generally "safe" according to UL and the input voltage is generally very steady. On the downside, the currents in a DC drive are higher, and the source impedance becomes an important factor in the design.

BASIC IDEAS

The fundamental determinant of what type of converter to use for a DC drive is set by the relative values of the supply voltage and the LED voltage. There are three cases.

1. The supply voltage is always higher than the LED voltage.
2. The supply voltage is always lower than the LED voltage.
3. Sometimes the supply voltage is higher than the LED voltage, and sometimes it's lower.

When we use the word "always," this includes variation of the supply voltage with time, as well as variation of the forward voltage of the LEDs due to temperature, binning, and so on. For example, the supply voltage is always higher than the LED voltage only if there is *never* a time in operation when it is lower.

In the first case the supply voltage is always higher than the LED voltage. A typical example of this would be a flashlight using four D cells, and running a single

Practical Lighting Design With LEDs, First Edition. Ron Lenk, Carol Lenk.

© 2011 the Institute of Electrical and Electronics Engineers, Inc. Published 2011 by John Wiley & Sons, Inc.

3 W white LED. The D cells are 1.5 V when new, and something like 0.9 V when at end of life, giving a battery voltage range of 3.6–6.0 V. The LED’s forward voltage is 3.2 V at 350 mA, and decreases with increasing temperature. In cases like this, a buck converter would normally be used.

In the second case the supply voltage is always lower than the LED voltage. A typical example of this would be a flashlight using two C cells, driving a set of seven 5 mm white LEDs. The C cells have the same voltage range as Ds, so the battery voltage range in this case is 1.8–3.0 V. The LEDs have a voltage of 3.2 V each, for a total string voltage of 22.4 V. In cases like this, a boost converter would normally be used.

Finally, in the third case the supply voltage is sometimes higher and sometimes lower than the LED voltage. An example of this would be the same as the example for number one, but using three “D” cells rather than four. Now the battery voltage range is 2.7–4.5 V. The LED voltage is right in the middle of this range. For a case like this, a buck-boost converter would typically be used.

BATTERY BASICS

All three of our examples used batteries. Therefore, we need some basic information about batteries. First let’s examine terminology. In common speech, we use the word “battery” to describe both a D-size battery that you buy in the supermarket and the big battery that powers a laptop computer. Technically, however, the packages the supermarket is selling contain cells, not batteries. A battery is the end device that powers something. It often consists of two or more cells, usually in series. But sometimes, as in children’s toys, it consists of a single cell, hence the confusion. Thus, a laptop really does have a battery. It has three or four 4.3 V lithium cells in series. Your flashlight with two D cells also is using a battery. We’ll try to consistently maintain the difference between the words.

The most basic piece of electrical information about a cell or battery is that it is a voltage source. To zero order approximation, an alkaline cell of any size (D, C, AA, AAA, etc.) produces 1.2 V. To the same approximation, a lithium cell produces 3.6 V. The reality is considerably more complicated. Cells do in fact change their output voltage, depending on load current, temperature, state of charge, and age. And different types of cells have different characteristics as well. Some understanding of all these aspects is necessary to be able to properly design a DC drive for LEDs.

Cells are actually complicated electrochemical systems. The “chemical” part of this word is referring to how the cell stores energy. Cells work differently from capacitors. Capacitors are purely electrical. They store energy in an electric field. Typical energy levels for a capacitor can be described by the image of a 1 μF capacitor charged to 5 V. The energy it stores is

$$E = \frac{1}{2} CV^2 = \frac{1}{2} 1 \mu\text{F} (5 \text{ V})^2 = 12.5 \mu\text{J}.$$

Cells store energy in chemicals. Typical energy for a cell can be seen by thinking of a 1.2V cell with 200mAh of charge. (Milliamp-hour is the usual rating for cell capacity.) The energy stored is approximately

$$E \approx V * Q = 1.2 \text{ V} * 200 \text{ mAh} * 3600 \text{ sec/hr} = 864 \text{ J.}$$

The cell stores some eight orders of magnitude more energy than the capacitor. This is why batteries are used rather than capacitors to run LEDs for lighting or for any device requiring significant power.

While we don't need to know exactly how the chemistry in the cell works, it's important to know that there are a couple of different kinds on the market. We've already mentioned the two most important, alkaline and lithium (nickel metal hydride [NiMH] has been largely superseded by lithium). The C and D cells sold in the store are typically alkaline. These are usually single-usage types. Their fully charged voltage is 1.5V, and end-of-life voltage is about 900mV. You can actually run them lower than this, but there is very little energy left in them below 900mV.

The other common type of cell is a lithium cell. These are almost always rechargeable. Their fully charged voltage is about 4.3V, and discharged is usually considered to be about 2.7–3.0V. Lithium cells are considerably more sensitive than alkaline cells. For example, they might explode if you continue to charge them after they're fully charged. And they can be permanently damaged if you discharge them below about 2.4V. The bottom line here is that practical lithium batteries almost always use a special IC to control them, to take care of all the problems they can present. For our purposes, we'll simply assume that this circuitry is already in place, and that the LED drive we're designing works with the output of this circuitry.

Now we can examine the performance of batteries. As we've already noted, the output voltage of a cell is constant only to zero order approximation. What really happens is that for small loads, the voltage is close to a constant. But as the load becomes bigger, the voltage drops, approximately linearly for medium loads, as though there were a resistor in series. When you get to heavy loads, the voltage drop is large and nonlinear, something like a hyperbolic tangent (see Figure 7.1). In a practical sense, batteries for LED lighting are almost never run at heavy loads. A little thought shows why. If you run the battery very hard, it will discharge quickly. For most lighting applications, you want the light to be on for a long time. Therefore,




Figure 7.1 12V Battery I-V Curve.

Source: Lenk (1998).

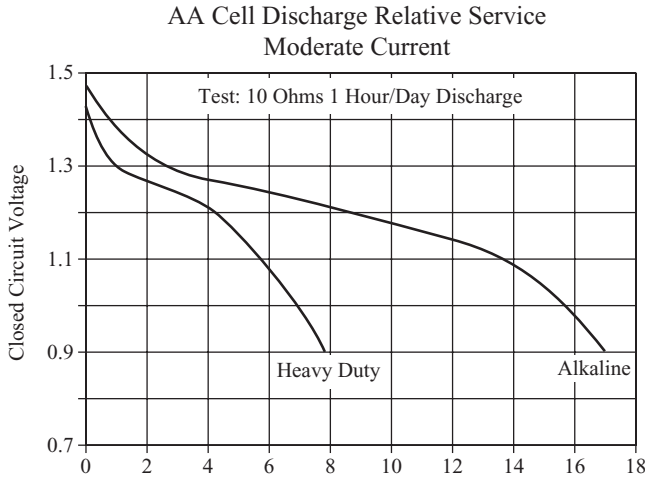


Figure 7.2 Alkaline-Cell Battery Voltage as a Function of Time with a Resistive Load.

Source: Rayovac, OEM 151 (R-3/99), “Application Notes & Product Data Sheet,” “Primary Batteries — Alkaline, & Heavy Duty,” Figure 1. Property of Spectrum Brands, Inc.

you almost always run the battery in its light or medium load ranges. And, thus, estimating the battery voltage as an ideal battery plus a resistor is reasonable for most of our applications.

But what is a light or medium load? And what resistor value should you assume? There’s no standard answer. It varies from type to type and to some extent, even from manufacturer to manufacturer for the same type. Let’s look at an example to see how to figure it out. Looking at Figure 7.2, the manufacturer supposes a constant 10Ω to constitute a moderate load for an AA alkaline, a 2.2Ahr cell (we’re looking at the “Alkaline” cell curve in Figure 7.2, not the “Heavy Duty” cell curve). Much of the time, the cell sits at 1.2 V, and so the current for a typical moderate load is around 120 mA for this cell. Note that as the cell discharges, the output voltage drops.

Now we can estimate the approximate series resistance in this cell. The data-sheet says that the fresh voltage of the cell is 1.55 V. Figure 7.2 shows the $10\text{-}\Omega$ load at the beginning of discharge seeing a cell voltage of 1.47 V. This is a drop inside the cell of 80 mV in response to a current of 120 mA. We thus estimate the internal cell resistance as $80\text{ mV}/120\text{ mA} = 667\text{ m}\Omega$. It’s unfortunate we can’t make the same estimate for the cell resistance when it’s almost finished discharging, but we’ll assume that it’s not radically different.

Another factor affecting batteries is ambient temperature. As cells get colder, their voltage drops and their impedance rises. This particular alkaline is rated from -30°C to $+55^\circ\text{C}$. Below -30°C , the chemicals inside the cell start to freeze. Practically, what this means is that the impedance becomes so high that the cell can’t source any current. And what happens above 55°C ? Unless you’re considerably above 55°C , probably nothing catastrophic will happen. But the self-discharge rate becomes so high that it seriously undermines operating life from the cell. (Batteries sitting on the shelf do discharge themselves gradually.)

Finally, there are aging effects on rechargeable cells. As you charge and discharge them repeatedly, they gradually lose their ability to store charge. After some hundreds of cycles, they will lose so much capacity that they are considered dead. In the old days, nickel-metal hydride batteries also had a memory effect. If they were only partially discharged before being recharged, they lost some of their storage capacity. Fortunately, lithium cells don't have this sort of problem.

OVERVIEW OF SMPS

Switch-mode power supplies (SMPS) are used almost universally to convert a source of power into a form suitable for a load. Common ones include the “wall-wart” that plugs into the wall to convert 120 VAC or 240 VAC into 5 VDC to charge your cell phone, and the one on the PCB of your computer that converts the 3.3 VDC from the silver box to 1 V to run the processor. Other converters deliver a constant current rather than a constant voltage, and this is the kind we will be designing for LED lighting. The fundamental distinction between all types of SMPS is whether the power source is AC or DC, and we will be following this distinction by treating the two in separate chapters. This chapter is about DC.

SMPS operate in switch-mode, hence the name (see Lenk, 1998). What this means specifically is that there are one or more transistors in the circuit, which switch on or off to control the operation of (usually) an inductor. Transistors *could* be operated in linear mode instead of switch-mode, but this is inefficient. Referring to Figure 7.3, we see the problem. The transistor is basically acting as a linear regulator. The input voltage of 12 V is dropped to 5 V across the transistor, which therefore has 7 V across it. If the output current is 1 A, then the power dissipation in the transistor is $7\text{ V} * 1\text{ A} = 7\text{ W}$, and the efficiency is $\eta = P_{\text{out}}/P_{\text{in}} = (5\text{ V} * 1\text{ A})/(12\text{ V} * 1\text{ A}) = 42\%$.

In switch-mode, the circuitry is much more efficient. Referring now to Figure 7.4, the transistor is shown as a switch, which is either open or shorted. The efficiency of the SMPS comes from operating the transistor only in these two modes, and trying to get from one to the other as fast as possible. When the switch is shorted (the transistor is on), as in (a), the input voltage is directly applied to the inductor. The diode is reverse-biased and so is off. Since the voltage on the switch side of the inductor is higher than the output voltage, the inductor current increases (remember that $V = Ldi/dt$).

In the next step, shown in (b), the switch is open (the transistor is off). The input voltage is disconnected from the rest of the circuitry. Now the voltage across the inductor is the other way. The output voltage is higher on the output than on the input of the inductor, so the current decreases. The current continues to flow in the inductor, and since it has to come from somewhere, it comes through the diode.




Figure 7.3 Operating a Transistor in Linear Mode Is Inefficient.

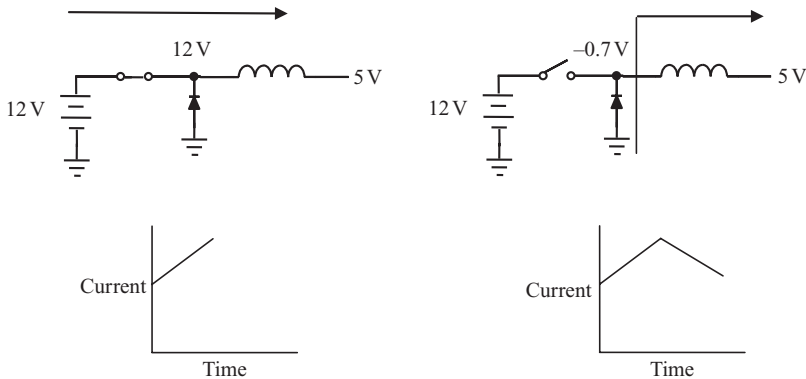


Figure 7.4 When the Transistor (t) Is On, Current in the Inductor (I) Increases; When the Transistor Is Off, Current in the Inductor Decreases.

The voltage on that node drops until the diode can conduct. The diode becomes forward-biased and turns on. That means the voltage on that side of the inductor drops to a diode drop below ground, about -0.7V .

Now here's the crucial step. The inductor current is thus forced to alternately increase and decrease each cycle. The inductor current goes into a capacitor, which is alternately slightly charged and discharged each cycle. By controlling how long the inductor current increases versus how long it decreases, the average current can be set. And it is set in such a way as to produce the desired average voltage on the output capacitor. The setting is done by a feedback circuit, the PWM (pulse-width modulation) controller, which in practice is always an IC. The controller measures the output voltage, and adjusts the amount of time the switch is on to keep the output voltage constant. The ratio of the on-time of the switch to the total time (on-time plus off-time) is called the duty cycle (DC; although this has to be distinguished from "direct current" by context).

There are quite a number of different ways of hooking up the transistor, inductor, diode, and capacitor, including methods that use multiple devices. All of the different configurations are generically called SMPS topologies. There are also a variety of modulation schemes, such as the constant-frequency control we've been discussing, as well as constant on-time control and others. But all of these systems work on the same basis, modulating the current in an inductor to regulate the output. In the next few sections, we will cover some of the basics of the topologies that will be suitable for driving LEDs from DC sources. We'll cover these designs in much more detail in Chapter 10.

BUCK

Let's cover the first case, where the source voltage is always higher than the LED voltage. The topology we will use is the buck converter, which is what we have been discussing in the previous section. Recall how this worked. To make the inductor




Figure 7.5 LM3405 Schematic for Buck.

Source: LM3405 datasheet, National Semiconductor, February 2007.

current increase, we needed to apply a voltage to it that was higher than the output voltage. If the voltage applied was lower than the output voltage, the current would decrease, not increase, and the output voltage would drop, no matter how long the transistor was on. Eventually, the output voltage would equal the input voltage.

Thus, a buck converter can only convert an input voltage to a lower output voltage. You can't get a higher voltage out of this topology than you put in. This is exactly the way the first case works. As an example circuit, we're going to look at National Semiconductor LM3405. This part is convenient to use because the switch is integrated inside the IC, and the whole thing fits in a 6-pin SOT, a tiny package. We will be using this circuit in a flashlight design in Chapter 10, converting the USB output of a computer to drive a single power LED.

Look at Figure 7.5, taken from the front page of the LM3405 datasheet. The input voltage can be as high as 20V on this device, and so the IC is powered directly from the input on the V_{IN} pin. The transistor for the buck is inside the IC, between the V_{IN} and the SW (switch) pins. The LM3405 switches the transistor on and off at a constant frequency of about 1.6MHz. This connects and disconnects the SW pin to the input voltage. When the transistor is on, the voltage is positive across L1, and when the transistor is off, the voltage on SW swings down to a diode drop below ground, turning on D1. The current in L1 thus increases and decreases. Note that the high switching frequency of the LM3405 means the value of L1 can be very small, typically around 10 μ H.

The current in L1 is smoothed by C2, and is fed to the LED. C4 is basically in parallel with C2, but also has some function for stability of the circuit. The current is sensed by R1, which produces a voltage proportional to the current and is fed back to the FB (feedback) pin. The IC controls the duty cycle of its internal switch to produce a voltage of 205 mV on the FB pin. This voltage is intentionally low in order to avoid significant power loss in R1 ($1A * 205mV = 205mW$, which can be handled by a 1/4W resistor). This may possibly require an RC filter between R1 and the FB pin, to reduce noise fed back to the IC. Other ICs have even lower feedback voltages. You should generally avoid using a normal IC used for producing a fixed

output voltage. For example, a 2.5V feedback at 1A would lose 2.5W in the resistor.

The LM3405 can be controlled on and off with a digital signal on the EN/DIM pin (1.8V on, 0.4V off). The final elements in Figure 7.5 are D2 and C3. These are used as a charge pump, to produce a voltage on the BOOST pin that is higher than the input voltage. What is this used for? Remember that there is a MOSFET inside the IC between V_{IN} and SW. When the MOSFET is turned on, we want the voltage on SW to be as close as possible to the input voltage, as otherwise there is a voltage drop across the transistor, leading to power loss. But to turn on an N-channel MOSFET, the gate voltage has to be higher than the source voltage. The BOOST pin provides that higher voltage.

BOOST

In the second case, the input voltage is always lower than the output voltage. The topology we will use is called a boost. As the name suggests, it takes the input voltage and boosts it up to a higher voltage. And just as the buck is incapable of producing a higher voltage output than its input, the boost is unable to produce an output that is (significantly) lower than its input. As an example circuit, we're going to use the Fairchild FAN5333. This part is convenient to use because the switch is again integrated inside the IC, and the whole thing fits in a 5-pin SOT. We will be using this circuit in a flashlight design in Chapter 10, boosting the output of 2 'D' alkaline cells to drive a 3W LED.

To understand how a boost works, look at Figure 7.6. Once again, the input voltage goes to the V_{IN} pin, this time limited by the manufacturer to 6V. Now in this case, the transistor, diode and inductor are still present, but are connected together differently. Here's how this works. The input voltage is always attached to the inductor. The reason the inductor current doesn't continue to increase is because the voltage on its other side is higher than the input voltage (that's why the boost has to have higher output than input voltage). The transistor is operated by periodically

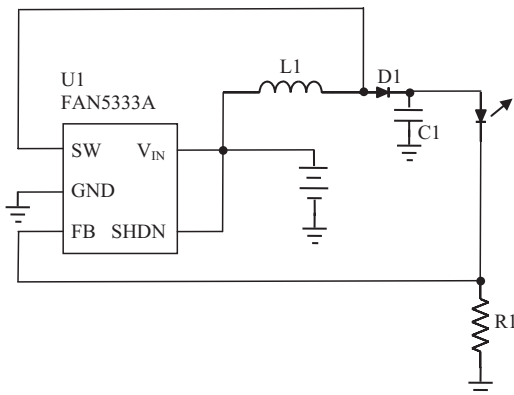


Figure 7.6 FAN5333A Schematic for Boost.

Source: FAN5333A datasheet, Fairchild Semiconductor, August 2005.

pulling the SW pin to ground. When the transistor is on, SW is at zero volts, and so the inductor current increases. When the transistor is off, SW goes up to a diode drop above the output voltage, and so the inductor current decreases.

The purpose of the diode D1 is thus clear. When the switch is on, the SW pin is pulled to ground. If the diode wasn't there, that would pull the output to ground as well. The current from the inductor goes through D1 and into C1 and the load. The current is sensed by R1, producing a feedback signal at the FB pin. The FAN5333A controls the duty cycle of its internal transistor to produce a regulated voltage at the FB pin (110mV), corresponding to the desired current through the LED. You should note that when the switch is on, the inductor current is going to ground. To keep the LED on during this time, there has to be enough energy stored in capacitor C1 to make sure that the voltage across the LED doesn't change significantly.

In this particular schematic, the shutdown pin, SHDN, is connected high, so that the IC is always on. The only other aspect to note in this design is that no charge pump circuit is required. The internal MOSFET goes from SW to ground, and so the input voltage is high enough to run its gate.

BUCK-BOOST

We've saved the best for last. The buck converter can only be used when the output voltage is less than the input voltage. Conversely, the boost converter can only be used when the output voltage is greater than the input. What do you use when the input is sometimes higher than, and at other times lower than, the output? This situation arises because battery voltages are dependent on a number of factors, including their state of charge and temperature. So a fresh battery may have a voltage greater than the LED it's driving, but when it's been mostly discharged, its voltage may be lower.

The traditional circuit for this power supply is called a buck-boost. As the name suggests, it is a buck and a boost stuck together. This involves two transistors and two diodes, one set for each. Conceptually, it works by turning off the buck and running the boost when the input voltage is low, and by turning off the boost and running the buck when the input is high.

For many LED drivers, however, this may be too expensive. Four external power devices eat up quite a bit of board space and money; using integrated switches costs even more. In this section, we're instead going to do something more clever. We will use the Supertex HV9910 to create a circuit that works regardless of the magnitude of the input or output voltages. We will be using this circuit in a flashlight design in Chapter 10, converting the output of a car battery to run an LED taillight.

To understand how this works, look at Figure 7.7. Once again, the input voltage goes to the V_{IN} pin, though this time it can be as high as 450V (the HV9910 is really designed for AC off-line application). Now in this case, the transistor, diode, and inductor are connected together exactly the same as they are in the boost converter.

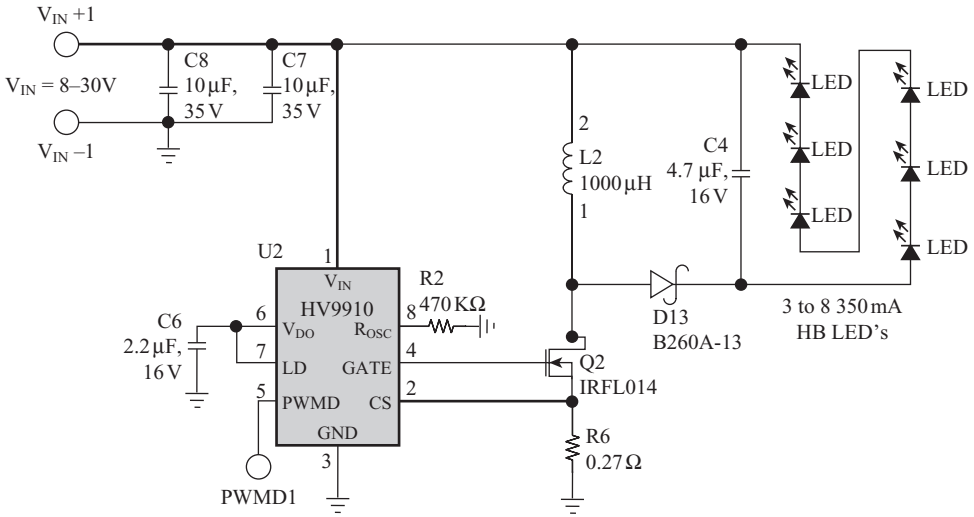


Figure 7.7 HV9910 Schematic for Buck-Boost.

Source: HV9910 datasheet, Supertex Inc., 2006.

The transistor is external, but this has no bearing on the operation. What's different about this circuit from both the LM3405 and the FAN5333A circuits is the position of the load, the LEDs. In the previous circuits, the LEDs were connected from the output to ground. In this circuit, by contrast, they are connected from the output *back to the input*. How does this work?

LEDs are current devices. As long as they have current flowing from their anode to cathode, they light up. It does not matter to their operation where ground is. So the idea of this circuit is to use a boost converter to produce a voltage higher than the input, and then run the LEDs from this. Since the output voltage is higher than the input, the LEDs can be referenced to the input rather than ground. The big advantage of this circuit is that it doesn't matter what the voltage across the LEDs is. It can be lower in magnitude than the input voltage is to ground, or higher. Since the top of the LED string is at higher voltage than the input, the boost operates properly no matter what.

To complete the circuit, capacitor C4 is in parallel with the diodes. It too is not ground referenced. Since this HV9910 circuit is indeed a boost, during Q2's on time the inductor current goes to ground, not through the LEDs. Thus C4 is there to provide current during this off-time.

Now while this circuit has notable advantages, it also has some downsides, which are fortunately taken care of by the IC. In the first place, the FAN5333A and the LM3405 measure the LED current by measuring the voltage across a ground-referenced resistor. Since the IC is also ground-referenced, this is convenient. The IC can just put the measured voltage directly into a comparator. But for this HV9910 circuit, the LEDs are above ground, and so sticking a resistor in series with them would require accurately measuring the difference between two high voltages, which

is difficult. The HV9910 instead measures the LED current by changing the way it controls the switch.

Both the FAN5333A and the LM3405 are standard SMPS (technically, they use current-mode control, but that doesn't matter to us here). They compare how much current there is with how much there is supposed to be. If those two are different, they send the difference to an amplifier that adjusts the duty cycle of the switch. If the current is the right amount, the duty cycle stays the same. That way, the duty cycle is adjusted to give exactly the right amount of current to the LEDs.

The HV9910 works differently (it uses peak current control). It turns on the switch and then measures the current as it is increasing through the inductor. When it reaches a certain level it turns off the switch. The average current is thus less than the set point; the value of the inductance determines just how much less. The point is that this still regulates the LED current, just in a different way than the other two ICs. And this difference is what allows the HV9910 to work in this topology. Since it is just looking for a peak current value, it can measure the current through the (external) switch. The inductor current is ramping up when the switch is on, and it goes through the switch and through a current sense resistor. The current sense resistor is referenced to ground, and so the control signal is once again straightforward. When the current through the inductor reaches the preset maximum level, the IC turns the switch off. Then the current flows through the LEDs. When the current is going through the LEDs, it decreases, and we have actually measured the maximum current. Thus we are again regulating something related to the LED current, and the difference between this value and the real current is controlled by the inductance value.

The other drawback of this circuit configuration is that the voltage that is seen by the transistor and diode is higher than the input voltage. In fact, they see the input voltage plus the LED string voltage. But unless the string is very long, this is not very much voltage and so parts are readily available to handle it at no extra cost or space. The high voltage also means somewhat higher losses in these two components, but again this may not be practically that significant in real applications.

INPUT VOLTAGE LIMIT

For the three ICs used in this chapter, we've mentioned the input voltage, that is, the maximum voltage that can be applied to the V_{IN} pin. But how do you practically decide what voltage IC you need?

The first thing to remember is that the maximum input voltage for the IC is not necessarily the same as the maximum input voltage the circuit can handle. Indeed, they are usually different. Although our selection of the FAN5333A and the LM3405 might suggest otherwise, most controller ICs do not have integrated power devices. In this case, then, the external power devices can be very high voltage, while the IC runs off a low supply voltage. The reason to use this method is that the high-voltage IC process is expensive. It is much cheaper to make the IC low-voltage. In this case,

then, the choice of IC input voltage is determined by where the supply voltage comes from. If the DC input voltage is low, say below 40 V, then you can use an IC whose input will directly tolerate that voltage. If the DC input voltage is higher, then you probably need to use a low voltage IC, and generate the supply voltage inside the circuit. In this case, a 12 V IC might be a good choice, 12 V being a common, inexpensive process. (Why not use a 5 V IC? Because 5 V is not always enough to run the gate of the MOSFET, and so the IC plus the transistor cost may be higher, even though the IC by itself is cheaper.)

To generate the supply voltage for the IC, should this be necessary, the easiest thing to do is just put a resistor in series with a zener (don't forget a bypass capacitor on the V_{IN} pin to counter noise). For modern ICs drawing low V_{IN} current this won't dissipate too much power in the resistor. If for some reason you need more than a couple of milliamps supply current, then you probably have to do something more complicated, such as adding a winding to a transformer to make an internal power supply.

It is only when the switch is integrated inside the IC that the IC needs to be able to handle the full input voltage. But even then, economics determines that the V_{IN} pin will usually want a low voltage, and the high voltage is reserved just for the pins with the power devices attached to them. Thus for example, the FAN5333A has a 35 V maximum on its SW pin, but 6 V maximum on V_{IN} . The LM3405 has both V_{IN} and SW rated at 20 V, but this is because the buck transistor is attached to V_{IN} . The HV9910 is the exception, in that has its V_{IN} pin rated at 450 V. But even this then regulates it down to 7.5 V on its V_{DD} pin, which is used for the internal circuitry. The high input voltage rating on the V_{IN} pin was an intentional choice to avoid needing the external resistor and zener, and to avoid certain start-up issues. Otherwise, it is unnecessary, as the transistor is external to the IC.

DIMMING

In Chapter 2 we've already mentioned the two popular ways of dimming LED circuits, PWMing and analog. We can now look at these methods more closely, and see how to apply them to our circuits.

Because the optical characteristics of the LEDs are (at present) dependent on drive current, it is desirable to operate them at their rated current. The way to dim LEDs while holding them at a constant current is to turn them on and off periodically, as shown schematically in Figure 7.8. This is called pulse width modulation (PWM), the same as for controlling ICs. During their on-time, the LEDs are run at




Figure 7.8 Pulse Width Modulation Turns the Current Rapidly On and Off to Get an Average Current.

their rated current. The percentage of time they are on is called their duty cycle (abbreviated DC, and not to be confused with the power source type). Since the rest of the time during the cycle (or period) they are off, their *average* current is lower. In fact, the average current is their on-time current multiplied by the duty cycle.

Now to avoid flickering of the LED, the frequency of the PWM has to be above about 60 Hz. The ICs we've used in this chapter easily accomplish this. For example, the LM3405 can be PWMed up to about 5 KHz. Since it's running at 1.6 MHz, you might suppose that it could be modulated much faster. But what is actually happening is that the IC is being turned on and off, and it takes about 100 μ sec to turn on. This limits the PWM frequency, because full LED current won't be present until the IC is fully on. The same thing limits the FAN5333A to only about 1 KHz. If you use ICs that were not designed with this type of modulation in mind, it can take so long for the IC to turn on that the modulation frequency will have to be in the visible range. This sort of IC is then unsuitable for dimming LEDs.

The other generally applicable method of dimming LEDs is to reduce their drive current. Many applications will not be affected by color shifts induced by reduced drive current. As shown in Figure 7.9, a resistor divider from the current sense resistor to the feedback of the IC is all this method takes, although you may want a capacitor to control noise. Bear in mind how this circuit works. As the variable resistor R3 is decreased in value, the voltage presented to the IC is lessened. This makes it generate more current through the LEDs. So in this case, the minimum LED current is set by R3, with the divider set at maximum. As the divider goes to zero, the LED current will try to go to infinity, and so R4 is added in series with the divider to provide a maximum for the current.




Figure 7.9 Dimming Circuit.

BALLAST LIFETIME

In Chapter 5 we talked about LED lifetime and how it depends on temperature and other factors. But in a lighting system, the lifetime of the entire system needs to be considered. Assuming that the LED lifetime is adequately under control, then the drive circuit may actually determine the lifetime of the system. Formally, the lifetimes of the individual parts add like paralleled resistors. Thus if one part has a much lower lifetime than the other, the system's lifetime is controlled by that part.

$$\text{SystemLifetime} = \frac{1}{\frac{1}{\text{Part \#1Lifetime}} + \frac{1}{\text{Part \#2Lifetime}}}$$

The lifetime of a properly designed drive circuit is usually very long. ICs in particular can last decades if they are not subjected to excessive voltages. And the power components, MOSFETs and diodes, are equally robust under the same conditions (although this is different in AC drives, see the next chapter). The real limiting factor on the driver lifetime is the input capacitor.

In real DC drive circuits, the input voltage is always buffered by a capacitor. The FAN5333A circuit in this chapter doesn't show it, but it is necessary. The reason, as outlined in the battery section of this chapter, is the fairly high impedance of batteries at typical switching frequency. *How much* capacitance is needed depends on the power level of the driver and the switching frequency. You don't want the input voltage to drop very much during the time when the switching transistor is on. So higher input current means you need more capacitance. Conversely, switching at a higher frequency means that the current is drawn for less time, and so you need less capacitance.

The problem for lifetime is the electrolytic capacitors. The factors to watch out for are the rated lifetime of the electrolytic capacitor, its rated temperature, rated voltage, and rated ripple current. For example, the HV9910 circuit shows two 10 μ F capacitors at the input. Let's suppose we use 35V electrolytic capacitors rated at 1000 hours at 85°C. These are probably the cheapest ones we can get. Now if the drive circuit is encased so that the air temperature runs at 65°C, then the capacitor is being run 20°C below its rating. For each 10°C drop below rating, the capacitor lifetime approximately doubles. So here, the lifetime of each capacitor is about 1000 hours * 2^(20°C/10°C) = 4000 hours, and the two capacitors in parallel are thus about 2000 hours. This is only about 3 months of continuous operation!

To increase the lifetime, we could switch to a 2000 hour 105°C rated capacitor. The lifetime is now 2000 hours * 2^(40°C/10°C) = 32,000 hours, and the two in parallel are thus about 16,000 hours. This is two years of continuous operation, and may be satisfactory for some applications. Of course, the increased specification of the capacitor comes at an increased cost as well. You can also get some additional lifetime by raising the voltage of the capacitor. For example, a 50V capacitor in a circuit whose average voltage is 30V will have 5/3 more lifetime. Of course, you still have

to include the lifetime of the LEDs and other drive components to calculate the real lifetime of the system.

The other factor that sometimes has to be considered is the ripple current rating of the capacitor. The rated life of the part is also specified at a certain ripple current. For example, the 10 μ F, 35 V electrolytic above has a rating of 1000 hours at 85°C at 36 mA. As temperature decreases the ripple current rating goes up, but only slightly. At 65°C, the ripple is rated at 45 mA. So generally, you don't get much of a bonus for running at lower ripple current than the rating. But you do have to make sure that the ripple current that the capacitor sees is not more than what it's rated for. The best way to check this is to actually put it in the circuit, and measure it with a current probe and an RMS meter rated for the frequency of operation.

A better choice for this and many LED DC drive circuits is to simply avoid the use of electrolytic capacitors entirely. For example, a 10 μ F, 35 V ceramic capacitor is available in a 1206 package for a 2000 piece price of 6¢. In fact, a 10 μ F, 35 V, 2000 hour at 105°C electrolytic runs 8¢ in a 5 × 6 mm package, and so the ceramic is a better choice anyway.

ARRAYS

Most LED drive circuits use a single series string of LEDs. Since all of the LED forward voltages are in series, this sets the output voltage of the circuit and together with the input voltage determines the topology. But there is a downside to having all of the LEDs in series. If any one of them fails to open, then the entire string is open and there is no light at all. This is evident from the calculation of the lifetime of the string. Since all of the lifetimes are in parallel, the lifetime of the string is the lifetime of an individual LED divided by the number of LEDs.

An alternative is to put multiple strings in parallel. If one LED fails, taking out its string, the other strings continue to produce light. Although the total light from the circuit with this failure may be reduced, the reduction may not be enough to call the system failed. For example, if there are 10 strings and one fails, there is a 10% reduction in light output. If total light is allowed to drop 30%, then this system still has a considerable amount of life left in it.

Simply putting multiple strings of LEDs in series doesn't work all that well. The problem is that the current sense resistor on the one string will generate a voltage mismatch with the other strings. A first pass at alleviating this is to put a resistor equal in value to the current sense resistor in series with each string, as shown in Figure 7.10. But even with 100 mV binning, a string of 5 LEDs can still have a variation of 500 mV from the next string. A way to deal with this is for each string to have a resistor in series that drops about 500 mV, as shown in Figure 7.11. The variation in voltage will then be partially taken up by the resistor, and the currents will be better balanced (confusingly, a resistor used this way is the original meaning of the word "ballast"). The resistor works better than the LED because a big change of current results in a big change of voltage across the resistor, unlike the LED where a big change of current is required to obtain just a small




Figure 7.10 The Effect of the Current Sense Resistor Is Compensated by Putting One in Series with Each String.




Figure 7.11 LED Forward Voltage Variation Can Be Compensated at the Cost of Additional Power.




Figure 7.12 Ballasting LED Strings with Total Current Sensing.

change in forward voltage. If the control IC requires less than 500mV, it can be divided down.

The problem with these schemes is that if the one string being used to sense the feedback voltage fails open, there is no feedback at all, and something unfortunate may happen to the circuit. You could instead add all of the currents into a single sense resistor, as in Figure 7.12. That way, if any of the strings fail open, the rest will remain powered. Of course, there is also additional power loss in this circuit, namely the string current times 500mV times the number of strings. Since this can be several watts, arrays are less common in practice than a single string of LEDs.

Chapter 8

Practical AC Drive Circuitry for LEDs

The previous chapter addressed DC drive circuitry for LEDs. This one is about AC drive. When your LED light needs to run off the electrical grid, AC drive circuitry is used. AC power conversion is considerably more complicated than DC. Not only is it technically harder, there are also government regulations concerning EMI emissions. This chapter will cover in depth what you need to know to be successful. The very first thing you have to know about is your *safety*.

SAFETY

Did your parents tell you not to stick that paperclip into the wall outlet? Or about not running a toaster in the bathtub? Your parents were on the right track. AC is potentially deadly. Before you do anything with AC power, you need to read this section and take it to heart.

One of the authors once watched as an engineer working on a 277 VAC lighting ballast accidentally touched the input line. The engineer was knocked backward off his chair, and lay twitching on the floor for seconds. This guy was lucky. He didn't remain connected to the power line, and this probably saved his life. The next day he was back at work. However, there's no reason to count on your being equally lucky! (As an aside, you might wonder what *we* were doing. At first everyone was so surprised that we just stared. Then we ran around turning off the power. The authors strongly recommend that you get a "kill switch" installed in your lab before you work on AC power. This is a big red button near the doorway that you activate by hitting it with your fist. It should turn off all electrical outlets in the lab, but not the light. That way if there is an accident, you don't have to scramble around trying to figure out where the power switches are.)

The most serious danger from AC power is fibrillation, that is, your heart muscles twitching uncontrollably. Suppose you touch one side of the AC line with one hand, and the other side with the other hand. Now there is a path for current to

flow through your body. And since your heart is in that path, it may receive enough of a shock to stop working.

Another path is through your feet to the ground. If you touch the AC line with your hand, current may flow through your body and your shoes to earth. This can also potentially stop your heart.

It is worthwhile noting that Underwriters' Laboratories (UL) has done tests on AC safety. They have determined that even 5 mA is enough to create an unpleasant shock. At their safety class, they let you try out 5 mA, and then 10 mA and finally 20 mA—apparently very few people dare go to 20 mA after feeling the 10 mA shock! At 120 VAC, even a 24 K Ω resistor between you and line voltage is not enough to protect you.

Here is a series of recommendations to avoid getting into trouble:

1. Switch off the AC line before doing anything on a circuit. Suppose you have a probe looking at the source of a transistor, and you want to look at the drain. Turn off the AC power first, then move the probe and then turn the circuit back on. Do this even with low voltage points in the circuit such as the source. You never know if you're going to accidentally bump the high voltage. This seems like a nuisance in the lab, but your safety should be important to you.
2. Don't take someone else's word that power is off. Sometimes you will be working with someone else on an AC circuit. You've been debugging, and he tells you to take a look at something. Stop! You need to personally look and see that the input voltage monitor reads zero or, better yet, that the plug is unplugged.
3. Use an isolation transformer. An isolation transformer is a big, heavy transformer made of steel. It provides isolation between the input AC line and its output. The output is usually 120 VAC, but sometimes it is double this. The important part is that if for some reason the output is shorted, its power is limited by saturation of the core. You can't get more than 50 W, or whatever its rating is, out of it. A 50 VA device, suitable for much LED work, weighs three or four pounds. If it is much lighter than this, it's probably not the right device.
4. Use a fuse in your circuit. While you can still get a lot of power through a fuse, it's time limited. You won't stay connected to the AC line forever. Of course, only components downstream of the fuse are protected.
5. Keep one hand behind your back if you have to touch a live circuit. Since the worst danger is current through your heart, if one hand is behind your back your body won't create a circuit from one arm, across your chest, and out through the other arm. It may be also convenient to grab hold of your belt.
6. No cheaters. We're not talking about copying someone else's design. A cheater is an AC adapter that converts a three prong plug to two prongs. They may give the impression of offering isolation, but they don't. That third prong is there for a reason, to give you a ground connection. Discarding it may allow the power to float up, presenting a safety hazard. We've also

seen an oscilloscope with a cheater have its probes literally burn when they were attached to power. Our suggestion: Ban cheaters from the lab.

7. Use insulation. This is the simplest and most obvious safety measure. After you finish wiring something with AC, wrap it up in electrical tape. (Ordinary masking or cellophane tape isn't good enough; it needs to withstand high voltage.) If you're going to leave a probe on a point for a while, wrap the tip of it up in tape too.
8. Put up a sign. Nobody except you really knows if a circuit is live or not. Put up a big "Danger! Live 120 VAC!" sign so that people coming through know to not touch.
9. Turn it off. If you're going out of the area for a bit, to the restroom or even to get a component, turn off the circuit. Maybe the vice-president of the company will come by showing a customer the work you're doing. Electrocuting the customer is bad for business.
10. Cordon off the area. If you need to leave an AC circuit on overnight, make sure that it is very safe. The cleaning crew at night needs to be safe too—try adding skull and crossbones to your sign, and a language or two other than English. Put up a plastic chain-link fence that someone has to actively push aside to get to the circuit.

Even with the best advice, there's no such thing as complete safety with AC power. Paying attention to what you and others are doing is your best safeguard. As a final note, you might wonder about that toaster. The reality is that in most modern bathrooms, there is a GFI (ground fault interrupter). This detects if there's any current going to ground rather than going to neutral. If there is, then it shuts off the AC power. So you're actually probably safe to make toast in the bathtub, if the house is new enough to have a GFI—but it's a good idea not to try it anyway.

WHICH AC?

We say "the AC line," but this is shorthand for a very complex electrical environment. To design in this environment, you have to know what it's like. To start with, what voltage is the AC line? If you're in the U.S., you probably will say 120 VAC. But in the rest of the world, it's usually 240 VAC. And 120 VAC isn't necessarily the right answer even in the United States. There's also 208 VAC, usually used in homes for dishwashers and the like, and 480 VAC used in industrial settings. While these two probably won't concern you as a lighting designer, there's also 277 VAC which is commonly used to power fluorescent ballast fixtures. This will be important if you're designing LED replacements for this type of lighting.

The 120 VAC line isn't really 120 VAC. That is its mean value. In fact, the power company usually (but not always) guarantees that the line voltage will be 120 VAC \pm 10%, which is 108–132 VAC. This still isn't good enough for a design. If your customer is at the end of a long transmission line, the range will be more like 85–135 VAC.

This still isn't good enough. Sometimes there are brownouts. For one reason or another, the line voltage goes lower than 85 VAC. If it sits there for any length of time, motors tend to burn out, but your lighting shouldn't. After all, incandescent light bulbs are just resistors and they don't fail with low voltage, they just become dim. So at the least, your LED circuit should be able to survive any line voltage from 0–135 VAC indefinitely. It doesn't have to work properly at voltages outside the normal range, but at least it shouldn't fail. Testing your light by letting it sit at 10 V intervals for an hour each from 0–135 VAC should be a required design test for your design.

Designing for the range of 0–135 VAC covers you for the U.S. market. But in the rest of the world, 240 VAC is widespread. And this has the same tolerance, so you need to design to 265 VAC. Incandescent light bulbs are country-specific, so you don't necessarily have to design for both 120 VAC and 240 VAC, normally called universal input. If you do, then you should have no problems, at least as long as the circuit produces the right light at both line voltages. If you don't design for universal input, you face a problem in Japan. There, half the country is on 240 VAC, the other half is 100 VAC (there are historical reasons for this).

In summary, deciding on the input voltage range of the driver is not just a technical decision. It also has market impact. The engineering and marketing departments have to decide:

1. Will there be universal input, which costs extra money for each unit made? Or country-specific input, which increases the number of different designs for both engineering and production?
2. Will the unit produce the same light over some range of input voltages, and, if so, what is that range? At what voltage will it turn off? Or will it imitate incandescent light bulbs and have light dependent on input voltage? If the latter, will it be linearly dependent, or quadratic? (Remember that an incandescent light bulb is a resistor, so its power is represented as V^2/R , quadratic in the input voltage to first approximation.)

There are yet more things to consider about the AC line. What is the voltage when you're connected to 120 VAC? (Yes, this is a trick question.) Normally, you can assume that the voltage that the power company provides is sinusoidal. In this case, the voltage varies each half-cycle from 0 V to $120\sqrt{2} = 168$ V, the peak value of the 120 Vrms line. But the factor of $\sqrt{2}$ is right only if the voltage is perfectly sinusoidal. In reality, there may be heavy loads on the line that change this.

How can the AC line be affected? Remember that what the power company generates is far away at a substation. It then goes through some long wires that may include other things, such as big motors. These long wires have impedance, and so their currents can change the waveform. Even inside a house, the wiring gives nonzero source impedance. For example, if you pull a lot of current (for example turning on the microwave) you get some drop in voltage. If you momentarily feed a lot of current back (turning off the motor in your vacuum) you get some surge in voltage. The 120 Vrms line can have both sags and spikes on it.

Then there is lightning. When lightning strikes near power lines, it can induce high voltages and currents on the AC line. This surge can produce voltage as high as 6000V or current as high as 3000A inside a house! The limits are set by arcing in the house wiring. Fortunately, these enormous voltages and currents are present only for very short times. We'll cover how to protect against lightning a little further on.

Before we close this section, let's briefly mention AC line frequency. In the United States, AC voltage runs at 60Hz; elsewhere it is frequently 50Hz. These frequencies are fairly precise, since they are related to the speed of the generator. Still, the precision in this measurement is for the long-term average frequency, not in an isolated moment. For most lighting applications this won't matter, but it is traditional to design AC circuitry to work from 47 to 63 Hz.

RECTIFICATION

A block diagram of an SMPS for running LEDs offline is shown in Figure 8.1. The first step in using AC power (aside from EMI and surge protection, which we will cover below) is rectification. Rectification converts the sinusoidal AC line to DC with some ripple (see Figure 8.2 for a bridge rectifier). When hot is positive with respect to neutral, current flows through diode D1, through the load (here a resistor), to ground, then through diode D3 from ground back to the neutral. When it's negative, neutral is positive with respect to hot. Then the current flows through D2, *through the load in the same direction*, through D4 and back to hot. Since current flows through the load in the same direction regardless of the polarity of the AC line, this is a DC supply. Of course, the actual voltage the load sees still runs from zero to peak each half-cycle.




Figure 8.1 Block Diagram of AC SMPS for LED Lighting.




Figure 8.2 A Bridge Rectifier.




Figure 8.3 Half-Wave Rectification.

Now there are a few options with the rectification. The one shown in Figure 8.2 is by far the most common. Note that the diodes are regular rectifiers, not fast or schottky diodes. At 60Hz, the recovery time of the diodes is unimportant (actually, regular diodes are better than fast because of EMI). And at 120 V, a 2 V forward drop (a pair of diodes conducting at the same time) is not a very significant power loss. The real choice is between discrete diodes and a bridge in a four pin package. This choice is usually based on cost and space, the integrated bridge having smaller footprint but costing a bit more.

There are two more choices for rectification that might be possible for LED lighting. For high power designs (in the multi-kilowatt range), the loss in the bridge, while a small percentage of the total power, may still be too large as a number of watts. Sometimes, MOSFETs are used to replace the diodes. When they're on, their loss is I^2R , which may be lower than $I \cdot V$, and when they're off they don't conduct. To accomplish this, the drain acts as the cathode, and the source as the anode. But driving the gate properly can be a little tricky. Since this high power supply is expected to be unusual for LED drives, we won't get into details.

The other possible choice is half-wave rectification (the bridge in Figure 8.2 does full-wave rectification). See Figure 8.3 for how this works. When hot is positive with respect to the neutral, it works just the same as does full-wave rectification described above. Current goes through diode D1, through the load to ground and then through diode D3 back to the neutral. But during the part of the cycle where the hot is negative with respect to neutral, there is no conduction. There is no output from the bridge. This system works in that you get DC. Current flows through the load in only one direction. But it is usually only applied at very low power, since it makes a mess of the AC line. It also may require more capacitance to support the load during the nonconduction time. And after all, for 2¢ more you can get the two extra diodes.

Now that we've rectified the AC line, the next question is whether we like the variation in voltage. Just rectifying the 120 VAC produces an output that goes between 0V and 168 V every 8.3 msec (1/60 Hz, and twice each cycle). This is undesirable because of visual flicker and potential effects on LED lifetime. The magnitude of the variation could be reduced by adding a capacitor, as shown in Figure 8.4. How much capacitance is needed? We can determine this by looking at energy stored in the capacitor.

As a side note, whether or not the light is flickering depends on who you ask. Some people can see 60Hz flashes from fluorescent tubes. Others won't notice anything even at 50Hz. What about the rectified 120Hz? It turns out that in the real world rectified AC is slightly asymmetric, causing a 60Hz modulation on the 120Hz




Figure 8.4 Reducing the Ripple from a Bridge Rectifier with a Capacitor.

flicker. So the solution is to reduce the amplitude, making the flicker even less noticeable.

Energy stored in a capacitor is $\frac{1}{2}CV^2$. Now at the peak of the line cycle, the capacitor will be charged up to this peak, 168V minus two diode drops, 166V. So the energy stored will be $(13.8 * C)$ in mJ when C is given in μF . Suppose that you want the minimum voltage to be 10% less, or 149V. The energy stored in this condition is $(11.1 * C)$ mJ. So the energy taken from the capacitor is $13.8C - 11.1C = (2.7 * C)$ in mJ when C is in μF .

Now we do a trick. If you look at the waveform in Figure 8.4, you see that after the peak is reached, the bridge input goes back down to zero, and then comes back up to that 149V point. But how long that takes involves the sine function. Rather than do the math, we will say that the capacitor has to provide power during the whole time from one peak to the next, 8.3msec. This is a good approximation for a capacitor whose voltage doesn't vary too much, and not a bad one even if it does. It's not off by more than a factor of two in the worst case.

How much energy is needed? If your supply draws constant power it's easy. Suppose you have a 4W input power. Then in 8.3msec you need $4\text{W} * 8.3\text{msec} = 33\text{mJ}$. We needed $(2.7 * C)$ in mJ, and so we need $33\text{mJ}/(2.7\text{mJ} * C) = 12.2\mu\text{F}$. We round it up to $15\mu\text{F}$ to account for tolerances.

That isn't all that much capacitance. There are problems with power factor doing this (see the section below), but the real problem is the voltage. You need a 200V capacitor at $15\mu\text{F}$. Your only choice is going to be an aluminum electrolytic capacitor. This capacitor at this voltage is relatively large (such as $21 \times 38\text{mm}$). And for a 240VAC system, a suitably rated capacitor may not even be available. Below, we'll give some guidelines on how to deal with this. As we'll also see below, this type of capacitor may also determine the lifetime of your system.

Now there's one more choice for a drive. Certain manufacturers, notably Seoul Semiconductor, make "AC LEDs." These are LED devices in a single package that have a forward voltage of something like line voltage in both directions. They are intended to be run directly from the AC line, without any other components, not even the bridge in Figure 8.2. Now, one could imagine there are going to be some problems. What happens when the line voltage increases? The current of course goes up, and the manufacturer says that's acceptable. But you still have to dissipate all that extra heat somehow. And what happens when there's a line surge or lightning? It would seem you need some protection circuitry in front. And then there's EMI. We haven't talked about it yet (see below), but fast diodes such as LEDs can generate a lot of electrical noise. Perhaps an EMI filter will be necessary. The authors

think there's more work to be done on this structure before it's ready for prime time. But if these issues can be worked out, then the AC drive circuitry discussed may well disappear. Ultimately, it's going to come down to cost.

TOPOLOGY SELECTION

Now that we have rectified the AC, we need to do something with it. The general structure (the “switching” block in Figure 8.1) of the power supply is referred to as its topology. There are a number of possibilities. Let us start by understanding the pitfalls of the simpler, nonswitching topologies.

Analogously to what we mentioned above, the easiest thing that you could do would be to simply attach a suitable number of LEDs in series to get a forward voltage approximately equal to the AC line voltage (see Figure 8.5). This is certainly very cheap. But it has the same practical problems as the integrated version. The line voltage goes from 0V up to 168V and back down to 0V 120 times per second. While this isn't visible to most of us, some people can see it (especially if it's 100 times per second rather than 120), and it's very hard on the LEDs. This variation can be fixed with a capacitor, as we discussed above. But worse yet, when the line voltage goes up 10% to 132V, there is a LOT more current into the LEDs. Remember that they have current that is exponential in the voltage, much faster than the quadratic rise of incandescent light bulbs. And when the line voltage goes down 10% to 108V, they're going to be quite dim. In short, while series connection of the LEDs may possibly be acceptable when you have a fixed DC voltage, the AC line varies too much for it to be practical in real off-line applications.

Definition of Ballast

The word “ballast” is used in two different ways. Traditionally, a ballast is a resistor used in the emitter of a transistor to compensate variations in voltage that affect gain. From there, it was straightforward to use the word ballast to signify a resistor in series with a string of LEDs. However, ballast is also the word used to mean the circuit that drives a fluorescent tube. This is because the original circuit used a capacitor to limit the amount of current into the tube. When electronic circuits for fluorescents became available, they were also called ballasts. In this book, we often call circuits for running LEDs “ballasts” also. But you should be aware that some people take exception to this, and call such circuits “drivers.” For our purposes there are no differences between the two terms.

A slightly more complex choice is to linearly ballast one or more strings of LEDs with resistors. Imagine putting a resistor in series with the LED string in Figure 8.5. Now this ballast resistor is supposed to take up the variation in the line voltage. If the voltage becomes too high, the current through the LEDs goes up. This causes the voltage across the resistor to also increase. Thus the resistor absorbs part of the increase in input voltage. But this comes at a cost. Suppose normal current is 40mA, which is $120\text{V} * 40\text{mA} = 5\text{W}$ into the LEDs. If you want to be able to take




Figure 8.5 Running LEDs Directly Off-Line.

up 10% = 12 V on the resistor, the resistor value should be $12\text{ V}/40\text{ mA} = 300\Omega$. In normal operation this resistor dissipates $(12\text{ V})^2/300\Omega = 480\text{ mW}$. But now suppose the voltage rises to 132 V. The LED voltage doesn't change much with current, so almost the whole increase in voltage goes across the resistor. Power in the resistor is now $(12\text{ V} + 12\text{ V})^2/300\Omega = 1.92\text{ W}$. Since the input voltage is capable of sitting at this voltage indefinitely, you need a 5 W resistor. This is a 1.5 cm \times 1.6 cm package. And you need to be able to dissipate an extra 2 W thermally. At an overall electrical efficiency of $(7\text{ W} - 2\text{ W})/7\text{ W} = 71\%$, a resistor is not a practical solution.

A better way to deal with line voltage variation is to use a linear regulator. You could imagine a voltage regulator that ran a string of LEDs at the same voltage regardless of input line voltage. You could even notice that as the LEDs warm up, their forward voltage and efficacy both decrease. So with a cleverly designed voltage regulator, this system would produce constant light output: As the forward voltage drops, the current increases, compensating for the efficacy decrease. Unfortunately, high voltage input regulators don't currently exist.

The other way to run LEDs with a regulator is to use one with a current output, producing the same current into the LEDs regardless of input voltage. There are such devices on the market. You could even make an additionally complex one that compensates for forward voltage drop by increasing the current.

The problem with both of these linear regulator concepts is that they still dissipate too much power. Any increase in line voltage is dropped across the linear regulator. So at 40 mA and 108 V across the LEDs, a 120 V input loses $(120\text{ V} - 108\text{ V}) * 40\text{ mA} = 500\text{ mW}$ in the regulator, and when line goes up to 132 V, 1 W. So the regulator has to be pretty strong to handle so much heat dissipation.

Since these linear schemes are out for most applications, we will use a switch-mode power supply. We have available the various topologies mentioned in Chapter 7 on DC drives, although there will be additional choices once we've added a transformer for isolation. Here we'll talk about the nonisolated types of converters for off-line use, and defer the question of isolation to the next section.

As in the previous chapter, we will be using a switch-mode power supply that regulates current. The easiest of these to work with offline is the buck type, as epitomized by the 9910 IC (parts by ST and others work much the same way). As you recall from the previous chapter, the buck is characterized by producing voltages that are strictly less than or equal to the input voltage. Since the input here is 168 V, and LEDs are only a few volts each, the buck is almost always a good choice.




Figure 8.6 How the Off-Line Buck Works.

Take a look at Figure 8.6 to understand conceptually how the off-line buck works. When MOSFET Q1 turns on, inductor L1 is effectively grounded. So the input voltage drops across the series LEDs, and then is dropped across the inductor, which starts to increase its current. Of course, since they're in series, the inductor current is the same as the LED current. When the current reaches a specified level (set by the IC and sensed by R1), the MOSFET turns off. Now the current in the inductor can't go through the MOSFET, and it has to go somewhere, so it goes through the diode D1, back to the input. The voltage on the end of the inductor that used to be grounded is now close to the input voltage, while the other end of the inductor remains lower by an amount equal to the LED voltage. So in this part of the cycle the inductor current and the LED current both decrease.

By controlling the point at which the MOSFET turns off, the current can be maintained at approximately a desired value. And if the inductor is relatively large, the minimum current will be close to the maximum current, so that an approximately constant current runs through the LEDs. This circuit also has the convenient feature that both the diode and the MOSFET never see a voltage higher than input voltage, just like other buck converters. This makes component selection relatively easy.

We might note that you could use the 9910 to generate an off-line buck-boost, just as in Chapter 7. However, there would be no point in this, unless you had an exceptionally long string of LEDs. Then a buck could only be driving them during the fraction of the time when the line voltage was higher than the string voltage. For most applications, the LED string voltage will be a fraction of the peak line voltage, and the buck version will work fine.

NONISOLATED CIRCUITRY

A real circuit using the 9910 is shown in Figure 8.7 (minus the input filter). We will be using this circuit in a BR40 design in Chapter 10. For the time being we'll do a quick overview. The components in the schematic that haven't been mentioned yet are related to control of the converter. Resistor R2 sets the current. When the voltage at R2 reaches 250mV, the MOSFET Q1 is turned off. Turn off is thus at $250\text{mV}/330\text{m}\Omega = 758\text{mA}$. Capacitors C1 and C2 are bypass for the switching (they both experience line voltage and must be rated for 200V). C3 is bypass for the internal linear regulator, which is 7.5V.

Resistor R1 sets the switching frequency of the converter. A higher frequency means a smaller inductor. So generally, higher is better. But a higher switching frequency may mean it's harder to filter the EMI (see below). A practical compromise between these conflicting desires is to pick a switching frequency just below 150KHz. Since EMI requirements start at 150KHz, this puts the main switching power where it doesn't need to be filtered. R1 at 249K Ω produces a switching frequency at 90KHz.

Now look at the power portion of the circuit. Transistor Q1 switches on and off at 90KHz under the control of U1. Again, U1 is controlling the peak current measured through R2. When Q1 is on, the current flows through the three LEDs, through the inductor, and then through the transistor and R2 to ground. Since there is a positive voltage across L1, its current increases during this time. When it has increased enough to reach the threshold set by R2, the transistor is turned off. Now the current through the LEDs and L1 is shunted through D1 back to the input. Since the voltage across L1 is now negative (one side is down three LEDs and the other side is up

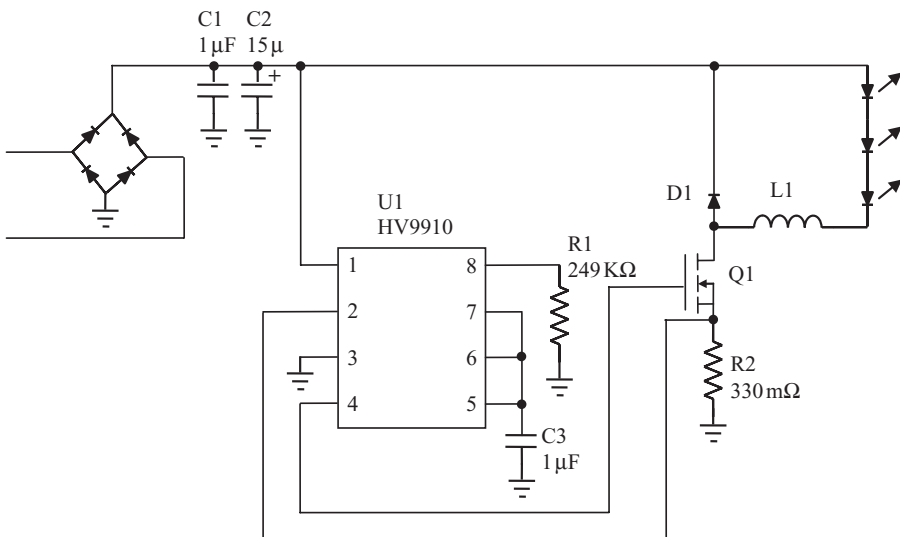


Figure 8.7 A Nonisolated Off-Line LED Driver.

one diode), the current decreases. Then after 11 μsec (one over 90 KHz), the whole cycle starts again.

This is an extremely simple power supply, and that is its advantage. There's no compensation circuitry design necessary, there's not much in the way of magnetic design and there are only eight pins on the IC. If a nonisolated supply is acceptable for your purposes, power supply design doesn't get any easier than this!

ISOLATION

You'll recall from this chapter's introductory section on safety that isolating the system from AC line with an isolation transformer is one of the important safety techniques to be used in the lab. Isolation is also very useful for products for the same reason, though it's not done the same way. A customer must not accidentally touch something that is connected to live AC. And Underwriters' Laboratories (UL) standards require isolated power supplies in many cases. Some further information on these requirements is in the section below on UL. (Note that we are using "UL" to refer to generic safety standards. Underwriters' Laboratories writes these standards with input from a variety of parties, but there are a number of companies that will test devices to these standards, UL being only one of them.)

The sort of isolation transformer you use in the lab is large. A 50 VA unit can weigh 5 pounds. Clearly this will not work for a lighting product. The reason that the isolation transformer is so big is because it is transforming 60 Hz (or 50 Hz). The low frequency means that the voltage is positive for a long time before it becomes negative. The inductor therefore has to be large to support all of those volt-seconds.

In an SMPS, we can employ a trick. Since it is switching at high frequency, at least 20 KHz and often ten to hundred times higher than that, the volt-seconds requirement for the transformer is proportionately reduced. You can transfer substantial power through an isolation transformer that is not substantially bigger than some capacitors. You simply put the transformer in place of the inductor that is there anyway, and separate off the load from the primary. With a properly designed transformer, everything on the secondary side is then UL safe.

That said, though, you should be aware that many designs use nonisolated converters. Nonisolated converters are both cheaper and easier to design than isolated ones. In a real product, after all, the converter is going to be inside an electrically insulated metal can or a plastic or glass enclosure of some kind. The customer shouldn't be able to access the inside of that. So then the only AC live parts that are potentially exposed are the LEDs and their PCB traces. If these are securely doubly insulated, then a nonisolated supply may be acceptable. Again, we recommend you follow UL standards and apply for UL certification.

Now how do we insert a transformer into Figure 8.7 in place of the inductor? There are two choices, shown in Figure 8.8 and Figure 8.9. Looking first at the schematic in Figure 8.8, we see that it is a forward converter.¹ You can tell that it is related to a buck because when the transistor turns on power is transferred to the

¹ A forward is a buck-derived converter. For more information on this converter, consult Lenk (1998).




Figure 8.8 Adding a Transformer Makes the Converter into a Forward.




Figure 8.9 Adding a Transformer Makes the Converter into a Flyback.

load, the same as a buck. It can therefore be controlled by the same IC that controls a buck, and so might be a good choice. But the downside is that whereas the buck had just one diode, this converter requires two diodes and an inductor on the secondary side. At the high currents required by LEDs, this inductor will be large, and so the forward converter is probably not the right choice.

Turning to the schematic of Figure 8.9, we have a flyback converter.² You can tell that it is related to a boost because when the transistor turns off, power is transferred to the load, the same as a boost. The difference between the flyback and the forward transformer is just the polarity of the transformer, signified by the dot. This converter has the same number of components as the boost, there is no extra inductor and only one diode on the secondary, and so is a good choice. Normally, special measures would have to be taken to have an IC control the flyback. But in this particular case we are controlling the peak current cycle-by-cycle. It will turn out that in this case, the same IC that we used for the original buck can also be used for this flyback, without raising concerns about loop stability.

Having chosen a flyback topology for our off-line converter, what we haven't yet said is that the next step is to decide whether this should be a continuous or

²This is a boost-derived converter.

discontinuous mode flyback (see Lenk 1998 for more information). There are pros and cons to both.

1. **Transformer size.** In a discontinuous mode flyback, the transformer acts as an inductor during the transistor on-time. All of the energy for the cycle is stored in the primary inductance. This results in a relatively large transformer. In a continuous mode flyback, however, the transformer is just a transformer. Therefore it is rather smaller than in discontinuous mode.
2. **Peak LED Current.** In the discontinuous mode flyback, no energy is going in to the capacitor while the transistor is on. Thus to get the same average current, the current, when it does flow, must be higher. This means that the capacitor must be adequately sized for the ripple current, which may add to the size problem. The continuous mode flyback has current going through the capacitor all the time and so peak and average currents are almost identical.
3. **Control Signal Isolation.** The LEDs are on the other side of the isolation barrier from the controller. Somehow the IC has to detect how much current the LEDs are getting, so the gate drive is turned off at the right time. Now in a discontinuous flyback, we get a lucky break. When the transistor is on, all of the energy is stored in the primary side of the transformer, which acts as an inductor. When the transistor turns off all of the stored energy goes to the secondary side, so that, on average, control of the turn-off point of the transistor determines the LED current.

The continuous mode flyback is more complex. There is always current coming out through the secondary of the transformer, and information on the amount of current is available only during the time the transistor is on. In short, the IC isn't taking into account the duty cycle. Information about the LED current has to be transferred across the isolation barrier. This requires additional circuitry, usually fairly complex, expensive, and prone to failure circuitry. For example, the common way to do secondary side current sensing is to use a 431 voltage reference IC as a comparison for the current sense voltage, and then use it to control an opto-isolator. The opto-isolator then controls the feedback of the IC. But the opto-isolator has a number of potential problems, including poorly controlled gain and aging.

Given this last problem with continuous-mode flyback converters, we usually choose a discontinuous mode flyback for our isolated converters. These have some limits on how much power they can easily transfer, but for a 120VAC line this is well over 50W. If you're going to build a light that uses more power than that for the LEDs, you probably can afford the size and cost of a larger ballast anyway.

COMPONENT SELECTION

One issue that hasn't been addressed in the previous two sections is how to select the values of the components. Some of it is obvious. If the LEDs are conducting 400 mA, you don't want to use a 500 mA diode, even though it would probably work.

You want to use a higher current diode, to take advantage of a lower forward voltage. Similarly for the MOSFET, you want to use one that has relatively low $R_{DS,on}$ so that your losses aren't too great. But what about the breakdown voltage, both for the MOSFET and diode, and for other components such as the input capacitor and the IC?

As a general rule, you want to obtain the highest voltage part you can afford. The peak voltage on a 120 VAC line running 10% high is 187 V, so a 400 V device gives plenty of margin. What do you need margin for? The input filter will include an MOV for lightning protection. But the MOV isn't like a zener diode that clamps at exactly its rated voltage. Its voltage increases substantially when lightning current is added to it. For example, the ROV07-241K sits at 240 V at 1 mA (where it is already dissipating 1/4W), but at 10A it's up to 395 V! Parts that are being protected by the MOV need quite a bit of voltage withstand capability above the rated MOV voltage. A 400 V part is adequate on a 120 VAC line. But 240 VAC is within normal bounds at a peak voltage of 373 V. Even 600 V is a little too close here; 1000 V would be desirable.

Looking first at the diodes (both the rectifier and the switching diode), there's very little difference in price between 400 V and 600 V. So you should choose at least 600 V. But the reality is that in production volumes, diodes are one of the least expensive components. So generally, 1000 V parts are the right choice. This will then cover you adequately both in the U.S. and around the world.

The MOSFET is somewhat more challenging. 600 V devices are common, and should always be used. This is high enough for use on 120 VAC. But that doesn't leave all that much margin for lightning let-through from the MOV on a 240 VAC line. So you have to make a money decision with marketing involvement. Do you pay the extra money for a 1000 V MOSFET and have a single design? Or do you use these only for the non-U.S. market, and have two different designs, with the costs associated with two different product codes?

The IC really ought to be rated at least 600 V. Although some are, in fact the 9910 is rated at only 450 V. This is probably adequate for the 120 VAC market, but is clearly unacceptable for the 240 VAC market. A suggested practical fix is shown in Figure 8.10. The resistor drops voltage equal to the current draw of the IC multiplied by the resistance. When the voltage surges above the transorb voltage, the current it has to absorb is limited by the resistor. A bypass capacitor is included. Shown are some typical values. This scheme doesn't cost much (the resistor and transorb are cheap) but does consume some power in steady-state. And this is true for all the 240 VAC components in this section: 240 VAC circuits are going to be slightly less efficient than 120 VAC circuits. The reduced losses in the bridge are usually not high enough to affect this conclusion.

The most trying question is the voltage rating of the input capacitor. Aluminum electrolytic capacitors are almost nonexistent above 450 VDC. This is adequate for 120 VAC, but clearly unacceptable for 240 VAC. What can be done? It's well known that putting two capacitors in series halves the capacitance and doubles the voltage (opposite of a resistor). But capacitors have a tolerance which affects this division, typically $\pm 20\%$ for this type. If one of the two capacitors in series is 20% and the




Figure 8.10 Protecting the HV9910 from High Voltages.




Figure 8.11 Resistors Balance Voltages for Series Capacitors.

other 20% low, the one that is 20% low will take 20% more of the total voltage than the other one. This could exceed the rating of the component. The solution is shown in Figure 8.11. Adding large value resistors in parallel with the capacitors helps to balance out the AC current. This ensures that voltages divide approximately evenly between the two caps even when their values are somewhat imbalanced. Having a resistor in parallel with the input capacitor may be a good idea anyway, since it ensures that the circuit will eventually dissipate its charge. Otherwise, the bridge blocks discharge of the capacitor in one direction, and the LEDs have too low a leakage current to discharge it in the other direction.

Still, a better plan is to not use an electrolytic capacitor at all, if that is possible. This is good for both power factor and lifetime (see below for both). The only capacitor needed is 10–100nF of high-frequency capacitance for the high-frequency switching. Such capacitors are available at 630V rating in an 1812 package, and at 1000V in an 1825 package.

EMI

One of the things that make design of AC ballasts much harder than DC is electromagnetic interference (EMI). Governments require that devices that attach to the AC line not produce more than a specific amount of electrical noise. And switch-mode power supplies, since they switch at high frequency, generate a lot of noise.

There are two types of noise, and one of the types has two varieties. Your design has to meet all of the noise requirements for all of these. The first type of noise to




Figure 8.12 Normal Mode EMI Filtering for a Two-Wire Input.




Figure 8.13 Common Mode EMI Filtering Added for a Three-Wire Input.

consider is conducted noise. This is noise that comes out of your design and is conducted along the input power wires. It is typically caused by devices that turn on and off quickly and carry power, in particular the MOSFET and the power diode. Devices that switch quickly but don't carry much current typically aren't as important for conducted EMI, although it's wise to create the best possible design just in case.

There are two types of conducted noise, normal mode and common mode. They are typically treated differently. Normal mode noise is due to different current in the hot and neutral power wires. Common mode noise is due to the current in the hot and neutral wires with respect to the ground wire. Note that for lighting purposes, many devices will have only two wires. In this case, only normal mode noise is possible, and that is the only type that can need filtering.³ Normal mode filtering starts with a high-frequency bypass capacitor at the input to the switching stage (see Figure 8.12), and will probably include a series inductor as well. Common mode filtering will require, in addition, a balun and capacitors from both hot and neutral to ground (Figure 8.13). The balun can have normal-mode inductance as well as common-mode, so that a separate normal-mode inductor isn't required. There are UL limits on how great a value the two caps to ground can be.

The second type of noise is radiated. This is typically caused by current loops, that is, current that runs in one direction and then returns, with a physical separation between the two paths. An example of this is shown in Figure 8.14. The MOSFET pulls current from the capacitor through a trace, and the source of the MOSFET returns the current to ground through another trace which is separated from the first trace. Current goes in a loop as shown, and this acts as a little antenna, radiating electromagnetic noise. The solution is to minimize the area of such loops, and try

³We see numerous designs that have only two-wire input and yet utilize a more expensive common-mode filter. This is seriously misguided.




Figure 8.14 Current Loops May Cause EMI Problems; Reducing Loop Area Helps.

to return currents either on top or bottom of the forward path or, in the case of a single-layer board, next to the forward path.

One more concept to note is that a nonisolated supply should have an isolated heat sink so that customer contact with the AC can be avoided. This potentially worsens the thermal interface, but its EMI consequences are not as dire as some have made them out to be. In a power supply, the main radiators are the MOSFET and the power diode, which couple capacitively through the heat sink isolation to make the heat sink an EMI radiator. In an LED light, however, the components dissipating the most power are the LEDs. Since these by design don't have much ripple, they are not very important to EMI. And since the ballast is a comparatively small part of the power, it may not need to be heatsunk to the heat sink. The conclusion is again the same: an isolated supply may be easier to deal with.

It is beyond the scope of this book to describe how to get your AC converter to pass EMI. Detailed information on how this can be done can be found in Lenk (1998). Here are some pointers that will help you get there with a minimum of confusion:

1. Put all of the components close together on the PCB.
2. Avoid having large currents go in loops.
3. Don't pick components that switch faster than you need them to. For example, choose a normal rectifier for the bridge, not an ultrafast.
4. Have currents on traces that are as close as possible to their return traces, either above or below, or on a single layer board next to the return trace.
5. Bypass the input to the converter with a small ($<1\ \mu\text{F}$) capacitor. Small value capacitors work at higher frequencies than large capacitors—the electrolytic in parallel with this small capacitor doesn't help the EMI at all.

We'll make a brief comment on EMI standards in the United States versus those in the European Union (EU). These standards are fundamentally very similar. The only significant difference at this time is that the EU also has a limitation on power line harmonics, EN61000-3-2. Chances are that if your design passes EMI, it probably will pass this standard as well.

Finally, let's comment on selecting the switching frequency of your SMPS. The reason for including these comments is that SMPS switching is often the most important source of EMI in a system. Now the fundamental consideration is that

EMI conducted limits start at 150 KHz. Below that frequency, you can have as much noise power as you like (subject to limits on harmonic content). Above 150 KHz, the amount of noise allowed decreases at 30 dB/decade, which is faster than a one-pole filter but slower than a two-pole filter. So the most sensible choices are either to switch below 150 KHz or else be considerably above that frequency. Switching at 200 KHz, for example, gives you the worst of both worlds.

You will also have to pass radiated emissions. The limits for this start at 30 MHz, so that the emissions don't directly impact switching frequency selection. However, switching at high frequencies requires transitions between on and off in the power elements to be fast. If your switching frequency is 1 MHz, the on-time of the transistor is less than 1 μ sec, and so it needs to turn on and off in less than 30 nsec to avoid big losses. These transitions are within the regulated radiated bands. If instead your switching frequency is 100 KHz, the transition times can be 10 times slower.

From an EMI perspective, a switching frequency below 150 KHz is going to be best. With lower frequencies, the downside is that component sizes get bigger. This is true particularly of inductors and capacitors. You can do a lot of math to find out that there's no one best answer; there are too many factors to consider and optimize. Our practical recommendation is this: pick a switching frequency just below 150 KHz (including component tolerances) or above 1 MHz. If size isn't an issue, the lower of these two is probably the better choice, and will be more efficient anyway.

POWER FACTOR CORRECTION

Many devices above 50 W are now required to be power factor corrected (PFC). Additionally, new regulations (Energy Star in the United States, in particular) seem to require PFC for lighting devices regardless of power level. So your design may well need to be power factor corrected. What is PFC?

Consider the current input to a switch-mode power supply when there is a large capacitor right after the bridge (see Figure 8.15). When the line voltage reaches its peak value, the capacitor is also charged to its peak. Now the line voltage decreases. But since the capacitor is large relative to the amount of energy the power supply is pulling from it, it discharges more slowly than the line voltage drops. Consequently, during this time the line is not providing any current to the capacitor. Once the line has gone to zero and come back up, it reaches the same voltage as the capacitor, and begins to charge it again. During this short time, all of the current needed for the whole cycle is delivered to the capacitor. As a result, the current from the line has a big spike close to the peak of the line voltage. This spike is one of the primary causes of bad power factor.

Power factor is in a practical sense how close the input current matches the sinusoidal input voltage. What you want is to have input current also be sinusoidal, and in phase with the voltage. The load should "look like" a resistor. The current to the big capacitor clearly doesn't meet these criteria, and as a result has poor power factor.




Figure 8.15 A Big Capacitor Maintains Constant Voltage During the Line Cycle, Generating Large Peak Currents and Bad Power Factor.




Figure 8.16 A Smaller and Cheaper PFC.

What's wrong with a poor power factor? The reality is that it wouldn't hurt your design very much to have a bad PF (usually your input capacitor is rated to take the RMS current). It's rather a problem for the electric power company. When you have a load that doesn't look like a resistor, it must look like a capacitor (or inductor). These pull current without dissipating it. So the power company supplies you current, but you don't pay for it, since you return it on the next part of the line cycle, and this is costly to the company. You have power losses in the power lines because of this current, and the company needs additional generating capacity for something that isn't billable. To minimize these problems, governments mandate minimum PF for many devices.

There are a couple of methods to correct the power factor of your design. The traditional method is to add a second converter to the input of your power supply. In this type of design, the first stage is a boost converter that takes the input voltage and boosts it to a higher DC voltage, typically 400V. It typically uses a special control IC that draws current from the line in phase with the input voltage, so that it has nearly perfect PF. Then the second converter takes this 400VDC and converts it to the output, in this case an LED current. This two stage converter works well, but does add substantially to the cost and size of the design.

A smaller and less costly alternative is shown in Figure 8.16. It basically replaces the single input electrolytic capacitor with two capacitors and adds three additional diodes. It can probably work up to about 60W. It is therefore a reasonable alternative to the two-stage design for many LED applications.

One of the authors devised an even cheaper scheme to do PFC with the 9910. This is shown in Figure 8.17. The brightness control pin of the IC is connected through a resistor divider to the AC line input. As the AC line goes up and down, so does the input current. The top resistor of the divider is selected to give low power dissipation at 120VAC; the bottom resistor is chosen to give full scale of the control pin (250mV) at line peak (168V). This circuit gives a PF of >0.9, and has essentially




Figure 8.17 Simple Power Factor Correction Circuit.

no cost and takes up no space. No electrolytic input capacitor is required, resulting in a cost savings and improving the lifetime of the power supply. The downside is that a sinusoidal current is now going through the LEDs, which as discussed below causes some loss of efficacy and potentially decreases their lifetime.

LIGHTNING

Anything attached to the AC line will occasionally face lightning on the line. Before considering what to do about this (if anything) let's first look at the phenomenon. Lightning is a natural phenomenon, and so can vary wildly from event to event. IEEE587 (now ANSI/IEEE C62.41) defines the types of lightning events that are generally considered to be worst-case. If you can survive them, then you're probably in good shape for the real world. A good overview with waveforms is given by Martzloff (1991), a pioneer in the field.

The first thing to know is that the type of lightning surge that occurs depends on location. Outside a house or a building, the surges are much worse than inside. Inside, the maximum voltage you can see is limited by the wiring. If the voltage goes too high, the wires will arc in the circuit box. Since most LED designs are done for indoor lighting, we'll concentrate the discussion on Class B.

Inside a building, the arcing limits the maximum voltage to 6000V, and the wiring impedance limits the maximum current to 3000A. These don't both occur at the same time. The source looks approximately like 6kV with 2Ω impedance. This still sounds like it's impossible to protect against. What saves you is that the time during which this occurs is very short.

There are basically three types of lightning surges you will need to protect against. The 6kV, 3kA type has a rise-time of 1.5μsec, and a decay time of 50μsec (and is thus called a 1.5/50 waveform). Then there is a 10/1000 waveform but it has a peak voltage of only 1300V and an impedance of 250mΩ. Finally there is a ring waveform, which rings up to 6kV, 500A with a frequency of 100KHz, decaying away within a few cycles. There are other types of waveforms as well, but these three will practically cover your design.




Figure 8.18 Adding an MOV to the Design Protects It Moderately Well from Lightning.

What can be done to protect your 600V design against 6000V lightning? The best approach is to add an MOV (metal-oxide varistor) to the input, as shown in Figure 8.18. An MOV resembles a large, power bidirectional zener. You can pump huge amounts of current into it for a short period of time, and it clamps the voltage. Because of this it is a good device for protection against lightning.

However, there are several aspects of MOVs that makes them less than perfect. For one thing, they have soft “knees”. That is, their voltage is not absolutely fixed, but climbs up quite a bit as you increase the current. In practical terms, this means that you have to have a fair amount of margin between the turn-on voltage of the MOV and the voltage at which your converter will break. Furthermore, MOVs perform differently as they age. After a number of lightning events, their clamp voltage starts to rise. After enough events, they can explode, leaving your circuitry unprotected against the next lightning strike.

So the best you can do is to put in the biggest MOV you can. They are relatively inexpensive, so the determining factor is how much space your design has. For a 120VAC design, a practical MOV is the –241, rated for 150VAC and for 240VAC the –431, rated for 275VAC.

Now if we are honest, we will recognize that many electronics designs out there don’t bother with an MOV. Why don’t they just blow up the first time there’s a storm outside? Part of the answer is that they do—bottom-end electronic devices just fail after a while, and lightning *is* one of the reasons. Another consideration is that if there is a device on the circuit that is protected, it may protect other devices on that line as well. So some devices leach off others’ better design practices. We recommend that you always include lightning protection, as LED lighting is advertised to last for a long time. It’s bad if the LEDs are still working great, but the lamp is dead for the lack of a 10-cent part. And it’s poor design practice to hope that someone else is protecting you. Ultimately, the bottom line is cost, which is a marketing decision.

DIMMERS

There are many homes with dimmers for their light bulbs. They are also fairly common in businesses—restaurants dim lights for ambiance, conference rooms have light dimmers for presentations, and so on. So at least for LED designs going into general service lighting the ability to work on a dimmer circuit is highly desirable.




Figure 8.19 Output Waveform of a Triac Dimmer.




Figure 8.20 Keeping an IC's Power Alive During the Off-Time of a Dimmer.

The trouble is that dimmers are a problem for most AC power supplies. Almost all of them work by turning off the AC power during a portion of each line cycle. This is almost always done with a triac circuit, because it's amazingly cheap. And since it's such an inexpensive way of doing it, you should expect triac dimmers to continue to dominate the market for the foreseeable future. There are also electronic dimmers, but we won't consider them.

The basic waveform generated by a triac dimmer is shown in Figure 8.19. The voltage is zero until some phase angle is reached, and then it abruptly resumes the normal sinusoidal voltage for the rest of the cycle. The missing part can be either at the beginning of each cycle or at the end, called respectively leading or trailing edge. You see both kinds in commercially available dimmers.

There are really two problems with this waveform. The first is that when the line voltage abruptly turns on, there is a large surge of current into the input capacitor, which tries to charge up very quickly. This can cause an electrolytic capacitor to blow up within minutes. The solution here is to be PFC. Without the big capacitor the rest of the circuitry doesn't mind the chopped AC waveform. It regulates the LED current during the portion of the cycle when the voltage is present, and the average LED brightness is reduced proportionally. This works as long as your converter's bandwidth is fast enough to respond to the missing part of the line voltage. (PFC does not require low bandwidth. For example, the gain circuit suggested in the section above will work just fine.) The only trick you may need to know is that if the IC takes a long time to turn on (some take milliseconds), then you may need to provide power to it during the line's off time. All this takes is a diode and a capacitor, as shown in Figure 8.20. The capacitor is chosen to provide current to the

IC during one cycle of 8.3 msec, and the diode blocks the capacitor from discharging back to the line while the input is zero.

The second problem is more serious: LEDs are so efficient that the dimmer may not work! The reason is that dimmers are meant for incandescent light bulbs, the smallest of which is typically a 40W device. Many LED light bulbs, however, are sub-10W. Below about 30W, the triac doesn't work right, you end up with abnormal AC waveforms, and the whole system doesn't dim properly. This problem is not solved by PFC, and indeed is made worse by it, because the current drawn by the supply is lower at the lower line voltages.

One solution that some IC manufacturers have been implementing is to increase the load current. This could be accomplished by adding a big resistor at the input. But then the LED efficiency is reduced to that of an incandescent. Instead, the IC people attach the resistor close to the zero-crossing of the AC line. This is the portion of the cycle where the triac needs to have enough current, so the resistor makes sure it fires properly. Additionally, since the line voltage is low during this time, the power dissipated in the resistor is relatively low. Doubtless there is room for further design improvements in this developing field.

RIPPLE CURRENT EFFECTS ON LEDs

The authors once built an array of LEDs and ran it with a lab supply set to current limit. The light was measured with an integrating sphere, and the current was set to produce sufficient light output. The current was then measured with a true-RMS meter. Next, an AC power supply was built and measured to have the same current. But when a light measurement was taken, the LEDs were dimmer. It wasn't a meter problem, because the meter had plenty of bandwidth. The problem, it was determined, was that LEDs lose efficacy when there's ripple current. So we had to use more power to obtain the same light output. We're going to examine the magnitude of this effect in this section. Some graphs showing approximate effects are shown in Figure 8.21.

There are actually two aspects of LED performance that are affected by ripple current that act to increase power dissipation. To get the same light output from a rippling current as from a constant current requires that the peak current be higher than the average. Now recall that the forward voltage of an LED is only approximately constant with current. In reality, the forward voltage increases somewhat as the current increases. Looking at Figure 8.22,⁴ we can see that increasing the current from 350mA to 500mA increases the forward voltage from about 3.18V to about 3.28, about a 3% increase. This increases the power dissipated while the current is this high.

A second loss mechanism is that the efficacy of the LEDs decreases with increasing current. Looking at Figure 8.23, the relative flux at 350mA is 1.00, while at 500mA it is approximately 1.30. The efficacy has thus decreased from approximately $1.00/0.35 = 2.86$ to approximately $1.30/0.5 = 2.60$, a 9% decrease.

⁴Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Co., 2007.




Figure 8.21 As Ripple Current Increases, Power Loss in the LED Also Increases.

Source: Betten and Kollman (2007). Used by permission of Electronics Technology, a Penton Media publication.




Figure 8.22 Forward Voltage Increases with Increasing Current.

Putting these two mechanisms together will result in a higher power required to get the same light output with ripple current. This leads to a higher temperature, shortening lifetime of the LEDs. There is a second-order effect as well. The increased temperature also decreases the efficacy. But for practical purposes, this effect is usually small enough to be ignored.

Let's look at a practical example to illustrate the effect. The numbers are summarized in Table 8.1 and Table 8.2. As a baseline, let's take a Luxeon Rebel being operated at a DC current of 350 mA, with a thermal resistance of 15°C/W to an ambient of 20°C. Suppose it has an output of 60 lumens at 20°C.

Now at 350 mA, the forward voltage of the LED is 3.15 V, so the power dissipation is $3.15\text{ V} \times 0.35\text{ A} = 1.103\text{ W}$. This gives a device temperature of $20^\circ\text{C} + 1.103$

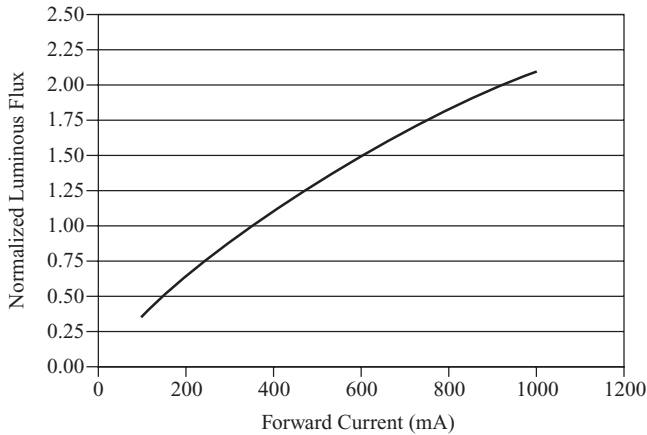


Figure 8.23 Increasing the Current Does Not Proportionally Increase the Light.
 Source: Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Co., 2007.

Table 8.1 Efficacy of LED at 350 mA, Including Temperature Effects

Parameter	Value	Units
I	350	mA
V _f	3.15	V
Power Dissipation	1.103	W
Pad Temperature	36.5	°C
Temperature Efficacy	96	%
Efficacy	57.6	Lm/W
Light	63.5	Lumens

Table 8.2 The Efficacy at an Average Current Is Less Than That at That DC Current

Parameter	Value at 250 mA	Value at 450 mA	Average	Units
I	250	450	350	mA
V _f	3.10	3.25	—	V
Power Dissipation	0.775	1.463	1.119	W
Pad Temperature	36.5	36.8	36.8	°C
Temperature Efficacy	96	96	96	%
Relative Light Output Due to Current	0.76	1.23	—	—
Efficacy	54.2	48.4	51.3	Lm/W
Light	42.0	70.8	56.4	Lumens




Figure 8.24 200mA_{pp} Ripple Current on a 350mA DC Drive.

$W * 15^{\circ}\text{C}/\text{W} = 36.5^{\circ}\text{C}$. At this temperature, the efficacy is reduced to 96%, which is $60\text{Lm}/\text{W} * 0.96 = 57.6\text{Lm}/\text{W}$. Light output is thus $57.6\text{Lm}/\text{W} * 1.103\text{W} = 63.5$ lumens.

Now let's examine one with a ripple current. We still want an average of 350 mA, but superimposed on top of this is 200 mA peak-to-peak ripple (see Figure 8.24). Peak current is thus 450 mA, and minimum current is 250 mA. Here's where the practical part comes in. Rather than trying to integrate the various parameters over the triangle wave, we're going to approximate the current as constant at 250 mA for half the cycle and at 450 mA for the other half.

Now at 450 mA, the forward voltage increases from 3.15 V to 3.25 V. At 250 mA, the forward voltage decreases from 3.15 V to 3.10 V. Power dissipated during the 250 mA part of the cycle is therefore $3.10\text{V} * 250\text{mA} = 0.775\text{W}$. During the 450 mA part of the cycle it is $3.25\text{V} * 450\text{mA} = 1.463\text{W}$. Note that 2/3 of the power is dissipated during the high part of the current cycle. But average power is essentially identical with the DC case, as was to be expected. Thus, device temperature is also almost identical, now $20^{\circ}\text{C} + 1.119\text{W} * 15^{\circ}\text{C}/\text{W} = 36.8^{\circ}\text{C}$. Since the temperature is the same, the temperature effect on the efficacy is the same, 96%. However, efficacy is also affected by the current level. At 250 mA, 76% of the light at 350 mA is produced, and at 450 mA 123% light comes out. Calculating efficacy at 250 mA, $60\text{Lm}/\text{W} * (0.76/0.775\text{W}) * 96\% = 54.2\text{Lm}/\text{W}$. At 450 mA, $60\text{Lm}/\text{W} * (1.23/1.463\text{W}) * 96\% = 48.4\text{Lm}/\text{W}$. Light emitted during each condition would be $54.2\text{Lm}/\text{W} * 0.775\text{W} = 42.0$ lumens and $48.4\text{Lm}/\text{W} * 1.463\text{W} = 70.8$ lumens. Total light output is the average of the two, 56.4 lumens.

Light output has therefore decreased 11% due to the ripple current. Looked at the opposite way, to get the same light output would require $1/(1 - 11\%) = 12\%$ more power—which would increase the losses even more. The conclusion is that ripple current can have a significant impact on the light output from your LEDs. If you want to get it right the first time, you have to account for ripple—or reduce it to a level where it's insignificant for the application.

LIFETIME

There are some additional factors affecting lifetime of your design, beyond those already mentioned in the chapter on DC drivers. These additional factors are 60Hz ripple in the input capacitor, lightning and ripple current in the LEDs.

Let's start with the input capacitor. In the previous chapter, we noted the need for a high-frequency capacitor at the input of switch-mode power supplies to provide low impedance for the high-frequency power pulses. This is still true in an off-line

converter. However, the input capacitor in AC converters in addition has 60Hz (or 50Hz) ripple. And it turns out that this is often the dominant ripple source. When you look at the AC current in the capacitor with a current probe, most of what you see is at line frequency.

There's really nothing new here, except that the capacitor ESR again has to be selected to ensure that you don't degrade the MTBF of the design. And off-line converters sometimes use large electrolytic capacitors. It is these that have limited lifetime. As discussed in previous chapters, there's usually not enough data on a capacitor datasheet to tell what the ESR is. So the usual method is to build the converter and measure the current. This is helpful for ensuring you're not exceeding absolute maximum ratings. But from a practical standpoint, lifetime is set by the capacitor temperature. So again, sticking a thermocouple on the capacitor tells you the lifetime. For most applications, it's possible to find a capacitor with a high enough lifetime at a high enough temperature to meet the MTBF. The problems usually arise when it also has to be physically small enough. The solution is to power factor correct the SMPS, so no electrolytic capacitor is needed.

There are frequent comments in the literature about limiting the ripple current to the LEDs. It is claimed that ripple current degrades the lifetime of the LEDs, but few specifics are given. It turns out that for reasonable levels of ripple, the dominant effect on LED lifetime is the increase in dissipated power, as we've seen above. This means that the normal calculation and measurement of LED temperature is enough to ensure their lifetime. But what is a reasonable level of ripple? The answer is that the LEDs' peak current should not exceed the absolute maximum rating of the device. As long as this is respected, the LEDs' lifetime will be reasonably well predicted by their temperature.

The final new factor affecting lifetime of the system is lightning. We've already seen that the best you can do about lightning is to put in the biggest MOV you can fit. But is that good enough? Because MOV clamp voltages degrade with usage in a nonspecified way, there's no analytical answer possible. Here's a practical approach to the answer. Following the guidelines given above for component selection, the voltage let through by the MOV will not be enough to damage the converter. Thus, lightning won't cause the unit to fail until the MOV fails. How long does this take? We see units in the field that are not very well protected (say, with a 7 or 10 mm MOV) that last for 10 years or more. While the frequency of lightning depends on the geographic area in which the unit is installed, in general a 10–14 mm MOV should be sufficient for most locations. Of course, if you can afford the space for a 20 mm MOV, that is preferable.

UL, ENERGY STAR, AND ALL THAT

In this final section we're going to survey some of the governmental and quasi-governmental regulations surrounding LED lamps. As noted in Chapter 2, governments have seen fit to regulate lighting, and there is a wide variety of safety and other agency requirements as well (see Table 8.3) (Dowling, 2009).

Table 8.3 A List of Standards**Photometry**

CIE 127-2007 (TC2-45) Measurement of LEDs
 IESNA LM-79 Electrical and Photometric Measurements of SSL Products
 IESNA LM-80 Method for Measuring Lumen Maintenance of LED Light Sources
 CIE TC2-46 CIE/ISO LED intensity measurements
 CIE TC2-50 Optical properties of LED arrays
 CIE TC2-58 Luminance and radiance of LEDs
 IESNA TM-21 Predicting Lumen Maintenance of LED Sources

Color

ANSI C78.377-2008 Chromaticity of SSL Products
 CIE 177-2007 (TC1-62) Color Rendering of White LED Light Sources
 CIE TC1-69 Color Quality Scale (new CRI)

Photo biological Safety

IES RP-27 Photo biological Safety
 IEC 60825-1-2001 Safety of Laser products (to be superseded)
 CIE S009 Photo biological Safety
 IEEE P1789 Recommended Practices of Modulating Current in High Brightness LEDs for Mitigating Health Risks to Viewers

Safety

ANSI C82.SS11 Power Supply
 ANSI C82.77-2002 Harmonic Emission Limits
 ANSI C78.09 82 Fixture Safety Specification
 FCC 47 CFR Part 15 Radio Frequency Devices
 IEC SC 34A 62031:2008 LED modules – Safety
 IEC SC 34C 61347-2-13:2006—Lamp Control Gear Part 2-13: DC or AC Control Gear for LED modules
 IEC SC 34A IEC 62560 Self-Ballasted LED Lamps
 IEC SC 34A [TBD] LED lamps >50V—Safety specs
 UL 8750 LED Light Sources for Use in Lighting Products

Performance

IEC SC 34C 62384-DC or AC supplied electronic control gear for LED modules
 EC SC 34A-Performance Standard for LED Lamps

Nomenclature

IES RP-16 Nomenclature and Definitions Addendum A: SSL Definitions
 IEC SC 34A-TS 62504. Terms and Definitions for LEDs and LED Modules

EMC & Other

IEC TC 34 EN 62547 LED EMC/Immunity
 IEC SC77A-EN 61000-3-2 LED EMC/Harmonics
 ANSI SSL2 LSD-45 Sockets & Interconnects
 ANSI C82.04 Driver Safety Circuitry

Source: Dowling, “LED lighting standards and guidelines are now building on a firm foundation,” *LEDs Magazine*, May/June 2009.

Indeed, some would say that governments have seen fit to over-regulate LED lighting. A prime example of this is the U.S. government's specification of power factor for LED bulbs. In the United States, the Department of Energy awards an Energy Star rating to devices that are in the top 10% at saving energy. To the extent that consumers look for the Energy Star insignia, this has helped to increase the average efficiency of items such as refrigerators and water heaters. These are all goods with well-established markets. With LED bulbs, however, the DOE decided to take the initiative, and put out requirements that LED bulbs need to meet Energy Star *before* the market was established. How do they know what the top 10% of bulbs do? For example, Energy Star LED bulbs will be required to have a minimum power factor of 0.7 (residential) or 0.9 (commercial). Adding power factor correction to an LED ballast costs money. In some cases, it may make it impossible to make the bulb (think of fitting the extra circuitry into an MR16 form factor). The government is thus mandating performance, without regard to the effect of this mandate on the market.

There are many other potentially onerous requirements for Energy Star. As a minor example, LED bulbs are required to meet FCC EMI requirements of part 15. While it seems reasonable that all electronic devices should do so, in fact CFLs are required to meet only part 18, which is somewhat easier. Why should LED lighting be singled out for different treatment? Our guess is that LEDs have been touted by enthusiasts as the perfect light source. Since they are (or rather will be, the enthusiasts claim) perfect sources, the government has decided to require them to be better in every way, regardless of whether or not the market is willing to pay for these improvements. It remains to be seen whether this hurts the market for LED bulbs.

UL also has quite a number of standards that will potentially affect LED lighting (see Table 8.4). Of course, the venerable UL1993 is still applicable. But they have now added UL8750, "Safety Standard for Light Emitting Diode (LED) Equipment for Use in Lighting Products." The most significant requirement, as always for AC powered equipment, is line isolation. A consumer handling the bulb in a reasonable way should not be able to come into contact with the AC line. But what is reasonable? Clearly, having live wires hanging out is too dangerous.

An easy solution is to make the ballast isolated. But this involves (at least) changing from an inductor to a transformer. This certainly adds to the cost, and probably the size and weight of the ballast. Are other solutions possible? While each UL evaluation is individualized, apparently bulbs in which the live electronics (and traces, solder pads, and other components) are covered with a plastic or glass cover are reasonably safe. Recently, individual LEDs have been recognized as being safe. After all, even if all of the electronics are covered up, the LEDs have internal bond wires that could be connected to AC line. No general formula can be given for passing UL, although a safe route is to use an isolated ballast. There are ways around this, and you have to talk to UL about them (preferably before getting very far into the design).

One other area of safety regulation that is almost unique to LED lighting is eye safety. There is a long history of limits on lasers and how it can be ensured that

Table 8.4 UL Standards Affecting LED Lighting

Standard	Product Type
Fixed Luminaires	
UL 1598 & CSA C22.2 #250	Luminaires
UL 1573	Stage and Studio Lighting
UL 1574	Track Lighting
Portable Luminaires	
UL 153	Portable Electric Lamps
UL 1993	Self-ballasted Lamps and Lamp Adapters
Specialty Luminaires	
UL 48	Electric Signs
UL 676	Underwater Lighting Fixtures
UL 844	Fixtures for use in Hazardous Locations
UL 924	Emergency Lighting and Power Equipment
UL 1786	Nightlights
UL 1838	Low-voltage Landscape Lighting Systems
UL 1994	Low-level Path Marking and Lighting Systems
UL 2108	Low-voltage Lighting Systems
UL 2388	Flexible Lighting Products
Power Supplies	
UL 1012	Power Units other than Class 2
UL 1310	Class 2 Power Units

Source: Straka (2009)

people aren't blinded by them. But with individual LEDs getting brighter and brighter, and with the limited light emission angle some devices give, LEDs are a potential threat to eyes as well. There are now some limits in place on how bright an individual device can be, without requiring some sort of diffuser or other means to limit the brightness. For example, there is an optical safety requirement in the EU's EN60825-1. While not an immediate problem for LED lighting, this is something to keep in mind for the future.

Chapter 9

Practical System Design With LEDs

In previous chapters, we have examined the components that make up a lighting design: the LEDs, the driver, their thermal performance, and so on. To make a complete design, however, these various components need to be integrated. We need to make a printed circuit board for the LEDs and for the ballast. We need to design optics to get the light from the LEDs to go where we want it to go and to have the characteristics we want. We need the design to survive in the customer's environment. And we also need to look forward to our next design, when the efficacy of the LEDs will have increased yet again. It would be wise to make our design flexible enough to accommodate this next generation of components without starting the design all over again. After all, those better components will be here in just a few months.

PCB DESIGN

There are specialists who do nothing but lay out printed circuit boards (PCBs.) And within this group of specialists are specialists who focus on power supply layouts. Given this degree of specialization, we won't be able to tell you in detail how to lay out PCBs for your LEDs and ballast. What we're going to do is to give some pointers on concepts we have found to work well based on long years of experience. We'll also advise on pitfalls to avoid, based on PCBs that came back and didn't work quite right.

The first step in creating a good PCB is to make a good schematic. A good schematic is more important than it at first seems, for two reasons. In the first place, a good schematic is easy to look at and understand. This is important for anyone else who has to look at the schematic, for example, for a consultant or for a design review. It's also important for you, when you have to go back and look at that old design 12 months from now.

A second reason a good schematic is important is because it helps generate a good layout. When you spend enough time to lay out a good schematic, it becomes

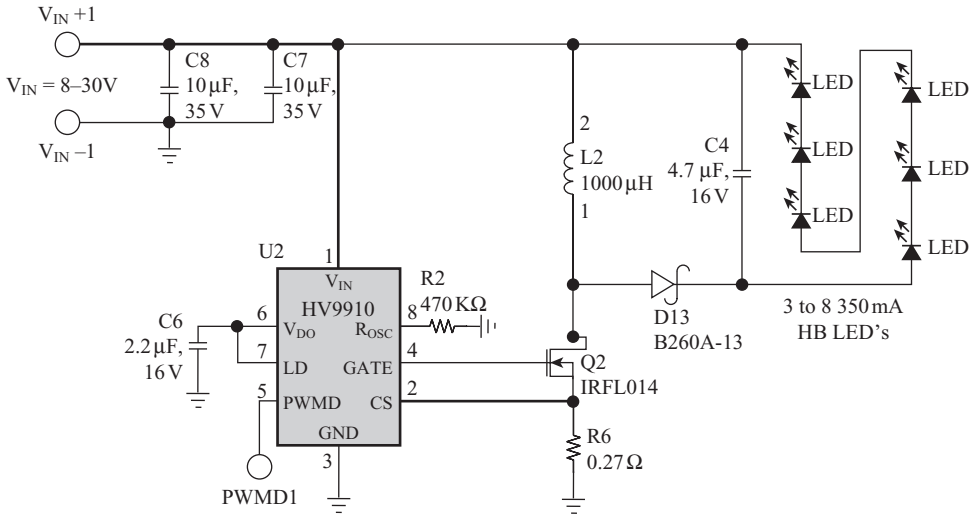


Figure 9.1 A Schematic which Could Be Improved.

Source: HV9910 datasheet, Supertex Inc., 2006.

a guide telling you where components should go on the PCB. Components that are “naturally” close together on the schematic often should be close together in the layout as well.

Look at Figure 9.2 for a sample of a good schematic versus Figure 9.1 for a schematic that could be improved.¹ The most immediately obvious difference is in the IC pin-out. The original schematic shows the pins out of numerical order and with some pins on all four sides. Probably this is to minimize the visual complexity of the schematic. But it’s not useful for laying out a board because the IC doesn’t come with pins that do that. The improved schematic shows the IC pins in their actual order. This makes it obvious that the IC connections to the gate and source of the MOSFET are on the same side as V_{IN} and GND. The pins thus have to either be routed around the IC (possibly a noise problem with the gate drive) or the IC could be rotated 180° or be mounted on the backside of the PCB. Conscientiously putting the pins in the right place on the schematic will make it obvious if connections need to cross over other connections, a frequent occurrence with ICs. This can be an EMI problem or even a routing problem if you’re trying to make a one-layer board.

The next aspect to notice about the improved schematic is the reference numbering. The original shows capacitors with reference designators C4, and C6 through C8. Where are C1 through C3, and C5? Reference designators should start from the number one and proceed sequentially without skipping numbers. Otherwise the

¹Don’t get us wrong, we think the 9910 is the best high-voltage LED driver on the market today. It’s just the schematic that could use some improvement.




Figure 9.2 A Good Schematic.

BOM will be confusing, and buyers and others will ask questions about where the missing parts are.²

A further reference designator problem is that the references of the original schematic have no order. C8 is to the left of C7, and C4 is way off to the right. If you have a hundred capacitors in the schematic, how can you find C47? We recommend that the proper way to draw the schematic is to start numbering components from the upper left, and proceed toward the lower right. Of course, there will be some components whose sequential numbering will be not as clear, but overall this makes finding a given part much easier.

There are still improvements that could be made to the improved schematic. These would be more important for a design with more components. For example, you could show the ground for power components separately from the small-signal component ground. They could use different ground symbols, or have a “p” and an “s” attached to the ground symbol respectively. In the current schematic, small-signal ground would be only for R1 and C2, and the IC ground would be where the two grounds connect. This is how we will actually lay out the board in Chapter 10.

In some cases, it may also be useful to color-code the schematic. For power supply schematics, we find it useful to indicate the high-current traces in one color (say, yellow) and the high-voltage traces in a second color (say, red). That way the PCB layout person knows that yellow traces have to be wide and red traces need to be well separated from other traces. The former is for current carrying capacity, the latter for arcing. You might also consider another color for signals with high-frequencies, to be careful with EMI and cross-talk.

²If you have two identical sections that you want to reference the same way, you could call the parts in one section 101, 102, etc., and the corresponding parts in the second section 201, 202, etc.

One ambiguous area of the schematic is the placement of R2. It measures the current from Q1, and it seems natural to put it there. On the other hand, the actual layout will have R2 right next to pins 2 and 3 of U1, to minimize noise in the feedback path. Which one should you pick for your schematic? Since there's no good solution, you just pick one. Maybe you can add a note about placement to the schematic.

Here is a checklist of design elements to aim for when putting together a schematic:

1. Draw the schematic with real pin-outs.
2. Order reference designators starting from 1, and don't skip numbers.
3. Assign reference designators starting from the upper left hand corner and working towards the lower right hand corner.
4. Have separate symbols for signal and power grounds.
5. Color code high-current, high-voltage and high-frequency traces.
6. Make notes about components that should have special attention during layout.

Now we have an easy-to-use schematic. Our PCB is basically going to look like our schematic. We place the components approximately in the relative positions they occupy on the schematic, and the connections between them route approximately the same as they do on the schematic. There are a few additional considerations. Let's talk about grounding first.

Grounding

Your PCB should have a good ground system. Just connecting all of the grounded components together with traces really doesn't work. What you want is to connect all of the power grounds together with thick traces or a plane, a copper pour. Power grounds are all the grounds carrying substantial currents, say more than 100mA. Use as much copper as space permits. Once you've connected all of these grounds together, it should connect back to the input capacitor's ground.

Don't make loops in the power ground. A loop means that there are two ways for the current to get back to the input. This creates an EMI problem. Of course, a via in the middle of the ground plane is acceptable, and usually necessary. Just make sure that the via isn't right at the edge of the plane, which would give a small trace on one side of it. It's better to retract the plane a little away to avoid having the small trace.

It's best to let each ground go directly back to the input via the plane, rather than trying to connect one trace to the next in a chain. All power grounds should go directly to the ground plane, directly if possible and otherwise through one or more vias.

The small-signal grounds should also be attached together using traces or a separate small plane. Don't just hook them into the power ground. This separate

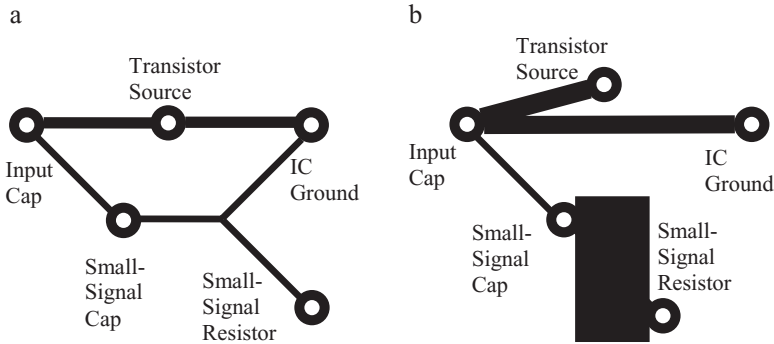


Figure 9.3 Poor Grounding Layout (Left) and Improved Layout (Right).

plane is called an island. Its purpose is to keep the high currents on power grounds from causing voltages on sensitive nodes. For example, when the MOSFET turns on there's a fast current spike from its source to ground. Since even ground planes have resistance and inductance, this can cause a fast voltage spike on the ground. This spike may be large enough to cause false-triggering of the IC. Connect the island to the power ground plane at a single point. We recommend the connection to be at the input capacitor, if possible. Multiple connections back to the power ground generate the noise you're trying to avoid.

Avoid chains if possible, make the signal ground into a star. A chain is a series connection of component grounds. For example, you have two capacitors and two resistors all of which are grounded, and you attach a trace from the small-signal ground to each of them in turn. The one at the end of the chain will be affected by the noise from all of the others. A star gives each component its own trace back to the ground, so that all small-signal grounds see the same noise simultaneously.

Figure 9.3 gives some idea of what should be avoided (a) and what should be done (b). On the left, the ground of the IC is returned to the input capacitor's ground by going through the transistor's source connection. The high frequency components of the two will mix, potentially causing the transistor—or the IC—to turn on or off at the wrong time. On the right, the IC ground and the transistor's source are returned to capacitor ground separately. On the left, the IC ground has *two* paths back to capacitor ground: one through the source, and another through a path through a small-signal ground. This loop splits the current from the IC ground, creating a ground loop that can affect EMI. On the right, the IC ground has a single path back to capacitor ground. On the left, the small-signal capacitor and resistor are in line with the current from the IC ground. Since the resistor joins this path before the capacitor does, they will see different ground voltages. They will also see different noise. This can potentially upset the operation of the IC, resulting in things such as missed switching cycles. On the right, the small-signal components' grounds are tied together on an island, and then returned by their own separate trace to capacitor ground.

Other Connections

Power connections should be treated much as power ground is. For example, the connection of the MOSFET drain, the inductor and the anode of the diode in the schematic of Figure 9.2 should be connected with a plane to the extent possible. Wide traces are acceptable, but make them as wide as will fit. Extra copper is free!

What constitutes a high-voltage trace is somewhat relative, but generally anything above 60V is appropriate. There aren't any on our schematic, although the drain will come close. You should leave extra spacing between this type of traces and others at low voltage, to avoid any possible arcing problems.

High-current traces, including power traces, should be wider than normal. In general, the wider the trace is, the better. This minimizes voltage drops and efficiency losses. On a related topic, we recommend that all ballasts and all PCBs with LEDs should use at least 1oz copper, not the standard ½oz. It is even better to use 2oz copper, if the cost is not prohibitive.

High-frequency traces need to avoid crossing other traces. Ideally, of course, they should be very short. But if it's necessary to run it some distance, running it above the ground plane is good. This minimizes the loop area for radiation. The gate drive from the IC to the MOSFET is particularly a problem. It has both high current and high frequency. This trace needs to be short, if possible. It is probably better to move the MOSFET than to have a long trace for this signal.

The following is a checklist of things to aim for in laying out a PCB:

1. Make ground a plane or a portion of a plane. Connect all of the high current grounds directly to it. Avoid loops and traces.
2. Make a separate ground island for small-signal grounds. Connect it at a single point to power ground.
3. Connect power connections with a plane. If you need to use traces, make them as wide as possible. The same is true for all high-current connections.
4. Separate high-voltage traces from others with extra space.
5. Use at least 1oz copper for the PCB.
6. Keep high-frequency traces short, and don't cross other traces with it. Avoid vias. The gate drive in particular needs to be short.

GETTING THE LIGHT OUT

The purpose of an LED light is to illuminate. The LEDs produce the light, but there is usually more to it than that. You want the light to have a certain distribution, bright in some places and dark in others. For example, a spotlight needs to have most of its light within a certain forward angle. You want the light to be of a specific color. For example, daylight incandescent bulbs use the same filament as normal incandescent bulbs do, but modify the light color by adding Neodymium to the glass. And

you want the light to be bright enough. If you've spent all that money on LEDs, you don't want to waste 30% of it in a piece of yellowing plastic.

All of these aspects are part of the system design of the optics. We'll start with lenses. Everyday experience gives us the idea of what lenses can do. In the form of a magnifying glass, they can bend and focus light. As a prism, they can split white light up into a rainbow of its constituent colors. A colored lens can absorb light of a particular color, removing it from the incoming light. And, while less obvious in daily experience, lenses also reflect light to some degree.

While lenses are used for some of these things in LED lights, some of the other things mentioned above we try to avoid. For example, the primary optics in an LED can be used to direct edge-emitted light from the LED forward, so that light comes out of the front rather than from the sides of the device. The silicone covering of an LED also acts as a lens and helps to reduce light loss. Since the index of refraction of the die and air are quite different, the silicone acts as an intermediary. By lowering the index of refraction in two steps rather than doing it all at once, reflections from the interface can be reduced. This increases the light output from an LED.

Of course, these lenses also have some undesirable characteristics. The primary optics in an LED don't redirect all of the light. Some of it is absorbed, some of it is directed in the wrong direction, and some just reflects back to the emitter and is lost as heat. The die-to-silicone interface and the silicon-to-air interface have similar problems. Additionally, silicone eventually yellows. This absorbs part of the light spectrum, reducing efficacy and changing the color of the emitted light.

Beyond lenses on the LEDs themselves, there are other optical devices, whether their optical properties are intentional or not. A flashlight may have a glass or plastic lens in front of the LED in order to obtain the right beam angle. A USB light may have a glass or plastic cover so that the user is not accidentally burned by the hot LED. And an LED light bulb may need to have a cover over it to avoid electrical shock hazard. It will probably also have overall structural elements to make it work mechanically and aesthetically, and these can reflect or absorb light as well.

As suggested by these examples, glass and plastic frequently form part of the interface between the LED light and the outside world. Looking first at glass, we note that dealing with glass is almost a forgotten art in the world of electronics. It takes some searching to find someone who knows enough about glass to make a design that is feasible with the material, mechanically robust, and cost-effective. To start with, there are a number of different common types of glass: soda-lime glass, borosilicate glass, and lead glass, among others. Soda-lime is the most common and least expensive type, used for bottles and jars. Lead glass is used in stemware and in incandescent light bulbs (it really does contain lead). Borosilicate glass is the most expensive, used in chemical test tubes and the like. It has great resistance to thermal shock, and low coefficient of thermal expansion.

In practice, your lights, if they use glass at all, will probably use soda-lime glass. Soda-lime glass has relatively good optical transmission in the visible region (see Figure 9.4). It's in the mid-90% range near the photopic peak in the green. It also transmits well throughout the IR, so that heat can escape by radiation.




Figure 9.4 Soda-Lime Glass Optical Transmission.

Source: http://en.wikipedia.org/wiki/File:Soda_Lime.jpg under license <http://creativecommons.org/licenses/by-sa/3.0/>.

Now while 92–94% transmission is very good for an optical device, as soon as you add glass to the optical path from the LEDs to the world, you’ve lost 6–8% of the light you struggled so hard to make.

Glass also has the downside of being easy to break. Maybe this isn’t a significant problem; incandescent bulbs have been made from glass for a century, and fluorescent tubes are made from glass and even contain a neurotoxin. On the other hand, UL thinks that glass is decorative, and doesn’t consider it to be a barrier against AC line voltage. Thus it may be a marketing decision as to whether plastic rather than glass is used.

Plastic doesn’t usually shatter like glass, but it presents other concerns. For one thing, there are many types of plastics, such as polyethylene, polyethylene terephthalate (PET), and polycarbonate. And they all have different properties, which will probably matter to your design. There are also many different grades of each type. How are you going to pick the right kind? The easy way is to pick the same kind that the competitors are already using. They must know what they’re doing, right?

The first thing to know about plastics is that they are typically less optically transmissive than glass. Glass will lose 6–8% of the light that strikes it. Even the best plastics are more like 10%, and some are much higher than that. So you’re going to have to put more power into the LEDs to obtain the same light output. You could pick a plastic with really good clarity, such as polycarbonate, but the cost of this is higher than other choices.

Plastics also have temperature limitations. Glass of course can be exposed to high temperatures without effect. But plastics can soften at very low temperatures, and completely melt not much above those. For example, polyethylene melts between 105°C and 130°C, depending on what type you use. PET starts softening at 75°C. Some polycarbonates on the other hand don’t melt until 265°C, and can be used well over 150°C. In fact, some plastics will burn if put into a fire, and may release toxic

gases. What will firefighters think about the presence of such a plastic in a burning house that contained your light bulb?

Plastics also yellow with age. Both heat and light are accelerators for this type of aging—naturally, just the characteristics you expect to see in an LED light. Of course, some plastics are better than others. In particular, polycarbonate is used in automobile headlights. They can sit in the sun for ten years and be right next to high-temperature incandescent bulbs without noticeable yellowing or hazing.

Some plastics diffuse light. Sometimes this is intentional, as when you want the light to be spread out more uniformly than when there are only a few point sources of light. The diffusion, however, is not loss-free. Diffusion results in increased optical loss, from 10% to 30%. So you have to crank up the LEDs even more to compensate. We should mention that the soft-white type incandescent bulb has very low loss in its diffusion. We have been told that it is only a few percent, so perhaps this would be a good model for manufacturers of LED bulbs to copy.

The conclusion is that almost anything you do to the emitted light is going to cause optical loss. If possible, it's better to get LEDs that have the right emission direction than to try to shape the light with lenses. Glass covers have optical losses. If the lack of mechanical robustness of glass is a problem you could go with plastics. As you can tell, we prefer polycarbonate. But plastics absorb more light than glass does, and adding diffusion absorbs even more light. It's probably best to get a consultant who knows all about plastics. (It's a good thing higher efficacy LEDs are coming out next month!)

LEDS IN HARSH ENVIRONMENTS

Most of the study we've done up till now has been for LEDs operating in a lab bench sort of environment. They get hot when you put power into them, but the environment is generally benign, 25°C or maybe 40°C at worst. The same assumption has been true for ballasts. But in a real system, the LED and the ballast heat each other, the mechanical part of the system may affect the LEDs, and the environment may not only be hot but unfriendly to LEDs and circuitry in a variety of other ways. For example, would LEDs work properly in a microwave oven? The authors don't know and would be glad to hear from someone with experience.

Some environments are going to be far hotter than others. An LED in a system left sitting on a car dashboard in summer in Florida will see an ambient temperature of 50°C. An LED in a lighting system in a traffic light in Alaska will see -40°C or lower. LEDs are unaffected by either of these extremes (although the electrolytic capacitors in the ballast won't work at -40°C). But they might be affected being in a display window in Alaska that is turned on only from 4 to 9 p.m. each night. They might be up at 85°C when they're on, and then drop to -40°C when they're turned off. Now they are being put through a wide ranging thermal cycle every day, 365 times a year. Thermal cycling tends to break things, because almost everything physically expands when hot and contracts when cold. Solder joints loosen, plastic

cracks, bond wires can fail. Any of these can potentially cause your light system to fail. In military designs, thermal cycle testing of the design is a requirement. It may also be a good idea for your design.

Another problem is also caused by thermal cycling which goes through 0°C. When a system is cooled below 0°C, humidity in the air tends to become frost. Then when it warms up again, water droplets may form. On a circuit board, these can cause electrical short circuits, and thus failure. On a ballast, it's common practice to seal the PCB with a conformal coat to avoid this problem. But on a PCB for LEDs, this may not work. What are the optical losses in the conformal coat? Will the conformal coat interact with the silicone on the LED?

Another potential environmental hazard for LED lighting systems is humidity. Humidity can degrade phosphor performance, leading to premature dimming. Companies routinely ship LED reels in sealed bags with silica absorbent to prevent moisture from compromising the parts. Some recommend baking LEDs from opened bags for 24 hours to evaporate the moisture. The fear is that when the LED is soldered, the moisture will form steam, potentially damaging the part. In any case, the effect of humidity on your system should be experimentally evaluated. In Malaysia, for example, it can be close to 100% humidity almost all the time.

A related problem is salt fog. In places close to the ocean, houses need new coats of paint every two years. This is because it is always humid there, and the air contains lots of salt. Salt corrodes paint and metal, among other things. So if, for example, you're making a light that can be used outdoors, you need to check that a salty, humid environment won't cause your light to fail prematurely.

These problems with temperature and humidity are exacerbated if your system contains plastic, for example, as the external case of your light. As we've said, there are many different kinds of plastic. Many aren't very good at high temperatures. They can become soft and lose mechanical strength, or they can become brittle and shatter easily. Many plastics will turn yellow with heat (or sunlight) and this may affect the color of the light that your system puts out. Again, if you plan to use a plastic in your product you should turn to a plastics expert to make sure you're doing it right.

LEDs can be blown up by ESD. But a lighting system can be also be affected. You have to make sure that the customer can't destroy it by touching the case or the electrical inputs with a static charge. Systems attached to the AC line are worse. Although we've shown protection for the ballast by adding in an MOV, *some* voltage always gets through. MOVs increase in voltage depending on how much current passes through them. If there's a lightning strike close enough to the light, it will be destroyed no matter what you do. Also remember that MOVs degrade with number of surges. If you have line surges ten times a day from a nearby motor starting up, this may damage the MOV enough that it no longer provides sufficient protection against lightning. The only practical solution is to put something that doesn't cost too much on the market, and evaluate returns to see if it is cost effective to increase the protection versus unit replacement and dissatisfied customers.

By the way, we'd like to offer this unpleasant thought. Incandescent bulbs are like resistors. If lightning comes along, the incandescent bulb absorbs a good deal of the energy. Of course, a direct hit will destroy anything, but the incandescent bulbs are suppressing some of the energy of some of the neighborhood's lightning. What happens when all those incandescent lightning protectors have been replaced with LED lights that are just electronics?

LEDs are semiconductors, and as such can be affected by vibration. When you put an LED light on an automobile, it experiences many bumps and jolts over its lifetime. Bond wires can fail in the LED from these vibrations, the same as they can fail in the controller in the ballast. Fortunately, in order to evaluate them, standardized tests have been developed for this sort of environment.

Think all those environments weren't enough? Lights can also be affected by dust. Dust settles on everything. Light bulbs are rarely cleaned, if ever, and thus the light output from the bulb dims with time, even if the light source itself stays constant. A device like an LED light bulb can collect even more dust than furniture; the presence of a high voltage on the surface, even though it is safe and passes safety regulations, may actively attract dust. Think of how the monitor on your laptop computer attracts dust. Although companies claim that this dust depreciation is "not our fault," the customer will still perceive that your bulbs have become dimmer with age, without thinking about why.

DESIGNING WITH THE NEXT GENERATION LED IN MIND

Way back in Chapter 1, we learned the benefit of Haitz's law. Every few months a brighter and cheaper LED comes out, and all of our lighting products become brighter and less expensive just by changing part numbers. Indeed, you can simply plan on this being true. But this comes with a downside too.

For one thing, probably we've designed with enough LEDs to obtain the required brightness in the first place. If we now substitute brighter LEDs, perhaps the light will be *too* bright. The customer wanted 630 lumens, not 750. In a traditional engineering environment, this would be viewed as a cost-cutting opportunity. We can now make a new design that uses fewer LEDs. We can reduce the cost of the heat sink, and the ballast can be redesigned to produce less power. Unfortunately, all of this takes time and engineering resources. Redesign of the ballast can take a month, and laying out new PCBs for the ballast and the LEDs can also take a couple of weeks. The analyses have to be redone. All the tests have to be run again: thermal, EMI, and aging. Maybe UL qualification has to be redone. But before this new design is implemented, the *next* new LEDs will be out!

While this may sound like assured employment for the engineering department, the reality is different. Management may decide to skip a generation to get adequate return on the engineering expenses already incurred, or you may have inventory that has to be used up before the next design can be produced. And then your competitors will come out with a better bulb while you're in between versions. The situation is

analogous to the microprocessor world, where some years ago a faster clock speed came out every month. Your designs simply have to plan on that.

The right way to deal with these problems is to plan for obsolescence. If your design needs 18 LEDs, it isn't too hard to figure out that in the near future it will only need 16, and not too far down the road it will only need 12. So when you design the ballast, verify that it works not just with the 18 LEDs you're using today, but also with the 12 you'll be using six months from now. That way you don't have to redesign and requalify the ballast, and you will save time and money.

Your LED PCB can be laid out to accommodate different numbers of LEDs. If you need the LEDs to be mechanically symmetrical, you could include extra footprints on the PCB so that it remains symmetric when you go to fewer devices. The extra footprints can be shorted out with a jumper when you don't need them. For that matter, the PCB can be laid out so that each device has multiple footprints, enabling you to switch LED vendors without laying out the board a second time.

LIGHTING CONTROL

One other aspect of lighting system design is the possibility of an interface of the light with a lighting controller. None of the designs in this book cover this possibility, because it is a specialty item and takes substantial time and money to implement in both hardware and software. Nonetheless, there may be circumstances in which you need to communicate with your light. As a particular example, large installations of lights in office buildings may need to be under central control. For example, they may all need to be turned off at 8 p.m. when everyone has gone home, and on again in the morning, so that a person doesn't have to physically walk around every floor turning them on and off every day.

We admit to not being experts in the field of lighting control. So here we're going to take some extracts from Wong and Zheng (2009), and leave it at that.

DALI Protocol

“Digital dimming using the DALI protocol would add more flexibility and the advantages of being able to individually control up to 64 luminaires from a single pair of signal cables. This approach offers significant advantages in larger installations such as offices, particularly where intelligent lighting systems are used that employ advanced energy management techniques such as presence detection and daylight harvesting.

“DALI is a digital addressable lighting interface designed to replace and enhance the traditional 0–10 V analog dimmer. It is an open-lighting control protocol backed by major lighting manufacturers, and it has been established as the de facto control standard in fluorescent ballasts.

“The two-wire DALI network supports up to 64 lighting devices. Since Manchester encoding is utilized, the polarity of the wires is interchangeable, resulting in fool-proof installation. Since the topology of the devices can be a star, a bus

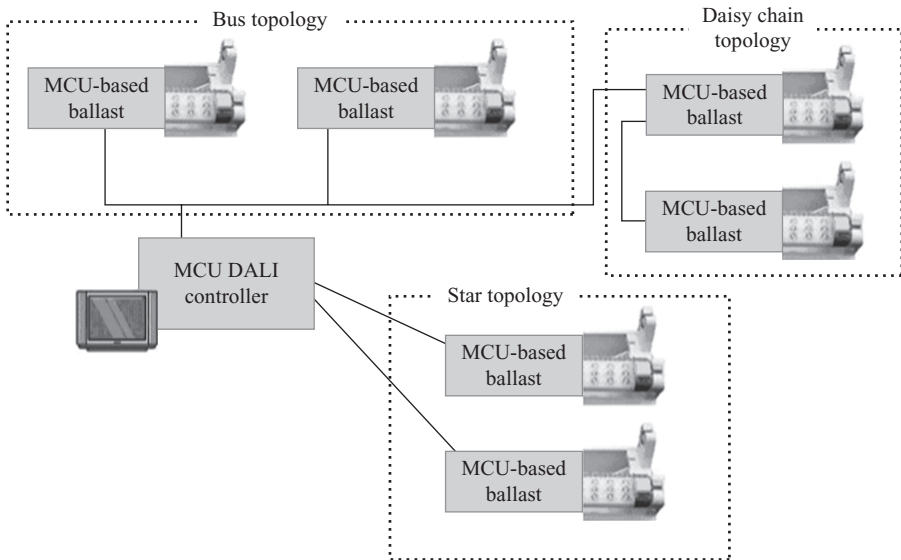


Figure 9.5 DALI Topologies.

or in a daisy-chain, retrofitting is simple when expanding a network (see Figure 9.5). The communication distance is less than 1000ft and is based on a master-slave arrangement, the master being the DALI controller and the slave being the ballast. A bidirectional communication enables the ballast to return its condition and setting back to the DALI controller. The DALI protocol has a relatively low baud rate of 1200 bits per second (bps), but is sufficient to turn on, off and dim a light. Since DALI is an addressable protocol, devices can be addressed individually or by group. A DALI master device can control a single light or a group of lights, and DALI command has been extended from typical level dimming and scene setting to incorporating environmental sensors for ambient light and motion sensing. The DALI standard also includes extensibility of manufacturer-specific commands.”

DMX512 Protocol

“DMX512 is a digital multiplexing lighting protocol that is popular in theatrical, architectural and entertainment lighting and can control up to 512 devices.”

802.15.4 and ZigBee Open-Standard Technology

“The ZigBee protocol has been spearheaded by utility companies to enable smart energy inside a home, where different appliances can communicate and control themselves. Currently, electricity usage is broadcast from an electricity meter to a portable in-home display, with the intent that when consumers can conveniently see

their electricity usage, they will turn off any unnecessary appliances. In the foreseeable future, the display panel could morph into a control panel to control lights and other appliances. To ensure interoperability, all ZigBee devices need to run a complete software stack.”

Powerline Communication

“Another wired communication technology, called powerline communication, requires no dedicated wire. Command and control functions are sent and received on 110 or 220V AC wires. Although many variations of powerline communications exist, frequency shift keying (FSK) is a popular choice because of the low data rate for command and control functions. In this method, high frequency is used for communication on a 60Hz power line. The quality of powerline communication inevitably varies, and the error correction, such as Forward Error Correction, reduces the error rate and improves the bandwidth of the medium. The data rate can achieve 30kbps or even higher. In comparison to a regular switch, powerline communication requires an MCU to run a network stack. However, the benefits are tremendous, because the software stack makes the network expandable, thereby connecting and controlling every smart appliance plugged to the powerline. Utility companies have started evaluating powerline communication as a complementary technology to a radio technology called ZigBee for the Smart Energy Home Area Network.”

Chapter 10

Practical Examples

In previous chapters we have reviewed all of the important properties of LEDs, their circuitry, and their physical and mechanical environment. Now it is time to put all of that information together. We are going to design four complete devices from scratch: A battery powered LED flashlight, a USB-powered reading light, an automotive tail light, and a 120 VAC powered LED light bulb. Of course, many others could have been added to the list. These four were chosen because they represent the four main types of drivers described in this book, and at the same time give opportunities to use much of the information that has been given.

The strategy of design we will use is typical of engineering projects. First we will lay out specifications according to what the marketing department thinks customers want. We'll then refine this into the realm of what is physically and cost-effectively possible. We will begin by making a first pass at the number and type of LEDs required, and estimating the power they require. Next, we'll make some models of the mechanical environment, and use this to get a handle on LED temperatures. Then we'll make a final pass through to tweak the values to get the optical output the marketing department specified. We'll look at the power conversion required, and select and design a power supply that will work and fit mechanically. This will include a PCB layout. Throughout, we'll be making choices based on cost, and our final step will be to estimate the cost of the total design.

EXAMPLE: DESIGNING AN LED FLASHLIGHT

Initial Marketing Input

We're going to start out by designing an LED flashlight. What the marketing department tells us they want in the first meeting is "a camping flashlight. Not one of those big police flashlights, something you can carry around. Get maybe 20–30 hours on a couple of batteries, how about C size? We want to sell it for about \$20, get it mass produced in China for \$2.50." They show us Table 10.1 to explain the cost structure they want.

Practical Lighting Design With LEDs, First Edition. Ron Lenk, Carol Lenk.

© 2011 the Institute of Electrical and Electronics Engineers, Inc. Published 2011 by John Wiley & Sons, Inc.

Table 10.1 Marketing Cost Structure for an LED Flashlight

	Price	Margin
Retail	\$20.00	50%
Distributor	\$10.00	35%
Us	\$6.50	60%
Manufacturer	\$2.60	

A little initial discussion reveals that what they have in mind is one of the newer flashlight types with a single, large-power LED in it. We point out that it would probably be cheaper to have six or seven T1 $\frac{3}{4}$ LEDs in it. They ask us to cost out both types.

Initial Analysis

The first thing to do here is to evaluate the battery, and see approximately how much light output you can get from “a couple of ‘C’ size” cells. We’ll assume that “a couple” means two, because that is a common number of cells in a flashlight battery. A quick web search¹ suggests that a C cell has an ampacity of about 6Ahr at an average voltage of 1.2V, which is 7.2Whr. Since there are two cells, we have 14.4Whr. At 25 hours of operation, we have power available of 576mW. This presumably has to go through a converter. Assuming we can afford to achieve 85% ballast efficiency on the \$2.50 budget, the power available to the LED is 490mW. We should be able to get an inexpensive power device in the 60Lm/W range, and we’ll lose maybe 10% in optics, which will generate in net 26 lumens. A little web searching suggests that 30 lumens is about right for a camping flashlight.

While we’re at it, let’s present them with an option to use two “D” cells instead. The “D” cell has ampacity of about 12.8Ahr, which at 1.2V is 15.4Whr. Two cells will be 30.7Whr. So with this bigger battery at 25 hours of operation, we have power available of 1.228W. Again assuming 85% efficiency, the power available to the LED is 1.044W. With a power device in the 60Lm/W range, and 10% loss in optics, we will generate 56 lumens.

Since this is more than enough, we could instead increase the run time by using the “D” cells. We could lower the light down to 30 lumens, which would run for $(56/30) * 25$ hours = 47 hours. As another alternative, we could reduce the cost of the LED, switching from 60Lm/W down to a $(30/56) * 60 = 32$ Lm/W device.

Let’s also take a quick look at multiple T1 $\frac{3}{4}$ devices. At www.mouser.com we can sort by ascending price. Looking at the first five devices that list both brightness and viewing angle, we construct Table 10.2.

The cost of 30 lumens appears to be completely impractical with T1 $\frac{3}{4}$ devices. Even if the cost were acceptable, you couldn’t fit 50 or 100 devices into a flashlight, even with a surface mount.

¹We used the following website: http://www.rayovac.com/technical/pdfs/pg_battery.pdf.

Table 10.2 A Few of the Possible 5 mm LEDs for the LED Flashlight

Device	Cost ea. (1 K pcs.)	Intensity (mcd)	Viewing Angle (°)	Lumens	Cost for 30 Lumens
LTW-2S3D7	\$0.27	13000	15	0.7	\$11.57
LTW-2E3C4	\$0.78	4200	15	0.2	\$117.00
LTW-2H7C5S	\$0.84	1800	56	1.3	\$19.38
LTW-2V3C5	\$0.90	2600	30	0.6	\$45.00
WP7114QWC/D	\$1.09	3200	20	0.3	\$109.00

Table 10.3 Battery Choices for the LED Flashlight

Option	2 “C” Cells	2 “D”’s Same Time	2 “D”’s Same Light
Ampacity (AHrs)	6	12.8	12.8
Mean Voltage (V)	2.4	2.4	2.4
Energy (Whrs)	14.4	30.7	30.7
Converter Efficiency	85%	85%	85%
Optical Efficiency	90%	90%	90%
LED Energy (Whrs)	11.0	23.5	23.5
LED Power (W)	0.58	1.04	0.58
Brightness (Lm)	26	56	26
Hours (Hrs)	25	25	47
Converter Cost (US\$)	\$0.50	\$0.50	\$0.50
Housing Cost (US\$)	\$0.50	\$0.50	\$0.50
LED Cost (US\$)			

Of course, these are just preliminary estimates. We haven’t accounted for temperature rise of the LED, which will reduce the light. And dissipating 1 W instead of $\frac{1}{2}$ W will make the LED hotter in the case of two “D” cells running at higher current, and thus dimmer. But these are secondary details. What the marketing department really cares about is cost. With 1 W or less we can use a single-power LED. Let’s put together a summary, Table 10.3, for the marketing department. We highlight the main difference, brightness versus run time.

When we start estimating the cost, we take a guess that the converter and housing will each be \$0.50. The difference between converting 500 mW and 1 W probably will make no difference to the cost of the converter. And heat-sinking an extra half watt isn’t going to matter to the aluminum housing. Thus, cost differences come down to which LED is used, not the difference in cell size or run time.

A web search on Lumileds parts shows that only the 3 W devices are available, even though a 1 W device would be sufficient for this application. We put together Table 10.4 for the LEDs’ cost for the two “C”-cell scenario.

Table 10.4 LED Selection Chart for the LED Flashlight

LED (Lm)	50	70	80	90	100
Typ. Efficacy (Lm/W)	63	76	90	100	91
25 Hr Light (Lm)	26	32	38	42	38
30Lm Runtime (Hrs)	22	27	32	35	32
CCT (K)	3000	3500	4000	5000	35–4000
CRI	90	85	85	70	?
Cost (US\$, 1 K pcs.)	\$3.35	\$2.98	\$2.98	\$2.98	\$2.98

The interesting thing about this table is the 1,000-piece pricing is the same for all the devices, independent of their brightness. The real difference is that you trade off CCT and brightness. It seems reasonable (since the marketing department didn't mention it) to use the highest brightness part and let the light be very cold white with a CRI of 70. This would be the 5000 K part in the chart.

It's also apparent that we can't meet the \$2.50 price target with an LED that costs \$3.00. We'll assume that at real volume this can be discounted by 30%, leaving us with \$2.00. With the \$1.00 for the converter and housing, we end up at a cost estimate of \$3.00. This is higher than what the marketing department wanted, but seems fairly reasonable. Of course, it will be up to us to find a \$0.50 converter!

Specification

The marketing department indicates that our choices are acceptable, including the high CCT. In particular, the operations department says they can produce this device for \$0.30 labor if the circuit board isn't too complicated. The marketing department decides to go with the "D" cell option since that doesn't affect the cost of the flashlight (batteries are not included). They also decide on longer, not brighter, since customers understand hours better than lumens. We thus draw up a specification, Table 10.5, for the marketing department to approve, before the design begins.

There are a number of issues that are indirectly addressed in this specification. The first thing to understand is the meaning of the terms in the header. In our convention, "typical" numbers refer to design goals. It does *not* mean that the average unit will perform this way. And it does not imply that units will be tested to this spec. Thus, if there is only a typical number in the specification, such as for MTBF, then the parameter is a design goal, and is guaranteed by the design, not by the production test. Conversely, any minimum or maximum specification is a guaranteed parameter, to be tested at production time unless otherwise noted. It is thus common in this sort of specification for there to be a lot of typical parameters, and few minima and maxima, as the operations department has to approve these latter parameters.

To start meeting the specification, we start with brightness. Lumileds's parts have a minimum guaranteed flux at 350 mA. While the minimum flux is guaranteed,

Table 10.5 Specifications for LED Flashlight

Parameter	Min.	Typ.	Max.	Units
Brightness (Vbat = 2.4 V)	27	30		Lumens
CCT		5000		K
CRI		70		—
Battery Size		2 “D” Cell Alkaline		
Battery Voltage (by design)	1.8	2.4	3.1	V
Touch Temperature			60	°C
Run Time		87		Hours
Ambient Temperature (by design)	10	25	40	°C
Reverse Battery Protection		No		
MTBF		50,000		Hours
Shelf Life		12		Months
BOM Cost			3.00	US\$
Total Cost			3.30	US\$

the typical and maximum flux are not. So we’re designing to the minimum, not the typical. Of course no one minds if it is brighter, so no maximum is specified. (In production there will be a maximum as well, to ensure that the converter is built correctly.) The other factor affecting the brightness is the converter. We expect its efficiency to go down as the battery voltage drops. To avoid problems in the production testing, brightness is specified at the typical battery voltage. CCT and CRI have already been mentioned; we’re maximizing run time.

Maximum battery voltage is set by the design having a two-cell alkaline battery. Note that this is an internal specification, not something that is tested in the design. Minimum battery voltage is not determined by the battery; it could theoretically be as close to zero as you like. But practically, below about 0.9 V/cell there isn’t much energy left. And converter design becomes more difficult and more expensive the lower the voltage at which it must operate. We consequently decide to not have the converter work below 1.8 V. Of course it should not blow up below that voltage, but operation is not guaranteed.

By touch temperature, we mean the maximum temperature on the flashlight that can be accessed by the user. At 40°C ambient temperature, this means only a 20°C rise. But the LED itself will be covered by a lens, and so the aluminum-to-air junction will meet this spec easily.

Run time is of course by design; there’s no way to test that in production. What we’ve done here is to use the 2 “D” cell spec, 30.7 Whr. Since our spec is only 30 lumens, with a 100 Lm/W part we need only 300 mW into the LED. Dividing by the converter’s 85% efficiency, that’s a power drain of 353 mW, corresponding to a run time of 87 hours.

At first we specify the ambient temperature to be at a minimum of –40°C, knowing that LEDs become more efficient at lower temperatures and that it might

get that cold when it is being used in the winter, for example, by a hiker. But a little consideration shows that there might be moisture condensation issues on the PCB when the flashlight is brought indoors after being in freezing conditions. Further, electrolytic caps (if we put in any) lose most of their capacitance below freezing. So we end up specifying minimum ambient as 10°C. The marketing department reluctantly agrees to this based on cost, but wants us to make our next flashlight not have this limitation.

Reverse battery protection seems reasonable; lots of people put in batteries backward. However, the only small and inexpensive way to implement this is with a series diode, which is going to drop at least half a volt. This would lower the minimum voltage the converter sees to 1.3V, and thus the converter would cost more. The marketing department decides to live without it.

MTBF has a design goal of 50,000 hours. Of course there is no way to test this, but a converter with only a few parts, plus an LED running not very hot, make this seem achievable. In the end we will meet this specification (the calculations aren't shown in this book).

Shelf life is something the marketing department wants to put on the label. When the flashlight is off, they don't want the battery to be drained as it is in a car or cell phone. But all electronic circuits have some leakage, including of course the battery itself. What we end up with is that the flashlight will have a mechanical switch, so that it won't draw any current at all. Shelf life is then set by the shelf life of the battery.

Power Conversion

Having selected the LED, the next task is to design its power supply. We've made an initial allocation of power to the LED of 300mW. The first thing we have to do is to translate this into current. The Luxeon has a typical forward voltage of 3.15V at 25°C, and decreases by about -3mV/°C. For our initial design, we'll assume the thermal pad in steady state is going to be about 60°C. We thus expect a forward voltage of $3.15\text{V} - 3\text{mV}/^\circ\text{C} * (60^\circ\text{C} - 25^\circ\text{C}) = 3.05\text{V}$. Dividing this into the power we get an LED design current of 98mA. We'll adjust this when we have a better estimate of the actual LED temperature.

The first and most critical item to select for the driver is the IC. This will be the most expensive component of the driver, and thus determines our ability to achieve our \$0.50 goal. To select the correctly priced IC, we introduce the authors' *divide-by-three* rule. It has turned out that over a very wide range of designs, the price of components in volume production can be reliably estimated. To do this, you take the 1,000-piece price at a distributor and divide it by three. For example, the 1N4004, a common bridge diode, is priced at a distributor at 4.3 cents for 1,000 pieces. When taken to volume production, it is a good estimate that the price will be 1.4 cents. As another example, a 100µF 100V aluminum electrolytic capacitor is priced at 12.2 cents at the distributor at 1,000 pieces; in volume this will cost 4.1 cents.

impedances at high frequencies. If the converter didn't have C1, it might see the battery voltage as less than 1.8V during the time its internal switch was on. To estimate the size of C1, we estimate the charge pulled from it during switching. The switching frequency is fixed at 1.5MHz, so the on-time is less than 700nsec. During that time, we're drawing an average of about 100mA. So the charge from the capacitor is $Q = I * t = 700\text{nec} * 100\text{mA} = 70\text{nC}$. To get no more than 100mV drop, we need $C = Q/V = 70\text{nC}/100\text{mV} = 700\text{nF}$. So we choose a 2.2 μF cap to account for tolerance. This is available at 16V in an 0805 package (to make it easy to place in the breadboard). The part costs \$0.01 in volume.

The purpose of C2 is to buffer the switching current from the LEDs, so they see DC. This calculation, to the level of accuracy needed, is the same as for C1. We pick the same capacitor in order to minimize the number of different parts needed (and to minimize stuffing errors).

D1 is the rectifier. Since we have 100mA output at some duty cycle, we pick a 1A device. Making the diode's rated current smaller wouldn't lower the price. Furthermore, at 1.5MHz, the recovery time would play a major role in efficiency, as would any forward voltage. We pick a Schottky diode to minimize both kinds of losses. We choose an SOD123 package because of its small size, and since any reverse voltage will do, we pick the device with the lowest forward voltage. This turns out to be the MBRX120, with a forward voltage of 320mV at 1A. Volume pricing looks to be about \$0.04.

The final item to select is L1. We've been rather cavalier in selecting the capacitors and diode, since there isn't much of a penalty to pay if they are oversized. But inductors become large and expensive quickly with size, and so we would do best to estimate the current through L1. The duty cycle for a boost converter is $DC = 1 - V_{in}/V_{out}$. The maximum duty cycle occurs for minimum input voltage, and the output voltage has to include the forward voltage of the rectifier. We find $DC_{max} = 1 - 1.8\text{V}/(3.15\text{V} + 0.32\text{V}) = 48\%$. Peak current through the inductor is then $I_{pk} = I_{DC} + V_{in} * t_{on}/L = 98\text{mA} + 1.8\text{V} * (48\%/1.5\text{MHz})/10\mu\text{H} = 156\text{mA}$. We used the 10 μH value from the recommendation of the datasheet.

To actually choose the inductor, we look for current ratings from 250mA to 500mA. We limit resistance to 1 Ω , to keep power losses down. The NLCV32T-100K seems a reasonable choice, with a current rating of 450mA, a resistance of 360m Ω , a tiny footprint, and an estimated cost of \$0.04 in production.

Our schematic is shown in Figure 10.2. Estimated cost is shown in Table 10.6. The cost for the power supply ends up at \$0.40 for the BOM. This leaves \$0.10 for manufacturer profit, labor costs, and errors in the cost estimate. The total cost, including the LED and PCB, is \$2.46. This leaves \$0.54 for labor costs, wires, and errors in the BOM cost estimate or manufacturer labor costs and profit. It also shows that the original desire for a \$2.50 flashlight probably couldn't have been satisfied. What is the most critical for this design (and for many designs) is to get the best possible price for the LED.

Let's do a couple of quick design checks. The battery voltage can be as high as 3.10V (= 2 * 1.55V/cell). The LED voltage can be as low as 3.05V. Are we violating the condition for a boost converter to work? Actually, we aren't. Remember




Figure 10.2 FAN5333A Final Schematic for Flashlight.

Table 10.6 Final BOM for LED Flashlight

Ref Des	Description	Part #	Mfr	1 Kpc Price	Est. Price
C1	2.2µF, 16 V, 0805	Any	Any	\$0.031	\$0.010
C2	2.2µF, 16 V, 0805	Any	Any	\$0.031	\$0.010
D1	1A, 20 V Schottky	MBRX120	MCC	\$0.120	\$0.040
L1	10µH, 360mΩ	NLCV32T-100K	TDK	\$0.121	\$0.040
R1	1.13Ω, 0805	Any	Any	\$0.102	\$0.034
U1	Controller	FAN5333A	Fairchild	\$0.585	\$0.195
PCB	1 in diameter				\$0.070
TOTAL					\$0.399
D2	LED	LXML-PW31	Lumileds	\$2.980	\$1.990
PCB	1 in diameter				\$0.070
TOTAL					\$2.46

that R1 has a voltage of 110mV across it, and D1 also has 360mV. Finally, L1 at 360mΩ at 100mA has another 36mV drop. Adding it all up, the duty cycle is $3.10V / (3.05V + 0.11V + 0.36V + 0.04V) = 87\%$, which is fortunately equal to the guaranteed duty cycle of the FAN5333A. If the duty cycle can't quite keep up, the battery will simply discharge directly into the LED for the little while the cells are at 1.55V. This won't be a problem; the LED will just be a little brighter for a while. The battery impedance and the I-V curves of the diode and LED will ensure that nothing catastrophic happens.

Thermal Model

Since we've got a power dissipation estimate for the LED (300mW), we can estimate the LED temperature. To do this, we need to create a thermal model for the

flashlight. What this means specifically is determining the thermal resistance of the path (or perhaps paths) from the LED to the ambient.

The first link in the path from the LED to the ambient temperature is the way the LED is mounted. For this design, we have opted not to use a metal-core PCB, as is customary, because of the cost. Instead, we're going to use a normal FR4 PCB with plated-through holes. We'll lay out the board to have a dozen such holes (artfully arranged in a circle around the LED, since the board is visible to users). Following the lead shown in the Lumileds' Application Brief,² we estimate the thermal resistance to the back side of the PCB to be (# holes in Lumileds' design/# holes in our design) * Lumileds' thermal resistance = (33/12) * 7 K/W = 19 K/W. We could decrease this by adding more holes, but the initial reaction of the marketing department was that that was aesthetically undesirable.

We've taken the LED thermal pad to the back side of the mounting PCB. For the next step in the thermal path, the PCB is glued to a small piece of aluminum. When we talk to manufacturing, they intend to use "some epoxy" for gluing. So we take the number for a typical epoxy in the same Lumileds document, 0.8 W/m-K for the thermal conductivity. Remember that the thermal resistance is one over the thermal conductivity multiplied by the length divided by the area. The thickness of the glue we presume will be the 100 μ m used by Lumileds. The PCB itself has a diameter of 1 in, and we're going to cover the back side of it with copper (no solder mask) so that this is the effective area in contact with the epoxy. We have $A = \pi (1 \text{ in})^2/4 = 0.239 \text{ in}^2 = 1.5 \times 10^{-4} \text{ m}^2$. We thus end up with thermal resistance through the epoxy of $10^{-4} \text{ m}/(0.8 \text{ W/m-K} * 1.5 \times 10^{-4} \text{ m}^2) = 0.8 \text{ K/W}$.

Thermal conductivity of aluminum is very high, something like 100 W/m-K. Further, all the aluminum pieces in this design are relatively thick. So we're just going to approximate the thermal resistance of the aluminum as zero. We have now progressed from the thermal pad of the LED to the aluminum case of the flashlight. But there's one more thermal resistance in the model: the resistance from the case to the air. As discussed in Chapter 6, the surrounding air is not really an infinite heat sink. It conducts heat through radiation and through convection. Which one dominates depends on the details of the arrangement and takes a lot of effort to compute. For practical purposes, we will use the estimation method of Figure 6.5. The power dissipation is 300 mW, and the surface area of the flashlight is $2 * \pi * (1/2 \text{ in})^2 + \pi * 1 \text{ in} * 8 \text{ in} = 26.7 \text{ in}^2$ (1 in diameter and 8 in long), so the power density is 0.01 W/in². Using a linear approximation, this is 1.0°C temperature rise, corresponding to a thermal resistance of 3.3°C/W.

Our completed thermal model is shown in Figure 10.3. We can now compute the typical LED temperature for lifetime calculation; the typical ambient temperature is 25°C. Thus LED temperature = ambient temperature + power * total thermal resistance = 25°C + 300 mW * (19 + 0.8 + 3.3 K/W) = 32°C. This very low temperature means that we don't have to worry about LED lifetime, and thus about MTBF. Clearly, if we ran the LED at its rated power of 3 W, it would be much hotter, around 94°C.

²Application Brief AB32, Lumileds, 10/08.




Figure 10.3 Thermal Model.

The other temperature we need is the maximum temperature of the aluminum. This occurs when the ambient temperature is maximal, 40°C. For this, the PCB and epoxy part of the string of thermal resistances don't matter. We have aluminum temperature = ambient temperature + power * air thermal resistance = 40°C + 300mW * 3.3 K/W = 41°C. Again, we are well below spec. (At 3 W this would be 50°C, still not a problem.)

Before we leave the thermal model, we should comment on two other thermal paths from the LED. Of course, for this flashlight, additional paths won't matter, since the temperature is low already. But in other designs they might matter, so a brief mention is in order. We already know from Chapter 6 that the air surrounding the LED is not very important for thermals. Indeed, for the Lumileds devices, the thermal resistance for this path is not even specified. And in any case, there is a lens on top of the LED, so any air circulating will not have a direct path to ambient temperature.

A more potentially interesting thermal path is through the wires that connect the PCB to the switch and battery. We can estimate the thermal resistance of this path the same way as all others. Thermal resistance = one over the thermal conductivity of copper times the length of the wire divided by the cross-sectional area of the wire. Using 2 in of AWG20 wire, we have $\Theta_R = 0.05\text{ m} / (400\text{ W/m-K} * 5 * 10^{-7}\text{ m}) = 250\text{ K/W}$. We have a positive and a negative wire, which gives 125 K/W, but clearly in this case this is insignificant compared with the main thermal path for the LEDs.

PCBs

Since all the numbers look good in this design, we won't need to go back and revise our assumptions. In particular, the typical thermal pad temperature of the LED is lower than we assumed (32°C rather than 60°C) and so the light output will be slightly higher than we estimated. If we were intent on squeezing out the last bit of performance from the batteries, we could re-estimate the light output and scale back the current from the driver. But this is likely to be a savings of a few percentage points at the most, and we forgo it in favor of getting the flashlight to market.

The last step in this design is to design the two PCBs, one for the driver and one for the LED. Since there is only one LED on the LED PCB, we will forgo showing it here, concentrating instead on the ballast PCB.

Our design is shown in Figure 10.4. The board is circular with a diameter of 1 in. (Yes, manufacturers make round boards.) It is obvious that with so few components in the design, the board could have been substantially smaller. The first thing

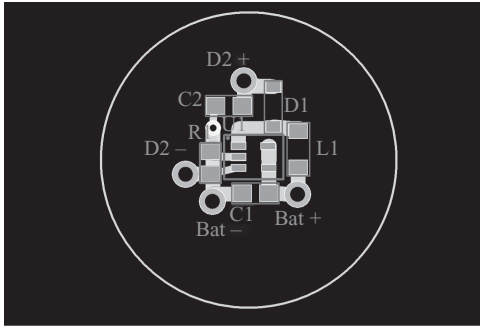


Figure 10.4 PCB for LED Flashlight Ballast. See color insert.

to note is that the connections to the battery and to the LED are holes. Since these items are located in a different part of the flashlight, manufacturing will run a pair of wires to each. The cost of doing this is very low.

C1, the input bypass capacitor, is located between the plus and minus battery inputs. These connections then go the shortest route possible to pins 4 and 5 of the IC for plus, and pin 2 for minus. We wanted the capacitor to be as close as possible to these pins, since this is where all the high-frequency components are located. The layout engineer had originally intended to do a single-layer board, but in order to get ground to pin 2 without a big loop, the engineer had to use two layers and add a via. There really isn't any cost difference between a one- and a two-layer board. It's really a boasting rights sort of thing: "I managed to do it on a single layer!" There's a cost difference only when four layers are needed.

The current sense resistor R1 is placed next to the current sense pin of the IC and ground. This minimizes the chance of noise pickup into the current sense circuit of the IC. The return connection from the LED is also right there, again to minimize noise pickup. L2 and D1 are also placed very close to the IC, again to minimize the possibility of interference. And finally, C2, the LED smoothing cap, is right next to the positive connection of the LED and as close as possible to ground.

The basic idea of this layout was to make it as compact as possible, since the IC switching frequency is so high. The sense resistor in particular was placed next to the IC pins, both signal and ground. In the final layout the power components were wrapped around more than the schematic showed in order to accomplish the tight packing. Maybe the schematic should have considered this, but we will leave it as is.

Final Design

In the end, our design didn't use any electrolytic capacitors. The concern about low temperature operation then comes down to possible moisture condensation on the PCB. We can take care of this by conformal coating the PCB. The final specification will thus go down to -40°C , as the marketing department requested in the first place.

It appears to be a good device; we'll have the marketing department get us one from the production line.

EXAMPLE: DESIGNING A USB LIGHT

Initial Marketing Input

For their next product, the marketing department wants to market a USB reading light. This is apparently inspired by the VP of marketing, who wants to see his keyboard on an airplane without turning on the overhead light. "I want to plug into the USB port on my laptop, and then have a flexible neck that leads up to a little LED. It doesn't have to be very bright. No, on second thought let's make it as bright as you can, and we'll put a control on it so I can adjust the brightness. I bet lots of people wish they had something like that."

Since there is only so much power a USB port can deliver, that sets the upper end of the brightness. If we use a resistor divider, the intensity will be nonlinear with the control setting, but, the VP says, "I don't think I care about that." How about an on/off switch? "Too much stuff. When the switch is at the bottom of its scale, you can turn it off then." And should it be totally off, or just really dim? "Oh, if you can't see it, it's okay. It can't be much power that way anyway, right? I guess I don't want to drain my battery." We promise to evaluate the standby current drain from the USB port.

Initial Analysis

The first thing is to look up the USB specification. Wikipedia³ informs us that "the USB 1.x and 2.0 specifications provide a 5V supply on a single wire from which connected USB devices may draw power. The specification provides for no more than 5.25V and no less than 4.75V ($5V \pm 5\%$) between the positive and negative bus power lines. For USB 2.0 the voltage supplied by low-powered hub ports is 4.4V to 5.25V.... A unit load is defined as 100mA in USB 2.0, and was raised to 150mA in USB 3.0.... All devices default as low-power but the device's software may request high-power."

To be backward-compatible, then, and to avoid having software, our light can't draw more than 100mA. And to be compatible with 2.0, the voltage can be as low as 4.4V. So our maximum input power shouldn't exceed 440mW. This means that for this design, we really don't need a 3W LED as we did for the flashlight; a 1W or even ½W LED would be enough. But again, Lumileds seems to have standard 3W LEDs, so that's what we're going to use.

Since there's a single LED, the output voltage is always going to be less than 4.4V. Therefore, we need a buck converter. We plan to use the LM3405, since its

³http://en.wikipedia.org/wiki/Universal_Serial_Bus#Mobile_device_charger_standards under license <http://creativecommons.org/licenses/by-sa/3.0/>

high switching frequency means the components will be small. Come to think of it, could it be made so small as to fit into the USB plug? A quick calculation suggests the power level is compatible with this idea. The datasheet suggests 84% efficiency at 5 V input; that's a loss in the power converter of only 70 mW. That shouldn't be a problem for something the size of a USB plug. Besides, it's plugged in to a USB port, which should offer some heat sinking.

Let's estimate light output. We have an input power of 440 mW and an efficiency of 84%, so power to the LED will be 370 mW. For this application, we imagine that the CCT doesn't matter that much and we will choose the same LXML-PW31 we used for the flashlight. (Remember to check this assumption with the marketing department.) With a device at about 100 Lm/W, we have maximum light output of about 37 lumens. (We'll add in the temperature effects in the more detailed analysis.) Using the estimate from Chapter 3, the illuminance on the keyboard will be about 176 lux. This is fairly bright and might be good for reading printed material.

We're also supposed to calculate the standby power. If we let the dimmer go down to 10%, the input power will be approximately 44 mW. The assumption for the approximation is that fixed overhead power draw, such as the IC operating current, is still a small contributor to total power. Since 44 mW at 4.4 V is 10 mA, and the IC draws 1.8 mA typically, this is a realistic approximation. A typical laptop battery has an ampacity of maybe 4000 mAh, so 10 mA current drain is a 0.25% capacity load. Another way of looking at it is that this load (even if it is linearly regulated down off the battery) would take 400 hrs or 2 weeks to drain the battery. We guess that this is okay, subject to marketing approval. After all, it will still be noticeable that the light is on, so you can unplug it; and USB ports typically are turned off when the laptop goes into hibernation.

The last thing is to estimate cost. The LED is priced at \$1.99. The LM3405 costs \$1.18 at 1,000 pieces, and we'll discount it 30% to estimate \$0.79. (We may really be able to get a substitute at the divide-by-three rule.) We'll guess the rest of the components to be \$0.30, giving a total of about \$3.00. The operations department hopes the mechanical parts are about \$1.50, and assembly cost for the PCB and the mechanical parts will be \$0.50. Thus the total cost will be about \$5.

Specification

We prepare a specification for this design, see Table 10.7.

The marketing department expresses some concern about the 10% minimum level. How dim is that going to be? Will it keep the VP awake? And is 10 mA really the best we can do? We point out that the USB port turns itself off when the computer is put to sleep, but agree to see how low we can go. Perhaps they'd like to turn the thing 100% off when the level is below 10%? This could be done with a comparator on the EN/DIM pin. "If you can do it without affecting the cost or footprint" is the answer. So that doesn't work out.

As for the CCT, they think that it should be more like a reading light. The color is somewhat more important for typing and reading than is the case with a camp flashlight. "What's the penalty for going down to 4000 K?" The answer is there's

Table 10.7 Initial Specification for USB Light

Parameter	Min.	Typ.	Max.	Units
Maximum Brightness		30		Lumens
CCT		5000		K
CRI		70		—
Illumination		175		Lux
Dim Range	10		100	%
Input Power at Maximum Brightness			440	mW
Input Power at Minimum Brightness		44		mW
Touch Temperature			60	°C
Ambient Temperature (by design)	10	25	40	°C
MTBF		50,000		Hours
BOM Cost			3.00	US\$
Total Cost			5.00	US\$

Table 10.8 Design Specification for USB Light

Parameter	Min.	Typ.	Max.	Units
Maximum Brightness		27		Lumens
CCT		4000		K
CRI		85		—
Illumination		160		Lux
Dim Range	10		100	%
Input Power at Maximum Brightness			440	mW
Input Power at Minimum Brightness		44		mW
Touch Temperature			60	°C
Ambient Temperature (by design)	10	25	40	°C
MTBF		50,000		Hours
BOM Cost			3.00	US\$
Total Cost			5.00	US\$

no cost penalty, and the CRI rises from 70 to 85, but the light is decreased 10%. Since it's pretty bright already, they like that better. The design spec is presented in Table 10.8.

Power Conversion

Now we can design the ballast. The first thing to figure out is the drive current. We approximated the LED power as 370 mW, which at 3 V is about 120 mA. The forward voltage of the LED at this current is about 3.12 V at 25°C, and drops by about $-3\text{mV}/^\circ\text{C}$. For our initial design, we'll assume the thermal pad in steady state is




Figure 10.5 First-Cut USB Light Schematic.

going to be about 60°C . We thus expect a forward voltage of $3.12\text{V} - 3\text{mV}/^{\circ}\text{C} * (60^{\circ}\text{C} - 25^{\circ}\text{C}) = 3.02\text{V}$. Dividing this into the power we get an LED target design current of 119 mA.

Our schematic is shown in Figure 10.5. We need to select values for the components L1, C1-5, D1, D2, and R1-3. Starting with the capacitors, the datasheet recommends $10\mu\text{F}$ ceramic for C1, with a X7R or X5R dielectric. A quick on-line check shows an 0603 size part for 8 cents at 1,000 pieces, so this will probably be suitable. We're using 0603 rather than 0805 because of the plan to fit the ballast into the USB plug. The manufacturer recommends $1\mu\text{F}$ in a 1206 package for the output capacitor C3. The same capacitance value is recommended for the compensation capacitor C4, but an 0603 will do since there isn't any power going through it. The boost cap C2 is recommended to be 10nF . Finally, C5 is a noise filter for our resistor divider. Since we expect to be using a $1\text{K}\Omega$ potentiometer for R2, we can't use a $1\mu\text{F}$ cap, as otherwise the LED will flash at turn-on while the FB pin voltage is rising. A cap of 1nF seems reasonable, which gives a time constant of $1\mu\text{sec}$.

Turning to the diodes, the datasheet says to use a 1N4148 for the boost diode D1. We initially choose a surface-mount version, but the layout department thinks they can do through-hole, which is slightly less expensive. For D2, the datasheet says to use a Schottky. We can pick one that is rated for the full output current of the LED, 120 mA. In fact, we can use the same 1A Schottky we used in the flashlight, the MBRX120, which is in a reasonably sized SOD-123.

To select the resistors, let's review how the divider works. When the potentiometer R2 is set to zero, the voltage on R3 will be the same as that on R1. As the value of R2 is increased, the voltage on R3 goes down, and to compensate, the LM3405 drives more current. So the value of R1 is determined by the minimum current. In our case, this is $1/10$ of the 119 mA, or 12 mA. The LM3405 is trying to hold the voltage at the FB pin to 205 mV. The value of R1 is thus going to be $205\text{mV}/11.9\text{mA} = 17.2\Omega$. There's insignificant power dissipated in the resistor, since $(205\text{mV})^2/17.2\Omega = 2\text{mW}$, and so we pick a 17.4Ω resistor in 0603.




Figure 10.6 LM3405 Schematic for USB Light.

The sourcing potentiometer R2 was assumed to be 1 K Ω . Presumably, what the people in the marketing department have in mind is a rotary wheel, not something that needs a “pot twaker”! We find one at a 1,000-piece price of 70 cents. The value of R2 in turn determines the value of R3. Since we want the dimming ratio to be 10:1, we pick 111 Ω in an 0603.

The final item is the inductor L1. The datasheet recommends 10 μ H for a 200 mA load. With this value, the datasheet suggests that peak-to-peak ripple current is 240 mA, so that our peak current would be 119 mA + (240 mA/2) = 239 mA. But we are worried. That’s double the current of the LED, which means it’s physically large. And if we approximate the triangular ripple current waveform as a sine wave, the RMS value would be $240 \text{ mA} / 2\sqrt{2} = 84 \text{ mA}$. The switcher is running at 1.6 MHz; how much skin-effect losses in the wire are there going to be?⁴ To minimize this, we’ll instead go with the highest inductance value the datasheet suggests, 22 μ H. This gives a 100 mA peak-to-peak ripple, giving peak current of 169 mA. We find a 4 mm \times 3 mm \times 4 mm 22 μ H inductor with a DC resistance of 2.1 Ω . Ignoring the skin effect, the losses are less than $(.169 \text{ A})^2 * 2.1 \Omega = 60 \text{ mW}$, a modest number. We expect the cost of such an inductor in volume to be about 10 cents.

The schematic is shown in Figure 10.6, with the BOM in Table 10.9. The total materials cost is \$3.39, gratifyingly close to our initial estimate of \$3.10.

Thermal Model

We need to estimate the temperature of the LED to make sure that we meet the lifetime, and also so we don’t burn it up. The people in the marketing department have decided they want a slick design for this light. The whole thing is going to be

⁴Skin effect is the tendency of high-frequency current to travel only on the outer skin of a wire. This increases the effective resistance and thus losses above that for a plain DC current. Refer to Lenk, 1998, for a discussion of skin effect.

Table 10.9 BOM for USB Light

Ref Des	Description	Part #	Mfr	1 Kpc Price	Est. Price
C1	10 μ F Ceramic, X5R, 0603	Any	Any	\$0.081	\$0.027
C2	10 nF	Any	Any	\$0.008	\$0.003
C3	1 μ F, X5R, 1206	Any	Any	\$0.132	\$0.044
C4	1 μ F, X5R, 0603	Any	Any	\$0.008	\$0.003
C5	1 nF	Any	Any	\$0.008	\$0.003
D1	1N4148, DO-35	Any	Any	\$0.045	\$0.015
D2	1A, 20V Schottky	MBRX120	MCC	\$0.120	\$0.040
L1	22 μ H, 2.1 Ω	4379-223KS	API Delevan	\$0.275	\$0.091
R1	17.4 Ω , 0603	Any	Any	\$0.004	\$0.001
R2	1000 Ω Pot, Rotary	3352T-102LF-ND	Bourns	\$0.698	\$0.233
R3	111 Ω , 0603	Any	Any	\$0.012	\$0.004
U1	LM3405	LM3405	National	\$1.180	\$0.790
PCB					\$0.070
TOTAL					\$1.238
D2	LED	LXML-PW51	Lumileds	\$2.980	\$1.990
PCB					\$0.070
TOTAL					\$3.39

encased in plastic, and they're putting a diffusion lens over the LED. The diffuser is going to lose about 10% of the light. This will reduce our brightness and illumination numbers by 10%, but since this was very bright already and the optical numbers were only typical, we'll just change the datasheet rather than changing the design.

More significantly for the design, encasing the whole thing in plastic gives the LED very high thermal resistance to the ambient temperature. In the case of the flashlight, the LED was attached to aluminum; plastic has about 1/500 the thermal conductivity that aluminum does.

In this case, it seems that most of the heat has to be transferred by the power wires down to the USB plug, where it will be dissipated into the metal case. Let's see how well we can do with these two wires. The distance from the LED to the output of the power supply is about 12 in. From our calculation in the flashlight example, we know that two strands of AWG20 wire running 2 inches have a thermal resistance of 125 K/W, so that 12 inches will be 750 K/W. With LED power dissipation of 370 mW, this gives a temperature rise of 278°C! We clearly want to cut this down by at least a factor of five, and 10 would be even better. How thick does the wire have to be? If we change from AWG20 to AWG18, the diameter increases from

0.032 in to 0.0403 in, a ratio of 1.259. The area ratio is the square of this, 1.586. The thermal resistance drops by this factor, as does the temperature rise, leaving us with a rise of 175°C, which is still too high.

Now, wire larger than AWG18 becomes hard to handle in manufacturing. Besides, it might make the neck of the light stiff. What would happen if we used two strands of AWG20 wire, one positive and one negative? With double the area, we get half the temperature, 139°C. This is still too high. How about using three wires? Now the temperature rise will be 93°C. At nominal 25°C, the LED will be at 118°C, and with a maximum ambient temperature of 40°C it will be at 133°C. The maximum is okay, since the rated temperature of the LED case is 135°C. As for the lifetime, at a typical case temperature of 118°C Lumileds's document⁵ shows 60,000 hours even at 350 mA, so this should be okay too. The real problem is that at this high temperature, the light output has dropped to 80%. We can compensate the forward voltage drop being lower than we assumed (another $-3 \text{ mV}/^\circ\text{C} * (120^\circ\text{C} - 60^\circ\text{C}) = -180 \text{ mV}$) by increasing the current to 130 mA, but we can't increase the power dissipated as that will send the temperature over the top. This is a 10% increase in current, and so the net light output is a further 10% less. We again can change the datasheet. Touch temperature will be unaffected because the plastic acts as an insulator.

PCB

This PCB is constrained by space considerably more than the previous one. The USB dongle is not very large (some of it is taken up with the connector), but the potentiometer is fairly large. Figure 10.7 shows a top view of the PCB on the left and a bottom view on the right. The bottom view is an "X-ray" view, that is, the view you would get by looking through the top of the board with X-ray vision. This makes it easy to see where connections go.

The outer two pins of the USB connector are 5 V and ground (the inner two are not used for this design). C1, the input bypass capacitor, is next to the IC. The LED return connection (marked "K" for cathode) has a very short run to R1, the current sense resistor, to minimize offset noise. This signal then routes around the board to get to the potentiometer R2. The presence of C5 next to R3 forms a good filter for the IC inputs. Note that the ground side of R1 goes right up to pin 2 of the IC to avoid offset voltages. For this design, since the current is very low and the switching frequency very high, we have concentrated on getting parts close together rather than on having a ground plane. The only other thing to notice is that the potentiometer R2 hangs out of the USB dongle slightly, allowing it to be rotated by the user.

Final Design

Our final spec and schematic are shown in Table 10.10 and Figure 10.8.

⁵Reliability Datasheet RD07, "Luxeon Rebel Reliability Data", 2007.




Figure 10.7 USB Light PCB, Top View and X-Ray Bottom View. See color insert.

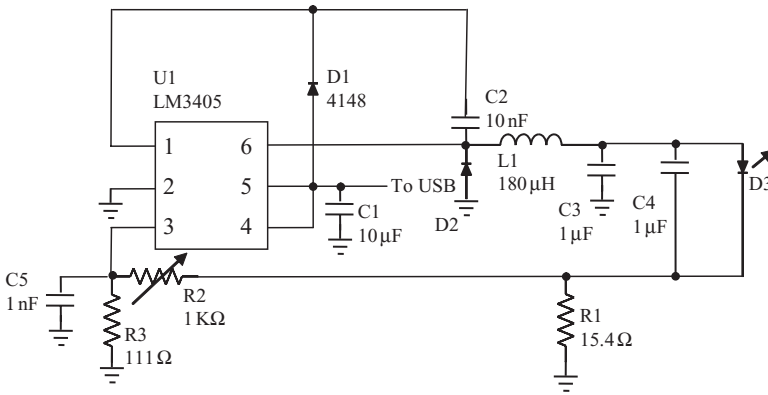


Figure 10.8 Final USB Light Schematic.

Table 10.10 Final Specification for USB Light

Parameter	Min.	Typ.	Max.	Units
Maximum Brightness		24		Lumens
CCT		4000		K
CRI		85		—
Illumination		140		Lux
Dim Range	10		100	%
Input Power at Maximum Brightness			440	mW
Input Power at Minimum Brightness		44		mW
Touch Temperature			60	°C
Ambient Temperature (by design)	10	25	40	°C
MTBF		50,000		Hours
BOM Cost			3.10	US\$
Total Cost			5.10	US\$

EXAMPLE: DESIGNING AN AUTOMOBILE TAILLIGHT

Initial Marketing Input

Our customer is an after-market provider of automobile lights, with a fairly good idea of what is wanted: “an inexpensive tail light made from 5 mm red LEDs. We want 160,000 mcd in stop mode, and 15% of that in taillight mode. It has to run over the normal automotive temperature range, -40 to $+65^{\circ}\text{C}$. We don’t want it to dim any when battery voltage is low, the battery can be as low as 7 V during cranking. Power comes directly from the battery, so it needs to be okay with a 60 V load dump. We want lifetime to be at least 10,000 hours. It needs to be a single string of LEDs; we can’t have some LED strings on and others off.”

This is all fairly specific. A rating in millicandela, though, doesn’t tell us how bright it has to be without specifying the solid angle over which the light is distributed. The customer doesn’t really understand the question; other companies are making tail lights this way, so that’s what they want. We take a guess that 60° is what will be available from the LEDs without optics.

Initial Analysis

We start by looking at 5 mm red LEDs in surface mount, limiting the search to those with ratings of at least 8000 mcd. All of these devices show fairly narrow angles, from 6° to 30° . We end up picking the device with the widest angle, a TOPLED from Osram, giving 10,600 mcd into a 30° angle at 50 mA. They cost \$0.23 at 1,000 pieces, so 16 of them (to get 160,000 mcd) would be \$3.68, and probably 1/3 of that in volume. This seems low enough that we won’t try driving the LEDs’ current harder than nominal.

These devices are 2.15 V forward voltage, and so 16 of them in series will be 34.4 V. At 50 mA, that’s 1.72 W. Since the input voltage ranges from 7 V to 60 V, the converter has to be a buck-boost. We’ll use the HV9910 circuit discussed in Chapter 7. This has the considerable advantage that it can take 60 V directly on its input pin, so no additional circuitry is required to handle load dump. On the other hand, it has a minimum operating voltage of 8.0 V, and there will also be some drop in the ORing diode (stop and taillight power are on separate lines). We will ask the customer if the cranking voltage requirement can be relaxed 1 V.

This circuit will provide a constant current to the LEDs, but how much variation will there be in their light output over the temperature range? The datasheet shows a curve that the light at -40°C is 1.4 times brighter, whereas at 65°C it is only 70% as bright. So we could either add $1/70\% = 1.4\times$ more LEDs, or we could sense the temperature and compensate for the efficacy.

Choice 1. The extra 40% LEDs would be seven devices; at \$0.23 each this would cost \$1.61. Since this is a 1,000-piece price this seems fine. What about board space? The devices are 3 mm \times 3.4 mm, so we’ll estimate total board space by allowing 4 mm \times 4 mm each, or a total area of 0.57 sq. in. That’s okay; the board will be

substantially larger than that anyway. The forward voltage will be increased to 49.4 V; we're still looking at buck-boost. The power is increased to 2.47 W, which should still be okay thermally.

What about reliability of 23 LEDs? If they're all in series, and even one fails, then the tail light is out. Still, what's the alternative? If we have two strings of 12 LEDs and one device goes out, the remaining string will produce only half the specified light. This doesn't meet regulations or customer spec, and may be detrimental if the driver is unaware that his tail light has dimmed. We'll just have to keep the LEDs cool enough that their MTBF can support the 10,000 hours.

Choice 2. A temperature compensation circuit can be added. We dislike thermistors because of their nonlinearity. While you can add resistors around them to linearize it, and the accuracy isn't that important, the real downfall of this method is consistent sourcing of the thermistors. As soon as the manufacturer substitutes a "direct replacement" part, the characteristics are completely different.

It is better to use an IC. A quick glance shows very inexpensive parts, \$0.19 at 1,000 pieces for a Microchip MCP9700. So this will certainly be cheaper than seven additional LEDs. On the other hand, it's going to take some circuitry around it. For example, the IC wants to run on 5 V. We could zener it down from the 7.5 V from the HV9910; but zeners need at least 1 mA to regulate, and that's the maximum the HV9910 can source. The other choice is to regulate it down from the battery. With a 7 V input, we need $(7\text{ V} - 5\text{ V})/1\text{ mA} = 2\text{ K}\Omega$. But now during load dump, this resistor has to dissipate $(60\text{ V} - 5\text{ V})^2/2\text{ K}\Omega = 1.5\text{ W}$, so we need a 2 W resistor. At 14 V nominal battery voltage, we are throwing away $14^2/2\text{ K}\Omega = 100\text{ mW}$, so it doesn't achieve average efficiency.

We also have to adapt the signal to the HV9910 FB pin requirements. The MCP9700 generates 500 mV at 0°C and increases by 10 mV/°C. Its range is then 100 mV at -40°C to 1.15 V at +65°C. The HV9910 wants to regulate at 250 mV. Since this is between these limits, we can probably design a resistor divider network that will properly scale the current with temperature.

One more consideration is that the MCP9700 takes 800 μsec to turn on. Before it does, the HV9910 will see 0 V in its circuit, which it will interpret as -50°C. Since this corresponds to very high efficacy for the LED, the driver will provide low current. So for 1 msec, the LEDs will be a little dim. This is okay; we were just checking that it didn't flash at turn on.

Output response to a 5 V step in the battery voltage is about 700 msec. This seems potentially more serious. When you take your foot off the gas pedal and apply the brake, the voltage tends to surge up a bit. This will produce a surge of current, as the HV9910 will think the temperature is very high. Thus, the tail light *will* flash.

This too can probably be fixed by adding a soft-start circuit to the HV9910 response, but if it takes 700 msec to turn on to full brightness, then you're actually slower than an incandescent bulb! One of the advantages to LED tail lights is supposed to be that they give the car behind you additional time to react. So a more complex system is needed.

In the end, we decide that the issues are complex enough that we will design with 23 LEDs. After the issues and costs are explained, the customer agrees.

Table 10.11 Specifications for LED Automobile Taillight

Parameter	Min.	Typ.	Max.	Units
Stop Brightness (V_{in} = full range, 25°C)		230,000		mcd
Stop Brightness (V_{in} = full range, Worst Case Temp)		160,000		mcd
Tail Brightness (V_{in} = full range, 25°C)		34,500		mcd
Emission Angle, ½ Brightness		+30		°
Dominant Wavelength		633		nm
LED Power		2.5		W
Input Power (V_{in} = 14 V)		3.0		W
Input Voltage	8.5	14	60	V
LED Temperature (T_{amb} = 40°C)			60	°C
Ambient Temperature (by design)	-40	25	65	°C
Board Width		5		Inches
Board Length		2		Inches
MTBF		10,000		Hours
LED Cost (1 K pcs.)			5.30	US\$
BOM Cost (1 K pcs.)			7.80	US\$

Specification

With this information in hand, we draw up a specification for the customer, shown in Table 10.11. We've taken a guess that 2.5 W won't raise the temperature of the LEDs more than 20°C. Since the part is rated at 100°C, we hope that when we estimate the lifetime, 60°C is enough to meet it. We've documented the desired size of the tail light. We've added \$2.50 for the price of the ballast components since we're showing them the 1,000-piece price. They have their own component sourcing.

The 8.5 V minimum input voltage is accepted. We're given the go-ahead to design this system.

Power Conversion

The circuit is shown in Figure 10.9. We've added a circuit that detects when the stop power is applied. In that case, the current sense resistor R3 is put in parallel with R5, generating the higher current required for stopping. When tail power is applied, R5 alone sets the current. We've included fuses because nothing should fail in such a way as to risk pulling large current from the battery. We've included an input inductor to form a basic EMI filter with C1. And we've included C3 to smooth out the current to the LEDs.

We need to select values for F1, C1, C3, L1, L2, R1-5, and D1-5. Let's start with R4 to set the frequency, as that determines the size of the inductor. For this design, we'll set a switching frequency of 50 KHz as a compromise between being far away from the 150 KHz edge of automotive EMI regulations and getting into large component sizes. We set $R4 = 475 \text{ K}\Omega$.




Figure 10.9 HV9910 Schematic for Taillight.

Given the frequency, we can determine the inductor value. First, we will use volt-second balance across the inductor to determine the duty cycle. The LEDs' forward voltage is about 49 V. Since the LEDs are in series with the inductor, this voltage is the voltage across the inductor while the transistor is off. When the transistor is on, the voltage across the inductor is the input voltage. The product of voltage and time has to be equal in these two conditions; otherwise the current in the inductor would change. Thus,

$$49V * t_{off} = V_{input} * t_{on}$$

By adding $49V * t_{on}$ to both sides, we can get a relation between the period, T , and t_{on} :

$$49V * t_{off} + 49V * t_{on} = V_{input} * t_{on} + 49V * t_{on}$$

which is

$$49V * T = (V_{input} + 49V) t_{on}$$

and since $t_{on}/T = \text{duty cycle}$,

$$DC = \frac{49V}{49V + V_{input}}$$

When $V_{input} = 8.5V$, $DC = 85\%$, and when $V_{input} = 60V$, $DC = 45\%$. T_{on} , which is D/f , is then $T_{on} = 17\mu\text{sec}$ at $V_{input} = 8.5V$, and $T_{on} = 9\mu\text{sec}$ at $V_{input} = 60V$. Since the inductance the datasheet recommends is proportional to T_{on} , the greatest value occurs at $V_{input} = 8.5V$. We end up with

$$L = \frac{V_{in} - V_{LED} * T_{on}}{0.3 * I_{LED}} = \frac{(60V - 8.5V) * 17 \mu\text{sec}}{0.3 * 50 \text{ mA}} = 5.8 \text{ mH}$$

The 0.3 in the denominator of the last equation came from the datasheet assumption that the peak-to-peak ripple current in the inductor is 30% of average. Thus peak current is 15% higher than average, which is to say 58 mA. Looking for inductors between 5 and 6 mH and from 50 to 75 mA rating, we see our choice for standard parts is just 5.6 mH. Choosing this, we see that they are all around 20 Ω DC. At 50 mA this gives power dissipation of 50 mW, a reasonable number. We choose the RL622-562 from Miller, with a circular footprint of 0.335 in and a height of 0.433 in. It has a 10% drop in inductance at 50 mA, so it should be fine at 58 mA. Since all of our parts are going to be automotive spec, we are going to use the straight distributor 1000 piece price, triple the normal volume pricing we assume. This inductor is \$0.315.

Next we turn to the transistor Q2. Maximum input voltage is 60 V, and the diodes add another 50 V on top of that. It's easy to pick a 200 V MOSFET, so we go for low price and small package. The ZVN3320 from Zetex lists at \$0.162 in 1000 pieces, and is in a SOT-23, nice and small. It's rated at 25 Ω , 60 mA continuous. With 50 mA of current and a 85% duty cycle, the on-state losses will be $(50 \text{ mA})^2 * 25 \Omega * 85\% = 53 \text{ mW}$.

We can also choose the diode D5 to have the same rating, 200 V. We look for a 1A part since those are common, and since they all pretty much have the same forward voltage, we pick an SMA package. We want something that isn't too fast, to avoid EMI problems. After all, the switching frequency is only 50 KHz. The RS1D seems suitable. Its power loss is only $700 \text{ mV} * 50 \text{ mA} * 88\% = 31 \text{ mW}$. Since the current is similar for D1 and D2, we pick the same part for these. The 200 V rating makes it good for protection against polarity reversal. Since D3 doesn't carry any current, we choose a surface-mount 4148 for it, giving enough voltage for polarity reversal protection.

We should also be ready now to pick the current sense resistors. We want 50 mA, with 15 mA peak-to-peak ripple, and so the value should be $250 \text{ mV} / (50 \text{ mA} + 0.5 * 15 \text{ mA}) = 4.35 \Omega$ when both resistors are in parallel. In taillight mode, the current should be $(50 \text{ mA} + 0.5 * 15 \text{ mA}) * 15\% = 8.6 \text{ mA}$. R5 should thus be $250 \text{ mV} / 8.6 \text{ mA} = 29.0 \Omega$. The closest standard value is R5 = 28.7 Ω . To get down to 4.35 Ω in stop mode, we need to parallel $(28.7 \Omega * 4.35 \Omega) / (28.7 \Omega - 4.35 \Omega) = 5.13 \Omega$. Now this resistance is the series combination of Q1 and R3. We want a Q1 resistance to be small compared to this, so that it doesn't affect the current much. Looking for something in the 20–30 V V_{DS} range, we find the DMG3420U-7, which is rated 25 m Ω with 4.5 V on its gate. So for R3 we can select 5.11 Ω . Since Q1 needs only 5 V on its gate, we choose D4 to be BZX84C5V1, and R1 and R2 to be 10 K Ω each.

As for power ratings for the resistors, with 250 mV present 85% of the time on R3, the power dissipation in this resistor is $(0.25 \text{ V})^2 * 85\% / 5.11 \Omega = 10 \text{ mW}$, so we can pick our standard 0805. R2 has a maximum power of $(5.1 \text{ V})^2 / 10 \text{ K}\Omega = 3 \text{ mW}$, and also gets an 0805. R1 has a maximum power of $(60 \text{ V} - 5 \text{ V})^2 / 10 \text{ K}\Omega = 300 \text{ mW}$ and gets a 1206.




Figure 10.10 LS E63F I-V Curve.

Source: Power TOPLED with Lens, datasheet, LS E63F, Osram, December 2009.

We still need to pick C3. This capacitor has to provide energy to the LEDs during the time the MOSFET is on, since during that time the inductor current is shunted to ground. The worst case scenario is when the MOSFET is on 85% of $20\mu\text{sec} = 17\mu\text{sec}$. During that time, the capacitor is providing 50 mA. The charge delivered is thus $Q = I t = 17\mu\text{sec} * 50\text{mA} = 850\text{nC}$. We want the LED current to remain approximately constant. Looking at the I-V curve in Figure 10.10, it looks like dropping from 50 mA to 40 mA corresponds to a forward voltage drop from 2.13 V to 2.10 V. So let's say the LED voltage shouldn't drop more than 10 mV per LED, or a total of 230 mV. Since $Q = CV$, we have $C = Q/V = 850\text{nC}/230\text{mV} = 3.7\mu\text{F}$. When we look for 4.7 μF at 63 V, it turns out there aren't any in a ceramic package. So we go to 100 V, a part in a 1210 package.

This leaves the input filter and the fuse. The purpose of C1 is to provide high-frequency bypass for the converter. It should probably be ceramic, even at 50 KHz. In fact, since it sees up to 60 V, it's convenient to simply pick the same 4.7 μF capacitor we just selected for the LEDs. Using the same component gives us some relief on pricing and minimizes stuffing errors.

By the same rationale, let's check if we could use the same inductor. We get a roll-off frequency of $1/(2\pi\sqrt{LC}) = 980\text{ Hz}$. This is a factor of $50 = 34\text{ dB}$ below the switching frequency, and so is a good choice for the filter inductor.

Finally for the fuse we find a problem. Our distributor carries automotive fuses, the smallest of which is rated for 20 A and costs over \$2.00! A quick call to the customer reveals that this component is part of the assembly driving the taillight, so we leave it off the BOM.

Our schematic and BOM are shown in Figure 10.11 and Table 10.12. We've ended up \$0.80 over budget, because, even though the LEDs are suitable, the con-

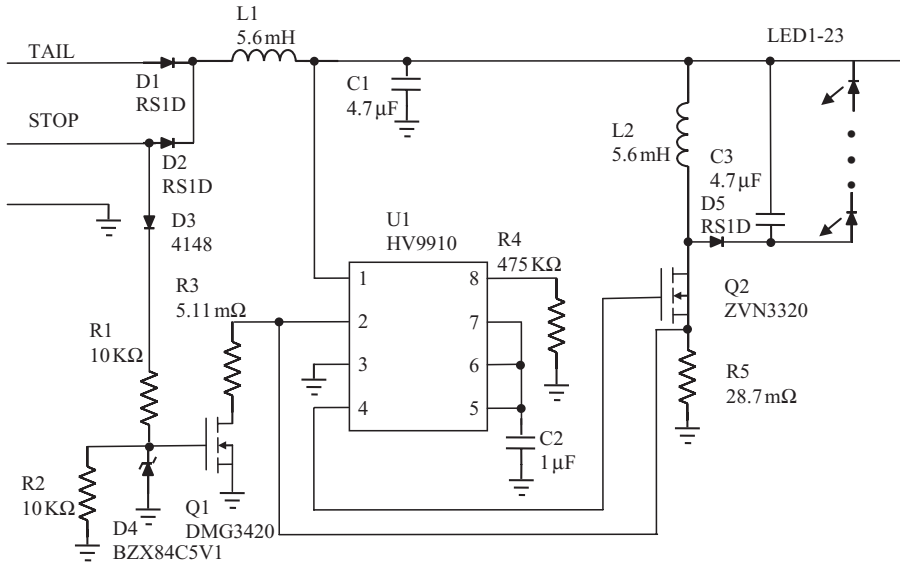


Figure 10.11 Final HV9910 LED Automobile Taillight Schematic.

Table 10.12 LED Automobile Taillight BOM

Ref Des	Description	Part #	Mfr	Est. Price
C1	4.7 μF, 100 V, 1206	Any	Any	\$0.286
C2	1 μF, 10 V, 0805	Any	Any	\$0.045
C3	4.7 μF, 100 V, 1206	Any	Any	\$0.286
D1	200 V, 1 A Ultrafast, SMA	RS1D	Any	\$0.189
D2	200 V, 1 A Ultrafast, SMA	RS1D	Any	\$0.189
D3	MMBD4148, SOT23	Any	Any	\$0.047
D4	5.1 V Zener, SOT-23	BZX84C5V1-TP	Any	\$0.051
D5	200 V, 1 A Ultrafast, SMA	RS1D	Any	\$0.189
L1	5.6 mH, 50 mA, 20Ω	RL622-562	J.W. Miller	\$0.315
L2	5.6 mH, 50 mA, 20Ω	RL622-562	J.W. Miller	\$0.315
Q1	20 V, 29 mΩ, SOT23	DMG3420U-7	Diodes Inc.	\$0.068
Q2	200 V, 25 Ω, SOT23	ZVN3320	Zetex	\$0.162
R1	10 KΩ, 1206	Any	Any	\$0.100
R2	10 KΩ, 0805	Any	Any	\$0.004
R3	5.11 Ω, 0805	Any	Any	\$0.004
R4	475 KΩ, 0805	Any	Any	\$0.004
R5	28.7 Ω, 0805	Any	Any	\$0.004
U1	Controller	HV9910	Supertex	\$0.725
PCB				\$0.210
TOTAL				\$3.335
LED1-23	633 nm LED	LS E63F	Osram	\$5.405
TOTAL				\$8.60

verter is \$0.80 more than expected. The most expensive items, in order, are the controller, the inductors and then the 4.7 μF caps. We recommend to the customer talking to the multiple vendors for the controller to get the price down, and checking with their Chinese manufacturer about cheaper replacement parts for the caps and inductors. They accept the design as is.

Thermal Model

For such a large board, there doesn't seem to be much power dissipation in our design, but on the other hand we've designed it so that the LED temperature rise should be small. We'll again use Figure 6.5 to form an estimate. We have 3.0 W of dissipation including the ballast. The surface area is 5 in \times 2 in = 10 sq. in. Power density is thus 0.3 W/in². Looking at the curve for a 40°C ambient, the temperature rise estimate is 40°C, which takes the temperature to 80°C, 20°C higher than the maximum temperature we had allocated.

We'll have to add some heat sinking to the design. The effective thermal resistance of the design right now is 40°C/3 W = 13.3°C/W. We need the temperature rise to be half of what it is with just the air, so we need an additional thermal path of 13.3°C/W. After going through some heat sink options, it looks like we're going to need a custom extrusion. This starts to sound worrisome, until it's pointed out that we forgot that we are already attached to a really big heat sink: the vehicle! All we need to do is to find an electrically insulating thermal interface, and bolt the board down through it.

We will go the Bergquist website (www.bergquistcompany.com), and look for thermally conductive insulators. The one called "economical," the Sil-Pad® 1100ST, is fiber glass-reinforced and suitable for automotive electronics. The thermal conductivity at rated contact pressure is 1.1 W/m-K. To find the thermal conductivity, we multiply by the area and divide by the thickness of the material, 1.1 W/m-K * (0.0065 m²)/0.00031 m = 23 W/K. The thermal resistance is the inverse of this, 0.043 K/W, considerably lower than needed. The LEDs will be at constant at ambient temperature. Similar materials at a distributor are priced at \$48.66 for a 10 in \times 12 in pad. This is big enough for 12 full boards, and using the divide-by-three rule we still end up with a cost addition of \$1.35 per board. This is half as much as the whole ballast! So in the end, we decide to only cover the ballast part of the board, since this is the only part with live connections exposed on the backside anyway. This is about 1.5 in \times 1.5 in = 2.25 in² = 1/5 of the board with this material, giving us five times the thermal resistance (still less than 1 W/K) and 1/5 the cost (\$0.27), which seems like an acceptable trade-off. The part of the board without the Sil-Pad can be insulated with fiberglass, or any other inexpensive electrical insulator.

PCB

The PCB for the entire assembly is shown in Figure 10.12. On the top, we see the LED layout. We wanted to make a symmetrical distribution of the light sources, and




Figure 10.12 LED Taillight. Top: Whole Board. Bottom Left: Topside Zoom of Ballast Area. Bottom Right: Bottom Side X-Ray Zoom of Ballast Area. See color insert.

also come close to the edge of the board. Of course the number of LEDs, 23, is a prime. What we've done here is to make an array of 3×8 LEDs, and replace the middle two with the single LED in the center. This is the most symmetric arrangement possible. There is a three-pin card edge connector on the bottom center for input power. The fuse isn't included, as decided above.

The traces connecting the LEDs are relatively large, but with only 50mA current, they don't completely use all the space available. Notice that the layout of the LED pads requires their connection from the lower-right hand corner to the upper-left hand corner, without connecting to the other two corner pads. It required some thought to organize the array in such a way.

In the middle you can see the ballast. The top side is shown on the bottom left of Figure 10.12, and the back side on the bottom right. As promised, it is smaller than 2×1 in. The positive input connector comes in through the two diodes and then to a via to the large trace on the backside to L1. Since L1 is through-hole, its leads are used as vias. The other end of L1 is power to the system, and goes by means of the copper pour on the topside to L2, C1, C3, and most importantly to LED1.

Ground is done much the same way. It comes in on the trace on the topside to the large pour on the bottom side. This then connects to the large trace on the topside

as the IC ground and the ground for the input capacitor C1 and the current sense resistor R5. Small-signal ground for R1 and C2 is made with a trace on the top-side and a single via down to the power ground plane. Again, this is to ensure that small-signal information to the IC isn't upset by the currents flowing through the power devices. Note that the ground plane on the bottom side is underneath the IC, again to minimize pickup into sensitive nodes and to help with EMI.

Final Design

The result is the BOM described in Table 10.13. The price is \$0.81 higher, because of the Sil-Pad, but the customer seems to expect this sort of price over-run. We make a note to ourselves to look up how much retail pricing is on the Internet. The volume pricing will be only \$0.27 higher anyway. (And we know that cheaper LEDs are on their way.)

Table 10.13 Final LED Automobile Taillight BOM

Ref Des	Description	Part #	Mfr	1 K Est. Price
C1	4.7 μ F, 100 V, 1206	Any	Any	\$0.286
C2	1 μ F, 10 V, 0805	Any	Any	\$0.045
C3	4.7 μ F, 100 V, 1206	Any	Any	\$0.286
D1	200 V, 1 A Ultrafast, SMA	RS1D	Any	\$0.189
D2	200 V, 1 A Ultrafast, SMA	RS1D	Any	\$0.189
D3	MMBD4148, SOT23	Any	Any	\$0.047
D4	5.1 V Zener, SOT-23	BZX84C5V1-TP	Any	\$0.051
D5	200 V, 1 A Ultrafast, SMA	RS1D	Any	\$0.189
L1	5.6mH, 50 mA, 20 Ω	RL622-562	J.W. Miller	\$0.315
L2	5.6mH, 50 mA, 20 Ω	RL622-562	J.W. Miller	\$0.315
Q1	20 V, 29 m Ω , SOT23	DMG3420U-7	Diodes Inc.	\$0.068
Q2	200 V, 25 Ω , SOT23	ZVN3320	Zetex	\$0.162
R1	10 K Ω , 1206	Any	Any	\$0.100
R2	10 K Ω , 0805	Any	Any	\$0.004
R3	5.11 Ω , 0805	Any	Any	\$0.004
R4	475 K Ω , 0805	Any	Any	\$0.004
R5	28.7 Ω , 0805	Any	Any	\$0.004
U1	Controller	HV9910	Supertex	\$0.725
PCB				\$0.210
<i>TOTAL</i>				\$3.335
LED1-23	633 nm LED	LS E63F	Osram	\$5.405
—	Electrical Insulator	Sil-Pad 1100ST	Bergquist	\$0.870
TOTAL				\$9.61

EXAMPLE: DESIGNING AN LED LIGHT BULB

Initial Marketing Input

A company has decided to move into the field of lighting. For their first product, the sales vice president says, they want “a floodlight light bulb.” The VP continues, “I have a 14-ft ceiling in my house, and it’s a bear to change out those light bulbs. So I want something that’s going to last as long as my house, so I don’t ever have to get out the ladder again. Make it as bright as you can, it’s a long way down to the floor. I don’t think we have to worry about the price, there’s lots of really expensive stuff out there right now, and we’ll make it up in volume.”

A short discussion about spotlights versus floodlights establishes that the VP does in fact want a floodlight, so we don’t need special optics to control the angle of emission. The discussion then turns to performance tradeoffs. The cost is not the only factor limiting the brightness. To make the light bulb brighter, the LEDs are also going to have to be hotter. And hotter LEDs means the MTTF (mean time to failure) is going to be less. Here the answer is that “it has to last longer than a CFL. But give me a table of brightness versus life so we can choose.”

Initial Analysis

There are going to be two sources of heat in this design, the ballast and the LEDs. Let’s estimate that we can put about 10W into this design. (If we’re off 50% either way it won’t affect our zero-order conclusions.) If we have 85% efficiency on the ballast, that’s about 1.5W in the ballast, and the rest in the LEDs. Since reflectors are relatively large, we’re going to assume that the only real problem is to get the heat out of the LEDs; the ballast isn’t going to overheat.

So how much heat can we really put in? The first concern to note, as usual, is that “floodlight” doesn’t really capture all the design information we need. We have to know the size of the bulb to determine the thermals. Looking at General Electric Lighting web page⁶ on indoor floodlights shows that the word floodlight encompasses bulb sizes of BR30, BR40, MR16, PAR16, PAR20, PAR30, PAR30L, PAR38, R20, R25, R30, and R40. We narrow this down by looking just at medium screw bases, which shortens the list to BR30, BR40, PAR30L, PAR38, R20, R25, R30, and R40.

Further examination shows that a lot of these sizes are available only as compact fluorescent bulbs. We’re tempted to say that the marketing department wants only the incandescent replacement sizes (BR30, BR40, and R20) because the CRI of the fluorescent bulbs is only 82, and he wants his house to look beautiful. Of course,

⁶http://genet.gelighting.com/LightProducts/Dispatcher?REQUEST=RESULTPAGE&CHANNEL=Consumer&FILTER=FT0010:Track+%26+Recessed_Indoor+Floodlight&CATEGORY=Lamps&BREADCRUMP=Track+%26+Recessed_Indoor+Floodlight%230



Figure 10.13 BR Bulbs.

this will be a problem for the LEDs also. On the other hand, perhaps energy-conscious consumers, the type who will spend a lot of money for an LED bulb, already are using CFLs and don't mind the CRI. Or maybe they're dissatisfied with the CRI of fluorescents, and that's why they want to change to LEDs.

The marketing department states authoritatively that they had the BR style in mind ("they don't have that odd shape the CFLs come in [R]"). We tell them that the brightest will be the largest, and so we settle on the BR40 size bulb.

This begs the question of CCT. The marketing department of course wants the bulb to look "as close to incandescent as you possibly can get it." After pointing out that we can't get a perfect CRI, we further mention that what they're asking is a CCT of 2750K, which will have significant impact on the brightness. At this point, they ask for a table detailing the brightness versus CCT choices.

Now we're ready to do some calculations. We need two tables, one showing the achievable brightness versus lifetime, the other showing the achievable brightness versus CCT. For the lifetime table, we can hold the CCT constant at 2750K, and increase the input power. For the CCT table, we'll hold the LED temperature constant, and take the power level from the previous table.

The key parameter in calculating the lifetime is the temperature of the LED. Fortunately, our company's preferred vendor, Lumileds, has provided graphs of lifetime versus temperature and drive current. All we have to do is to look up the answer once we've determined these two. And the temperature isn't the die temperature, which would depend on the power dissipation in the LED. The lifetime is based on the temperature of the case. This is independent of the current into the LEDs, as long as the total power into the system is constant. For other LED vendors, whose device lifetime depends on die temperature, an iterative calculation can be used to determine die temperature after case temperature has been determined.

To determine the case temperature what is needed is a thermal model. Since a BR40 goes into a can, we'd better assume that the light bulb ambient goes up to 50°C. We'll mount the LEDs onto a MCPCB, which Lumileds⁷ suggests can achieve 5 K/W. The PCB will in turn be mounted to a large piece of aluminum, which for our purposes we'll assume to have zero thermal resistance. Finally, there is thermal resistance to the ambient. This depends on a number of factors, such as bulb size versus can size, and whether the can is insulated or not. We can't use the estimates

⁷Lumileds, Application Brief AB32, 2008.




Figure 10.14 Initial Thermal Model for BR40.

Table 10.14 Decreasing the LED Temperature Requires Decreasing Light Output

Thermal Pad Temp (°C)	Lifetime (Hrs.)	Efficacy (%)	Power (W)	Light (Lumens)
140	9,000	75	6.0	450
135	12,000	75	5.7	428
130	20,000	76	5.3	403
125	25,000	77	5.0	385
120	47,000	78	4.7	367
115	60,000	80	4.3	344
110	60,000	81	4.0	324
100	60,000	83	3.3	274
90	60,000	85	2.7	230

of Figure 6.5 because the can doesn't allow free access to the air for either convection or radiation. We'll take an initial, hopefully conservative, guess of 10 K/W, and measure it later. We therefore estimate the total thermal resistance from LED case to ambient to be 5 K/W + 10 K/W = 15 K/W, with an ambient temperature of 50°C. Since the maximum temperature on the lifetime curves is 140°C regardless of current, this model immediately shows that we won't be putting more than $(140^\circ\text{C} - 50^\circ\text{C})/15 \text{ K/W} = 6 \text{ W}$ into the design (see Figure 10.14).

At what current should we drive the LEDs? Although this doesn't affect the case temperature, it does still affect the lifetime. Let's just try all three currents shown in the curve, 350 mA, 700 mA and 1 A. These correspond to 1 W, 2 W, and 3 W at the accuracy we're trying to achieve for this initial calculation.

At 350 mA, the maximum case temperature to reach 10,000 hours (like a fluorescent) is 140°C. This corresponds, as we just calculated, to 6 W. At 700 mA, 10,000 hours is at 120°C, so power is $(120^\circ\text{C} - 50^\circ\text{C})/15 \text{ K/W} = 5 \text{ W}$. And at 1 A, we can dissipate $(110^\circ\text{C} - 50^\circ\text{C})/15 \text{ K/W} = 4 \text{ W}$. So as usual, the choice is between cost and performance: 4 W at 1 A is just two devices, so this is minimum cost. 6 W at 350 mA is six devices, but this will be the brightest design. Presumably, we could run the LEDs at an even lower current and allow even higher power input. But the absolute maximum junction temperature is 150°C, so the best that we could possibly obtain would be $(150^\circ\text{C} - 50^\circ\text{C})/15 \text{ K/W} = 7 \text{ W}$, which is not that much brighter than the 6 W we can get at 350 mA, and considerably more expensive.

So we're going to look at the table of brightness versus lifetime at 350 mA drive current. We use 6 W power and an efficacy of 100 Lm/W at 25°C. It is no surprise that, as shown in Table 10.14, the brightest will be at the highest device temperature.

Table 10.15 Sensitivity of Light Output versus LED Temperature to Thermal Resistance

Thermal Pad Temp (°C)	Lifetime (Hrs.)	Efficacy (%)	Power (W)	Light (Lumens)
140	9,000	75	9.0	675
135	12,000	75	8.5	638
130	20,000	76	8.0	608
125	25,000	77	7.5	578
120	47,000	78	7.0	546
115	60,000	80	6.5	520
110	60,000	81	6.0	486
100	60,000	83	5.5	457
90	60,000	85	5.0	425

Table 10.16 The Higher the CCT, the Brighter the Bulb

LED Type (Lm)	Typ. Efficacy (Lm/W)	CCT (K)	Brightness (Lm)	Cost (US)
50	63	3000	378	\$20.10
70	76	3500	456	\$17.88
80	90	4000	540	\$17.88
90	100	5000	600	\$17.88
100	91	3500–4000	546	\$17.88

The increase of efficacy with decreasing temperature isn't enough to compensate for the decrease in power to achieve that temperature.

Since we just took a guess at the thermal resistance to the ambient of this light bulb, it's reasonable to ask what the sensitivity to 15 K/W is. Let's rewrite this table with a thermal resistance of 10 K/W. The only change will be in the last two columns, as shown in Table 10.15, since we can put more power in with lower thermal resistance (9W instead of 6W).

The result is the same: higher temperature produces more light. The same insensitivity will be shown if the ambient temperature is lower or higher than the 50°C we assumed.

The other thing that the marketing department wanted to know was brightness versus CCT. Using the 6W number, we obtain the results shown Table 10.16. As expected, the higher the CCT, the brighter the bulb. We've included the price curve for the LEDs, showing that the cost is decoupled from the brightness-CCT consideration.

Specification

We present Tables 10.14–10.16 to the marketing department. Their first comment is that 6W isn't much power, and the second is that 378 lumens isn't even as bright

Table 10.17 Specifications for LED BR40 Light Bulb

Parameter	Min.	Typ.	Max.	Units
3000 K Bulb				
Target Brightness		400		Lumens
CCT	2950	3000	3150	K
Production Cost Target		25		US\$
5000 K Bulb				
Target Brightness		600		Lumens
CCT	4750	5000	5300	K
Production Cost Target		23		US\$
Common Parameters				
Production Brightness Tolerance	-10		10	%
Input Voltage	0	120	135	VAC
Lightning Protection		Yes		
Input Power		7		W
AC Conversion Efficiency Target		85		%
PFC	0.7			
Dimmer Compatibility		No		
Ambient Temperature (by design)	10		50	°C
MTBF		20,000		Hours
Dimensions		BR40		

as a 40 W incandescent! We explain that these are just estimates; we may be able to improve the figures in the actual design. Since they have to pick a number for life-time, the team goes for 20,000 hours, since this is only a 10% decrease in brightness.

As for CCT, this is something of a revelation to the people in the marketing department. A 40 W incandescent bulb costs the same as a 75 W bulb, and a 9 W CFL costs the same as a 23 W. But with LEDs, the brightness depends on the cost. More light, more money. This is actually what normal market pricing for commodities ought to be like, but it's going to take some getting used to. Those in the marketing department still want their incandescent color bulb, but they also want the brightest bulb possible. The decision is to make two light bulbs, one at 3000 K, the other brighter at 5000 K.

We thus draw up a specification, Table 10.17, for marketing approval before the design begins.

There are a number of issues which are indirectly addressed in this specification. The target brightness is in reality the minimum that the marketing department will accept. We need to improve this if at all possible. In production, we need to hold the tolerance of the light output, so we specify this as a percentage of nominal. The CCT tolerance is from the Lumileds binning information. The production cost target

is the cost of the LEDs, plus \$1 for the ballast, \$2 for the materials, and \$2 for the labor and shipping to the United States.

The input voltage can go as high as 135VAC, which is standard. Of course, the bulb needs to be protected from lightning. Producing light when the line voltage goes down to zero is of course not possible. What is meant is that the unit won't be damaged by any line voltage less than nominal. The problem is that supplying the LEDs with constant power means the input current has to increase as the input voltage goes decreases. Many components will dissipate more power as current increases, potentially damaging them. This could in theory be prevented by turning off the ballast when the line voltage goes below a previously determined value. But the other way to accomplish this is to have the ballast be corrected for power factor. That way, the power decreases as the voltage decreases, so the current doesn't increase. Since we know how to do this inexpensively, we decide to include PFC. The marketing department likes this, anyway, since that means they can add an Energy Star rating on the bulb. And decreasing brightness is what incandescent bulbs do anyway, so the marketing department won't expect customers to object.

We target an AC conversion efficiency of 85%, which should be acceptable as there is a fair amount of room inside a BR40 for the ballast. The input power is the 6W for the LEDs divided by this efficiency. We opt out on dimmer compatibility, since this is difficult to achieve. Perhaps it can go in the next generation, or maybe a premium can be charged for it. No typical ambient temperature is specified. There are many different environments for the bulb, and it needs to have at least its rated MTBF for all of them.

The bulb shape is specified to be a BR40. We discuss some of the options here with those in marketing before deciding this. There are of course the usual trade-off arguments. On one hand, the marketing department thinks the bulb should be exactly the same as the incandescent shape, so that consumers don't have to get used to something new. On the other hand, that shape was certainly not designed to optimize the thermal performance of LEDs, so that maybe it should be changed. After all, once it's in the ceiling in a can, no one pays any attention to what shape it is.

There's an additional consideration here. It might be supposed that the BR40 is so large (5in at its widest) that it will block up the possibility of air flow in a 6-in can. Maybe air flow can be used to increase the power dissipation capability? Unfortunately, this idea doesn't work. Since the can and its lamp are pointing downward, the hot air goes up into the can. Air circulation is probably minimal inside the can. Further, the LEDs are presumably at the top, flat end of the bulb. Thus there really isn't much reason to deviate from the standard shape, and a really good reason not to make it smaller: That would reduce the top surface area which is the primary site for heat dissipation. One further choice would be a PAR38, which is almost the same size as the BR40, but with a flatter top. If we can somehow shorten the path from the top to the ambient, that would help. We hold this idea in our back pocket.

The one factor left out of our calculations and the spec is how much light loss to expect in the glass cover over the front of the bulb. We hope to hold this down to 10%, and tell those in marketing that we hope to make it up elsewhere.

Power Conversion

An important question not addressed in the specifications is the status of line isolation for the light bulb. If the ballast is not isolated from the AC line, the LEDs will have to be. This is both for safety and for meeting UL requirements. Nothing with a galvanic connection to the AC line can be exposed to touch (see Chapter 8). On the other hand, if the ballast is isolated, the LEDs can be directly exposed, which will help their thermals. But then the ballast is substantially harder to design. In the end, we compromise. The main heat dissipation mechanism in the system is going to be the aluminum plate in the front. The LEDs from Lumileds are already rated for UL safety. Insulation is necessary, then, only in the pads and traces. Since these occupy a relatively small percentage of the plate, we will insulate just these, leaving the rest of the plate galvanically isolated from power. This accomplishes the best of both worlds. The ballast can be nonisolated, and the LEDs have good thermal connection to ambient. And in addition, there is no glass blocking the light. This saves maybe 8–10% of the light output from the LEDs, as we promised the marketing department!

With the decision to use a nonisolated system, the biggest challenge in this design is still the conversion from AC line voltage to LED drive current. We have to provide well regulated output over a variety of line voltages, have it power-factor corrected, meet EMI, and so on. The first concern to address is how much room we have to accomplish all of this. A look at the drawing in Figure 10.13 shows that there is a fairly cylindrical neck space available for the ballast. Using a ruler, we estimate that the width is about $1\frac{7}{8}$ there. We don't expect to be using very heavy glass there, so we can make the ballast circuit board with a diameter of 1.8 in. The neck is long, and so we don't expect to have a significant height limitation on components.

We're going to start with an HV9910 design, as shown in Figure 8.7, reproduced here as Figure 10.15. We've changed the schematic to show six LEDs in series. We've added in the power factor correction circuit of Figure 8.17, and removed the input electrolytic capacitor. Finally, we've added an input inductor for EMI, an MOV for lightning and a fuse. Note that the fuse is closer to the line than the bridge and the MOV: if the diodes short or the MOV fails, we want to make sure the circuit isn't a fire hazard.

Let's start with the easy things first. We're only pulling 7W from the AC line, so the current at 120VAC is only 60mA. Remember that since this circuit is PFC, the current decreases as line voltage decreases. So for the fuse we can select something with about double that current rating, a 1/8A fuse. To keep the size small, we'll go with a through-hole device, the MQ125. For the same reason, the diodes D1-4 are nothing special. We can choose 1A devices; there isn't much difference in size or cost to go lower in current. And in line with the principles discussed in Chapter 8, we select 600V parts, the 1N4004.

The most critical decision for the rest of the components is the switching frequency of U1, the HV9910. Knowing this allows us to design everything else. We follow the guide shown in Chapter 8 and settle for a switching frequency of just less




Figure 10.15 First Attempt at a BR Design.

than 150 KHz. The tolerance on the switching frequency of the IC is $\pm 20\%$ at 25°C . Let's add in another 5% for temperature variation, and 1% for the resistor tolerance. To be sure we don't hit 150 KHz, then, we need to choose the switching frequency to be $150\text{ KHz} * (0.8) * (0.95) * (0.99) = 113\text{ KHz}$. And since the datasheet shows that the IC is apparently tested at 100 KHz, we just round down to this frequency, and use R1 as 226 KΩ.

With the switching frequency in place, we can select the inductor. Following the information in the HV9910 datasheet, we have the duty cycle as a function of input voltage $D = V_{\text{LEDs}}/V_{\text{in}} = 6 * 3.15\text{ V}/V_{\text{in}}$. The on-time of the switch is $T_{\text{on}} = D/F_s = 18.9/(V_{\text{in}} * 100\text{ KHz})$. The recommended inductance is then $L = (V_{\text{in}} - V_{\text{LEDs}}) * T_{\text{on}} / (0.3 * I_{\text{LED}}) = (V_{\text{in}} - 18.9) * 18.9 / (V_{\text{in}} * 100\text{ KHz}) / (0.3 * 0.35\text{ A}) = 1.8\text{ mH} - (34\text{ VmH}/V_{\text{in}})$. Clearly, the greatest inductance is needed at the highest line voltage. Since this design is PFC, the peak voltage is $120\sqrt{2} = 168\text{ V}$, giving an inductance of 1.6 mH.

Actually, this calculation has an error in it. Did you spot it? We want 350 mA *average* through the LEDs. But the ballast is PFC, and so the current is a function of the line voltage. In particular, to get 350 mA at the average voltage (120 V) we need $350\text{ mA} * \sqrt{2} = 500\text{ mA}$ current at line peak. Substituting this corrected current into the inductance calculation, we find that the actual inductance required is 1.1 mH.

Now the datasheet and inductance formula assumed that the ripple current was 30% peak-to-peak of the current. At line peak we have 500 mA, and so the peak current is $I + (30\%/2) * I = 575\text{ mA}$. We need to find an L2 that is 1.1 mH at 575 mA. Looking at some options on the web between 1.0 and 1.5 mH and between 570 mA

and 700 mA, we find that as usual there is a tradeoff between size and resistance. A very small inductor ($6.6 \text{ mm} \times 4.5 \text{ mm} \times 2.9 \text{ mm}$) is available with a resistance of 13.8Ω . This would yield a power loss of $(350 \text{ mA})^2 * 13.8 \Omega = 1.69 \text{ W}$, definitely too high. The next lowest resistance one is 6.0Ω , which would be 735 mW , still 10% of input power. In the end, we end up with one that is 1.0 mH at 1.54Ω , which is 189 mW , in a $12 \text{ mm} \times 12 \text{ mm} \times 8 \text{ mm}$ package, the SPD127R-105M from API Delevan.

Knowing the current, we can now select the current sense resistor, R4. At line peak we're going to set the PFC circuit to produce the 250 mV that is full gain for the HV9910. So we want the line peak current in the transistor to be 575 mA . (Remember that the HV9910 trips off at a current, so we have to include the ripple to make sure we get enough current.) Since the current sense pin trips at 250 mV , we want $R4 = 250 \text{ mV}/575 \text{ mA} = 435 \text{ m}\Omega$. The closest common value is $430 \text{ m}\Omega$.

With knowledge of the current, we can also select Q1 and D5. For the transistor, the square of the average current is only 0.12 A^2 . So we pick something like a 1Ω $R_{\text{DS,on}}$ which only dissipates 120 mW . Even at line peak, it's only dissipating 250 mW . It's not sensible to pick a transistor lower in resistance, because it would cost more money. Besides, lower $R_{\text{DS,on}}$ also is accompanied by greater gate capacitance, which would increase our transition losses. In the end, we select an STD5NM60T4, a 600 V device.

D5 is also carrying 350 mA average. It's being switched at 100 KHz , so it needs to be ultra-fast. A 1 A diode would certainly work, but since we have the room, we select a 3 A . That way the forward voltage loss will be lessened. We also prefer slower recovery rather than faster, to minimize EMI. We select an RS3J-13-F.

Let's do the PFC circuit next. We want a resistor divider that produces 250 mV when there is 169 V input. At the same time, we don't want to dissipate too much power in the top resistor either. Picking the largest practical value, $1 \text{ M}\Omega$, for R2, we need R3 to be $1.48 \text{ K}\Omega$. The closest common value is $1.5 \text{ K}\Omega$. R2 dissipates only $(120 \text{ V})^2/1 \text{ M}\Omega = 14 \text{ mW}$.

Some experience suggests that a $1 \text{ M}\Omega$ input impedance to a circuit tends to be noisy. So we'll add a capacitor in parallel with R3. With a rectified period of 8.3 msec , we want to track the line voltage with, say, not more than $800 \mu\text{sec}$ delay. With the $1 \text{ M}\Omega$ resistor, this time constant implies an 800 pF capacitor. We'll round down to the common value 680 pF .

The only components left to design are the input L and C, and the MOV. Turning to the EMI filter first, the input capacitor should not be too large, so that we don't have power factor problems. We limit ourselves to 100 nF , which we find at 630 V , available at reasonable cost with 10% tolerance in an 1812 package.

Given the capacitor, the inductor is determined by how much roll-off we want at the switching frequency. For example, if we place the filter (double) pole at $1/10$ the switching frequency, 4.5 KHz , we get 40 dB of attenuation at the switching frequency. The inductor for 4.5 KHz has to be 12.5 mH . With input current at 60 mA at 120 VAC , it will be root two higher at line peak, 85 mA . Looking in the range of $12\text{--}18 \text{ mH}$ and $85\text{--}120 \text{ mA}$, we find inductors have typical resistance of 24Ω . Since

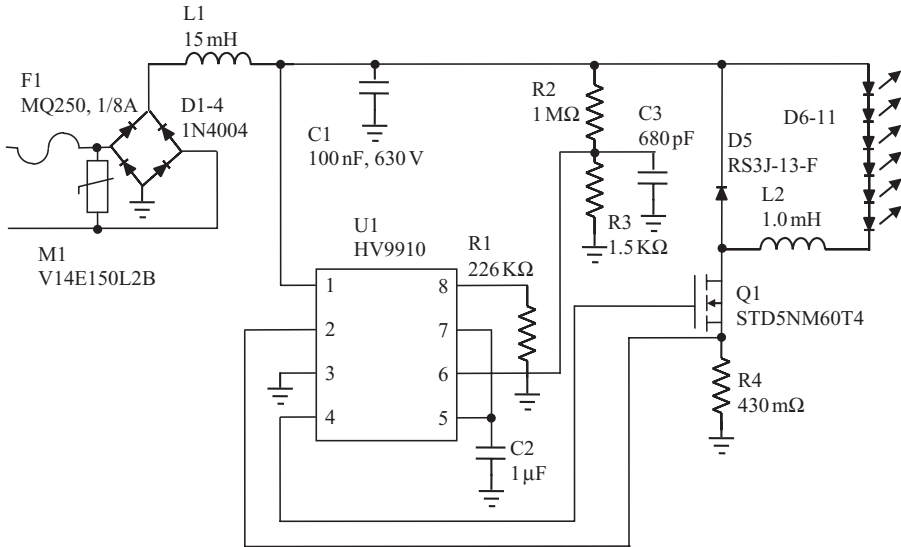


Figure 10.16 Final HV9910 Schematic for LED BR40.

at 60mA this is only 85mW, we pick the TSL1112RA-153JR12-PF, which has 15mH at a rated current of 120mA.

We've left the MOV for last because it's not determinate. You just pick the biggest one that will fit. We'll choose a 20mm disk, and let the layout person complain if it doesn't fit. The completed schematic is shown in Figure 10.16, with the BOM in Table 10.18. \$13.67 for volume cost of the electronics seems pretty good; we should be able to make a profit at that even after the mechanicals are added in.

PCBs

The PCB layout for the ballast is shown in Figure 10.17. The diameter of the board is 1.8in. The thing that immediately jumps out is that the MOV is by far the largest component on the board. If the light didn't have a huge neck, we would have had to shrink the MOV down in order to make the design fit.

The input stage is all at the top of the PCB. "N" and "H" are the two inputs (neutral and hot). The traces have at least 50mil spacing, to avoid the possibility of arcing with the high AC voltage. The only awkward part of the design is that the output of L1, which is the power bus for the converter, has a long run to get up to where all the components are. It's probably difficult to avoid at least one such run in a compact board like this.

Table 10.18 Final BOM for LED BR40, 3000 K Version

Ref Des	Description	Part #	Mfr	1 Kpc Price	Est. Price
C1	100 nF, 630 V, 1812	Any	Any	\$0.168	\$0.056
C2	1 μ F, 10 V, 0805	Any	Any	\$0.008	\$0.003
C3	680 pF, 0805	Any	Any	\$0.008	\$0.003
D1-4	600 V, 1 A Regular	1N4004	Any	\$0.172	\$0.057
D5	600 V, 3 A Fast, SMC	RS3J-13-F	Diodes, Inc.	\$0.273	\$0.091
F1	125 V, 1/8 A Fast Fuse	MQ125	Bel Fuse	\$0.248	\$0.083
L1	15 mH, 120 mA, 24 Ω	TSL1112RA- 153JR12-PF	TDK	\$0.875	\$0.292
L2	1.0 mH, 620 mA, 1.5 Ω	SPD127R- 105M	API Delevan	\$1.078	\$0.359
M1	150VAC, 20mm MOV	V14E150L2B	Littelfuse	\$0.246	\$0.082
Q1	600 V, 1 Ω , DPAK	STD5NM60T4	STM	\$1.040	\$0.347
R1	226 K Ω , 0805	Any	Any	\$0.004	\$0.001
R2	1 M Ω , 0805	Any	Any	\$0.004	\$0.001
R3	1.5 K Ω , 0805	Any	Any	\$0.004	\$0.001
R4	430 m Ω , 0805	Any	Any	\$0.102	\$0.034
U1	Controller	HV9910	Supertex	\$0.725	\$0.241
PCB1				\$0.210	\$0.007
TOTAL				\$5.165	\$1.722
D6-11	LED	LXML-PW71	Lumileds	\$17.880	\$11.940
PCB2				\$0.210	\$0.007
TOTAL				\$23.26	\$13.67

The input bypass capacitor C1 is right next to the output to the LEDs. This proximity will help keep the line noise from getting to them, keeping down radiated EMI. The IC is right close up to this node as well, but is buffered by having its small-signal components (R1, R3, C3, and C2) grounded through a dedicated trace with a single connection to the IC ground. Observe that the current sense resistor, R4, is as close as physically possible to the IC pins. Note too that the gate and source of Q1 are also immediately next to the IC, minimizing the run for these high-frequency lines.

L2 is an outlier, because it is relatively very large. Still, the layout has a short fat trace between it and the drain of Q1, and a very short trace to the wire that leads




Figure 10.17 BR40 PCB. Right: Topside. Left: Bottom Side with X-Ray Vision. See color insert.

to the anode of the LED string. We may consider twisting the pair of wires coming off the board to the LED string to further reduce radiative area.

The ground is large and underneath the IC and L2. It really only has four connections: the two anodes of the bridge, C1 and a via to the IC ground. All the other grounds in the system come to the IC ground, including the current sense resistor.

Chapter 11

Practical Measurement of LEDs and Lighting

Most of this book up to this point has been design-oriented: how to calculate how hot the LEDs become, how to design to obtain the desired number of lumens. But by the end of the design cycle, at the latest (but hopefully before), you have to check what is really happening. Maybe the LEDs are running hotter than you expected. Measuring LED temperature is not as simple as just sticking a thermocouple onto one. The light output may be less than it ought to be, or the CRI may be low. You can't just hold a lux meter 3 ft away and tell if it's bright enough, and guessing CRI is unreliable too. This chapter is going to look in detail how to make these sorts of measurements, accurately and in a repeatable fashion.

MEASURING LIGHT OUTPUT

Lux Meter

The cheap and easy way to measure the light output of a system is by using a lux meter. A lux meter measures the luminous intensity on a surface. Figure 11.1 shows a typical unit, a Mastech MS6610. It has a detector, the white circular area, which is the sensor part. In a good unit this detector will be "cosine corrected," compensating for the angle that the light is coming in to the sensor. The meter should also be calibrated, so that its measurement takes into account the CIE human eye response curve.

Now while this device is easy to use, you have to be cautious with this meter. It is very sensitive to the distance from the light source. Remember that lux is proportional to the square of the distance from the light source. Suppose that at 1 m away you measure 3500 lux. If for your next measurement you are 1 centimeter further away, the measurement will be reduced to $3500 \text{ lux} / (1.01 \text{ m} / 1 \text{ m})^2 = 3431 \text{ lux}$. If these two measurements are taken on two consecutive days, you might conclude that your light source has gotten dimmer. The same problem occurs if you are slightly off center. The cosine correction fixes the light coming in at an angle to the sensor,



Figure 11.1 Typical Lux Meter.

Source: http://www.p-mastech.com/products/06_ifep/ms6610.html.

but the illuminance actually *is* different when you're off-center. You don't want to come to the conclusion that your light source has dimmed (or gotten brighter) because you varied the measurement setup.

To use this kind of meter, we recommend a setup that fixes the measuring distance. This could be as simple as using a 3-ft cardboard mailing tube. Place the tube over the light source and the lux meter sensor at the other, open end of the tube. This eliminates most of the distance variation in measurement. To eliminate angular variation, try making this tube into a semi-permanent fixture. Glue it down to a piece of cardboard, and mark the light source position on the cardboard for repeatability. You could also mount the sensor to a second piece of cardboard, with a ring on the cardboard showing where to position it on the tube.

When using a Lux meter, one must be careful to have no background light. Also be aware that if a tube or enclosed box is used, the sides of the box or tube will reflect the light and add to the measured numbers. If you're just taking a relative measurement, such as in the case of percentage light drop over a certain time period, then the effects of the sidewalls do not matter. If, however, the actual luminous intensity is to be measured, you should either use a very large box with dark walls, or you should line the walls with nonreflective flocking paper. The edges of the box or tube are a source of light leaking in and should have a drape of black material to cover them. A dark measurement, that is, one that is taken with the light source turned off, should be made to determine the background light level. You should do this every time you take a measurement, and subtract it from the number you have measured.

It is tempting to take a lux measurement and multiply by the surface area of an equivalent sphere to calculate the luminous flux. If the light source is uniform, then the equivalent lumens can be calculated this way. However, very few light sources are uniform, including incandescent bulbs. If the light source is uniform but only over a limited angle such as in a reflector lamp, then calculating lumens might give

a grossly incorrect answer. In almost all cases, whether the light comes from an LED or from an incandescent light bulb, the light source is not uniform. Therefore, a lux meter cannot be used to determine lumens. Lux meters are properly used for measuring illuminance, that is, light on an area, and not luminous flux.

Integrating Sphere

To properly measure luminous flux, an integrating sphere must be used. This is a sphere, preferably four to five times larger in diameter than the light source, into which the light source is placed. They are available in sizes ranging from a few centimeters to those that are big enough for a person to walk into. The inside of the sphere is coated with highly (>98%) reflective material. Light from the source emitting in any direction is reflected inside the sphere and is measured at a collector (see Figure 11.2 and Figure 11.3). Thus, the luminous flux is physically integrated across all directions, hence the name. One can see this effect by peering in a thin crack after the sphere is closed. The inside of the sphere seems to emanate a uniform glow.

It's important that light from your source doesn't shine directly into the integrating sphere's collector. The collector is behind a baffle, a light-blocking element that keeps light from directly reaching the collector. Baffles are designed to be a minimum size. Larger baffles have a negative effect on sphere accuracy. Care should be taken to ensure that the light source does not leak around the baffle. It is difficult to see leakage by looking for shadows of the baffle created by the light source, since the sphere is so highly reflective. A piece of paper held behind the baffle shows the

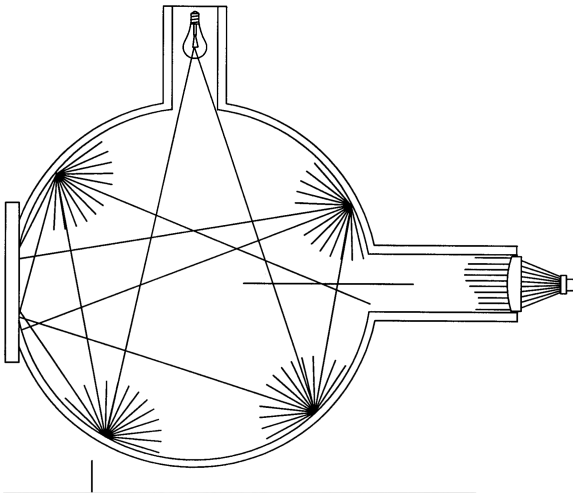


Figure 11.2 Sketch of Operation of an Integrating Sphere.

Source: <http://www.aaccuracy.com/optical.htm>.




Figure 11.3 Walk-In Integrating Sphere.

shadows more clearly. If the shadow does not completely cover the collector, then the light source should be moved until the baffle completely covers the collector. In the case of a directional light, the light source may be pointed away from the collector. When all else fails, a quick fix is to tape a piece of white paper, as small as possible, to the stage (mount) of the light source to aid the baffle.

The integrating sphere is calibrated to a standard light source, a tungsten light bulb. The standard should be an NIST traceable standard, which means that it is calibrated against another standard that is calibrated to a standard at NIST with a known lumen output. The calibration is said to be transferred from the NIST standard to the standard at a lab, and then transferred to the standard that you buy. Calibration of the sphere will last for a year or for 50hrs of actual on-time, whichever is reached first.

A calibration bulb comes with a calibration file that relates how that standard compares with the vendor's standard. At your lab, you will transfer the calibration from your standard to an auxiliary bulb that stays permanently inside the sphere. The calibration bulb with a known luminous flux is placed on the stage inside the sphere. A high precision current source is used to drive both the calibration bulb and the auxiliary. The calibration bulb is measured first. The amount of light measured at the collector reflects the sensitivity response of the integrating sphere. Then the auxiliary bulb is measured and a luminous flux value can be assigned to it, based on its output compared with the output of the calibration bulb. This "calibration transfer" should be done once a month to maintain the auxiliary bulb calibration. The auxiliary bulb is the daily workhorse and performs the absorption correction for the sphere. That way your expensive calibration bulb doesn't need frequent replacement.

When a light source is placed inside the integrating sphere, it changes the system's response. In other words, the surface of the light source is not 100% reflective and it will absorb some light and affect the calibration, and in turn the measurement.



Figure 11.4 Gooch & Housego 6-Inch Integrating Sphere with OL770-LED Spectroradiometer.
 Source: http://www.olinet.com//content/file/1210967318B220_770-LED_Brochure.pdf.

This self-absorption is corrected by turning off the light source under test and turning on the auxiliary light. A measurement now will show the effect of the absorption and can be added to the calibration. After all this setup, the unit under test can finally be measured. It's best to follow this proper procedure every time, to eliminate questions about measurement accuracy or drift. It seems like a lot of work, but it becomes routine with time.

The collector of the integrating sphere can be connected to any light collecting meter. The preferred kind is the spectroradiometer (see Figure 11.4). This meter separates the light by wavelength bins and measures the amount of light at each wavelength. The resulting spectral distribution is necessary for calculations of CCT, chromaticity coordinates, and CRI. Alternatively, a photometer can be connected to the integrating sphere. This would be similar to the lux meter, able to measure only lumens and not the spectral information, CCT, chromaticity coordinates, or CRI. In fact, a photometer is basically a lux meter measuring the irradiance at the collector port. The difference is that the integrating sphere reflects the light evenly inside and allows the irradiance to be converted to flux.

Standard sphere size for most labs is 20 in to measure light bulb-size sources. Spheres measuring 6 in are fine for measuring LED-sized sources. The surface area of the light source should be less than 2% of the surface area of the sphere. A general rule is that the diameter of the sphere should be at least four times larger than the diameter of the light source. The 6-in spheres have a port against which to place the light source. These are referred to as 2π measurements since light is being emitted only from the front side of the source and only half of the solid angle needs to be measured. To measure light bulbs inside recessed can lights, a 6-in port may be added to a 20-in sphere to perform the same 2π measurement. If the recessed light is only 4 in in diameter, a white cover plate should be used to cover up the gap.

Goniophotometer

Measuring angular light distribution is important for bulbs such as floodlights and spotlights that have prescribed beam angles. To measure angular light distribution, a goniometer is used. When combined with the light measurement device, the system is called a goniophotometer. These are room-sized contraptions with a sensor mounted on a swing arm. Because of the size and cost of this process, these measurements are usually performed in testing labs.

One can obtain an approximate measurement by using a 6-in integrating sphere to measure irradiance. The light source is mounted on a rotation stage at a distance at least five times that of the largest dimension of the light source. Care is taken to ensure that no other light sources are interfering. Then relative light distribution can be measured by turning the light source.

Special Considerations in Measuring Light Output of LEDs

Measuring light output of LEDs requires special considerations beyond those for measuring end-product light sources. The light-capturing aspect is straightforward; a 6-in integrating sphere will suffice. A larger sphere also may be used as long as it is sensitive enough to accurately measure the amount of light put out by a single LED.

The tricky part is due to the self-heating of the LEDs. This leads to continually dropping lumens until thermal equilibrium is reached. Even if the LED is mounted on a large heat sink, the junction will likely rise 10–20°C in the second that it takes to make the measurement. This can result in a difference of several percentage points in luminous flux from a measurement taken at room temperature. LED manufacturers use a 25 msec pulse of current to the LED to ensure that the junction temperature remains at room temperature. Short of using special equipment, one is unlikely to obtain the same luminous flux measurement as that shown on data sheets. An alternative method has been proposed by NIST (Zhong and Ohno, 2008) wherein the LED is mounted on an electric cooler and allowed to settle to steady state before a measurement is taken.

LED MEASUREMENT STANDARDS

Luminaire Light Output (LM-79)

The Illuminating Engineering Society of North America (IESNA) has a standard, IES LM-79-2008, specifying that the light output of LED-based luminaires must be measured inside an integrating sphere at thermal steady state. It reflects the best practices in light measurement and prevents ill-founded or unscrupulous claims of

light output. The old practice with incandescent bulbs was to use a relative method of taking the bulb light output and to multiply it by the optical efficiency of the rest of the fixture. With LED-based luminaires, the fixture design is integral to the thermal performance and would affect the light output in terms of both optical loss and thermal roll-off. The net light output must be measured in an integrating sphere, the absolute method. LM-79 also specifies that the power into the driver for the LEDs is the net input power. Net light output divided by the net input power gives the luminaire efficacy.

LED Lifetime (LM-80)

The other IESNA standard that concerns us is LM-80-2008. The intent of this is to standardize LED manufacturer lifetime testing procedures. As was mentioned before, LEDs don't burn out, they just get dimmer over time. LM-80 defines LED end of life as the time when the LED reaches 70% of its initial lumens (L_{70}). It also specifies that the LEDs should be tested for 6000 hr (8.3 months) at three different junction temperatures.

However, LM-80 stops short of specifying how to extrapolate this information to a useful lifetime, because there isn't a definitive model yet. So besides verifying that the LED manufacturer follows LM-80 lifetime testing guidelines, one should also verify that the manufacturer has measured dimming for a reasonable amount of time to allow reasonable extrapolation to lifetime. For example, 1,000 hours of testing should not lead to claims of a 50,000-hour lifetime.

The lifetime issue becomes important when qualifying for Energy Star certification. For that, you must use LEDs from a manufacturer that has LM-80 compliant test data. You must measure the LED temperature following the manufacturer's recommendation and verify that it is within the LM-80 test temperatures. Then the LM-80 data must be greater than 91.8% of initial lumens after 6,000 hours of operation to qualify for a 25,000-hour life rating. It must be greater than 94.1% for a 35,000-hour life rating.

ASSIST

The Alliance for Solid-State Illumination Systems and Technologies (ASSIST), at the Lighting Research Center at Rensselaer Polytechnic Institute in New York, has put together a number of recommended test procedures for LED lights, focusing on tests that emulate real-life operation. A current list of such recommendations can be found at <http://www.lrc.rpi.edu/programs/solidstate/assist/recommends.asp>. Particularly important is "LED Life for General Lighting," which recommends 6000 hours of testing for LED lights. They recommend then "a functional fit (e.g., exponential decay) to the data" in order to predict the lifetime of the light.

MEASURING LED TEMPERATURE

Recall that LEDs are current devices. You can't set their optical output with a voltage. The optical output depends on the current. Furthermore, a small change in the voltage will result in a large change in the current. Therefore, as discussed in Chapter 7, control of the LED requires a current. This current is often a constant with time.

Diodes generically change their forward voltage with temperature. This is true of LEDs as well. As the temperature increases, the forward voltage of the LED decreases. However, fully understanding this concept requires some more work.

Think back to the introductory chapter and recall that the LED consists of a semiconductor die, some phosphor, and a package. Of these, the electrical characteristics are determined exclusively by the semiconductor die. When we talk about the temperature affecting the forward voltage, we are necessarily talking about the die temperature.

Unfortunately, the die is inside the package, and its temperature cannot be directly measured. Think of how you might try to measure the die temperature. The standard method for measuring temperature is to attach a thermocouple and use a thermometer. This won't work in this case for at least two reasons. For one, the die is an electrical circuit with very tiny elements. A thermocouple consists of two pieces of metal, and the metal might short out the circuit elements. In addition, since the die is very small, the added mass of the thermocouple would change the thermal performance of the LED. In any case, attaching a thermocouple requires opening up the package of the LED, which is destructive to any other tests you might be trying to run on the device.

Another method of measuring the die temperature is with an infrared (IR) thermometer. While this is routine with normal semiconductors, it might be a problem with a source emitting intense light. And again, this would involve opening the package of the LED.

There are two standard methods of measuring the die temperature of an LED. Both have some problems. The easier of the two methods is to measure the LED package temperature, and infer the die temperature. In the example shown in Figure 11.5, a thermocouple is attached to the package. Simultaneously, the power into the LED is measured. The die temperature is inferred from the equation

$$T_{DIE} = T_{PKG} + P * \Theta_{JC}$$

Of course, you have to hope that the manufacturer's number for the thermal resistance from the junction to the case, Θ_{JC} , is correct.

There are a number of practical considerations to note before making this measurement. To start with, the size of the thermocouple should be as small as possible. This helps to avoid any chance of the measurement changing the results. In practice, 36 gauge wire works reasonably well.

Next, the position of the thermocouple is quite important. The better manufacturers will specify the point at which the package is to be measured. Failing this,



Figure 11.5 Location of Thermocouple to Measure LED Temperature.

Source: “Thermal Design Using Luxeon® Power Light Sources,” Lumileds AB05, 2006.

attach the thermocouple right next to the anode or cathode metal of the device. (Caution! Don’t let the thermocouple actually contact the metal. If it does it will have the same voltage as the LED has. This may affect the meter reading, or may actually be hazardous if the LED is connected to the AC line.)

Attaching the thermocouple properly is something of an art. You shouldn’t just place it against the device. This results in poor contact and gives temperature measurements which are artificially low. You certainly shouldn’t hold it in place with your finger. The added thermal mass of your finger completely changes the results. Taping the thermocouple down is better, but it is annoying when the adhesive on the tape comes loose with temperature (you might try Kapton tape). A drop of glue is best. Use the smallest drop you can manage, to avoid the thermal mass issue. Use fast-setting glue, such as a 5-minute epoxy or cyanoacrylate (Superglue). Also make sure that the bead of the thermocouple and the first 3 mm of wire are flush against the surface.

There is a second possibility for measuring the die temperature. Since forward voltage is temperature dependent, you could use the forward voltage to measure the die temperature. In fact, this is what is done for microprocessors. The microprocessor has a diode inside it. This diode is run at the same current as another, external diode of the same type. By measuring the forward voltage of both diodes, the difference in temperature can be computed.

In theory, this same method could be used for LEDs. You could measure the forward voltage of the LED when it turns on and again when it is in steady-state, and the difference will be the temperature increase. Unfortunately, it is that first step that is hard. Since LEDs are tiny, they have tiny thermal time constants. Unless you have the equipment to measure the forward voltage very fast, say within the first 25 msec of turning it on, you don’t know the initial die temperature. In practical applications, that makes this method unusable for absolute temperature measurement, though it works reasonably well for differential temperature measurements.

MEASURING THERMAL RESISTANCE

Having measured the temperature of the LED, we can measure the temperatures of other parts of the system using similar methods. For example, you will almost certainly want to measure both the temperature of the outer case of your light and the

ambient temperature. So now, knowing the difference in temperature between two points, and the power dissipated, it's straightforward to calculate the thermal resistance of that path, because thermal resistance equals the temperature difference divided by the power.

Not so fast! We still don't know the power in the LEDs. We have accurately measured the electrical power going in to them, but only a part of that is then lost as heat power. The rest is emitted as optical power (there's also some IR power, but it is usually negligible for visible light LEDs). So to accurately calculate the thermal resistance, we need to know how much of the electrical input power is being emitted as heat power.

The wrong way to go about this is to try to look up the optical conversion efficiency of the LED. You get bogged down in all sorts of different numbers and differing measurement techniques by different people, and in the end you're confused. But you have actually already measured all the numbers you need. The spectroradiometer measured the optical power in each wavelength bin, and if you look for it you will find the total optical power in the data. The correct power to use for the thermal resistance, then, is the total electrical power minus the optical power (in watts, of course).

$$\Theta = \frac{\Delta T}{P_{in} - P_{optical}}$$

Our cautions about accuracy and precision apply to this calculation (see below). You are subtracting two almost equal numbers, and then dividing into a third number. Be careful. You should cross-measure the two thermocouples at a single point to verify that they give the same reading. Since for the measurement of thermal resistance they are being subtracted from one another, the absolute accuracy is not so critical. If one thermocouple reads 35.4°C and the other reads 34.4°C, the temperature difference is only 1.0°C, which gives you only two digits of precision for the difference. This limits your thermal resistance precision also to two digits, even though your temperature probe has three digits displayed. If you put 10 times more power into the system, you could get an additional digit, but in practice that probably won't work. Two digits accuracy is good enough anyway.

If you decide to determine the thermal capacitance, the same cautions apply again. You measure time (presumably) in seconds, so if you reach 63% of steady-state in 10 seconds, you can't know the thermal capacitance to more than two digits. Slower systems are easier, although again you probably don't need a lot of precision for thermal capacitance.

The one difficult aspect of measuring thermal resistance is the thermal resistance to ambient. A glance back at Figure 6.5 shows that the effective thermal resistance, which is the y-axis value divided by the x-axis value, is fairly nonlinear in power, and nonlinear again in ambient temperature. The practical thing to do here is to measure the temperature of the case of your light and the temperature of the ambient as close as possible to the actual operating environment as you can. For example, measure a floodlight in a can, not in open space. If you now increase or

decrease the input power level (say) by a small amount, you will be able to use the effective thermal resistance as an estimate of what the new temperatures are going to be. Ultimately, of course, there's no substitute for real measurements of temperature.

MEASURING POWER, POWER FACTOR, AND EFFICIENCY

In addition to optical and thermal measurements, electrical measurements of lights also have to be taken. You need to measure the amount of power going into the light to put on the label, and in order to calculate efficacy. You also need to measure power out of the ballast, to measure its efficiency. And you have to measure power factor, again for labeling purposes. None of these measurements are especially difficult once you know how to take them. It's just that there are many wrong ways to do it that you need to avoid.

Accuracy versus Precision

Before getting started, some important terminology needs to be discussed. In conversation, we often use accuracy and precision as synonyms, but they're not. The accuracy of a meter, for example, is whether the number it is showing you is correct. Suppose that your meter says the voltage is 5.011 V. If you measured that same voltage with a calibrated standard and it also said 5.011 V, then your meter is accurate to four decimal places. But suppose instead that the calibrated standard says the voltage is really 5.013 V. Then your meter measurement is accurate to only three decimal places. The fourth decimal place isn't right. Writing down that fourth decimal place in your notebook gives false data. Using instead a meter with six decimal places probably gives you better accuracy, but you don't know that unless it's been calibrated.

Precision is reflected in the number of decimal places displayed. In the example above, the precision is four decimal places. The six decimal place display meter has greater precision than the four decimal place display meter. But remember, that doesn't necessarily mean that the accuracy is any better. Both accuracy and precision are important, but in different ways. You need both, especially when you measure efficiency, as discussed below.

Measuring DC Power

In a DC power system, power is “just” current times voltage. As the scare quotes signify, however, this is not a trivial measurement to take. To start with, voltage has to be measured at some particular place. For example, suppose you're building a driver to run an LED from a USB port, as in Chapter 10. As a test, you hook up the driver to your 5 V lab supply. If you now read the voltage from the front panel meter

of the supply, there are several problems. To start with, how accurate is the front panel meter? Chances are that no one bothered to calibrate a bench meter, because it's intended to be just an indicator. It's not intended for accuracy. It also probably has low precision. So you need to use a dedicated digital volt meter (DVM). Of course, the DVM ought to be calibrated. If you measure the voltage today as 5.011 V and tomorrow it's 5.013 V, has the power supply changed, or is there a meter problem?

Next you need to consider where you're going to put the DVM leads. If you hook them up to the terminals of the power supply, you're reading the power supply voltage. This may not be the same as the voltage at the input pins of your device. If there is significant resistance in the wires, or significant current to the device, there will be an IR drop along both the positive and the negative wires. This will make the voltage applied to the device lower than that being sourced by the supply. As an example, suppose you have 1-meter-long cables for the positive and negative, and the device is drawing 40 mA. Using a standard 24AWG wire for the cabling, we have a resistance of 84 m Ω /m. With one meter for each polarity, we have a total of 168 m Ω . At 40 mA, this corresponds to a drop of 67 mV between the supply and the device. At 5 V, that's an error of more than 1%. The right thing to do is to put the DVM probes as close as possible to the inputs of the device you're measuring, both positive and negative.

Here is another thing to watch out for. Sometimes you only have one meter, so you measure the voltage, and then disconnect the probes so you can use the same meter to measure the current. This has a problem, because the ammeter drops some voltage; that's how it measures current. So you need to leave the voltmeter in place and find a second meter to measure the current. And again for the same reason, the voltmeter's probes should be closer to the device being tested than the ammeter.

Another problem with disconnecting the voltmeter is that the voltage may change with load current. Most power supplies won't change very much, but they all have finite output impedance. Now if you turn the dial on the power supply, of course the output voltage changes. More subtly, if you are running from a battery the battery voltage changes all the time. If you change the load current coming out of the battery, its output voltage will change. If you run the battery for more than a minute, it discharges enough that its voltage changes. Again, the conclusion is that you have to have a voltmeter monitoring the voltage continuously.

Now you are measuring the voltage and the current into the device. Power is the product of the two, right? Yes and no. As an example, suppose you've measured the voltage to be 5.011 V and the current to be 252 mA. It's easy to whip out your calculator and find that the product of the two is 1.262772 W. But that answer is *not right*. Your measurement of the voltage was accurate to four decimal places, and your measurement of the current was accurate to only three decimal places. Your power number can't be accurate to seven places when your measurement is only accurate to three! The rule to follow here is that the product of the two measurements can be no more accurate than the least accurate of the individual measurements. In this case, the current and therefore the power are only known to three decimal places. The power is 1.26 W. This is the right number to record in

your notebook. If this isn't accurate enough (and it may well not be when you're measuring efficiency) then you need to increase the accuracy of the current measurement.

A brief word about rounding and truncation is in place here. If your calculator shows you 1.2685 W, truncating this number means just dropping the numbers after the significant digits. In this case, truncating to three significant figures gives 1.26 W. Rounding on the other hand means that you change the last digit to the nearest number, based on the following digits. In this case, the six gets rounded up to seven, because the number 68 is closer to 70 than it is to 60. For all of the measurements you take, you should be rounding instead of truncating. This power should be recorded as 1.27 W.

Let's turn to some other problems commonly encountered in making electrical measurements. A very common one is that the numbers on the DVMs are jittery. While you're watching, the voltage changes from 5.011 V to 5.013 V, then to 5.010 V and then to 5.012 V. There's no pattern. You could write down just some random number from this set, but that doesn't give you that fourth decimal place of precision.

While there are several possible causes for this, the most common is that your meter is being confused by noise. Your SMPS is switching on and off at 45 KHz or 1.6 MHz, and this pulls current at the same frequency. It's almost certain that your input filter still leaves some degree of ripple on the input. This ripple voltage can upset your DVM. At frequencies like 45 KHz, the DVM may alias the ripple into the measurement, resulting in what appears to be random fluctuations. At frequencies like 1.6 MHz, although the meter has no response, the ripple may be rectified by the meter's input circuitry, resulting in an offset that may fluctuate with time.

The solution to the problem of jittery numbers is to eliminate the ripple. An easy way to do this is to attach a bunch of capacitance at the input. Of course, this is not for production. The capacitance you add is only to enable accurate measurements. For example, try adding ten times the amount of input capacitance you would normally use. If your board has a 10 μ F capacitor, add 100 μ F. Then add some high frequency bypass capacitors as well. Add a 1 μ F ceramic in parallel together with a 100 nF ceramic. This usually cleans up the jitter problem. If not, there may be more subtle issues involved, but it goes beyond the scope of this book to attempt to detail them and their fixes.

A similar but unrelated problem is temperature drift. When you turn the supply on, the current is 145 mA, but then it starts drifting up to 148 mA, 152 mA, and so on. The problem here is that the device being tested is warming up. As SMPS warm up, their reference voltages change, IR drops change, and so on. What you do about this depends on how accurate your reading needs to be. An easy but low accuracy fix is to make the measurement the moment you turn on the power. If the drift happens over a period of minutes, this may be good enough to make repeatable comparisons between versions of the supply. If you need greater accuracy, or if the drift is quick, you have no choice, you need to let the unit thermally stabilize before you take a measurement. This may take quite a bit of time, but has the advantage that the measurement will be taken under real operating conditions.

We should also briefly talk about meter ranges. When you're drawing 190 mA, it's natural to put the meter on the 200 mA scale rather than the 2 A scale, so you get more significant digits. Then when you increase the load to 210 mA, you switch the scale up, or it switches automatically. But how do you know that you get the same reading on the two different scales? They need not have exactly the same gain. It could be that 190 mA on the 200 mA scale reads 0.21 A on the 2 A scale. Potentially, this makes a graph of current have a spurious jump in it. In short, if you're going to be making measurements that are close to the range transition of a meter, you should set the meter to stay fixed at the higher range.

Measuring AC Power

To start off with, we want to remind you that AC power can be very dangerous. Please make sure you've read the section in Chapter 8 on safety before you start taking measurements on the AC line.

Measuring AC power is considerably more difficult than measuring DC power. This is because the AC voltage and current are constantly changing, going positive and negative, and their voltage and current waveforms need not match up with each other. The fact that they do not match up is related to power factor, which we'll discuss in a moment. The alternating and differing waveforms are related to the power.

It's relatively straightforward to measure the root mean square (RMS) AC voltage, or the RMS AC current. This is the proper computation for AC measurements, because the square of the voltage or current is what dissipates power (V^2/R or I^2R). You square the voltage or current, take its average over time, and then take the square root of this number to get back to the right units.

$$V_{rms} = \sqrt{\frac{1}{T} \int_0^T V^2(t) dt}$$

To measure AC voltage or current, you should use a meter that's rated for that purpose. Since AC line is usually 50 or 60 Hz, meters are typically prepared to measure this frequency (however, a typical meter will not measure the AC voltage or current at high frequency, such as the current through the input capacitor of your SMPS). The difficulty with these meters comes about when you have a SMPS. The SMPS may have a very high crest factor, beyond the rating of your meter (see Figure 11.6). The crest factor is the ratio of the peak value to the RMS value of the waveform. For a sine wave, such as the AC line voltage, the crest factor is $\sqrt{2} = 1.4$. But the current in Figure 11.6 has a much higher crest factor, maybe 2.0. The high peak current is caused by a large capacitor right after the bridge rectifier. The capacitor's voltage drops only a little bit during the 8.3 msec cycle, and so it charges back up only at line peak. All of the energy the SMPS used during the rest of the 8.3 msec has to be put back into the capacitor during this short time.




Figure 11.6 High Current Drawn at AC Line Peak Gives a High Crest Factor.

Many meters cannot accurately measure RMS values when the crest factor exceeds 1.7. So if you have a SMPS with a large input capacitor (i.e., without power factor correction), you need to find a meter that will handle the crest factor in order to accurately measure the RMS current. (RMS voltage is usually, but not always, easier to measure, because the AC line source impedance is usually, but not always, much lower than the load impedance.)

We now can measure RMS current and voltage. You might suppose that RMS power is the product of these two, but it isn't! The trouble is that the integral of the product may not equal the product of the integrals.

$$P_{rms} = \sqrt{\frac{1}{T} \int_0^T [V(t)I(t)]^2 dt} \neq \sqrt{\frac{1}{T} \int_0^T V^2(t) dt} * \sqrt{\frac{1}{T} \int_0^T I^2(t) dt}$$

Without getting into complicated mathematics, let's jump to the answer. Proper measurement of AC power requires a special meter designed for that purpose. The one we have successfully used is a Yokogawa WT210. It simultaneously measures RMS current, voltage, and power. It also reads power factor. It's a little pricey, but we think an instrument similar to this one is the only practical way to take real measurements of AC power.

Measuring Power Factor

We mentioned that power factor relates to way that the voltage and current waveforms do not match up. The spike in the current waveform of Figure 11.6 certainly gives rise to a poor power factor. But any lag or lead in the current may also result in poor PF (see Figure 11.7). Again, a suitable meter is requisite for obtaining the right measurement.

Measuring Ballast Efficiency

Efficiency is the ratio of output power to input power, and, as per by the Second Law of Thermodynamics, is always less than one. You've seen in Chapter 10 that we always have to make assumptions about the efficiency of a ballast during the initial design of an LED light. So once the ballast has been built, we have to verify its efficiency with a measurement.




Figure 11.7 Phase Shift between Current and Voltage Also Gives Low Power Factor.

Measuring the efficiency is as easy as just measuring the input power and the output power—if the efficiency is low, say less than 85%. Of course you need several meters, but that’s to be expected. The tricky thing comes about if your ballast efficiency is high. If you measure 10.0W in and 9.0W out, the efficiency is 90%. But if the input meter is off by 1%, you get 9.0W/9.9W = 91%. Is that 1% due to the new transistor you put in or a meter difference? You can’t tell.

The way to avoid questions like this is to take what we call “cross-measurements.” Before you take efficiency measurements, put both voltmeters on the same exact node, and verify that they read the same number. It doesn’t have to be the *right* number in the sense of absolute accuracy; it just has to be the *same* number. If they both have the same error, it will cancel out when you do the division to determine the efficiency. Do the same thing for the current. Take both ammeters and place them in series with the same exact line. If they both have the same reading, then you will be able to rely on your power measurement. You can do the same if you’re using integrated power meters, such as the WT210 mentioned above. Before measuring efficiency, use them both to measure an identical load, and verify that the wattages are identical.

EMI and Lighting

We offer a few quick words about measurement of EMI and lightning. There are a variety of home-built methods of measuring EMI. For example, people put a 1x oscilloscope probe on the input and run it into a spectrum analyzer to determine the conducted emissions. To measure radiated emissions, people make a one turn loop of wire (a sniffer) and attach it to an oscilloscope probe and run that into their spectrum analyzer. While these methods may have some value for determining relative improvement of performance, generically they just aren’t accurate enough. When you’re ready, pay the money and go to an EMI lab. There’s really no substitute. Usually you can measure lightning there as well.

As to what to actually do at the EMI lab, preparation is the key to success. Whenever we go to the EMI lab, we take along a box of all sorts of different sizes and values of inductors and capacitors. We do a quick scan (a prescan), and if the measurements are out we try adding in or swapping out a capacitor or inductor. You can always find something that works. Then when you go back to your own lab, you can find a design that fits onto the board and uses those values.

Determination of lightning goes by much the same concepts. Go with three of four different sized MOVs. Use the biggest one first. If that passes your lightning requirements, go to the next smaller size. Continue until the unit blows up—then you know that that last size was too small.

Oh, and always take along at least two identical units for your EMI testing. There's nothing more frustrating than getting there and finding out that the unit that worked perfectly that morning has stopped working that afternoon.

ACCELERATED LIFE TESTS

We've already talked about LM-80, which gives guidance on how to measure LED lifetime. But LED lifetime is not necessarily the only factor governing product lifetime. If the ballast fails in a month, or the plastic yellows in a year, these other components rather than the LEDs will have determined the lifetime of the product.

Let's suppose that you have already built units, and verified that the temperature of the LEDs is such that their lifetime is going to be adequate for your product per their manufacturer's tests. If your lighting product is guaranteed for 10,000 hours, you can't afford to wait 15 months to collect lifetime data. And if you're getting an Energy Star rating that requires a 3 year warranty, management is going to want to be very sure that there aren't any failures during that time. What can you do?

The thing to do is to recognize that the lifetime and warranty issues are usually the very last step in testing. The reality is that most products coming from the development cycle have plenty of problems that show up right away, for example, in the first 1000 hours ("infant mortality"). You want to progressively test longer and longer periods, to avoid wasted time. The first step in testing your new product should be to identify and fix these, as they are relatively easy and quick to find.

To get started, just try running the product for a while. You don't need a fancy oven. Plug in a few units to the AC outlet (or 5 V USB port, or whatever), and let them run overnight. When you come in next morning, there may be one or two units that don't work anymore. Debugging the failures will immediately strengthen your design, without your having to set up complicated resources to do a full life test.

How many units should you use? Six or ten is probably the right sort of number. You're not looking to collect statistically significant data; you just want to catch any significant design problems. In reality, this test may be done with hand-built units, rather than production units. That's okay, do it with just three. But make sure all three have identical components, so that you're catching a life issue, not still doing design work.

As a next step for lighting products, try a longer life test together with a cycle test. For example, get ten units from a preproduction run and run them on the AC line for a week (this doesn't preclude you from continuing to run tests on the hand-built units). At the same time, switch on and off three units a thousand times. A cycle test like this can catch start-up problems that a continuously on test won't see. You don't have to turn the switch on and off one thousand times (although this can be

done, we've done it). For under \$100 you can buy an AC line timer that switches power on and off at a fixed period. Two seconds on and then two seconds off is a period of four seconds, so in one day you can run 20,000 cycles!

After you've passed these preliminary tests you can move on to life tests. Of course, you still can't afford to wait 10,000 hours. And, in fact, even 10,000 hours of testing isn't long enough to give you good statistical data on the MTTF (mean-time-to-failure) of your units. MTTF is what management really wants to know: Over the 10,000 hour guaranteed life, how many units will actually fail?

In order to know how many failures there are going to be before a given time, you have to actually have some failures. Suppose, for example, that you test your product for 10,000 hours and there are no failures. Perhaps if you had tested them for 10,100 hours 35% would have failed (the electrolytic capacitor wears out, etc.). The problem here is that to reliably estimate how many are going to fail in the 10,000 hours you need to know both the mean time to failure and the standard deviation around this mean. In this example, suppose the MTTF is something like 10,200 hours, and the standard deviation is 100 hours. That means that 10,000 hours is only two standard deviations away from the mean. Then you should expect that in the first 10,000 hours you will get $\exp(-2) = 13.5\%$ of the units to fail—clearly unacceptable.

The goal is to accelerate the failure rate so you can collect the MTTF data. As already discussed in Chapter 5, most failures are governed by the Arrhenius law. They can thus be accelerated by increasing temperature. So what we want to do is to put the products into a thermal chamber and heat them past normal operating conditions to get them to fail prematurely.

If only it were that simple! Although temperature accelerates the failure rate, you don't know *how much* it accelerates it. So to use an accelerator, you also have to collect enough information to know how much acceleration a given temperature change results in. Of course, this information will be useful to you anyway. Presumably the failure rate in Phoenix, Arizona, is going to be higher than that in Juneau, Alaska.

Here's what to do practically. You need two temperature controlled ovens. Set one oven to run 40°C hotter than the average ambient temperature in which the product is going to be set up. Set the second oven to be 30°C hotter. Now into each oven place 30 units, and put a final 30 units in the ambient temperature in which the product will be used.

The hottest oven will presumably cause the units to fail first. Carefully record the time of each failure to the nearest hour. (This can be done, for example, with a recording power meter; when a unit fails, the power drops.) When half of the units have failed, that's the MTTF. And it's straightforward on Excel to calculate the standard deviation of the failure rate.

Now you must wait a while longer (usually about twice as long) for half of the units in the 30°C oven to fail. You collect the data on these as well, and find the MTTF. Since you know the mean at both 40°C and 30°C above actual operation, you can compute what the mean will be in the actual environment by assuming an exponential relationship. For example, suppose the mean failure at 40°C takes

600 hours, and at 30°C it took 1400 hours. For each factor of 10°C increase, the mean failure rate increases by $1400/600 = 2.33$. Therefore, at 20°C above normal, the MTTF will be 1400 hours * 2.33 = 3250 hours, at 10°C above normal the MTTF will be 1400 hours * $2.33^2 = 7620$ hours. And at normal operating temperature, the MTTF will be 1400 hours * $2.33^3 = 17,800$ hours. Since you also know the standard deviation of these numbers, you can now reliably estimate the failure rate at 10,000 hours.

These instructions begged a few questions. Why did we pick 40°C and 30°C? The reason is that many components have roughly a “power of two” rule dependence on temperature: raising the temperature by 10°C causes the MTTF to halve. So when we heard that we needed 10,000 hour data, we estimated that 30°C would accelerate the failure rate by $2^{(30^\circ\text{C}/10^\circ\text{C})} = 2^3 = 8$, giving about 1200 hours equivalent, seven weeks. This seemed like a reasonable amount of time to wait. If the data were needed for 20,000 hours, we would have picked 50°C and 40°C instead.

In that case, why not pick 50°C and 40°C right away? The problem is absolute maximum component temperatures. Some devices, such as capacitors, just have shorter lifetimes when overheated. This after all is the effect we’re looking for. But some just fail right away when overheated. For example, if the plastic case melts, you’re out of the temperature regime where the Arrhenius law applies. So you can’t go above the maximum temperature at which catastrophic failure starts. This puts an upper bound on how much acceleration can be achieved. Beyond this, you just have to test the product for long enough, even if that means months of testing.

Why did we pick 30 units? Wouldn’t 10 be enough? The number has to do with the statistical certainty of the measurement. The uncertainty decreases as the square root of the number of units tested. 100 units will give you $\sqrt{10} = 3$ times better accuracy than 10 units. As a compromise between using a large numbers of units and suffering from poor accuracy, 30 units has been generally agreed upon. You may be able to get by with fewer units, but you need to talk to an expert.

A few other issues should be mentioned. You don’t really have to wait for half of the units to fail. You can get a number for the MTTF from just a few units’ failure. This will again affect the accuracy, and needs expert guidance. What if none of them fail in what ought to be a reasonable time? Good news! This provides you at least a lower bound on the failure rate, and if this is long enough, maybe you don’t care how long the actual MTTF is. Of course, knowing what the MTTF actually is will provide more comfort to the finance department, which has to budget money for returns.

You might note that all of these calculations have assumed a Gaussian failure rate. One circumstance you might face where this is a faulty assumption is if there are two failure mechanisms with roughly the same rate. This will show up as a bimodal distribution of failures. An easy test for this is just to plot the failure time with a bar graph. If you can see that there are two lumps, this is diagnostic. You could also use Excel to look for nonzero skew. In either case, consult with an expert.

What were those 30 units running at ambient good for? Well, it’s fine and good that we estimated the life time from accelerated tests. Chances are that the test has yielded a good number. But there’s nothing like a real measurement. You should just

keep running those 30 units, basically forever. After a year or two, you'll start to see failures. That will confirm your calculations. Even if you change the design as a result of your accelerated testing, you should consider leaving those 30 original design units running. Those units will always have the most hours on them, even when you start the next 30 new-design units.

Finally, some types of failure aren't accelerated by temperature. An example is mechanical flexing. If your product has to be folded back and forth in normal usage, such as the USB light is, you need to make sure that the unit doesn't mechanically fail. There are special vibration tables built to test the strength of mechanical design. But for something like the USB light we have another suggestion. Give it to your child to play with. If it's still intact the next day, it's probably a reasonably good design.

Chapter 12

Practical Modeling of LEDs

In this final chapter, we take a look at building computer models of LEDs and of the thermal systems to which they are attached. Computer simulations can take a wide variety of forms. At their simplest, they might just be an Excel spreadsheet that calculates the operating point of an LED. At their most complex they might be a multi-physics simulation that simultaneously models the electrical and optical performance of the LEDs, their driver, the thermal system in which they operate, and the dynamic interaction between all of these.

Practical models tend to fall between these two extremes. You want sufficient detail to get some design estimates, but you can't spend a month doing tedious calculations to set up the model, either. We'll be using Excel to find equations that fit our data, and Spice software to do simulations. Multiphysics type simulations will be accomplished with Spice models emulating both thermal and optical performance of LEDs and their systems.

Why make a model at all? One reason is that there are so many LED choices out there; new ones appear every month. It wouldn't be practical to lay out a breadboard and test each one, not to mention the time and expense of getting samples. A model can quickly take information from the datasheet and accurately predict which LED is best for your application.

Models can be useful tools; the trick is to know when to use them. Some cases are so simple that simulation wouldn't be worthwhile; for example, nobody simulates an RC time constant. In other cases, simulation may not give reliable answers. An example of this is the performance of a MOSFET in a switch-mode power supply. Rise and fall times are very sensitive to the model parameters, and yet determine much of the power loss in the transistor. In this case, trying it out in the lab is the only reliable choice. Finally, some cases are sufficiently complicated that you can't guess the answer and lab tests take too much time. LEDs fall into this last category.

Many ICs these days have Spice models available on the Internet. Unfortunately, such is not the case for LEDs. In part this is doubtless due to their complexity. Unlike normal silicon ICs, they have not only electrical, but also thermal and optical characteristics, all of which play an important role in their performance. This chapter aims to show how to derive models from datasheet specifications. We'll also try to give some ideas about when a model is "good enough."

Practical Lighting Design With LEDs, First Edition. Ron Lenk, Carol Lenk.

© 2011 the Institute of Electrical and Electronics Engineers, Inc. Published 2011 by John Wiley & Sons, Inc.

PRELIMINARIES

What is the benefit of computer simulation? The brief answer is that simulation is useful for things you can't easily measure in the lab. An example of this is production yield. You can build one or maybe even 10 devices in your lab. If all 10 work, then you could hope for a yield between 90% and 100%. But 10% fallout could wreck a business. And 90% yield is a bogus conclusion from the data, anyway. All 10 of your LEDs (and resistors, capacitors, etc.) came from the same lot. What happens when you get a new reel and the parts' values are different?

The appropriate method for determining yield is computer simulation. The two common methods are worst case analysis (WCA) and Monte Carlo (MC) simulations. WCA is the more conservative method for determining production yield. The method is detailed in the author's book on power supply design. Briefly, you vary the value of each component or parameter between its guaranteed maximum and minimum and see how that affects performance of the circuit to specification. It isn't easy in the lab to find an 80nF capacitor and a 120nF capacitor to determine worst-case performance with your 100nF capacitor, but it's easy in your simulation. If some component causes the circuit to fall out of spec, you can either adjust the parameter or estimate how much yield loss it will produce.

The more common method of determining production yield is MC simulation. You again tell the computer highest and lowest values of all the parameters, and it randomly picks values over that range for each. If you run this 1000 times, and only once fail specifications, you get a reasonable guess at the yield being 99.9%. (Note: This flat distribution with equal probability for any value in range may not hold for some devices. For example, LED efficacy seems to always be at the bottom of the bin. This is because manufacturers are cherry-picking devices. Make sure you watch out for this sort of nonflat distribution; it can affect yield dramatically.)

A related area in which simulation is useful is optimization. As part of the design process the circuit has to be optimized, which usually means meeting specifications at minimum cost. If you remove one LED from the circuit, the remaining ones have to run at higher current to produce the same light. But then they get hotter, which reduces the efficacy; the higher current also reduces the efficacy. Plus, the losses in your ballast may be higher. In cases like this, simulation is good because there are so many interacting factors. You need to be able to scan the parameters very widely to find the optimum. This is best done with a computer.

A few words are in line about the authors' preference in simulation tools. For simple calculations, we always use Excel. "Simple" here means that one line follows from the previous lines; there aren't any feedback loops. Excel is good, for example, for evaluating expected light output from an LED given its current and temperature.

For anything more complex, we use Spice software. This is in keeping with our strategy of keeping simulations simple and transparent (more on this later). There are very fancy simulators that are based on Spice and have large component libraries and cost thousands of dollars. But we prefer a very basic simulator with just the Spice primitives and a low cost. We always use AIM-Spice (AIM-Software,

Trondheim, Norway; www.aimspice.com). A free version is offered that permits enough components for many applications; the full version is inexpensive. This is the simulator used for all the examples in this book. Its text-based input should be readily understandable and translatable to whatever simulator you are using.

Even the most basic Spice simulator contains fairly complicated models. For example, in AIM-Spice the transistor model, Q, has 41 specifiable parameters. The authors bet that not many people know what they all do—certainly we don’t! As a result, some of these models we just don’t use, unless truly “any device will do.” So when it comes to modeling LEDs, we don’t try to adjust the 18 parameters of the “D” diode model. We’ll build a model from scratch, using a “B” source.

So what we’re going to do, specifically, is to take datasheet parameters and curves, and fit functions to them. Once we have the function, we can write it as a “B” source, and this then is the model we’ll be using. This curve-fitting approach requires some care and recognition of its limitations. For example, it is known that any five data points can be exactly fit by a fourth-order polynomial ($a_0 + a_1x + a_2x^2 + a_3x^3 + a_4x^4$). So if you pick five points on an I-V curve, you can get exactly those values with this equation. But if you now try to find the value at a sixth point, you’ll usually find that the equation’s prediction is wildly off.

A better plan is to stick with the simplest equation possible. In practice, this is almost always a linear or quadratic equation, or an exponential of one. It doesn’t fit the curve exactly, but who says the graph was measured exactly? It will be good enough in the sense of least-squares.

Once you have this equation, it’s usually trustworthy from the minimum to the maximum value for which you selected data. It doesn’t go wildly out of control at in-between points. And it is usually a pretty reasonable guess for values somewhat beyond the limits—just remember that the validity at these values is unknown. You can’t base a design on the hope that your equation is right beyond the limits for which it’s been designed.

PRACTICAL OVERVIEW OF SPICE MODELING

In this section, we’re going to give a brief overview of how to use Spice. Of course, there are entire books written on the subject. But this is going to be more like a review of the main things you need to know, along with some practical tips how to get things to run smoothly. The simulator used for the examples will be AIM-Spice.

Let’s start by taking a look at the structure of an AIM-Spice listing. A simple model is shown in Table 12.1. It consists of a 1 mΩ resistor being fed by a 100 mA

Table 12.1 A Resistor Model

```
RESISTOR
R1 IN 0 0.001
I1 IN 0 -0.1
```

test source. The first line of the listing is a description of the model, in this case “RESISTOR.” It would also be typical to use this as the file name.

The next line is the actual resistor model. A listing consists of the part reference number, its connections, and its values. Spice already has a model of a resistor, which is called “R.” Since this is the first resistor in this model, the reference number is “R1.” The next resistor, if there were one, would be “R2,” and so on. The “R” part is necessary to tell Spice what model to use, but the numbering is arbitrary. The next part could just as well be “R100,” although we recommend consecutive numbering, as three-digit numbers will be used for submodels below.

Next to be added to the list are the connections of the resistor “R1.” Since a resistor has two terminals, it requires two connections, here “IN” and “0” (zero). The names of the connections are arbitrary, but should be selected to convey useful information. When you have a hundred connections, node 89 isn’t going to remind you what it does. The node “0” is special; zero signifies ground.

Finally, the last element of “R1” is its value, in ohms. Here we choose 1 m Ω . You can also use common prefixes, such as “1K” to signify 1000 Ω . But note that you don’t type in “ohms”; the simulator already knows what the units are.

The third line in our “Resistor” model is a current sink, “I1.” Note that the “I” model is a sink, not a source, so if current is to come out of it, it must have a negative value. Current sources have the same structure as resistors. “I1” is connected to “R1” at the node “IN” and has the value 100 mA.

The model is now complete. We expect to get a voltage of 100 μ V on “IN,” corresponding to 100 mA through a 1 m Ω resistor. To run the model, we push the “OP” button in the AIM-Spice window. As expected, V(IN)=100 μ V, the text signifying “the voltage at node IN equals 100 μ V.”

Since we’ve already introduced “I,” we should mention some other capabilities of it and its voltage source equivalent, “V.” Putting a number after either one’s connections makes it a DC source. But there are other choices in addition. The most useful is “PULSE,” which makes the source change value at some time and then change back later. You can consult the help section of AIM-Spice for details. The only thing to be careful of is not to make the initial value of “PULSE” different from the DC value. You might get something different in the transient simulation than in the operating point.

One more useful thing to do with an “I” or a “V” source is to sweep the value. This tells the simulator to repeatedly run a DC operating point on the circuit. Each sweep uses a different value of the source. You get to this operation by pushing the “DC” button on the AIM-Spice simulator. (Remember that the operating point is “OP.”) You select the minimum and maximum values, as well as the step size, from the window that pops up.

To turn to other circuit elements, a capacitor is the same as a resistor, using “C” instead of “R.” A typical listing would be “C1 IN 0 IN IC=1.” We have the part reference designator followed by the two connections (no polarity). Next is the value, in this case 1 nF. Finally, the “IC=1” tells the simulator that the initial condition (voltage) on the capacitor is 1 V.

Table 12.2 A Diode, Including a Diode Model

```

DIODE
D1 ANODE 0 DNORMAL
I1 ANODE 0 -0.1
.MODEL DNORMAL D

```

The next circuit element we need is “D,” a diode model. As mentioned above, it has a huge number of values potentially to be set. But we don’t need to. Practically, there are really only two uses for diodes in simulations. One use, obviously, is as an actual diode model. In this case we usually are not concerned with the exact details of the model; we just want it to conduct in one direction and not the other. We can use the model as is, without setting any parameters. Such a use is shown in Table 12.2.

Unlike for “R” and “C,” the diode model has a polarity. The anode connection is first, the cathode is second. The last item on the line, “DNORMAL,” tells the simulator what model to use for the diode (in this case it is a fictitious model name). The last line of the listing is the model for the diode, “.MODEL DNORMAL D.” First comes the word .MODEL (don’t forget the period in front!). Next is the model name, here again the name we made up, DNORMAL. Then follows the name of the Spice model to be used, “D.” Any nonstandard parameters would be listed last. Running the model, we see that the voltage at the anode is 774 mV at 100 mA.

The other use of “D” is when an ideal diode is needed, which means that it shouldn’t have the forward drop of a normal diode. An ideal diode could be used, for example, when a voltage source is required that only sources and doesn’t sink current, and we don’t want the voltage-dependent drop across a normal diode. To make this, we can change the last line to “.MODEL DIDEAL D N=0.01.” With this model, the voltage at the anode is 7.7 mV.

As mentioned above, we won’t use “D” for modeling an LED. With so many parameters available, there’s no telling that you’re getting the performance you want. Instead, we will directly build the performance we want using a “B” source. “B” is a “nonlinear dependent source.” What this means practically is that, unlike the other models, you can write equations to govern the output voltage (or current). In AIM-Spice, the allowed functions are basic arithmetic, powers, trigonometric and inverse trigonometric, logarithmic and exponential, absolute value, and the Heaviside step function. This is more than enough to make any model needed.

One more model that will be used is “X.” This is not properly a model in itself. It refers to a subcircuit that is defined at the end of the listing. It is useful when you want to have the same circuit appear in the listing multiple times. Rather than type in the details each time, you can call it “X” and just define it once. The details of usage will be explained below, when we use it to model LEDs.

Once in a while, you need to have a measurement of time in your Spice circuit. This can be easily done with a C-I timer. Table 12.3 shows a circuit that generates 1 V/sec, so that the voltage at “TIME” can be used as a surrogate for the time.

Table 12.3 A Timer Model

TIMER
 I1 TIME 0 -1
 C1 TIME 0 1

WHAT NOT TO DO

After all these good things being said about simulations, are you ready to jump into the chapter and type the models here into your simulator? Not just yet! Simulations have pitfalls waiting to trap those who rush into them.

Let's start off with a rule for simulation models: *Simple models are best*. Many people (particularly simulation tool vendors) might disagree with this, but listen to this sad story. One of the authors was building a simulation model of a complex power system, the one aboard what is now called the International Space Station. The model worked fine, correctly showing the bandwidth and phase margin of the power converters. The strange thing was, telling the system to turn off didn't work. An hour's worth of viewing waveforms showed that the problem was that an LM319 comparator model didn't go low when its inputs told it to. Was the problem in the rest of the circuit or in the LM319 model? It didn't seem that the problem could be the IC model because it was provided as part of the simulation program. But extensive probing couldn't find anything relevant wrong with the rest of the circuit.

A special test simulation was run, with nothing but the LM319, a couple of function generators, and a pull-up resistor to a +15 V supply on the output. The pull-up resistor was there because the LM319 is an open collector comparator. This means that it can sink current but not source it, so it needs a pull-up resistor to a voltage to go high. The result was the same: the output was always high. Changing the pull-up resistor value didn't affect it. Changing the voltage supply didn't affect it. In frustration, the supply was removed entirely. The output was *still* high! In the end, the vendor who sold the simulator had modeled the LM319 as having a high output—the inputs didn't have any connection to the output.

Of course, they fixed the problem right away when I told them. But the message is that you're dependent on the people who modeled the devices in your simulation. You trust that they correctly modeled all of the devices you're using. You trust that they modeled the devices not only according to how they are normally operated, but also according to the operating region in which you're using them (did they include reverse breakdown in the diode?) When you upgrade from version 20.01 to 20.02, you're also hoping that the people who patched the software didn't break a model that was working before. Anyone who's used a computer knows how remote the chances are that the software is perfect.

As a result, we recommend using the most basic models possible. Our models rely exclusively on Spice primitives. Does this sound like taking a step back in time? It does, but the visibility into what you're actually getting is worth it. And hours wasted finding someone else's software bugs will convince you of this.

Many potential problems with computer simulation are avoided by staying simple, but not all. A problem that nobody likes to acknowledge are simple misspellings. If you type a “C” instead of a “D,” the simulation will doubtless show strange things right away. But suppose instead of $10\text{K}\Omega$ you accidentally type $100\text{K}\Omega$? Or $19\text{K}\Omega$? It’s easy to type an extra zero, or hit the nine instead of the zero. With a hundred lines of code, or a large schematic page, this sort of problem can go unnoticed for a long time. If it’s a pull-up resistor, maybe it won’t be a big deal, but what if it’s part of a timing circuit?

A related problem is connecting components to the wrong nodes. Of course, this is particularly easy with a text editor, but it can happen with circuit schematics as well. Making clean schematics (no lines crossing each other unnecessarily) helps this somewhat. And you must make a schematic for any text editor anyway, with all the nodes and components labeled. We’ll show an example below.

WHAT TO DO

The solution to many of these problems is anchoring. You anchor a model by building it in the lab, and making sure that key parameters match your simulation. For example, suppose you are building a model of a current source running an LED. Go into the lab with a current source and measure the current with an ammeter, measure the temperature with a thermometer, and measure the light output with an integrating sphere. Then go to your model and set the current and temperature to the values you measured, and verify that the light output predicted by the model is correct (to some accuracy, perhaps 5%). If not, then fix the model—the data are right by definition! To reiterate: *Don’t spend all your time on modeling. You have to go to the lab for a reality check.*

You need to do the same anchoring for each important output of your model. If your model is supposed to predict die temperature, then you need to program in just the current and verify that the temperature predicted matches the measured value. Most important, your measurement should be relatively close to the actual operating point of the circuit. Otherwise, your model might be working in a regime where something gives incorrect results and you won’t know it. The best plan is to run several measurement/simulation pairs over a range of operating conditions. This enhances your confidence that the model will correctly predict operation. The one thing you don’t want to do is to tweak the model so it matches the data at only the one point measured. Such models invariably give bad predictions at any other operating point.

We’ve already mentioned this in passing, but it’s worth repeating. *You can’t get out what you didn’t put in.* As an easy example, suppose you’ve modeled an LED’s I-V curve so that it correctly predicts forward voltage at a current of 350mA. Will it give the right answer at the absolute maximum current of 1A? Perhaps. Will it give the right answer at a peak pulse current of 14A? No chance! You didn’t program in things such as bond wire resistance or the formation of hot spots on the die.

Chances are the LED will explode if you try it. That's why you need to anchor your model in the lab close to its real operating point.

MODELING FORWARD VOLTAGE

Let's start by modeling the forward voltage of an LED as a function of current. Figure 12.1 shows the curve the vendor provides for this parameter, presumably from measured data. The data runs from 100 mA to 1 A. Note that this curve is at 25°C; the effects of temperature will be added later.

The curve looks fairly smooth. The scale of both the current and the voltage are linear, so a quadratic function might be a reasonable guess to fit the data. In this case, however, we know from engineering theory that a diode's voltage and current are exponentially related. So we're going to do a fit to $I = ae^{bV}$.

Let's start by collecting data points. The number of data points we collect isn't strictly determined. We should have more than three for reasonable accuracy and less than 10; it doesn't produce additional accuracy to have more. In practice, the number of points is usually determined by how many points are easy to read from the graph without having to interpolate.

We have entered the data in Excel in Table 12.4. We have selected 100 mA and 1 A because these are the endpoints of the graph. We picked 350 mA rather than 400 mA because the datasheet specifies forward voltage at that point. We'll use the specification later, but for now we just read the values from the graph. The other three points are chosen because they're spaced apart a bit and are easy to read. The point at 400 mA was skipped because it's so close to 350 mA.

Now we're going to fit the exponential function to this data. Excel has a function that does a least-squares fit to linear data, so we'll take the log of the current




Figure 12.1 Luxeon Rebel I-V Curve.

Source: Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Co., 2007.

Table 12.4 Values from an I-V Curve for an LED

	A	B
1	I (mA)	V (V)
2	100	2.87
3	200	3.02
4	350	3.18
5	600	3.32
6	800	3.42
7	1000	3.50

Table 12.5 LED I-V Curve with Logarithm of Current

	A	B	C
1	I (mA)	V (V)	ln(I)
2	100	2.87	4.60517
3	200	3.02	5.298317
4	350	3.18	5.857933
5	600	3.32	6.39693
6	800	3.42	6.684612
7	1000	3.50	6.907755

(because $I = ae^{bV}$ implies $\ln(I) = \ln(a) + bV$). In the next column, the formulas should read “=LN(A2)” and so on. Table 12.5 shows the result.

Now we can perform the linear least-squares fit. In cell B9 we use “SLOPE(C2:C7, B2:B7),” and in B10 we use “INTERCEPT(C2:C7, B2:B7).” See Table 12.6.

The equation we want is $I = \exp(\text{intercept} + \text{slope} * V)$. In D2 we type “=EXP(B\$10+ B2*B\$9).” The dollar signs are to keep the auto-incrementing of equations in Excel turned off for this value; we don’t want the values sliding down when we paste the function into the other cells. We also check the error in E2 with “=(D2-A2)/D2.” The final result, presented in Table 12.7, shows that we have fit the data points to better than 10% accuracy over the entire range. This is good enough for most practical purposes. We show all the formulas in Table 12.8.

So our model is $\ln(I) = -5.756 + 3.6398 * V$ with I in mA. Notice that when $V = 0$, $I = 3 \mu\text{A}$. Dark current isn’t specified for this LED, but this number is practically zero for our purposes. Thus the model is reasonable from 0 to 1A. Note however, that the model is *not* reasonable for $V < 0$. The exponential is always positive, and so the model predicts that with reverse voltage applied to the LED, current is still positive! If you want to add in breakdown voltage, this has to be done separately. We’ll talk about this below.

Now that we have the equation, we turn to Spice to create a model. Now a “B” source can be either a current or a voltage source. Of course, we want to drive the

Table 12.6 Excel Fitting of LED I-V Curve

	A	B	C
1	I (mA)	V (V)	ln(I)
2	100	2.87	4.60517
3	200	3.02	5.298317
4	350	3.18	5.857933
5	600	3.32	6.39693
6	800	3.42	6.684612
7	1000	3.50	6.907755
8			
9	Slope	3.6398	
10	Intercept	-5.75564	

Table 12.7 LED I-V Curve Model Showing Goodness of Fit of Equation

	A	B	C	D	E
1	I (mA)	V (V)	ln(I)	Eqn.	Error
2	100	2.87	4.60517	108.91732	9%
3	200	3.02	5.298317	188.02203	-6%
4	350	3.18	5.857933	336.61075	-4%
5	600	3.32	6.39693	560.31508	-7%
6	800	3.42	6.684612	806.31889	1%
7	1000	3.50	6.907755	1078.8634	8%
8					
9	Slope	3.6398			
10	Intercept	-5.75564			

Table 12.8 Equations of Excel Model of LED I-V Curve


	A	B	C	D	E
1	I (mA)	V (V)	ln(I)	Eqn.	Error
2	100	2.87	=LN(A2)	=EXP(B\$10+B2*B\$9)	=(D2-A2)/A2
3	200	3.02	=LN(A3)	=EXP(B\$10+B3*B\$9)	=(D3-A3)/A3
4	350	3.18	=LN(A4)	=EXP(B\$10+B4*B\$9)	=(D4-A4)/A4
5	600	3.32	=LN(A5)	=EXP(B\$10+B5*B\$9)	=(D5-A5)/A5
6	800	3.42	=LN(A6)	=EXP(B\$10+B6*B\$9)	=(D6-A6)/A6
7	1000	3.5	=LN(A7)	=EXP(B\$10+B7*B\$9)	=(D7-A7)/A7
8				=EXP(B\$10+B8*B\$9)	
9	Slope	=SLOPE (C2:C7,B2:B7)			
10	Intercept	=INTERCEPT (C2:C7,B2:B7)			

Table 12.9 First Model of Luxeon Rebel

```

LUXEON REBEL
B1 AN 0 V=0.27474*(LN(V(ANODE)-V(AN))+19.5716)
R1 ANODE AN 0.001
I1 ANODE 0 -0.1

```

**Figure 12.2** First Model of Luxeon Rebel.

LED with a current source. Two current sources in series would compete with each other in a simulation. So we use the voltage source form of “B,” and thus the logarithmic form of the equation $V = 0.27474 * (\ln(I) + 5.756)$, with current in mA and voltage in volts.

Let’s turn now to implementing the forward voltage equation in Spice. A listing is shown in Table 12.9. The model is of an LED being fed by a 100mA current test source. The LED itself will be a compound device, consisting of a “B” source and a resistor. The “B” source is generating the LED’s equation, while the resistor is a dummy in series with the “B” and is used to measure the current into the LED. We need to measure the current for the equation we’re using for the LED.

The LED consists of a node “ANODE” and has its cathode connected to ground. “ANODE” actually goes to the current-sense resistor. The resistor is then connected to the “B” source at an internal node “AN” (see Figure 12.2; we will use the convention that node names are in boxes, reference designators are not). The “B” source goes from “AN” to ground. Its value is the equation $V = 0.27474 * (\ln(V(\text{ANODE}) - V(\text{AN})) + 19.5716)$. Here, $V(\text{ANODE}) - V(\text{AN})$ is a measure of the current; it’s the voltage across the resistor. The expression “ $V(\text{ANODE})$ ” means “use the value of the voltage at the node ‘ANODE.’” Since this equation uses current in mA, we have used a 1 mΩ resistor. A value of 1A through it will generate 1mV, and so this is the correct value for the equation. We’ve multiplied the current in amps by 1000 to get the current in mA—that’s why the resistor is 1 mΩ.

Now we can test our model. Pushing “OP” shows that the voltage at “ANODE” is 2.84675 V, comparing well with our graph value of 2.87 V. As expected, the voltage at “AN” is 100 μV lower than that at “ANODE.” This small value doesn’t matter practically, which is why we picked a small value resistor to measure the current in the first place.

Table 12.10 A Model of Multiple Luxeon Rebels in Series

MULTIPLE LUXEON REBELS

X1 IN LED1 LED

X2 LED1 0 LED

I1 IN 0 -0.1

.SUBCKT LED ANODE CATHODE

B100 AN CATHODE V=0.27474*(LN(V(ANODE)-V(AN))+19.5716)

R100 ANODE AN 0.001

.ENDS

We can now change the value of the current source to verify that the model generates the correct forward voltage. For example, at 1A the model shows 3.48 V, compared with 3.50 V on the graph. In short, this model adequately reproduces the forward voltage of the LED. You might note that AIM-Spice also allows the current in “I1” to be swept, using the “DC” button. However, the plot is not very useful, because there’s no straightforward way to plot current versus voltage, rather than the other way around.

Now consider what happens if we want seven LEDs in series. We could have just copied “B” and “R” seven times, but it isn’t very efficient. It’s better to build a model of the LED, to which we can then refer seven times. That’s what we’ll do next, as shown in Table 12.10.

In the model in Table 12.10, the LEDs are represented by the circuit model “X.” “X” always refers to a subcircuit, the name of which is at the end of the line. Since LEDs have two connections, “X” does also. For other subcircuits, “X” could have any number of connections. “X1 IN LED1 LED” tells the simulator that the LED subcircuit is to be connected from nodes “IN” and “LED1” and is to use the model named “LED.” Similarly, “X2 LED1 0 LED” tells it that the LED subcircuit is to be connected from “LED1” to ground and to use the same model, “LED.”

The “X” model tells the simulator to look for a subcircuit called “LED.” This is found in the last four lines of the model. The first of these says that it is a subcircuit called “LED” and that it has two connections, “ANODE” and “CATHODE.” Note that there is a period in front of this line. The next two lines are the same as we used before. The only difference is that we have used reference designators beginning with 100 rather than 1, to signify that these models are inside a subcircuit. Finally, the subcircuit’s last line is “.ENDS,” again with a period in front. This statement is necessary to separate the subcircuit from (potentially) other subcircuits.

Running “OP” shows what we would expect. Node “LED1” is at the same 2.85 V that resulted before, and node “IN” is at twice that, 5.69 V. So now we have a model of an LED that can be used with a single line. We could now combine it with, for example, a model of a power supply. Instead, we will further improve our LED model. Next we will turn to reverse breakdown of the LED, before getting to optical output.

REVERSE BREAKDOWN

The reader will have noticed that we have not addressed reverse breakdown voltage of the LED, except to note that the I-V model doesn't calculate I properly with negative V. There are two reasons for this. The first, practical reason is that LED manufacturers recommend that you do not operate LEDs this way. Some of them don't even specify what the reverse breakdown voltage is. It's easy to break the LED by putting current through it in the wrong direction. Thus, for most applications, having a model for the reverse breakdown should be unnecessary.

A second reason is that the discussion about modeling reverse breakdown is going to get complicated. So the reader may skip this section, unless there is a particular need for a model to work with negative voltages. (The model we generate won't be used in the rest of the chapter.)

Those caveats out of the way, let's see what happens if we just apply a voltage from cathode to anode. The listing is shown in Table 12.11. By running a DC sweep on this model, we can see that we get results that are not only implausible (the device probably doesn't conduct 11KA at 5V reverse), but also wrong: At 0V the current should be 0A, not -5000A!

Let's fix the wrong result at 0V first. This turns out to be a problem with the model, not with convergence. The LN function can't deal with zero voltage from "ANODE" to "AN," because the log of zero is negative infinity. Changing the value of V1 from -5V to 0V and running the DC operating point shows what really is going on. The current is 3 μ A. This isn't the real value; it's set by the GMIN option, minimum conductance. When the voltage on V1 is less than zero, the current is programmed to be even smaller, but it can't because of the GMIN. A little experimentation leads us to choose GMIN=1.0E-20. Now the current at 0V is 0A, and the sweep shows that the current is zero from 0V down to -5V, to the resolution of the simulator.

Now that the current is zero in the reverse direction, we turn to making the device conduct at 5V reverse. What we're going to do is fairly simple conceptually. In parallel with the model we've built so far, we're going to put a 5V voltage source in backward. That way, when the voltage applied to the LED goes negative nothing happens until it becomes more negative than 5V, and then it starts conducting.

Table 12.11 First Model of LED's Reverse Breakdown Characteristics

LED REVERSE BREAKDOWN

X1 ANODE 0 LED

V1 ANODE 0 -5

.SUBCKT LED ANODE CATHODE

B100 AN CATHODE V=0.27474*(LN(V(ANODE)-V(AN))+19.5716)

R100 ANODE AN 0.001

.ENDS

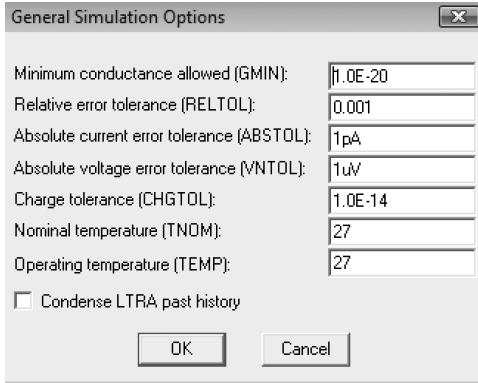


Figure 12.3 Setting the GMIN Option to 1.0E-20 Fixes the Current at 0V.

Table 12.12 Improved Model of LED's Reverse Breakdown Characteristics

LED REVERSE BREAKDOWN

X1 ANODE 0 LED

V1 ANODE 0 -5

.SUBCKT LED ANODE CATHODE

B100 AN CATHODE V=0.27474*(LN(V(ANODE)-V(AN))+19.5716)

R100 ANODE AN 0.001

V100 CATHODE CLAMP 4.3

D100 CLAMP ANODE DMODEL

.ENDS

.MODEL DMODEL D

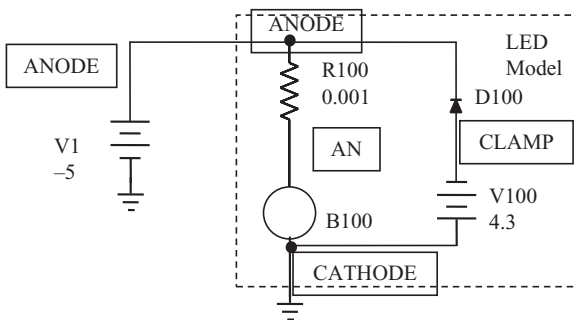


Figure 12.4 LED Model Including Reverse Breakdown.

We can't do literally this, because a voltage source will always try to source current. So in series with the voltage source we'll have to add a diode. This prevents the voltage source from sinking current when it's not supposed to, and additionally gives the benefit of a realistic I-V curve for the reverse breakdown. The model is shown in Table 12.12 and Figure 12.4.




Figure 12.5 With the Breakdown Modeled, the Complete I-V Curve Is Shown. The y-Axis Is the Negative of Anode-to-Cathode Current.

We've set the "V100" source to 4.3 V rather than 5 V because the diode has a forward voltage of about 0.7 V. The diode is modeled as a "DMODEL," which we define at the bottom as just being the normal "D" model. This is done with the statement ".MODEL DMODEL D."

Now when we run DC on this, we see that the current with 5 V reverse on the LED is 5 mA, and at 5.1 V it's up to 270 mA. Doing a DC sweep from -5.3 V to $+3.2$ V shows both the forward voltage and reverse breakdown characteristics expected, as shown in Figure 12.5.

MODELING OPTICAL OUTPUT

We've built a model of an LED that correctly determines its forward voltage as a function of the current through it. As a next step, we want to determine the light output from the LED. Again, we're going to be working at 25°C , and will add thermal effects later. We're starting with the model shown in Table 12.10 (not including the breakdown effects of the last section).

Just as for forward voltage, we're going to fit some curves. The graph in Figure 12.6 shows (normalized) light output as a function of current. To obtain real lumens from this, we have to know exactly which device we're using. Let's assume that our part number produces 100 lumens at 350 mA. Thus the graph's y-axis can be multiplied by 100 lumens to get light output. Our values extracted from the graph are shown in Table 12.13. We've added in 0 lumens at 0 mA to enhance the low current data.

It's tempting now to just do a linear fit of the equation, since the curve looks fairly straight. The only drawback is that at low currents, the light seems to drop off faster than it would linearly as current decreases. So we're going to model this as a quadratic, in order to better capture the low current light output.

For this purpose, we use an Excel graph, as shown in Figure 12.7.¹ We've graphed the data, and then added a trendline. We picked a polynomial trendline of order two, so that it's a quadratic fit, and checked the boxes to display the equation and the R^2 value. R^2 value is 0.999, indicating an outstandingly good fit.

¹Note that this is a scatter plot, not a line plot. To get the trendline equation right, it *has* to be a scatter plot. Excel gives the wrong equation for a line plot!

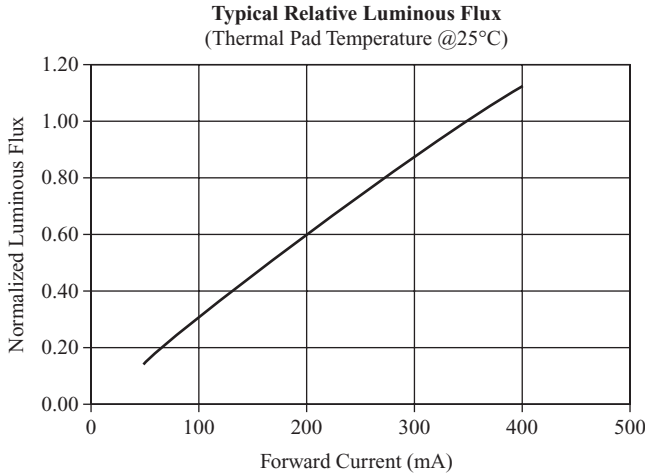


Figure 12.6 Normalized Light Output versus Current.

Source: Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Co., 2007.

Table 12.13 LED Light Output as a Function of Drive Current

	A	B
1	I (mA)	Light (Lm)
2	0	0
3	100	37
4	200	63
5	350	100
6	600	150
7	800	185
8	1000	212




Figure 12.7 Data versus Equation for Optical Output versus Drive Current.

Table 12.14 Spreadsheet Showing Goodness of Fit of Equation for LED Light Output as a Function of Drive Current

	A	B	C	D
1	I (mA)	Light (Lm)	Eqn.	Error
2	0	0	3.647	#DIV/0!
3	100	37	33.577	-9%
4	200	63	61.507	-2%
5	350	100	99.652	0%
6	600	150	153.227	2%
7	800	185	187.087	1%
8	1000	212	212.947	0%

Table 12.15 A Model of Two LEDs Including Optical Output, at Fixed Temperature

TWO REBELS WITH OPTICAL OUTPUT AT FIXED TEMP

X1 IN LED1 OPTIC1 LED

X2 LED1 0 OPTIC2 LED

I1 IN 0 -1

.SUBCKT LED ANODE CATHODE OPTICAL

B100 AN CATHODE V=0.27474*(LN(V(ANODE)-V(AN))+19.5716)

B101 OPTICAL 0 V=-0.0001*((V(ANODE)-V(AN))*10^6)^2+0.3093*((V(ANODE)-V(AN))*10^6)+3.647

R100 ANODE AN 0.001

R1000 OPTICAL 0 1MEG

.ENDS

As an aside, you might wonder why we didn't check the Excel box forcing the equation to be zero at zero. If you try this, you'll see that the correlation coefficient actually gets worse (as does the error). This seems counterintuitive, until you realize that by removing the constant term, you've reduced the number of parameters from three to two. And fewer parameters make a worse fit, just as having enough parameters makes a perfect fit.

We can now apply the equation to our data, and verify that the error is less than 10% over the entire range (see Table 12.14). The error at 0 mA is of course a division by zero, but the error is actually only 4 lumens.

To include this equation in our Spice model is simple. Just as the voltage of the LED is a value at a node (the anode), the light output is going to be a value at a new node. We'll call it "OPTICAL." As shown in Table 12.15 and Figure 12.8 we add a "B101" source into the subcircuit with the equation in it. The voltage from "ANODE" to "AN" is 1 mV for 1 A, and we need it to be 1000, as that is the value

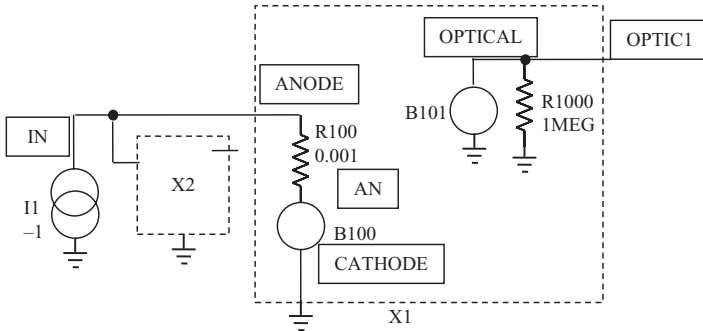


Figure 12.8 LED Model with Light Output, Temperature Not Yet Included.

in our equation, so we multiply by a million. The subcircuit now has another output, “OPTICAL,” and it is listed as “OPTIC1” and “OPTIC2” in the main listing. The light output is obtained by looking at the voltage on these nodes. And indeed, at 1A through the LED, the optic nodes show 213 V, which is to say 213 lumens, just as the graph shows. Note that inside the model we have added dummy resistors, “R1000,” because the simulator does not recognize nodes that lack connections. Dummies are usually $1\text{M}\Omega$, to keep the simulator from registering very small currents.

MODELING TEMPERATURE EFFECTS

The LED model now accounts for both the forward voltage and the optical output as a function of current—but only at 25°C . Next we need to add in the effects of temperature. Spice has a global setting for the temperature. You can consider this to be the ambient temperature. Since the LEDs will be dissipating power, they will be hotter than the ambient temperature, and thus need to have their own temperature parameter. We’ll call this node “THERM.”

The effect of temperature on forward voltage is straightforward. The only data provided is that it is $-3\text{mV}/^\circ\text{C}$. What we’re going to do for our model is create a voltage “AMB” corresponding numerically to the ambient temperature. This is done with a “V” source. This is then passed in to the LED model on one of its nodes.

Internal to the LED, we create a “THERM” node. For the moment, this is going to just be a settable value, using “V100.” When we build a thermal model for the LED, it will be determined by the thermal contact of the LED to the ambient temperature. Finally, we need to add dummy resistors as before to ensure convergence. The model including thermal effects on the forward voltage is shown in Table 12.16. Running it as shown shows that the forward voltage at “ANODE” has dropped 100 mV, as expected for a 33°C temperature rise, $33^\circ\text{C} * -3\text{mV}/^\circ\text{C} = -99\text{mV}$.

Finally, we need to add the effect of temperature on the light output. Looking at Figure 12.9, we see that this is linear also. At 20°C the factor is 1.00, while at

Table 12.16 Model of a Luxeon Rebel Including Temperature Effect on Light Output

```

MULTIPLE LEDS WITH OPTICAL OUTPUT AND TEMP EFFECTS

X1 IN LED1 OPTIC1 THERM1 AMB LED
X2 LED1 0 OPTIC2 THERM2 AMB LED

V1 AMB 0 27

R1000 AMB 0 1MEG

I1 IN 0 -1

.SUBCKT LED ANODE CATHODE OPTICAL THERM AMB
B100 AN CATHODE V=0.27474*(LN(V(ANODE)-V(AN))+19.5716)-0.003*(V(THERM)-
V(AMB))
B101 OPTICAL 0 V=-0.0001*((V(ANODE)-V(AN))*10^6)^2+0.3093*((V(ANODE)-
V(AN))*10^6)+3.647

R100 ANODE AN 0.001
R1000 OPTICAL 0 1MEG
R1001 THERM 0 1MEG

V100 THERM 0 60
.ENDS
    
```

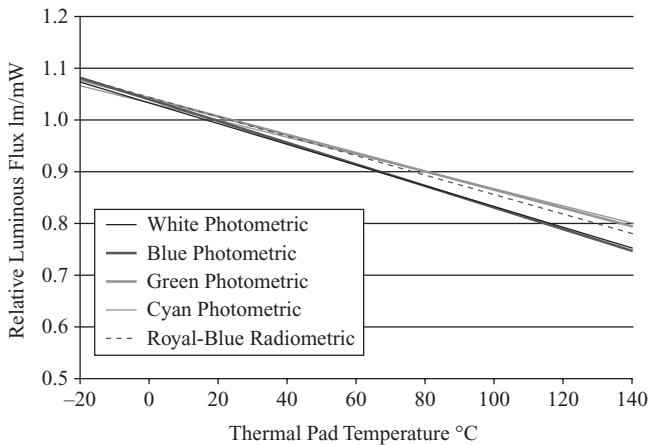


Figure 12.9 The Effect of Temperature on Light Output.

Source: Technical Datasheet DS56, Power Light Source Luxeon Rebel, Philips Lumileds Lighting Co., 2007.

140°C it is 0.75. The slope is thus $(0.75 - 1.00)/(140^\circ\text{C} - 20^\circ\text{C}) = -0.0021/^\circ\text{C}$, a decrease in light of about 1% for every 5°C increase in temperature.

The model including this is shown in Table 12.17, with the schematic shown in Figure 12.10. We have multiplied the normal light output by the factor above times the temperature of the LED minus 20°C.

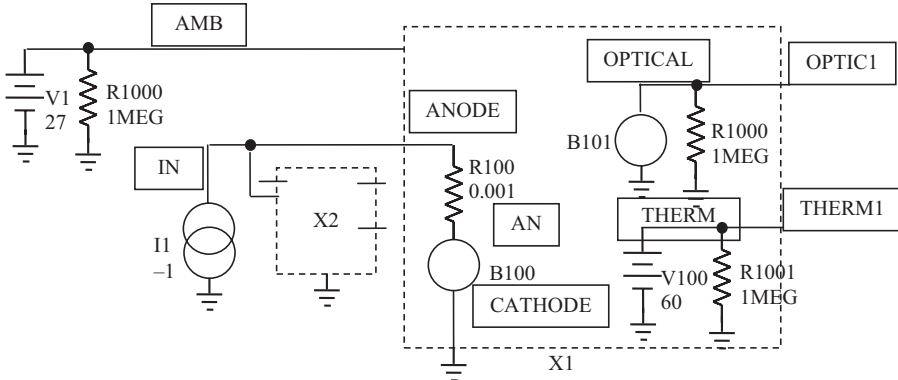


Figure 12.10 LED Model with Temperature Effect on Both Forward Voltage and Optical Output.

Table 12.17 A Model of Two Rebels Including Temperature Effects on Forward Voltage

TWO REBELS WITH TEMPERATURE EFFECT ON FORWARD VOLTAGE

X1 IN LED1 OPTIC1 THERM1 AMB LED

X2 LED1 0 OPTIC2 THERM2 AMB LED

V1 AMB 0 27

R1000 AMB 0 1MEG

I1 IN 0 -1

.SUBCKT LED ANODE CATHODE OPTICAL THERM AMB

B100 AN CATHODE V=0.27474*(LN(V(ANODE)-V(AN))+19.5716)-0.003*(V(THERM)-V(AMB))

B101 OPTICAL 0 V=(-0.0001*((V(ANODE)-V(AN))*10^6)^2+0.3093*((V(ANODE)-V(AN))*10^6)+3.647)*(1-0.0021*(V(THERM)-20))

R100 ANODE AN 0.001

R1000 OPTICAL 0 1MEG

R1001 THERM 0 1MEG

V100 THERM 0 60

.ENDS

MODELING THE THERMAL ENVIRONMENT

We've built a model of an LED, which was the main goal of this chapter. At the moment, however, this model has the LED temperature as an internally set parameter. What actually happens physically, of course, is that the power dissipated by the LED sets the LED temperature (along with the ambient). Our goal in this section is to model this.

Table 12.18 Complete LED Model

```

MULTIPLE LUXEON REBELS
X1 IN LED1 OPTIC1 THERM1 AMB LED
X2 LED1 0 OPTIC2 THERM2 AMB LED
V1 AMB 0 27
R1 THERM1 AMB 10
R2 THERM2 AMB 20
II IN 0 -1
.SUBCKT LED ANODE CATHODE OPTICAL THERM AMB
B100 AN CATHODE V=0.27474*(LN(V(ANODE)-V(AN))+19.5716)-0.003*(V(THERM)-
  V(AMB))
B101 OPTICAL 0 V=(-0.0001*((V(ANODE)-V(AN))*10^6)^2+0.3093*((V(ANODE)-
  V(AN))*10^6)+3.647)*(1-0.0021*(V(THERM)-20))
B102 THERM 0 I=-((V(ANODE)-V(AN))*10^3)*(V(ANODE)-V(CATHODE))
R100 ANODE AN 0.001
R1000 OPTICAL 0 1MEG
.ENDS

```

To complete our LED model, we will add a thermal resistance from the device to the ambient temperature. This will replace the fixed temperature at “THERM” with the calculated value based on the power dissipation of the LED. The power is of course the current through the LED (as measured by “R100”) multiplied by the voltage across the LED. We choose one of the thermal resistances to be 10°C/W, and the other 20°C/W. Note that this thermal resistance is not in the LED model itself. This is because it isn’t a property of the LED. Since it is part of the external world, for example, the same as the drive circuit, it is included in the main listing.

We add in another “B” in the model that calculates the power dissipated in the LED as the current through it multiplied by the voltage across it. This is the expression $(V(\text{ANODE})-V(\text{AN})) \cdot 10^3 \cdot (V(\text{ANODE})-V(\text{CATHODE}))$, with the first term being the current. Since we now have electrical paths to ground for the “AMB” and “THERM” nodes, we can remove their dummy resistors. “OPTICAL” is the only node that still needs a dummy.

The completed LED model is shown in Table 12.18 with the schematic shown in Figure 12.11. Running it, we see that the one LED is at a temperature of 61°C with a light output of 195 lumens, while the other is at 93°C with a correspondingly reduced light output of 180 lumens. This model can now be attached to electrical components, for example, those modeling the power supply. But this is beyond the scope of this book.




Figure 12.12 The Thermal Response of Two LEDs to Current Pulse.




Figure 12.13 The Optical Response of Two LEDs to Current Pulse.

shows that the LED with the shorter thermal time constant reaches its new steady-state temperature in about 5 seconds. The other LED takes nearly 20 seconds. The shorter time constant LED also has lower thermal resistance to the ambient temperature, and so doesn't get as hot (50°C vs. 80°C). The optical output shown in Figure 12.13 shows that corresponding to the hotter temperature, the longer time constant LED puts out less light (185 lumens vs. 200 lumens).

SOME COMMENTS ON MODELING

The foregoing sections epitomize our approach to modeling. You build up one section at a time, and check that it works before adding the next piece. Trying to debug it all at once is too difficult. Of course, the same thing could have been done with Spice with a schematic entry system. It's just one extra interface with potential problems of which you have to be aware.

You might notice that none of the models developed had any convergence issues. The Spice software is actually fairly robust. Convergence problems tend to arise only when there are abrupt transitions. An example would be using an ideal switch, "SW," for a transistor. In this case, minimizing the stiffness of the problem

is usually helpful. The stiffness is the ratio of the biggest number to the smallest number. For example, with the switch, the ratio of the off-resistance to the on-resistance is the stiffness to be minimized. If $1\ \Omega$ is a good model for the transistor, don't model the switch as $1\ \text{m}\Omega$. There are global parameters in Spice for helping along convergence, such as RELTOL, but our general feeling is that it shouldn't be necessary to tinker with these. If you're having convergence problems, you should work on your model further.

References

- Alliance for Solid-State Illumination Systems and Technologies (ASSIST). 2006. "ASSIST Recommends ... LED Life for General Lighting: Recommendations for the Definition and Specification of Useful Life for Light-Emitting Diode Light Sources." 1(7), February 2005, revised April 2006.
- BETTEN, JOHN, and ROBERT KOLLMANT. 2007. "LED Lighting Illuminates Buck Regulator Design." *Power Electronics Technology*, October, p. 38. Available online at <http://powerelectronics.com/mag/710PET24.pdf>.
- DOWLING, KEVIN. 2009. "LED Lighting Standards and Guidelines Are Now Building on a Firm Foundation." *LEDs Magazine*, May/June.
- HAITZ, ROLAND (Editorial Board). 2007. Editorial on Haitz's Law, *Nature Photonics*, 1(23).
- JAMESON, KIMBERLY A., SUSAN M. HIGHNOTE, and LINDA M. WASSERMAN. 2001. "Richer Color Experience in Observers with Multiple Photopigment Opsin Genes." *Psychonomic Bulletin & Review*, 9(2): 244–261.
- KALLONIATIS, MICHAEL, and CHARLES LUU. 2007. *Principles of Vision*. Webvision. <http://www.ncbi.nlm.nih.gov/bookshelf/br.fcgi?book=webvision&part=ch24psych1>.
- LENK, RON. 1998. *Practical Design of Power Supplies*. Hoboken, N.J.: IEEE Press/Wiley.
- MACADAM, DAVID L. 1942. "Visual Sensitivities to Color Differences in Daylight." *Journal of the Optical Society of America*, 32: 247.
- MARTZLOFF, FRANÇOISE. 1991. "A Standard for the 90s: IEEE C62.41 Surges Ahead." <http://www.eeel.nist.gov/817/pubs/spd-anthology/files/Standard%20for%2090s.pdf>.
- OHNO, YOSHI. 2004. "Color Rendering and Luminous Efficacy of White LED Spectra." *Proceedings of SPIE*, 5530: 88.
- SCHEIDT, PAUL. 2010. "Innovative Packaging Improves LEDs' Light Output, Lifetime and Reliability." *Electronic Design News*, January.
- STRAKA, TODD. 2009. "Navigating the Product Safety Certification Process for Solid-State Lighting Products." *LEDs Magazine*, November/December, pp. 37–38.
- TSAO, JEFFREY Y., et al. 2010. "Solid-State Lighting: An Energy-Economics Perspective." *Journal of Physics D: Applied Physics*, 43: 354001.
- TURNER, MICHAEL. 1996. http://www.electronics-cooling.com/articles/1996/may/may96_01.php.
- WONG, BOBBY, and ZHENG, LANCE. 2009. "Intelligent LED Lighting Systems and Network Technology Choices." In *Designing with LEDs*, e-book. Electronic Design News, April. http://www.edn.com/article/459201-Designing_with_LEDs_E_Book.php.
- ZONG, YUQIN, and YOSHI OHNO. 2008. "New Practical Method for Measurement of High-Power LEDs." In CIE Expert Symposium 2008 on Advances in Photometry and Colorimetry, pp. 102–106.

Index

Symbols

802.15.4, 143

A

Absorption, lenses, 137

AC

277VAC, 101

bridge rectifier, 103

brownout, 102

frequency, 103

Japan, 102

peak voltage, 102

rectification, 103

capacitance, 104

half-wave, 104

safety, 99

recommendations, 100

source impedance, 102

voltage, 101

voltage range, 101, 102

Accelerated life tests, 203

Accuracy, 197

AC LED, 105

AC power meter, 200

AIM-Spice, 208

Alkaline battery, 83

Ambient, definition, 74

ANSI C78.377, 40, 59

ANSI/IEEE C62.41, 119

ASSIST, 40, 193

B

Ballast

AC, nonisolated, 109

boost, introduction, 88

buck

charge pump, 88

details, 108

buck-boost, introduction, 89

component temperature ratings, 80

DC vs. AC, 81

dimmers, triac, 121

dimming, 120

IC keep-alive circuit, 121

efficiency, 79

EMI

common mode filter, 115

current loops, 115

heatsink, 116

normal mode filter, 115

recommendations, 116

fluorescent, 23

flyback, 111

continuous vs. discontinuous, 111

forward, 111

inductor current in, 86

isolation, 110

isolation transformer, 110

lifetime, AC factors, 125

lightning, 119

nonswitching, 106

opto-isolator, 112

overview, 85

PFC

definition, 117

simple circuits, 118

two-stage converter, 118

second meaning, 95, 106

switching frequency, 109

thermal analysis, 79

- Ballast (*cont'd*)
 - three types, 81
 - topologies, 86
- Bases, 21
 - pin base, 21
- Battery
 - aging, 85
 - alkaline, 83
 - comparison example, 146
 - energy storage capability, 83
 - impedance example, 151
 - introduction to, 82
 - I-V, 83
 - lithium, 83
 - circuitry, 83
 - model, 84
 - self-discharge, 84
 - shelf life, 150
 - temperature effects, 84
- Batwing, 56
- Biax, 21
- Binning
 - can't buy just one bin, 59
 - human perception, 59
- Black body, 39
- Boost, when used, 82
- Brownout, 102
- Buck, when it should be used, 82
- Buck-boost, when it should be used, 82

- C**
- Candela, definition, 32
- Candlepower, definition, 34
- Capacitor
 - lifetime, 94, 125
 - ripple current rating, 95
 - series connection for higher voltage, 114
 - UL limit, 115
 - voltage rating, 113
- CCT. *See* correlated color temperature
- Cell, battery, 82
- Center beam candle power, definition, 34
- CFLs
 - dimmers, 18
 - flicker, 17
 - problems
 - high temperature, 24
 - turn-on delay, 17
- Cheater, 100
- Chromaticity coordinates, 37
- CIE, 28
- Circadian rhythm, 44
- Color rendition index (CRI), 42, 43
- Color temperature, 39
- Conduction, 75
- Convection, 75
- Correlated color temperature (CTT), 16, 39
- Cosine corrected, 187
- Crest factor, definition, 200
- CRI. *See* color rendition index
- Current-mode control, 91

- D**
- DALI, 142
- DC drive circuit, 81, 82
- Dimmers, 18, 120
 - minimum wattage, 18
 - problems, 18
- Divide-by-Three Rule, 150
- DMX512, 143
- Drive circuit, 81
- Driver, 24
- Dust depreciation, 25, 140

- E**
- Efficacy
 - change of, with thermal aging, 62
 - definition, 51
 - drive current, 53
 - human perception, 51
 - limit, 9
 - Ohno, 9
 - real limit, 9
 - roll-off, current, 52
 - Stokes' law, 9
 - temperature effects, 63
- EMI, measurement of, 202. *See also under* ballast
- Encapsulant, 62
- Energy Star, 193
- Epoxy, 62
- Examples, 145
 - AC EMI filter design, 183
 - flashlight, 145
 - LED light bulb, 175

taillight, 165
 USB light, 157
 Excel, 207

F

Fan, 78
 issues, 78
 noise, 79
 power, 79
 Fixture, definition, 13
 Flicker, 17
 Floodlights, 20
 Fluorescent bulb(s)
 aging, 19
 ballast, 14
 electronic, 24
 switching frequency, 24
 instant start, 23
 magnetic, 24
 rapid start, 23
 CCT and CRI, 16
 dimmers, 18
 efficacy
 system, 17
 tube, 17
 efficiency, 4
 failure modes, 15
 filament, 23
 dimming, 24
 power, 24
 flicker, 17, 104
 lifetime, 19
 numbering system, 21
 plasma temperature, 14
 pre-heat, 17
 problems
 dust, 25
 fixture losses, 25
 shapes, 20
 start time, 24
 temperature, 15
 Foot candle, definition, 35
 Fuse, 100

G

Glass
 optical transmission, 137
 types of, 137

Goniometer, 192
 Goniophotometer, 192

H

Haitz's Law
 data, 8
 deviations from, 8
 Halogen, vs. incandescent, 14
 Harsh environments, 139
 Heatsink, 76
 attachment, 77
 EMI, 116
 HID
 definition, 15
 restart time, 17
 turn-on time, 17
 Human eye response, 555 nm,
 51

I

IC
 supply current, 92
 voltage rating, 91
 IEEE587, 119
 Illuminance, 32
 Incandescent bulb(s)
 aging, 19
 daylight white CCT, 16
 dimmers, 18
 dimming, 14
 efficacy, 17, 30
 failure mode, 14
 filament, 13
 flicker, 17
 gas mixture, 14
 lifetime, 19
 safety, 14
 temperature, 14
 turn-on time, 17
 Included angle, 57
 Induction lighting, 15
 Integrating sphere, 189
 IR, 28
 Isolation transformer, 100
 Isolation, 110
 Isotemperature, definition,
 39
 Isotropic light, 30

L

- Lambertian, 35, 56
- Lamp, definition, 13
- LED(s)
 - 5 mm, 2
 - 435 nm, 3
 - absolute maximum temperature, 64
 - AC, 105
 - aging, 19
 - Arrhenius, 66
 - application checklist, 7
 - array, 95
 - binning, 58
 - brightness, temperature effects, 63
 - constant current, 62
 - constant power, 63
 - cost example, 148
 - CRI, 54
 - current
 - absolute max, 50
 - overdrive, 50, 57
 - x,y variation with, 55
 - definition, 1
 - designing for the future, 141
 - dimmers, 18
 - dimming
 - analog, 92
 - duty cycle, 93
 - flicker, 93
 - PWM, 92
 - dust depreciation, 140
 - efficacy, 10, 51
 - temperature effects, 63
 - Energy Star, 128
 - eye safety, 128
 - flicker, 18
 - floating above ground, 90
 - forward voltage, 48
 - temperature effects, 62
 - typical, 49
 - future of, 9
 - harsh environments, 139
 - humidity, 139
 - included angle, 57
 - industry, 10
 - I-V equation, 47
 - I-V model parameters, 47
 - lifetime, 19
 - catastrophic failure, 67
 - lumen degradation, 66
 - light emission distribution, 56
 - materials, 5
 - measurement accuracy, 60
 - optical emission angle, 56
 - optical spectrum, 54
 - blue spike, 54
 - optics, 6
 - paralleling, 67
 - parameters, 58
 - phosphor, saturation, 50
 - power, 2
 - problems, 10
 - high temperature, 24
 - traffic light, 21
 - R9, 54
 - reverse breakdown, 49
 - fixes, 49, 50
 - ripple current effects, 122
 - salt fog, 139
 - selection, 58
 - size of die, 5
 - small, vs. power, 2
 - standards, 126
 - UL, 129
 - $T1^{3/4}$, 2
 - thermal aging
 - die, 62
 - epoxy, 62
 - package, 62
 - silicone, 62
 - thermal cycle, 139
 - thermal effects, forward voltage, 62
 - thermal shift
 - encapsulant, 62
 - mechanisms, 61
 - thermal time constant, 65
 - tolerance, 60
 - unsuitability, 22
 - usages, 2
 - vibration, 140
 - viewing angle, 57
- Lenses, 137
 - absorption, 137
- Life tests, 203
- Lifetime, system, 94
 - determined by input capacitor, 94

Light

- binning, 40
- blackbody, 39
 - candela, definition, 32
 - candlepower, definition, 34
- CCT, definition, 39
- center beam candle power, definition, 34
- chromaticity coordinates, definition, 37
- circadian rhythm, 44
- color names, 39
- color rendition, 41
- color spaces, 41
 - RGB, 37
- color temperature, 39
- cones, 43
- CRI
 - color samples, 43
 - definition, 42
- eye response curve, 29
- foot candle, definition, 35
- illuminance, 34
 - definition, 32
 - examples, 37
- IR, 28
- isotropic, definition, 30
- Lambertian, 35
- luminance, 36
 - definition, 32
- luminous flux, definition, 28, 32
- luminous intensity, definition, 32
- lux, definition, 35
- MacAdam ellipses, 39
 - steps, 40
- mesoscopic vision, definition, 29
- monochromatic, 28
- nit, definition, 37
- perceived vs. measured, 43
- photometric, definition, 27
- photopic vision
 - definition, 28
 - efficacy, 29
- radiant and luminous flux conversion, 29
- radiant flux, definition, 28
- radiometric, definition, 27
- radiometric vs. photometric, 28
- rods, 43
- scotopic vision
 - definition, 28
 - efficacy, 29

- solid angle formula, 33
- UV, 28
 - view angle, definition, 33
 - visible wavelength, 27
- Light bulbs; *see also* fluorescent
 - bulb(s); incandescent bulb(s); LED(s)
 - A, 19
 - bases, 21
 - BR, 20
 - brightening, 66
 - definition, 13
 - economics, 23
 - energy usage, 23
 - floodlight, 20
 - G, 20
 - halogen, 20
 - most common type, 19
 - numbering system, 19
 - number of, 1, 23
 - PAR, 20
 - specialty, 21
 - spotlight, 20
 - types, 19
- Lighting
 - control, 142
 - future of, 22
 - industry, 10
 - systems, 23
 - upward drift in brightness, 22
- Lightning, 103, 119, 140
 - measurement of, 202
- Lithium battery, 83
- LM-79, 192
- LM-80, 66, 193
- Lumens
 - 555 nm, 9
 - 683 lumens, 9
 - definition, 9
- Luminance, 32
- Luminous flux, 28
- Luminous intensity, 32
- Lux, definition, 35
- Lux meter, 187

M

- MacAdam ellipses, 39
- MCPCBs, 77

- Measurement, 187
 - accelerated life test, 203
 - accuracy and precision, 197
 - AC power meter, 200
 - angular light distribution, 192
 - ASSIST, 193
 - ballast efficiency, 201
 - bi-modal life data, 205
 - crest factor, 200
 - cross-measurements, 202
 - EMI, 202
 - Energy Star, 193
 - goniometer, 192
 - goniophotometer, 192
 - integrating sphere, 189
 - baffle, 189
 - calibration, 190
 - collector, 189
 - sizing, 191
 - jitter, 199
 - LED(s)
 - lifetime, 193
 - self-heating, 192
 - standards, 192
 - temperature, 194
 - life test, 203
 - estimating life, 204
 - number of units, 203
 - lightning, 202
 - LM-79, 192
 - LM-80, 193
 - luminaire light output, 192
 - luminous flux, 189
 - lux meter, 187
 - background light, 188
 - distance sensitivity, 187
 - mechanical life test, 206
 - meter ranges, 200
 - MTTF, 204
 - power
 - AC, 200
 - DC, 197
 - power factor, 201
 - power-of-two rule, 205
 - rounding and truncation, 199
 - spectroradiometer, 191
 - thermal resistance, 195
 - Mesoscopic vision, 29
 - Metal-Core PCBs, 77
 - Model/modeling, 207
 - “B,” 211
 - convergence, 229
 - diode, 211
 - LED
 - complete model, 225
 - forward voltage, 214
 - optical output, 221
 - resistor, 209
 - reverse breakdown, 219
 - temperature effects, 224
 - thermal transient, 228
 - overview, 209
 - problems, 212, 213
 - subcircuit, 211
 - thermal environment, 227
 - timer, 212
 - “X,” 211
 - Monte Carlo, 208
 - MOV, 120
 - aging, 120
 - MTTF, 204
- N**
- Nit, definition, 37
- O**
- OLEDs, 11
 - Optimization, of design, 208
 - Opto-isolator, 112
- P**
- Package, aging, 6, 62, 66
 - PCB(s) (printed circuit boards), 131, 134
 - checklist, 136
 - circular, 155
 - good schematic = good PCB, 131
 - grounding, 134
 - high-frequency traces, 136
 - high-voltage traces, 136
 - power connections, 136
 - Peak current control, 91
 - Phosphor, 3
 - efficiency, 3
 - fluorescent light, 4
 - red, 62
 - saturation, 50

thermal effects, 62
 tri-band, 3
 Photometric light, 27
 Photopic vision, 28
 Planck line, 39, 59
 Planck, 39
 Plastics
 optical transmissivity, 138
 problems, 140
 properties, 138
 temperature limit, 138
 yellowing, 139
 Power factor
 correction, 117
 measurement of, 201
 Powerline communication, 143
 Precision, 197
 Price, estimating, 150
 Printed circuit boards. *See* PCB(s)
 Production yield, 208

R

Radiant flux, 28
 Radiation, 75
 Radiometric light, 27
 Rectification, 103
 Reverse battery protection, 150
 Rounding, 199

S

Safety, 99, 128
 recommendations, 100
 Schematic, 131
 checklist, 134
 color coding, 133
 connection with good layout, 131
 reference numbering, 132
 Scotopic vision, 28
 Side-emitting, 56
 Silicon/silicone, 62
 Sil-Pad, 172
 SMPS, definition, 85
 Solar radiation spectrum, 31
 Solid angle, formula, 33
 Specification
 guaranteed, 148
 typical, 148
 Spectroradiometer, 191

Spice, 207
 Spotlights, 20
 Stefan-Boltzmann law, 75
 Steradians, 32
 Switch-mode power supplies, 85

T

Temperature
 electrical analogy, 72
 units, 72
 Temperature compensation circuit, 166
 Thermal analysis
 ambient, 74
 analogy with electrical, 72
 ballast, 79
 conduction, 75
 convection, 75
 fan, 78
 heatsink, 76
 introduction to, 71
 MCPCBs, 77
 radiation, 75
 radiation enhancement, 79
 temperature estimate, 75
 Thermal capacitance
 definition, 71
 electrical analogy, 72
 estimating, 74
 measurement of, 196
 units, 72
 Thermal cycle, 139
 condensation, 139
 Thermally conductive insulators, 172
 Thermal power
 electrical analogy, 72
 units, 72
 Thermal resistance
 definition, 71
 electrical analogy, 72
 measurement of, 195
 parallel, 72
 units, 72
 wires, 162
 Thermal time constant, definition, 72
 Title 24, 21, 23
 Touch temperature, 149
 Track lighting, 20
 Triax, 21

240 Index

Truncation, 199

Tsao, 22

Typical, definition of in specification, 148

U

UL

AC safety tests, 100

capacitor limit, 115

glass, 138

isolation, 110

safe voltage, 81

standards list, 129

USB specification, 157

UV light, 28

V

View/viewing angle, 57

definition, 33

W

White light, creating, 3

Worst Case Analysis,
208

X

X-ray, PCB, 163

Z

ZigBee, 143