

# Praise for Practical Common Lisp

*“Finally, a Lisp book for the rest of us. If you want to learn how to write a factorial function, this is not your book. Seibel writes for the practical programmer, emphasizing the engineer/artist over the scientist and subtly and gracefully implying the power of the language while solving understandable real-world problems.*

*“In most chapters, the reading of the chapter feels just like the experience of writing a program, starting with a little understanding and then having that understanding grow, like building the shoulders upon which you can then stand. When Seibel introduced macros as an aside while building a test framework, I was shocked at how such a simple example made me really ‘get’ them. Narrative context is extremely powerful, and the technical books that use it are a cut above. Congrats!”* —**Keith Irwin, Lisp programmer**

*“While learning Lisp, one is often referred to the CL HyperSpec if they do not know what a particular function does; however, I found that I often did not ‘get it’ just by reading the HyperSpec. When I had a problem of this manner, I turned to Practical Common Lisp every single time—it is by far the most readable source on the subject that shows you how to program, not just tells you.”*

—**Philip Haddad, Lisp programmer**

*“With the IT world evolving at an ever-increasing pace, professionals need the most powerful tools available. This is why Common Lisp—the most powerful, flexible, and stable programming language ever—is seeing such a rise in popularity. Practical Common Lisp is the long-awaited book that will help you harness the power of Common Lisp to tackle today’s complex real-world problems.”* —**Marc Battyani, author of CL-PDF, CL-TYPESETTING, and mod\_lisp**

*“Please don’t assume Common Lisp is useful only for databases, unit test frameworks, spam filters, ID3 parsers, Web programming, Shoutcast servers, HTML generation interpreters, and HTML generation compilers just because these are the only things that happen to be implemented in the book Practical Common Lisp.”* —**Tobias C. Rittweiler, Lisp programmer**

*“When I met Peter, who just started writing this book, I asked myself (not him, of course), ‘Why yet another book on Common Lisp, when there are many nice introductory books?’ One year later, I found a draft of the new book and recognized I was wrong. This book is not ‘yet another’ one. The author focuses on practical aspects rather than on technical details of the language. When I first studied Lisp by reading an introductory book, I felt I understood the language, but I also had the impression, ‘so what?’—meaning I had no idea about how to use it. In contrast, this book leaps into a ‘practical’ chapter after the first few chapters that explains the very basic notions of the language. Then the readers are expected to learn more about the language while they are following the ‘practical’ projects, which are combined to form a product of significant size. After reading this book, the readers will feel they are expert programmers on Common Lisp since they have ‘finished’ a big project already. I think Lisp is the only language that allows this type of practical introduction. Peter makes use of this feature of the language in building up a fancy introduction to Common Lisp.”* —**Taiichi Yuasa, Professor, Department of Communications and Computer Engineering, Kyoto University**

# Practical Common Lisp

PETER SEIBEL

## **Practical Common Lisp**

**Copyright © 2005 by Peter Seibel**

Lead Editor: Gary Cornell

Technical Reviewers: Mikel Evins, Steven Haflich, Barry Margolin

Editorial Board: Steve Anglin, Dan Appleman, Ewan Buckingham, Gary Cornell, Tony Davis, Jason Gilmore, Jonathan Hassell, Chris Mills, Dominic Shakeshaft, Jim Sumser

Assistant Publisher: Grace Wong

Project Manager: Beth Christmas

Copy Edit Manager: Nicole LeClerc

Copy Editor: Kim Wimpsett

Production Manager: Kari Brooks-Copony

Production Editor: Ellie Fountain

Compositor: Susan Glinert

Proofreaders: Katie Stence, Liz Welch

Indexer: Kevin Broccoli

Cover Designer: Kurt Krames

Manufacturing Manager: Tom Debolski

### Library of Congress Cataloging-in-Publication Data

Seibel, Peter.

Practical COMMON LISP / Peter Seibel.

p. cm.

Includes index.

ISBN 1-59059-239-5 (hc. : alk. paper)

1. COMMON LISP (Computer program language) I. Title.

QA76.73.L23S45 2005

005.13'3--dc22

2005005859

All rights reserved. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage or retrieval system, without the prior written permission of the copyright owner and the publisher.

Printed and bound in the United States of America 9 8 7 6 5 4 3 2 1

Trademarked names may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, we use the names only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

Distributed to the book trade in the United States by Springer-Verlag New York, Inc., 233 Spring Street, 6th Floor, New York, NY 10013, and outside the United States by Springer-Verlag GmbH & Co. KG, Tiergartenstr. 17, 69112 Heidelberg, Germany.

In the United States: phone 1-800-SPRINGER, fax 201-348-4505, e-mail [orders@springer-ny.com](mailto:orders@springer-ny.com), or visit <http://www.springer-ny.com>. Outside the United States: fax +49 6221 345229, e-mail [orders@springer.de](mailto:orders@springer.de), or visit <http://www.springer.de>.

For information on translations, please contact Apress directly at 2560 Ninth Street, Suite 219, Berkeley, CA 94710. Phone 510-549-5930, fax 510-549-5939, e-mail [info@apress.com](mailto:info@apress.com), or visit <http://www.apress.com>.

The information in this book is distributed on an “as is” basis, without warranty. Although every precaution has been taken in the preparation of this work, neither the author(s) nor Apress shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the information contained in this work.

The source code for this book is available to readers at <http://www.apress.com> in the Downloads section and also at <http://www.gigamonkeys.com/book/>.

*For Lily, Mom, and Dad*

# Contents at a Glance

About the Author .....	xix
About the Technical Reviewer .....	xxi
Acknowledgments .....	xxiii
Typographical Conventions .....	xxv
<b>CHAPTER 1</b>	<b>Introduction: Why Lisp? ..... 1</b>
<b>CHAPTER 2</b>	<b>Lather, Rinse, Repeat: A Tour of the REPL ..... 9</b>
<b>CHAPTER 3</b>	<b>Practical: A Simple Database ..... 19</b>
<b>CHAPTER 4</b>	<b>Syntax and Semantics ..... 37</b>
<b>CHAPTER 5</b>	<b>Functions ..... 51</b>
<b>CHAPTER 6</b>	<b>Variables ..... 65</b>
<b>CHAPTER 7</b>	<b>Macros: Standard Control Constructs ..... 79</b>
<b>CHAPTER 8</b>	<b>Macros: Defining Your Own ..... 89</b>
<b>CHAPTER 9</b>	<b>Practical: Building a Unit Test Framework ..... 103</b>
<b>CHAPTER 10</b>	<b>Numbers, Characters, and Strings ..... 115</b>
<b>CHAPTER 11</b>	<b>Collections ..... 127</b>
<b>CHAPTER 12</b>	<b>They Called It LISP for a Reason: List Processing ..... 141</b>
<b>CHAPTER 13</b>	<b>Beyond Lists: Other Uses for Cons Cells ..... 153</b>
<b>CHAPTER 14</b>	<b>Files and File I/O ..... 163</b>
<b>CHAPTER 15</b>	<b>Practical: A Portable Pathname Library ..... 179</b>
<b>CHAPTER 16</b>	<b>Object Reorientation: Generic Functions ..... 189</b>
<b>CHAPTER 17</b>	<b>Object Reorientation: Classes ..... 203</b>
<b>CHAPTER 18</b>	<b>A Few FORMAT Recipes ..... 219</b>
<b>CHAPTER 19</b>	<b>Beyond Exception Handling: Conditions and Restarts ..... 233</b>
<b>CHAPTER 20</b>	<b>The Special Operators ..... 245</b>
<b>CHAPTER 21</b>	<b>Programming in the Large: Packages and Symbols ..... 263</b>
<b>CHAPTER 22</b>	<b>LOOP for Black Belts ..... 277</b>

<b>CHAPTER 23</b>	Practical: A Spam Filter .....	291
<b>CHAPTER 24</b>	Practical: Parsing Binary Files .....	311
<b>CHAPTER 25</b>	Practical: An ID3 Parser .....	335
<b>CHAPTER 26</b>	Practical: Web Programming with AllegroServe .....	363
<b>CHAPTER 27</b>	Practical: An MP3 Database .....	385
<b>CHAPTER 28</b>	Practical: A Shoutcast Server .....	401
<b>CHAPTER 29</b>	Practical: An MP3 Browser .....	411
<b>CHAPTER 30</b>	Practical: An HTML Generation Library, the Interpreter .....	431
<b>CHAPTER 31</b>	Practical: An HTML Generation Library, the Compiler .....	449
<b>CHAPTER 32</b>	Conclusion: What's Next? .....	465
<b>INDEX</b>	.....	481

# Contents

About the Author .....	xix
About the Technical Reviewer .....	xxi
Acknowledgments .....	xxiii
Typographical Conventions .....	xxv
■ <b>CHAPTER 1</b> <b>Introduction: Why Lisp?</b> .....	1
Why Lisp? .....	2
Where It Began. ....	4
Who This Book Is For. ....	7
■ <b>CHAPTER 2</b> <b>Lather, Rinse, Repeat: A Tour of the REPL</b> .....	9
Choosing a Lisp Implementation .....	9
Getting Up and Running with Lisp in a Box. ....	11
Free Your Mind: Interactive Programming .....	12
Experimenting in the REPL .....	12
“Hello, World,” Lisp Style .....	13
Saving Your Work. ....	15
■ <b>CHAPTER 3</b> <b>Practical: A Simple Database</b> .....	19
CDs and Records .....	19
Filing CDs .....	21
Looking at the Database Contents. ....	21
Improving the User Interaction .....	23
Saving and Loading the Database .....	25
Querying the Database .....	27
Updating Existing Records—Another Use for WHERE .....	31
Removing Duplication and Winning Big .....	32
Wrapping Up. ....	36



<b>CHAPTER 4</b>	<b>Syntax and Semantics</b>	37
	What's with All the Parentheses?	37
	Breaking Open the Black Box	38
	S-expressions	39
	S-expressions As Lisp Forms	41
	Function Calls	42
	Special Operators	43
	Macros	44
	Truth, Falsehood, and Equality	45
	Formatting Lisp Code	47
<b>CHAPTER 5</b>	<b>Functions</b>	51
	Defining New Functions	51
	Function Parameter Lists	53
	Optional Parameters	53
	Rest Parameters	55
	Keyword Parameters	56
	Mixing Different Parameter Types	57
	Function Return Values	58
	Functions As Data, a.k.a. Higher-Order Functions	59
	Anonymous Functions	61
<b>CHAPTER 6</b>	<b>Variables</b>	65
	Variable Basics	65
	Lexical Variables and Closures	68
	Dynamic, a.k.a. Special, Variables	69
	Constants	74
	Assignment	74
	Generalized Assignment	75
	Other Ways to Modify Places	76
<b>CHAPTER 7</b>	<b>Macros: Standard Control Constructs</b>	79
	WHEN and UNLESS	80
	COND	82
	AND, OR, and NOT	82

Looping .....	83
DOLIST and DOTIMES .....	84
DO .....	85
The Mighty LOOP .....	87
<b>CHAPTER 8    Macros: Defining Your Own .....</b>	<b>89</b>
The Story of Mac: A Just-So Story .....	89
Macro Expansion Time vs. Runtime .....	90
DEFMACRO .....	91
A Sample Macro: do-primers .....	92
Macro Parameters .....	93
Generating the Expansion .....	95
Plugging the Leaks .....	96
Macro-Writing Macros .....	100
Beyond Simple Macros .....	102
<b>CHAPTER 9    Practical: Building a Unit Test Framework .....</b>	<b>103</b>
Two First Tries .....	103
Refactoring .....	105
Fixing the Return Value .....	106
Better Result Reporting .....	108
An Abstraction Emerges .....	109
A Hierarchy of Tests .....	110
Wrapping Up .....	112
<b>CHAPTER 10   Numbers, Characters, and Strings .....</b>	<b>115</b>
Numbers .....	116
Numeric Literals .....	117
Basic Math .....	119
Numeric Comparisons .....	121
Higher Math .....	122
Characters .....	122
Character Comparisons .....	122
Strings .....	123
String Comparisons .....	124

<b>CHAPTER 11</b>	<b>Collections</b>	127
	Vectors	127
	Subtypes of Vector	129
	Vectors As Sequences	130
	Sequence Iterating Functions	130
	Higher-Order Function Variants	133
	Whole Sequence Manipulations	134
	Sorting and Merging	135
	Subsequence Manipulations	136
	Sequence Predicates	137
	Sequence Mapping Functions	137
	Hash Tables	138
	Hash Table Iteration	140
<b>CHAPTER 12</b>	<b>They Called It LISP for a Reason: List Processing</b>	141
	“There Is No List”	141
	Functional Programming and Lists	144
	“Destructive” Operations	145
	Combining Recycling with Shared Structure	147
	List-Manipulation Functions	149
	Mapping	151
	Other Structures	152
<b>CHAPTER 13</b>	<b>Beyond Lists: Other Uses for Cons Cells</b>	153
	Trees	153
	Sets	155
	Lookup Tables: Alists and Plists	157
	DESTRUCTURING-BIND	161
<b>CHAPTER 14</b>	<b>Files and File I/O</b>	163
	Reading File Data	163
	Reading Binary Data	165
	Bulk Reads	165
	File Output	165
	Closing Files	167

Filenames . . . . .	168
How Pathnames Represent Filenames . . . . .	169
Constructing New Pathnames . . . . .	171
Two Representations of Directory Names . . . . .	173
Interacting with the File System . . . . .	173
Other Kinds of I/O . . . . .	175
<b>CHAPTER 15 Practical: A Portable Pathname Library . . . . .</b>	<b>179</b>
The API . . . . .	179
*FEATURES* and Read-Time Conditionalization . . . . .	180
Listing a Directory . . . . .	182
Testing a File's Existence . . . . .	185
Walking a Directory Tree . . . . .	187
<b>CHAPTER 16 Object Reorientation: Generic Functions . . . . .</b>	<b>189</b>
Generic Functions and Classes . . . . .	190
Generic Functions and Methods . . . . .	191
DEFGeneric . . . . .	193
DEFMethod . . . . .	194
Method Combination . . . . .	196
The Standard Method Combination . . . . .	197
Other Method Combinations . . . . .	198
Multimethods . . . . .	200
To Be Continued ... . . . .	202
<b>CHAPTER 17 Object Reorientation: Classes . . . . .</b>	<b>203</b>
DEFCLASS . . . . .	203
Slot Specifiers . . . . .	205
Object Initialization . . . . .	206
Accessor Functions . . . . .	209
WITH-SLOTS and WITH-ACCESSORS . . . . .	212
Class-Allocated Slots . . . . .	213
Slots and Inheritance . . . . .	214
Multiple Inheritance . . . . .	215
Good Object-Oriented Design . . . . .	218

<b>CHAPTER 18</b>	<b>A Few FORMAT Recipes</b>	219
	The FORMAT Function	220
	FORMAT Directives	221
	Basic Formatting	222
	Character and Integer Directives	223
	Floating-Point Directives	225
	English-Language Directives	226
	Conditional Formatting	227
	Iteration	228
	Hop, Skip, Jump	230
	And More	231
<b>CHAPTER 19</b>	<b>Beyond Exception Handling: Conditions and Restarts</b>	233
	The Lisp Way	234
	Conditions	235
	Condition Handlers	235
	Restarts	238
	Providing Multiple Restarts	240
	Other Uses for Conditions	241
<b>CHAPTER 20</b>	<b>The Special Operators</b>	245
	Controlling Evaluation	245
	Manipulating the Lexical Environment	246
	Local Flow of Control	248
	Unwinding the Stack	252
	Multiple Values	256
	EVAL-WHEN	258
	Other Special Operators	260
<b>CHAPTER 21</b>	<b>Programming in the Large: Packages and Symbols</b>	263
	How the Reader Uses Packages	263
	A Bit of Package and Symbol Vocabulary	265
	Three Standard Packages	266
	Defining Your Own Packages	267
	Packaging Reusable Libraries	270

Importing Individual Names .....	271
Packaging Mechanics .....	272
Package Gotchas .....	273
<b>CHAPTER 22 LOOP for Black Belts .....</b>	<b>277</b>
The Parts of a LOOP .....	277
Iteration Control .....	278
Counting Loops .....	278
Looping Over Collections and Packages .....	280
Equals-Then Iteration .....	281
Local Variables .....	282
Deconstructing Variables .....	282
Value Accumulation .....	283
Unconditional Execution .....	285
Conditional Execution .....	285
Setting Up and Tearing Down .....	287
Termination Tests .....	288
Putting It All Together .....	290
<b>CHAPTER 23 Practical: A Spam Filter .....</b>	<b>291</b>
The Heart of a Spam Filter .....	291
Training the Filter .....	295
Per-Word Statistics .....	297
Combining Probabilities .....	299
Inverse Chi Square .....	301
Training the Filter .....	302
Testing the Filter .....	303
A Couple of Utility Functions .....	305
Analyzing the Results .....	306
What's Next .....	309
<b>CHAPTER 24 Practical: Parsing Binary Files .....</b>	<b>311</b>
Binary Files .....	311
Binary Format Basics .....	312
Strings in Binary Files .....	314
Composite Structures .....	316

Designing the Macros .....	317
Making the Dream a Reality .....	318
Reading Binary Objects .....	320
Writing Binary Objects .....	322
Adding Inheritance and Tagged Structures .....	323
Keeping Track of Inherited Slots .....	325
Tagged Structures .....	327
Primitive Binary Types .....	329
The Current Object Stack .....	332
 <b>CHAPTER 25 Practical: An ID3 Parser .....</b>	 335
Structure of an ID3v2 Tag .....	336
Defining a Package .....	337
Integer Types .....	338
String Types .....	339
ID3 Tag Header .....	343
ID3 Frames .....	344
Detecting Tag Padding .....	346
Supporting Multiple Versions of ID3 .....	348
Versioned Frame Base Classes .....	350
Versioned Concrete Frame Classes .....	351
What Frames Do You Actually Need? .....	352
Text Information Frames .....	354
Comment Frames .....	356
Extracting Information from an ID3 Tag .....	357
 <b>CHAPTER 26 Practical: Web Programming with AllegroServe .....</b>	 363
A 30-Second Intro to Server-Side Web Programming .....	363
AllegroServe .....	365
Generating Dynamic Content with AllegroServe .....	368
Generating HTML .....	370
HTML Macros .....	373
Query Parameters .....	374
Cookies .....	377
A Small Application Framework .....	379
The Implementation .....	380

<b>CHAPTER 27</b>	<b>Practical: An MP3 Database</b>	385
	The Database	385
	Defining a Schema	388
	Inserting Values	390
	Querying the Database	392
	Matching Functions	394
	Getting at the Results	397
	Other Database Operations	398
<b>CHAPTER 28</b>	<b>Practical: A Shoutcast Server</b>	401
	The Shoutcast Protocol	401
	Song Sources	402
	Implementing Shoutcast	405
<b>CHAPTER 29</b>	<b>Practical: An MP3 Browser</b>	411
	Playlists	411
	Playlists As Song Sources	413
	Manipulating the Playlist	417
	Query Parameter Types	420
	Boilerplate HTML	422
	The Browse Page	423
	The Playlist	426
	Finding a Playlist	429
	Running the App	430
<b>CHAPTER 30</b>	<b>Practical: An HTML Generation Library, the Interpreter</b>	431
	Designing a Domain-Specific Language	431
	The FOO Language	433
	Character Escaping	435
	Indenting Printer	437
	HTML Processor Interface	438
	The Pretty Printer Backend	439
	The Basic Evaluation Rule	443
	What's Next?	447



- CHAPTER 31 Practical: An HTML Generation Library, the Compiler ..... 449
  - The Compiler ..... 449
  - FOO Special Operators ..... 454
  - FOO Macros ..... 459
  - The Public API ..... 462
  - The End of the Line ..... 463
- CHAPTER 32 Conclusion: What's Next? ..... 465
  - Finding Lisp Libraries ..... 465
  - Interfacing with Other Languages ..... 467
  - Make It Work, Make It Right, Make It Fast ..... 467
  - Delivering Applications ..... 475
  - Where to Go Next..... 477
- INDEX ..... 481

# About the Author



■ **Peter Seibel** is either a writer-turned-programmer or a programmer-turned-writer. After picking up an undergraduate degree in English and working briefly as a journalist, he was seduced by the Web. In the early '90s he hacked Perl for *Mother Jones* magazine and Organic Online. He participated in the Java revolution as an early employee at WebLogic and later taught Java programming at the University of California–Berkeley Extension. He's also one of the few second-generation Lisp programmers on the planet and was a childhood shareholder in Symbolics. He lives in Oakland with his wife, Lily, and their dog, Mahlanie.

# About the Technical Reviewer

■ **Barry Margolin** taught himself computer programming in high school in the late '70s, first on DEC PDP-8 time-sharing systems and then on Radio Shack TRS-80 personal computers, and he learned operating system design by reverse engineering these systems. He went to M.I.T., where he learned Lisp programming from Bernie Greenberg, author of the Multics MacLisp Compiler and Multics Emacs (the first Emacs clone to be written in Lisp); David Moon (one of the implementers of ITS Maclisp and a founder of Symbolics); and Alan Bawden (perhaps one of the best Lisp macrologists). After getting his computer science degree, he went to work for the Honeywell Multics development group, maintaining Emacs. When Honeywell discontinued Multics development, he went to Thinking Machines Corporation to maintain their Lisp Machine development environment. Since then, he has worked for Bolt, Beranek, and Newman—which became BBN Planet, then GTE Internetworking, and then Genuity, until being acquired by Level(3)—providing technical support for their Internet services. He's now working for Symantec providing level-two customer technical support for its enterprise firewall products.

# Acknowledgments

This book wouldn't have been written, at least not by me, if not for a few happy coincidences. So, I have to start by thanking Steven Haflich of Franz, who, after we met at a get-together of Bay Area Lispniks, invited me to lunch with some Franz salespeople where, among other things, we discussed the need for a new Lisp book. Then I have to thank Steve Sears, one of the sales guys at that lunch, who put me in touch with Franz's president, Fritz Kunze, after Fritz mentioned he was looking for someone to write a Lisp book. And, of course, many thanks to Fritz for convincing Apress to publish a new Lisp book, for deciding I was the right guy to write it, and for providing encouragement and assistance along the way. Thanks also to Sheng-Chuang Wu of Franz, the instrument of much of that assistance.

One of my most indispensable resources while working on the book was the newsgroup `comp.lang.lisp`. The `comp.lang.lisp` regulars answered what must have seemed to them an endless stream of questions about various aspects of Lisp and its history. I also turned frequently to the Google archives for the group, a treasure trove of technical expertise. So, thanks to Google for making them available and to all `comp.lang.lisp` participants past and present for providing the content. In particular, I'd like to recognize two long-time `comp.lang.lisp` contributors—Barry Margolin, who has been providing tidbits of Lisp history and his own brand of quiet wisdom for as long as I've been reading the group; and Kent Pitman, who, in addition to having been one of the principal technical editors of the language standard and the author of the Common Lisp HyperSpec, has written hundreds of thousands, if not millions, of words in `comp.lang.lisp` postings elucidating various aspects of the language and how it came to be.

Other indispensable resources while working on the book were the Common Lisp libraries for PDF generation and typesetting, CL-PDF and CL-TYPESETTING, written by Marc Battyani. I used CL-TYPESETTING to generate handsome PDFs for my own red-pen editing and CL-PDF as the basis for the Common Lisp program I used to generate the line art that appears in this book.

I also want to thank the many people who reviewed draft chapters on the Web and sent me e-mails pointing out typos, asking questions, or simply wishing me well. While there were too many to mention them all by name, a few deserve special mention for their extensive feedback: J. P. Massar (a fellow Bay Area Lispnik who also bucked up my spirits several times with well-timed pizza lunches), Gareth McCaughan, Chris Riesbeck, Bulent Murtezaoglu, Emre Sevinc, Chris Perkins, and Tayssir John Gabbour. Several of my non-Lisping buddies also got roped into looking at some chapters: thanks to Marc Hedlund, Jolly Chen, Steve Harris, Sam Pullara, Sriram Srinivasan, and William Grosso for their feedback. Thanks also to Scott Whitten for permission to use the photo that appears in Figure 26-1.

My technical reviewers, Steven Haflich, Mikel Evins, and Barry Margolin, and my copy editor, Kim Wimpsett, improved this book in innumerable ways. Any errors that remain are, of course, my own. And thanks to everyone else at Apress who participated in getting this book out the door.

Finally, and most of all, I want to thank my family: Mom and Dad, for everything, and Lily, for always believing I could do it.

# Typographical Conventions

Inline text set like this is code, usually the names of functions, variables, classes, and so on, that either I've just introduced or I'm about to introduce. Names defined by the language standard are set like this: DEFUN. Larger bits of example code are set like this:

```
(defun foo (x y z)
  (+ x y z))
```

Since Common Lisp's syntax is notable for its regularity and simplicity, I use simple templates to describe the syntax of various Lisp forms. For instance, the following describes the syntax of DEFUN, the standard function-defining macro:

```
(defun name (parameter*)
  [documentation-string]
  body-form*)
```

Names in *italic* in those templates are meant to be filled in with specific names or forms that I'll describe in the text. An italicized name followed by an asterisk (\*) represents zero or more occurrences of whatever the name represents, and a name enclosed in brackets ([ ]) represents an optional element. Occasionally, alternatives will be separated by a bar (|). Everything else in the template—usually just some names and parentheses—is literal text that will appear in the form.

Finally, because much of your interaction with Common Lisp happens at the interactive read-eval-print loop, or REPL, I'll frequently show the result of evaluating Lisp forms at the REPL like this:

```
CL-USER> (+ 1 2)
3
```

The CL-USER> is the Lisp prompt and is always followed by the expression to be evaluated, (+ 1 2), in this case. The result and any other output generated are shown on the following lines. I'll also sometimes show the result of evaluating an expression by writing the expression followed by an  $\rightarrow$ , which is followed by the result, like this:

```
(+ 1 2)  $\rightarrow$  3
```

Occasionally, I'll use an equivalence sign ( $\equiv$ ) to express that two Lisp forms are equivalent, like this:

```
(+ 1 2 3)  $\equiv$  (+ (+ 1 2) 3)
```



# Introduction: Why Lisp?

If you think the greatest pleasure in programming comes from getting a lot done with code that simply and clearly expresses your intention, then programming in Common Lisp is likely to be about the most fun you can have with a computer. You'll get more done, faster, using it than you would using pretty much any other language.

That's a bold claim. Can I justify it? Not in a just a few pages in this chapter—you're going to have to learn some Lisp and see for yourself—thus the rest of this book. For now, let me start with some anecdotal evidence, the story of my own road to Lisp. Then, in the next section, I'll explain the payoff I think you'll get from learning Common Lisp.

I'm one of what must be a fairly small number of second-generation Lisp hackers. My father got his start in computers writing an operating system in assembly for the machine he used to gather data for his doctoral dissertation in physics. After running computer systems at various physics labs, by the 1980s he had left physics altogether and was working at a large pharmaceutical company. That company had a project under way to develop software to model production processes in its chemical plants—if you increase the size of this vessel, how does it affect annual production? The original team, writing in FORTRAN, had burned through half the money and almost all the time allotted to the project with nothing to show for their efforts. This being the 1980s and the middle of the artificial intelligence (AI) boom, Lisp was in the air. So my dad—at that point not a Lisper—went to Carnegie Mellon University (CMU) to talk to some of the folks working on what was to become Common Lisp about whether Lisp might be a good language for this project.

The CMU folks showed him some demos of stuff they were working on, and he was convinced. He in turn convinced his bosses to let his team take over the failing project and do it in Lisp. A year later, and using only what was left of the original budget, his team delivered a working application with features that the original team had given up any hope of delivering. My dad credits his team's success to their decision to use Lisp.

Now, that's just one anecdote. And maybe my dad is wrong about why they succeeded. Or maybe Lisp was better only in comparison to other languages of the day. These days we have lots of fancy new languages, many of which have incorporated features from Lisp. Am I really saying Lisp can offer you the same benefits today as it offered my dad in the 1980s? Read on.

Despite my father's best efforts, I didn't learn any Lisp in high school. After a college career that didn't involve much programming in any language, I was seduced by the Web and back into computers. I worked first in Perl, learning enough to be dangerous while building an online discussion forum for *Mother Jones* magazine's Web site and then moving to a Web shop, Organic Online, where I worked on big—for the time—Web sites such as the one Nike put up

during the 1996 Olympics. Later I moved onto Java as an early developer at WebLogic, now part of BEA. After WebLogic, I joined another startup where I was the lead programmer on a team building a transactional messaging system in Java. Along the way, my general interest in programming languages led me to explore such mainstream languages as C, C++, and Python, as well as less well-known ones such as Smalltalk, Eiffel, and Beta.

So I knew two languages inside and out and was familiar with another half dozen. Eventually, however, I realized my interest in programming languages was really rooted in the idea planted by my father's tales of Lisp—that different languages really are different, and that, despite the formal Turing equivalence of all programming languages, you really can get more done more quickly in some languages than others and have more fun doing it. Yet, ironically, I had never spent that much time with Lisp itself. So, I started doing some Lisp hacking in my free time. And whenever I did, it was exhilarating how quickly I was able to go from idea to working code.

For example, one vacation, having a week or so to hack Lisp, I decided to try writing a version of a program—a system for breeding genetic algorithms to play the game of Go—that I had written early in my career as a Java programmer. Even handicapped by my then rudimentary knowledge of Common Lisp and having to look up even basic functions, it still felt more productive than it would have been to rewrite the same program in Java, even with several extra years of Java experience acquired since writing the first version.

A similar experiment led to the library I'll discuss in Chapter 24. Early in my time at WebLogic I had written a library, in Java, for taking apart Java class files. It worked, but the code was a bit of a mess and hard to modify or extend. I had tried several times, over the years, to rewrite that library, thinking that with my ever-improving Java chops I'd find some way to do it that didn't bog down in piles of duplicated code. I never found a way. But when I tried to do it in Common Lisp, it took me only two days, and I ended up not only with a Java class file parser but with a general-purpose library for taking apart any kind of binary file. You'll see how that library works in Chapter 24 and use it in Chapter 25 to write a parser for the ID3 tags embedded in MP3 files.

## Why Lisp?

It's hard, in only a few pages of an introductory chapter, to explain why users of a language like it, and it's even harder to make the case for why you should invest your time in learning a certain language. Personal history only gets us so far. Perhaps I like Lisp because of some quirk in the way my brain is wired. It could even be genetic, since my dad has it too. So before you dive into learning Lisp, it's reasonable to want to know what the payoff is going to be.

For some languages, the payoff is relatively obvious. For instance, if you want to write low-level code on Unix, you should learn C. Or if you want to write certain kinds of cross-platform applications, you should learn Java. And any of a number companies still use a lot of C++, so if you want to get a job at one of them, you should learn C++.

For most languages, however, they payoff isn't so easily categorized; it has to do with subjective criteria such as how it feels to use the language. Perl advocates like to say that Perl “makes easy things easy and hard things possible” and revel in the fact that, as the Perl motto has it, “There's more than one way to do it.”<sup>1</sup> Python's fans, on the other hand, think Python is clean and simple and think Python code is easier to understand because, as *their* motto says, “There's only one way to do it.”

---

1. Perl is also worth learning as “the duct tape of the Internet.”

So, why Common Lisp? There's no immediately obvious payoff for adopting Common Lisp the way there is for C, Java, and C++ (unless, of course, you happen to own a Lisp Machine). The benefits of using Lisp have much more to do with the experience of using it. I'll spend the rest of this book showing you the specific features of Common Lisp and how to use them so you can see for yourself what it's like. For now I'll try to give you a sense of Lisp's philosophy.

The nearest thing Common Lisp has to a motto is the koan-like description, "the programmable programming language." While cryptic, that description gets at the root of the biggest advantage Common Lisp still has over other languages. More than any other language, Common Lisp follows the philosophy that what's good for the language's designer is good for the language's users. Thus, when you're programming in Common Lisp, you almost never find yourself wishing the language supported some feature that would make your program easier to write, because, as you'll see throughout this book, you can just add the feature yourself.

Consequently, a Common Lisp program tends to provide a much clearer mapping between your ideas about how the program works and the code you actually write. Your ideas aren't obscured by boilerplate code and endlessly repeated idioms. This makes your code easier to maintain because you don't have to wade through reams of code every time you need to make a change. Even systemic changes to a program's behavior can often be achieved with relatively small changes to the actual code. This also means you'll develop code more quickly; there's less code to write, and you don't waste time thrashing around trying to find a clean way to express yourself within the limitations of the language.<sup>2</sup>

Common Lisp is also an excellent language for exploratory programming—if you don't know exactly how your program is going to work when you first sit down to write it, Common Lisp provides several features to help you develop your code incrementally and interactively.

For starters, the interactive read-eval-print loop, which I'll introduce in the next chapter, lets you continually interact with your program as you develop it. Write a new function. Test it. Change it. Try a different approach. You never have to stop for a lengthy compilation cycle.<sup>3</sup>

- 
2. Unfortunately, there's little actual research on the productivity of different languages. One report that shows Lisp coming out well compared to C++ and Java in the combination of programmer and program efficiency is discussed at <http://www.norvig.com/java-lisp.html>.
  3. Psychologists have identified a state of mind called *flow* in which we're capable of incredible concentration and productivity. The importance of flow to programming has been recognized for nearly two decades since it was discussed in the classic book about human factors in programming *Peopleware: Productive Projects and Teams* by Tom DeMarco and Timothy Lister (Dorset House, 1987). The two key facts about flow are that it takes around 15 minutes to get into a state of flow and that even brief interruptions can break you right out of it, requiring another 15-minute immersion to reenter. DeMarco and Lister, like most subsequent authors, concerned themselves mostly with flow-destroying interruptions such as ringing telephones and inopportune visits from the boss. Less frequently considered but probably just as important to programmers are the interruptions caused by our tools. Languages that require, for instance, a lengthy compilation before you can try your latest code can be just as inimical to flow as a noisy phone or a nosy boss. So, one way to look at Lisp is as a language designed to keep you in a state of flow.



Other features that support a flowing, interactive programming style are Lisp's dynamic typing and the Common Lisp condition system. Because of the former, you spend less time convincing the compiler you should be allowed to run your code and more time actually running it and working on it,<sup>4</sup> and the latter lets you develop even your error handling code interactively.

Another consequence of being “a programmable programming language” is that Common Lisp, in addition to incorporating small changes that make particular programs easier to write, can easily adopt big new ideas about how programming languages should work. For instance, the original implementation of the Common Lisp Object System (CLOS), Common Lisp's powerful object system, was as a library written in portable Common Lisp. This allowed Lisp programmers to gain actual experience with the facilities it provided before it was officially incorporated into the language.

Whatever new paradigm comes down the pike next, it's extremely likely that Common Lisp will be able to absorb it without requiring any changes to the core language. For example, a Lisper has recently written a library, AspectL, that adds support for aspect-oriented programming (AOP) to Common Lisp.<sup>5</sup> If AOP turns out to be the next big thing, Common Lisp will be able to support it without any changes to the base language and without extra preprocessors and extra compilers.<sup>6</sup>

## Where It Began

Common Lisp is the modern descendant of the Lisp language first conceived by John McCarthy in 1956. Lisp circa 1956 was designed for “symbolic data processing”<sup>7</sup> and derived its name from one of the things it was quite good at: LISt Processing. We've come a long way since then: Common Lisp sports as fine an array of modern data types as you can ask for: a condition system

- 
4. This point is bound to be somewhat controversial, at least with some folks. Static versus dynamic typing is one of the classic religious wars in programming. If you're coming from C++ and Java (or from statically typed functional languages such as Haskell and ML) and refuse to consider living without static type checks, you might as well put this book down now. However, before you do, you might first want to check out what self-described “statically typed bigot” Robert Martin (author of *Designing Object Oriented C++ Applications Using the Booch Method* [Prentice Hall, 1995]) and C++ and Java author Bruce Eckel (author of *Thinking in C++* [Prentice Hall, 1995] and *Thinking in Java* [Prentice Hall, 1998]) have had to say about dynamic typing on their weblogs (<http://www.artima.com/weblogs/viewpost.jsp?thread=4639> and <http://www.mindview.net/WebLog/log-0025>). On the other hand, folks coming from Smalltalk, Python, Perl, or Ruby should feel right at home with this aspect of Common Lisp.
  5. AspectL is an interesting project insofar as AspectJ, its Java-based predecessor, was written by Gregor Kiczales, one of the designers of Common Lisp's object and metaobject systems. To many Lispers, AspectJ seems like Kiczales's attempt to backport his ideas from Common Lisp into Java. However, Pascal Costanza, the author of AspectL, thinks there are interesting ideas in AOP that could be useful in Common Lisp. Of course, the reason he's able to implement AspectL as a library is because of the incredible flexibility of the Common Lisp Meta Object Protocol Kiczales designed. To implement AspectJ, Kiczales had to write what was essentially a separate compiler that compiles a new language into Java source code. The AspectL project page is at <http://common-lisp.net/project/aspectl/>.
  6. Or to look at it another, more technically accurate, way, Common Lisp comes with a built-in facility for integrating compilers for embedded languages.
  7. *Lisp 1.5 Programmer's Manual* (M.I.T. Press, 1962)

that, as you'll see in Chapter 19, provides a whole level of flexibility missing from the exception systems of languages such as Java, Python, and C++; powerful facilities for doing object-oriented programming; and several language facilities that just don't exist in other programming languages. How is this possible? What on Earth would provoke the evolution of such a well-equipped language?

Well, McCarthy was (and still is) an artificial intelligence (AI) researcher, and many of the features he built into his initial version of the language made it an excellent language for AI programming. During the AI boom of the 1980s, Lisp remained a favorite tool for programmers writing software to solve hard problems such as automated theorem proving, planning and scheduling, and computer vision. These were problems that required a lot of hard-to-write software; to make a dent in them, AI programmers needed a powerful language, and they grew Lisp into the language they needed. And the Cold War helped—as the Pentagon poured money into the Defense Advanced Research Projects Agency (DARPA), a lot of it went to folks working on problems such as large-scale battlefield simulations, automated planning, and natural language interfaces. These folks also used Lisp and continued pushing it to do what they needed.

The same forces that drove Lisp's feature evolution also pushed the envelope along other dimensions—big AI problems eat up a lot of computing resources however you code them, and if you run Moore's law in reverse for 20 years, you can imagine how scarce computing resources were on circa-80s hardware. The Lisp guys had to find all kinds of ways to squeeze performance out of their implementations. Modern Common Lisp implementations are the heirs to those early efforts and often include quite sophisticated, native machine code-generating compilers. While today, thanks to Moore's law, it's possible to get usable performance from a purely interpreted language, that's no longer an issue for Common Lisp. As I'll show in Chapter 32, with proper (optional) declarations, a good Lisp compiler can generate machine code quite similar to what might be generated by a C compiler.

The 1980s were also the era of the Lisp Machines, with several companies, most famously Symbolics, producing computers that ran Lisp natively from the chips up. Thus, Lisp became a systems programming language, used for writing the operating system, editors, compilers, and pretty much everything else that ran on the Lisp Machines.

In fact, by the early 1980s, with various AI labs and the Lisp machine vendors all providing their own Lisp implementations, there was such a proliferation of Lisp systems and dialects that the folks at DARPA began to express concern about the Lisp community splintering. To address this concern, a grassroots group of Lisp hackers got together in 1981 and began the process of standardizing a new language called Common Lisp that combined the best features from the existing Lisp dialects. Their work was documented in the book *Common Lisp the Language* by Guy Steele (Digital Press, 1984)—CLtL to the Lisp-cognoscenti.

By 1986 the first Common Lisp implementations were available, and the writing was on the wall for the dialects it was intended to replace. In 1996, the American National Standards Institute (ANSI) released a standard for Common Lisp that built on and extended the language specified in CLtL, adding some major new features such as the CLOS and the condition system. And even that wasn't the last word: like CLtL before it, the ANSI standard intentionally leaves room for implementers to experiment with the best way to do things: a full Lisp implementation provides a rich runtime environment with access to GUI widgets, multiple threads of control, TCP/IP sockets, and more. These days Common Lisp is evolving much like other open-source languages—the folks who use it write the libraries they need and often make them available to others. In the last few years, in particular, there has been a spurt of activity in open-source Lisp libraries.

So, on one hand, Lisp is one of computer science's "classical" languages, based on ideas that have stood the test of time.<sup>8</sup> On the other, it's a thoroughly modern, general-purpose language whose design reflects a deeply pragmatic approach to solving real problems as efficiently and robustly as possible. The only downside of Lisp's "classical" heritage is that lots of folks are still walking around with ideas about Lisp based on some particular flavor of Lisp they were exposed to at some particular time in the nearly half century since McCarthy invented Lisp. If someone tells you Lisp is only interpreted, that it's slow, or that you have to use recursion for everything, ask them what dialect of Lisp they're talking about and whether people were wearing bell-bottoms when they learned it.<sup>9</sup>

### BUT I LEARNED LISP ONCE, AND IT WASN'T LIKE WHAT YOU'RE DESCRIBING

If you've used Lisp in the past, you may have ideas about what "Lisp" is that have little to do with Common Lisp. While Common Lisp supplanted most of the dialects it's descended from, it isn't the only remaining Lisp dialect, and depending on where and when you were exposed to Lisp, you may very well have learned one of these other dialects.

Other than Common Lisp, the one general-purpose Lisp dialect that still has an active user community is Scheme. Common Lisp borrowed a few important features from Scheme but never intended to replace it.

Originally designed at M.I.T., where it was quickly put to use as a teaching language for undergraduate computer science courses, Scheme has always been aimed at a different language niche than Common Lisp. In particular, Scheme's designers have focused on keeping the core language as small and as simple as possible. This has obvious benefits for a teaching language and also for programming language researchers who like to be able to formally prove things about languages.

It also has the benefit of making it relatively easy to understand the whole language as specified in the standard. But, it does so at the cost of omitting many useful features that are standardized in Common Lisp. Individual Scheme implementations may provide these features in implementation-specific ways, but their omission from the standard makes it harder to write portable Scheme code than to write portable Common Lisp code.

*Continued*

- 
8. Ideas first introduced in Lisp include the if/then/else construct, recursive function calls, dynamic memory allocation, garbage collection, first-class functions, lexical closures, interactive programming, incremental compilation, and dynamic typing.
  9. One of the most commonly repeated myths about Lisp is that it's "dead." While it's true that Common Lisp isn't as widely used as, say, Visual Basic or Java, it seems strange to describe a language that continues to be used for new development and that continues to attract new users as "dead." Some recent Lisp success stories include Paul Graham's Viaweb, which became Yahoo Store when Yahoo bought his company; ITA Software's airfare pricing and shopping system, QPX, used by the online ticket seller Orbitz and others; Naughty Dog's game for the PlayStation 2, Jak and Daxter, which is largely written in a domain-specific Lisp dialect Naughty Dog invented called GOAL, whose compiler is itself written in Common Lisp; and the Roomba, the autonomous robotic vacuum cleaner, whose software is written in L, a downwardly compatible subset of Common Lisp. Perhaps even more telling is the growth of the Common-Lisp.net Web site, which hosts open-source Common Lisp projects, and the number of local Lisp user groups that have sprung up in the past couple of years.

Scheme also emphasizes a functional programming style and the use of recursion much more than Common Lisp does. If you studied Lisp in college and came away with the impression that it was only an academic language with no real-world application, chances are you learned Scheme. This isn't to say that's a particularly fair characterization of Scheme, but it's even less applicable to Common Lisp, which was expressly designed to be a real-world engineering language rather than a theoretically "pure" language.

If you've learned Scheme, you should also be aware that a number of subtle differences between Scheme and Common Lisp may trip you up. These differences are also the basis for several perennial religious wars between the hotheads in the Common Lisp and Scheme communities. I'll try to point out some of the more important differences as we go along.

Two other Lisp dialects still in widespread use are *Elisp*, the extension language for the Emacs editor, and *Autolisp*, the extension language for Autodesk's AutoCAD computer-aided design tool. Although it's possible more lines of Elisp and Autolisp have been written than of any other dialect of Lisp, neither can be used outside their host application, and both are quite old-fashioned Lisps compared to either Scheme or Common Lisp. If you've used one of these dialects, prepare to hop in the Lisp time machine and jump forward several decades.

## Who This Book Is For

This book is for you if you're curious about Common Lisp, regardless of whether you're already convinced you want to use it or if you just want to know what all the fuss is about.

If you've learned some Lisp already but have had trouble making the leap from academic exercises to real programs, this book should get you on your way. On the other hand, you don't have to be already convinced that you want to use Lisp to get something out of this book.

If you're a hard-nosed pragmatist who wants to know what advantages Common Lisp has over languages such as Perl, Python, Java, C, or C#, this book should give you some ideas. Or maybe you don't even care about using Lisp—maybe you're already sure Lisp isn't really any better than other languages you know but are annoyed by some Lisper telling you that's because you just don't "get it." If so, this book will give you a straight-to-the-point introduction to Common Lisp. If, after reading this book, you still think Common Lisp is no better than your current favorite languages, you'll be in an excellent position to explain exactly why.

I cover not only the syntax and semantics of the language but also how you can use it to write software that does useful stuff. In the first part of the book, I'll cover the language itself, mixing in a few "practical" chapters, where I'll show you how to write real code. Then, after I've covered most of the language, including several parts that other books leave for you to figure out on your own, the remainder of the book consists of nine more practical chapters where I'll help you write several medium-sized programs that actually do things you might find useful: filter spam, parse binary files, catalog MP3s, stream MP3s over a network, and provide a Web interface for the MP3 catalog and server.

After you finish this book, you'll be familiar with all the most important features of the language and how they fit together, you'll have used Common Lisp to write several nontrivial programs, and you'll be well prepared to continue exploring the language on your own. While everyone's road to Lisp is different, I hope this book will help smooth the way for you. So, let's begin.



# Lather, Rinse, Repeat: A Tour of the REPL

In this chapter you'll set up your programming environment and write your first Common Lisp programs. We'll use the easy-to-install Lisp in a Box developed by Matthew Danish and Mikel Evins, which packages a Common Lisp implementation with Emacs, a powerful Lisp-aware text editor, and SLIME,<sup>1</sup> a Common Lisp development environment built on top of Emacs.

This combo provides a state-of-the-art Common Lisp development environment that supports the incremental, interactive development style that characterizes Lisp programming. The SLIME environment has the added advantage of providing a fairly uniform user interface regardless of the operating system and Common Lisp implementation you choose. I'll use the Lisp in a Box environment in order to have a specific development environment to talk about; folks who want to explore other development environments such as the graphical integrated development environments (IDEs) provided by some of the commercial Lisp vendors or environments based on other editors shouldn't have too much trouble translating the basics.<sup>2</sup>

## Choosing a Lisp Implementation

The first thing you have to do is to choose a Lisp implementation. This may seem like a strange thing to have to do for folks used to languages such as Perl, Python, Visual Basic (VB), C#, and Java. The difference between Common Lisp and these languages is that Common Lisp is defined by its standard—there is neither a single implementation controlled by a benevolent dictator, as with Perl and Python, nor a canonical implementation controlled by a single company, as with

- 
1. Superior Lisp Interaction Mode for Emacs
  2. If you've had a bad experience with Emacs previously, you should treat Lisp in a Box as an IDE that happens to use an Emacs-like editor as its text editor; there will be no need to become an Emacs guru to program Lisp. It is, however, orders of magnitude more enjoyable to program Lisp with an editor that has some basic Lisp awareness. At a minimum, you'll want an editor that can automatically match `()`s for you and knows how to automatically indent Lisp code. Because Emacs is itself largely written in a Lisp dialect, Elisp, it has quite a bit of support for editing Lisp code. Emacs is also deeply embedded into the history of Lisp and the culture of Lisp hackers: the original Emacs and its immediate predecessors, TECMACS and TMACS, were written by Lispsers at the Massachusetts Institute of Technology (MIT). The editors on the Lisp Machines were versions of Emacs written entirely in Lisp. The first two Lisp Machine Emacs, following the hacker tradition of recursive acronyms, were EINE and ZWEI, which stood for EINE Is Not Emacs and ZWEI Was EINE Initially. Later ones used a descendant of ZWEI, named, more prosaically, ZMACS.

VB, C#, and Java. Anyone who wants to read the standard and implement the language is free to do so. Furthermore, changes to the standard have to be made in accordance with a process controlled by the standards body American National Standards Institute (ANSI). That process is designed to keep any one entity, such as a single vendor, from being able to arbitrarily change the standard.<sup>3</sup> Thus, the Common Lisp standard is a contract between any Common Lisp vendor and Common Lisp programmers. The contract tells you that if you write a program that uses the features of the language the way they're described in the standard, you can count on your program behaving the same in any conforming implementation.

On the other hand, the standard may not cover everything you may want to do in your programs—some things were intentionally left unspecified in order to allow continuing experimentation by implementers in areas where there wasn't consensus about the best way for the language to support certain features. So every implementation offers some features above and beyond what's specified in the standard. Depending on what kind of programming you're going to be doing, it may make sense to just pick one implementation that has the extra features you need and use that. On the other hand, if we're delivering Lisp source to be used by others, such as libraries, you'll want—as far as possible—to write portable Common Lisp. For writing code that should be mostly portable but that needs facilities not defined by the standard, Common Lisp provides a flexible way to write code “conditionalized” on the features available in a particular implementation. You'll see an example of this kind of code in Chapter 15 when we develop a simple library that smoothes over some differences between how different Lisp implementations deal with filenames.

For the moment, however, the most important characteristic of an implementation is whether it runs on our favorite operating system. The folks at Franz, makers of Allegro Common Lisp, are making available a trial version of their product for use with this book that runs on Linux, Windows, and OS X. Folks looking for an open-source implementation have several options. SBCL<sup>4</sup> is a high-quality open-source implementation that compiles to native code and runs on a wide variety of Unixes, including Linux and OS X. SBCL is derived from CMUCL,<sup>5</sup> which is a Common Lisp developed at Carnegie Mellon University, and, like CMUCL, is largely in the public domain, except a few sections licensed under Berkeley Software Distribution (BSD) style licenses. CMUCL itself is another fine choice, though SBCL tends to be easier to install and now supports 21-bit Unicode.<sup>6</sup> For OS X users, OpenMCL is an excellent choice—it compiles to machine code, supports threads, and has quite good integration with OS X's

- 
3. Practically speaking, there's very little likelihood of the language standard itself being revised—while there are a small handful of warts that folks might like to clean up, the ANSI process isn't amenable to opening an existing standard for minor tweaks, and none of the warts that might be cleaned up actually cause anyone any serious difficulty. The future of Common Lisp standardization is likely to proceed via de facto standards, much like the “standardization” of Perl and Python—as different implementers experiment with application programming interfaces (APIs) and libraries for doing things not specified in the language standard, other implementers may adopt them or people will develop portability libraries to smooth over the differences between implementations for features not specified in the language standard.
  4. Steel Bank Common Lisp
  5. CMU Common Lisp
  6. SBCL forked from CMUCL in order to focus on cleaning up the internals and making it easier to maintain. But the fork has been amiable; bug fixes tend to propagate between the two projects, and there's talk that someday they will merge back together.