

The Art of UNIX Programming

Eric S. Raymond



With contributions
from UNIX legends
Thompson, Kernighan,
McIlroy, Arnold,
Bellovin, Korn,
Gettys, Packard,
Lesk, Feldman,
McKusick,
and Spencer



ADDISON-WESLEY PROFESSIONAL COMPUTING SERIES

The photograph of Ken and Dennis in Chapter 2 appears courtesy of Bell Labs/Lucent Technologies.

The epigraph on the Portability chapter is from the Bell System Technical Journal, v57 #6 part 2 (July–Aug. 1978) pp. 2021–2048 and is reproduced with the permission of Bell Labs/Lucent Technologies.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers discounts on this book when ordered in quantity for bulk purchases and special sales. For more information, please contact:

U.S. Corporate and Government Sales
(800) 382-3419
corpsales@pearsontechgroup.com

For sales outside of the U.S., please contact:

International Sales
(317) 581-3793
international@pearsontechgroup.com

Visit Addison-Wesley on the Web: www.awprofessional.com

Library of Congress Cataloging-in-Publication Data

Raymond, Eric S.

The art of UNIX programming / Eric Steven Raymond.

p. cm.

ISBN 0-13-142901-9 (pbk.)

1. UNIX (Computer file) 2. Operating systems (Computers) I. Title.

QA76.76.O63R395 2003

005.4'32--dc22

2003058264

Copyright © 2004 by Pearson Education, Inc.

Printed in the United States of America. Published simultaneously in Canada.

For information on obtaining permission for use of material from this work, please submit a written request to: Pearson Education, Inc., Rights and Contracts Department, 75 Arlington Street, Suite 300, Boston, MA 02116, Fax: (617) 848-7047.

ISBN 0-13-142901-9

Text printed in the United States on recycled paper at RR Donnelley Crawfordsville in Crawfordsville, Indiana.

6th Printing September 2008

Preface

Unix is not so much an operating system as an oral history.

—Neal Stephenson

There is a vast difference between knowledge and expertise. Knowledge lets you deduce the right thing to do; expertise makes the right thing a reflex, hardly requiring conscious thought at all.

This book has a lot of knowledge in it, but it is mainly about expertise. It is going to try to teach you the things about Unix development that Unix experts know, but aren't aware that they know. It is therefore less about technicalia and more about *shared culture* than most Unix books—both explicit and implicit culture, both conscious and unconscious traditions. It is not a 'how-to' book, it is a 'why-to' book.

The why-to has great practical importance, because far too much software is poorly designed. Much of it suffers from bloat, is exceedingly hard to maintain, and is too difficult to port to new platforms or extend in ways the original programmers didn't anticipate. These problems are symptoms of bad design. We hope that readers of this book will learn something of what Unix has to teach about good design.

This book is divided into four parts: Context, Design, Tools, and Community. The first part (Context) is philosophy and history, to help provide foundation and motivation for what follows. The second part (Design) unfolds the principles of the Unix philosophy into more specific advice about design and implementation. The third part (Tools) focuses on the software Unix provides for helping you solve problems. The fourth part (Community) is about the human-to-human transactions and agreements that make the Unix culture so effective at what it does.

Because this is a book about shared culture, I never planned to write it alone. You will notice that the text includes guest appearances by prominent Unix developers,

the shapers of the Unix tradition. The book went through an extended public review process during which I invited these luminaries to comment on and argue with the text. Rather than submerging the results of that review process in the final version, these guests were encouraged to speak with their own voices, amplifying and developing and even disagreeing with the main line of the text.

In this book, when I use the editorial ‘we’ it is not to pretend omniscience but to reflect the fact that it attempts to articulate the expertise of an entire community.

Because this book is aimed at transmitting culture, it includes much more in the way of history and folklore and asides than is normal for a technical book. Enjoy; these things, too, are part of your education as a Unix programmer. No single one of the historical details is vital, but the gestalt of them all is important. We think it makes a more interesting story this way. More importantly, understanding where Unix came from and how it got the way it is will help you develop an intuitive feel for the Unix style.

For the same reason, we refuse to write as if history is over. You will find an unusually large number of references to the time of writing in this book. We do not wish to pretend that current practice reflects some sort of timeless and perfectly logical outcome of preordained destiny. References to time of writing are meant as an alert to the reader two or three or five years hence that the associated statements of fact may have become dated and should be double-checked.

Other things this book is not is neither a C tutorial, nor a guide to the Unix commands and API. It is not a reference for *sed* or *yacc* or Perl or Python. It’s not a network programming primer, nor an exhaustive guide to the mysteries of X. It’s not a tour of Unix’s internals and architecture, either. Other books cover these specifics better, and this book points you at them as appropriate.

Beyond all these technical specifics, the Unix culture has an unwritten engineering tradition that has developed over literally millions of man-years¹ of skilled effort. This book is written in the belief that understanding that tradition, and adding its design patterns to your toolkit, will help you become a better programmer and designer.

Cultures consist of people, and the traditional way to learn Unix culture is from other people and through the folklore, by osmosis. This book is not a substitute for person-to-person acculturation, but it can help accelerate the process by allowing you to tap the experience of others.

1. The three and a half decades between 1969 and 2003 is a long time. Going by the historical trend curve in number of Unix sites during that period, probably somewhere upwards of fifty million man-years have been plowed into Unix development worldwide.

Who Should Read This Book

You should read this book if you are an experienced Unix programmer who is often in the position of either educating novice programmers or debating partisans of other operating systems, and you find it hard to articulate the benefits of the Unix approach.

You should read this book if you are a C, C++, or Java programmer with experience on other operating systems and you are about to start a Unix-based project.

You should read this book if you are a Unix user with novice-level up to middle-level skills in the operating system, but little development experience, and want to learn how to design software effectively under Unix.

You should read this book if you are a non-Unix programmer who has figured out that the Unix tradition might have something to teach you. We believe you're right, and that the Unix philosophy can be exported to other operating systems. So we will pay more attention to non-Unix environments (especially Microsoft operating systems) than is usual in a Unix book; and when tools and case studies are portable, we say so.

You should read this book if you are an application architect considering platforms or implementation strategies for a major general-market or vertical application. It will help you understand the strengths of Unix as a development platform, and of the Unix tradition of open source as a development method.

You should *not* read this book if what you are looking for is the details of C coding or how to use the Unix kernel API. There are many good books on these topics; *Advanced Programming in the Unix Environment* [Stevens92] is classic among explorations of the Unix API, and *The Practice of Programming* [Kernighan-Pike99] is recommended reading for all C programmers (indeed for all programmers in any language).

How to Use This Book

This book is both practical and philosophical. Some parts are aphoristic and general, others will examine specific case studies in Unix development. We will precede or follow general principles and aphorisms with examples that illustrate them: examples drawn not from toy demonstration programs but rather from real working code that is in use every day.

We have deliberately avoided filling the book with lots of code or specification-file examples, even though in many places this might have made it easier to write (and in some places perhaps easier to read!). Most books about programming give too many low-level details and examples, but fail at giving the reader a high-level feel for what is really going on. In this book, we prefer to err in the opposite direction.

Therefore, while you will often be invited to read code and specification files, relatively few are actually included in the book. Instead, we point you at examples on the Web.

Absorbing these examples will help solidify the principles you learn into semi-instinctive working knowledge. Ideally, you should read this book near the console of a running Unix system, with a Web browser handy. Any Unix will do, but the software case studies are more likely to be preinstalled and immediately available for inspection on a Linux system. The pointers in the book are invitations to browse and experiment. Introduction of these pointers is paced so that wandering off to explore for a while won't break up exposition that has to be continuous.

Note: While we have made every effort to cite URLs that should remain stable and usable, there is no way we can guarantee this. If you find that a cited link has gone stale, use common sense and do a phrase search with your favorite Web search engine. Where possible we suggest ways to do this near the URLs we cite.

Most abbreviations used in this book are expanded at first use. For convenience, we have also provided a glossary in an appendix.

References are usually by author name. Numbered footnotes are for URLs that would intrude on the text or that we suspect might be perishable; also for asides, war stories, and jokes.²

To make this book more accessible to less technical readers, we invited some non-programmers to read it and identify terms that seemed both obscure and necessary to the flow of exposition. We also use footnotes for definitions of elementary terms that an experienced programmer is unlikely to need.

Related References

Some famous papers and a few books by Unix's early developers have mined this territory before. Kernighan and Pike's *The Unix Programming Environment* [Kernighan-Pike84] stands out among these and is rightly considered a classic. But today it shows its age a bit; it doesn't cover the Internet, and the World Wide Web or the new wave of interpreted languages like Perl, Tcl, and Python.

About halfway into the composition of this book, we learned of Mike Gancarz's *The Unix Philosophy* [Gancarz]. This book is excellent within its range, but did not attempt to cover the full spectrum of topics we felt needed to be addressed. Nevertheless we are grateful to the author for the reminder that the very simplest Unix design patterns have been the most persistent and successful ones.

2. This particular footnote is dedicated to Terry Pratchett, whose use of footnotes is quite...inspiring.

The Pragmatic Programmer [Hunt-Thomas] is a witty and wise disquisition on good design practice pitched at a slightly different level of the software-design craft (more about coding, less about higher-level partitioning of problems) than this book. The authors' philosophy is an outgrowth of Unix experience, and it is an excellent complement to this book.

The Practice of Programming [Kernighan-Pike99] covers some of the same ground as *The Pragmatic Programmer* from a position deep within the Unix tradition.

Finally (and with admitted intent to provoke) we recommend *Zen Flesh, Zen Bones* [Reps-Senzaki], an important collection of Zen Buddhist primary sources. References to Zen are scattered throughout this book. They are included because Zen provides a vocabulary for addressing some ideas that turn out to be very important for software design but are otherwise very difficult to hold in the mind. Readers with religious attachments are invited to consider Zen not as a religion but as a therapeutic form of mental discipline—which, in its purest non-theistic forms, is exactly what Zen is.

Conventions Used in This Book

The term “UNIX” is technically and legally a trademark of The Open Group, and should formally be used only for operating systems which are certified to have passed The Open Group's elaborate standards-conformance tests. In this book we use “Unix” in the looser sense widely current among programmers, to refer to any operating system (whether formally Unix-branded or not) that is either genetically descended from Bell Labs's ancestral Unix code or written in close imitation of its descendants. In particular, Linux (from which we draw most of our examples) is a Unix under this definition.

This book employs the Unix manual page convention of tagging Unix facilities with a following manual section in parentheses, usually on first introduction when we want to emphasize that this is a Unix command. Thus, for example, read “`munger(1)`” as “the ‘munger’ program, which will be documented in section 1 (user tools) of the Unix manual pages, if it's present on your system”. Section 2 is C system calls, section 3 is C library calls, section 5 is file formats and protocols, section 8 is system administration tools. Other sections vary among Unixes but are not cited in this book. For more, type **man 1 man** at your Unix shell prompt (older System V Unixes may require **man -s 1 man**).

Sometimes we mention a Unix application (such as *Emacs*), without a manual-section suffix and capitalized. This is a clue that the name actually represents a well-established family of Unix programs with essentially the same function, and we are discussing generic properties of all of them. *Emacs*, for example, includes *xemacs*.

At various points later in this book we refer to ‘old school’ and ‘new school’ methods. As with rap music, new-school starts about 1990. In this context, it's

associated with the rise of scripting languages, GUIs, open-source Unixes, and the Web. Old-school refers to the pre-1990 (and especially pre-1985) world of expensive (shared) computers, proprietary Unixes, scripting in shell, and C everywhere. This difference is worth pointing out because cheaper and less memory-constrained machines have wrought some significant changes on the Unix programming style.

Our Case Studies

A lot of books on programming rely on toy examples constructed specifically to prove a point. This one won't. Our case studies will be real, pre-existing pieces of software that are in production use every day. Here are some of the major ones:

cdrttools/xcdroast

These two separate projects are usually used together. The `cdrttools` package is a set of CLI tools for writing CD-ROMs; Web search for “`cdrttools`”. The `xcdroast` application is a GUI front end for `cdrttools`; see the `xcdroast` project site <<http://www.xcdroast.org/>>.

fetchmail

The *fetchmail* program retrieves mail from remote-mail servers using the POP3 or IMAP post-office protocols. See the `fetchmail` home page <<http://www.catb.org/~esr/fetchmail>> (or search for “`fetchmail`” on the Web).

GIMP

The *GIMP* (GNU Image Manipulation Program) is a full-featured paint, draw, and image-manipulation program that can edit a huge variety of graphical formats in sophisticated ways. Sources are available from the GIMP home page <<http://www.gimp.org/>> (or search for “GIMP” on the Web).

mutt

The *mutt* mail user agent is the current best-of-breed among text-based Unix electronic mail agents, with notably good support for MIME (Multipurpose Internet Mail Extensions) and the use of privacy aids such as PGP (Pretty Good Privacy) and GPG (GNU Privacy Guard). Source code and executable binaries are available at the Mutt project site <<http://www.mutt.org/>>.

xm1to

The *xm1to* command renders DocBook and other XML documents in various output formats, including HTML and text and PostScript. For sources and

documentation, see the xmlto project site <<http://cyberelk.net/tim/xmlto/>>.

To minimize the amount of code the user needs to read to understand the examples, we have tried to choose case studies that can be used more than once, ideally to illustrate several different design principles and practices. For this same reason, many of the examples are from my projects. No claim that these are the best possible ones is implied, merely that I find them sufficiently familiar to be useful for multiple expository purposes.

Author's Acknowledgements

The guest contributors (Ken Arnold, Steven M. Bellovin, Stuart Feldman, Jim Gettys, Steve Johnson, Brian Kernighan, David Korn, Mike Lesk, Doug McIlroy, Marshall Kirk McKusick, Keith Packard, Henry Spencer, and Ken Thompson) added a great deal of value to this book. Doug McIlroy, in particular, went far beyond the call of duty in the thoroughness of his critique and the depth of his contributions, displaying the same care and dedication to excellence which he brought to managing the original Unix research group thirty years ago.

Special thanks go to Rob Landley and to my wife Catherine Raymond, both of whom delivered intensive line-by-line critiques of manuscript drafts. Rob's insightful and attentive commentary actually inspired more than one entire chapter in the final manuscript, and he had a lot to do with its present organization and range; if he had written all the text he pushed me to improve, I would have to call him a co-author. Cathy was my test audience representing non-technical readers; to the extent this book is accessible to people who aren't already programmers, that's largely her doing.

This book benefited from discussions with many other people over the five years it took me to write it. Mark M. Miller helped me achieve enlightenment about threads. John Cowan supplied some insights about interface design patterns and drafted the case studies of *wily* and VM/CMS, and Jef Raskin showed me where the Rule of Least Surprise comes from. The UIUC System Architecture Group contributed useful feedback on early chapters. The sections on *What Unix Gets Wrong* and *Flexibility in Depth* were directly inspired by their review. Russell J. Nelson contributed the material on Bernstein chaining in Chapter 7. Jay Maynard contributed most of the material in the MVS case study in Chapter 3. Les Hatton provided many helpful comments on the *Languages* chapter and motivated the portion of Chapter 4 on *Optimal Module Size*. David A. Wheeler contributed many perceptive criticisms and some case-study material, especially in the Design part. Russ Cox helped develop the survey of Plan 9. Dennis Ritchie corrected me on some historical points about C.

Hundreds of Unix programmers, far too many to list here, contributed advice and comments during the book's public review period between January and June of 2003. As always, I found the process of open peer review over the Web both intensely challenging and intensely rewarding. Also as always, responsibility for any errors in the resulting work remains my own.

The expository style and some of the concerns of this book have been influenced by the design patterns movement; indeed, I flirted with the idea of titling the book *Unix Design Patterns*. I didn't, because I disagree with some of the implicit central dogmas of the movement and don't feel the need to use all its formal apparatus or accept its cultural baggage. Nevertheless, my approach has certainly been influenced by Christopher Alexander's work³ (especially *The Timeless Way of Building* and *A Pattern Language*), and I owe the Gang of Four and other members of their school a large debt of gratitude for showing me how it is possible to use Alexander's insights to talk about software design at a high level without merely uttering vague and useless generalities. Interested readers should see *Design Patterns: Elements of Reusable Object-Oriented Software* [GangOfFour] for an introduction to design patterns.

The title of this book is, of course, a reference to Donald Knuth's *The Art of Computer Programming*. While not specifically associated with the Unix tradition, Knuth has been an influence on us all.

Editors with vision and imagination aren't as common as they should be. Mark Taub is one; he saw merit in a stalled project and skillfully nudged me into finishing it. Copy editors with a good ear for prose style and enough ability to improve writing that isn't like theirs are even less common, but Mary Lou Nohr makes that grade. Jerry Votta seized on my concept for the cover and made it look better than I had imagined. The whole crew at Addison-Wesley gets high marks for making the editorial and production process as painless as possible, and for cheerfully accommodating my control-freak tendencies not just over the text but deep into the details of the book's visual design, art, and marketing.

3. An appreciation of Alexander's work, with links to on-line versions of significant portions, may be found at Some Notes on Christopher Alexander <<http://www.math.utsa.edu/sphere/salingar/Chris.text.html>>.

1

Philosophy: Philosophy Matters

Those who do not understand Unix are condemned to reinvent it, poorly.

Usenet signature, November 1987
—Henry Spencer

1.1 Culture? What Culture?

This is a book about Unix programming, but in it we're going to toss around the words 'culture', 'art', and 'philosophy' a lot. If you are not a programmer, or you are a programmer who has had little contact with the Unix world, this may seem strange. But Unix has a culture; it has a distinctive art of programming; and it carries with it a powerful design philosophy. Understanding these traditions will help you build better software, even if you're developing for a non-Unix platform.

Every branch of engineering and design has technical cultures. In most kinds of engineering, the unwritten traditions of the field are parts of a working practitioner's education as important as (and, as experience grows, often more important than) the official handbooks and textbooks. Senior engineers develop huge bodies of implicit knowledge, which they pass to their juniors by (as Zen Buddhists put it) "a special transmission, outside the scriptures".

Software engineering is generally an exception to this rule; technology has changed so rapidly, software environments have come and gone so quickly, that technical

cultures have been weak and ephemeral. There are, however, exceptions to this exception. A very few software technologies have proved durable enough to evolve strong technical cultures, distinctive arts, and an associated design philosophy transmitted across generations of engineers.

The Unix culture is one of these. The Internet culture is another—or, in the twenty-first century, arguably the same one. The two have grown increasingly difficult to separate since the early 1980s, and in this book we won't try particularly hard.

1.2 The Durability of Unix

Unix was born in 1969 and has been in continuous production use ever since. That's several geologic eras by computer-industry standards—older than the PC or workstations or microprocessors or even video display terminals, and contemporaneous with the first semiconductor memories. Of all production timesharing systems today, only IBM's VM/CMS can claim to have existed longer, and Unix machines have provided hundreds of thousands of times more service hours; indeed, Unix has probably supported more computing than all other timesharing systems put together.

Unix has found use on a wider variety of machines than any other operating system can claim. From supercomputers to handhelds and embedded networking hardware, through workstations and servers and PCs and minicomputers, Unix has probably seen more architectures and more odd hardware than any three other operating systems combined.

Unix has supported a mind-bogglingly wide spectrum of uses. No other operating system has shone simultaneously as a research vehicle, a friendly host for technical custom applications, a platform for commercial-off-the-shelf business software, and a vital component technology of the Internet.

Confident predictions that Unix would wither away, or be crowded out by other operating systems, have been made yearly since its infancy. And yet Unix, in its present-day avatars as Linux and BSD and Solaris and MacOS X and half a dozen other variants, seems stronger than ever today.

Robert Metcalf [the inventor of Ethernet] says that if something comes along to replace Ethernet, it will be called "Ethernet", so therefore Ethernet will never die.¹ Unix has already undergone several such transformations.

—Ken Thompson

1. In fact, Ethernet has already been replaced by a different technology with the same name—twice. Once when coax was replaced with twisted pair, and a second time when gigabit Ethernet came in.

At least one of Unix's central technologies—the C language—has been widely naturalized elsewhere. Indeed it is now hard to imagine doing software engineering without C as a ubiquitous common language of systems programming. Unix also introduced both the now-ubiquitous tree-shaped file namespace with directory nodes and the pipeline for connecting programs.

Unix's durability and adaptability have been nothing short of astonishing. Other technologies have come and gone like mayflies. Machines have increased a thousand-fold in power, languages have mutated, industry practice has gone through multiple revolutions—and Unix hangs in there, still producing, still paying the bills, and still commanding loyalty from many of the best and brightest software technologists on the planet.

One of the many consequences of the exponential power-versus-time curve in computing, and the corresponding pace of software development, is that 50% of what one knows becomes obsolete over every 18 months. Unix does not abolish this phenomenon, but does do a good job of containing it. There's a bedrock of unchanging basics—languages, system calls, and tool invocations—that one can actually keep using for years, even decades. Elsewhere it is impossible to predict what will be stable; even entire operating systems cycle out of use. Under Unix, there is a fairly sharp distinction between transient knowledge and lasting knowledge, and one can know ahead of time (with about 90% certainty) which category something is likely to fall in when one learns it. Thus the loyalty Unix commands.

Much of Unix's stability and success has to be attributed to its inherent strengths, to design decisions Ken Thompson, Dennis Ritchie, Brian Kernighan, Doug McIlroy, Rob Pike and other early Unix developers made back at the beginning; decisions that have been proven sound over and over. But just as much is due to the design philosophy, art of programming, and technical culture that grew up around Unix in the early days. This tradition has continuously and successfully propagated itself in symbiosis with Unix ever since.

1.3 The Case against Learning Unix Culture

Unix's durability and its technical culture are certainly of interest to people who already like Unix, and perhaps to historians of technology. But Unix's original application as a general-purpose timesharing system for mid-sized and larger computers is rapidly receding into the mists of history, killed off by personal workstations. And there is certainly room for doubt that it will ever achieve success in the mainstream business-desktop market now dominated by Microsoft.

Outsiders have frequently dismissed Unix as an academic toy or a hacker's sandbox. One well-known polemic, the *Unix Hater's Handbook* [Garfinkel], follows an antagonistic line nearly as old as Unix itself in writing its devotees off as a cult religion of

freaks and losers. Certainly the colossal and repeated blunders of AT&T, Sun, Novell, and other commercial vendors and standards consortia in mispositioning and mis-marketing Unix have become legendary.

Even from within the Unix world, Unix has seemed to be teetering on the brink of universality for so long as to raise the suspicion that it will never actually get there. A skeptical outside observer's conclusion might be that Unix is too useful to die but too awkward to break out of the back room; a perpetual niche operating system.

What confounds the skeptics' case is, more than anything else, the rise of Linux and other open-source Unixes (such as the modern BSD variants). Unix's culture proved too vital to be smothered even by a decade of vendor mismanagement. Today the Unix community itself has taken control of the technology and marketing, and is rapidly and visibly solving Unix's problems (in ways we'll examine in more detail in Chapter 20).

1.4 What Unix Gets Wrong

For a design that dates from 1969, it is remarkably difficult to identify design choices in Unix that are unequivocally wrong. There are several popular candidates, but each is still a subject of spirited debate not merely among Unix fans but across the wider community of people who think about and design operating systems.

Unix files have no structure above byte level. File deletion is irrevocable. The Unix security model is arguably too primitive. Job control is botched. There are too many different kinds of names for things. Having a file system at all may have been the wrong choice. We will discuss these technical issues in Chapter 20.

But perhaps the most enduring objections to Unix are consequences of a feature of its philosophy first made explicit by the designers of the X windowing system. X strives to provide "mechanism, not policy", supporting an extremely general set of graphics operations and deferring decisions about toolkits and interface look-and-feel (the policy) up to application level. Unix's other system-level services display similar tendencies; final choices about behavior are pushed as far toward the user as possible. Unix users can choose among multiple shells. Unix programs normally provide many behavior options and sport elaborate preference facilities.

This tendency reflects Unix's heritage as an operating system designed primarily for technical users, and a consequent belief that users know better than operating-system designers what their own needs are.

This tenet was firmly established at Bell Labs by Dick Hamming² who insisted in the 1950s when computers were rare and expensive, that open-shop computing, where customers wrote their own programs, was imperative, because “it is better to solve the right problem the wrong way than the wrong problem the right way”.

—Doug McIlroy

But the cost of the mechanism-not-policy approach is that when the user *can* set policy, the user *must* set policy. Nontechnical end-users frequently find Unix’s profusion of options and interface styles overwhelming and retreat to systems that at least pretend to offer them simplicity.

In the short term, Unix’s laissez-faire approach may lose it a good many nontechnical users. In the long term, however, it may turn out that this ‘mistake’ confers a critical advantage—because policy tends to have a short lifetime, mechanism a long one. Today’s fashion in interface look-and-feel too often becomes tomorrow’s evolutionary dead end (as people using obsolete X toolkits will tell you with some feeling!). So the flip side of the flip side is that the “mechanism, not policy” philosophy may enable Unix to renew its relevance long after competitors more tied to one set of policy or interface choices have faded from view.³

1.5 What Unix Gets Right

The explosive recent growth of Linux, and the increasing importance of the Internet, give us good reasons to suppose that the skeptics’ case is wrong. But even supposing the skeptical assessment is true, Unix culture is worth learning because there are some things that Unix and its surrounding culture clearly do better than any competitors.

1.5.1 Open-Source Software

Though the term “open source” and the Open Source Definition were not invented until 1998, peer-review-intensive development of freely shared source code was a key feature of the Unix culture from its beginnings.

2. Yes, the Hamming of ‘Hamming distance’ and ‘Hamming code’.

3. Jim Gettys, one of the architects of X (and a contributor to this book), has meditated in depth on how X’s laissez-faire style might be productively carried forward in *The Two-Edged Sword* [Gettys]. This essay is well worth reading, both for its specific proposals and for its expression of the Unix mindset.

For its first ten years AT&T's original Unix, and its primary variant Berkeley Unix, were normally distributed with source code. This enabled most of the other good things that follow here.

1.5.2 Cross-Platform Portability and Open Standards

Unix is still the only operating system that can present a consistent, documented application programming interface (API) across a heterogeneous mix of computers, vendors, and special-purpose hardware. It is the only operating system that can scale from embedded chips and handhelds, up through desktop machines, through servers, and all the way to special-purpose number-crunching behemoths and database back ends.

The Unix API is the closest thing to a hardware-independent standard for writing truly portable software that exists. It is no accident that what the IEEE originally called the *Portable Operating System Standard* quickly got a suffix added to its acronym and became POSIX. A Unix-equivalent API was the only credible model for such a standard.

Binary-only applications for other operating systems die with their birth environments, but Unix sources are forever. Forever, at least, given a Unix technical culture that polishes and maintains them across decades.

1.5.3 The Internet and the World Wide Web

The Defense Department's contract for the first production TCP/IP stack went to a Unix development group because the Unix in question was largely open source. Besides TCP/IP, Unix has become the one indispensable core technology of the Internet Service Provider industry. Ever since the demise of the TOPS family of operating systems in the mid-1980s, most Internet server machines (and effectively all above the PC level) have relied on Unix.

Not even Microsoft's awesome marketing clout has been able to dent Unix's lock on the Internet. While the TCP/IP standards (on which the Internet is based) evolved under TOPS-10 and are theoretically separable from Unix, attempts to make them work on other operating systems have been bedeviled by incompatibilities, instabilities, and bugs. The theory and specifications are available to anyone, but the engineering tradition to make them into a solid and working reality exists only in the Unix world.⁴

4. Other operating systems have generally copied or cloned Unix TCP/IP implementations. It is their loss that they have not generally adopted the robust tradition of peer review that goes with it, exemplified by documents like RFC 1025 (*TCP and IP Bake Off*).

The Internet technical culture and the Unix culture began to merge in the early 1980s, and are now inseparably symbiotic. The design of the World Wide Web, the modern face of the Internet, owes as much to Unix as it does to the ancestral ARPANET. In particular, the concept of the Uniform Resource Locator (URL) so central to the Web is a generalization of the Unix idea of one uniform file namespace everywhere. To function effectively as an Internet expert, an understanding of Unix and its culture are indispensable.

1.5.4 The Open-Source Community

The community that originally formed around the early Unix source distributions never went away—after the great Internet explosion of the early 1990s, it recruited an entire new generation of eager hackers on home machines.

Today, that community is a powerful support group for all kinds of software development. High-quality open-source development tools abound in the Unix world (we'll examine many in this book). Open-source Unix applications are usually equal to, and are often superior to, their proprietary equivalents [Fuzz]. Entire Unix operating systems, with complete toolkits and basic applications suites, are available for free over the Internet. Why code from scratch when you can adapt, reuse, recycle, and save yourself 90% of the work?

This tradition of code-sharing depends heavily on hard-won expertise about how to make programs cooperative and reusable. And not by abstract theory, but through a lot of engineering practice—unobvious design rules that allow programs to function not just as isolated one-shot solutions but as synergistic parts of a toolkit. A major purpose of this book is to elucidate those rules.

Today, a burgeoning open-source movement is bringing new vitality, new technical approaches, and an entire generation of bright young programmers into the Unix tradition. Open-source projects including the Linux operating system and symbionts such as Apache and Mozilla have brought the Unix tradition an unprecedented level of mainstream visibility and success. The open-source movement seems on the verge of winning its bid to define the computing infrastructure of tomorrow—and the core of that infrastructure will be Unix machines running on the Internet.

1.5.5 Flexibility All the Way Down

Many operating systems touted as more 'modern' or 'user friendly' than Unix achieve their surface glossiness by locking users and developers into one interface policy, and offer an application-programming interface that for all its elaborateness is rather narrow and rigid. On such systems, tasks the designers have anticipated are very

easy—but tasks they have not anticipated are often impossible or at best extremely painful.

Unix, on the other hand, has flexibility in depth. The many ways Unix provides to glue together programs mean that components of its basic toolkit can be combined to produce useful effects that the designers of the individual toolkit parts never anticipated.

Unix’s support of multiple styles of program interface (often seen as a weakness because it increases the perceived complexity of the system to end users) also contributes to flexibility; no program that wants to be a simple piece of data plumbing is forced to carry the complexity overhead of an elaborate GUI.

Unix tradition lays heavy emphasis on keeping programming interfaces relatively small, clean, and orthogonal—another trait that produces flexibility in depth. Throughout a Unix system, easy things are easy and hard things are at least possible.

1.5.6 Unix Is Fun to Hack

People who pontificate about Unix’s technical superiority often don’t mention what may ultimately be its most important strength, the one that underlies all its successes. Unix is fun to hack.

Unix boosters seem almost ashamed to acknowledge this sometimes, as though admitting they’re having fun might damage their legitimacy somehow. But it’s true; Unix is fun to play with and develop for, and always has been.

There are not many operating systems that anyone has ever described as ‘fun’. Indeed, the friction and labor of development under most other environments has been aptly compared to kicking a dead whale down the beach.⁵ The kindest adjectives one normally hears are on the order of “tolerable” or “not too painful”. In the Unix world, by contrast, the operating system rewards effort rather than frustrating it. People programming under Unix usually come to see it not as an adversary to be clubbed into doing one’s bidding by main effort but rather as an actual positive help.

This has real economic significance. The fun factor started a virtuous circle early in Unix’s history. People liked Unix, so they built more programs for it that made it nicer to use. Today people build entire, production-quality open-source Unix systems as a hobby. To understand how remarkable this is, ask yourself when you last heard of anybody cloning OS/360 or VAX/VMS or Microsoft Windows for fun.

The ‘fun’ factor is not trivial from a design point of view, either. The kind of people who become programmers and developers have ‘fun’ when the effort they have to put out to do a task challenges them, but is just within their capabilities. ‘Fun’ is therefore

5. This was originally said of the IBM MVS TSO facility by Stephen C. Johnson, perhaps better known as the author of *yacc*.

a sign of peak efficiency. Painful development environments waste labor and creativity; they extract huge hidden costs in time, money, and opportunity.

If Unix were a failure in every other way, the Unix engineering culture would be worth studying for the ways it keeps the fun in development—because that fun is a sign that it makes developers efficient, effective, and productive.

1.5.7 The Lessons of Unix Can Be Applied Elsewhere

Unix programmers have accumulated decades of experience while pioneering operating-system features we now take for granted. Even non-Unix programmers can benefit from studying that Unix experience. Because Unix makes it relatively easy to apply good design principles and development methods, it is an excellent place to learn them.

Other operating systems generally make good practice rather more difficult, but even so some of the Unix culture's lessons can transfer. Much Unix code (including all its filters, its major scripting languages, and many of its code generators) will port directly to any operating system supporting ANSI C (for the excellent reason that C itself was a Unix invention and the ANSI C library embodies a substantial chunk of Unix's services!).

1.6 Basics of the Unix Philosophy

The 'Unix philosophy' originated with Ken Thompson's early meditations on how to design a small but capable operating system with a clean service interface. It grew as the Unix culture learned things about how to get maximum leverage out of Thompson's design. It absorbed lessons from many sources along the way.

The Unix philosophy is not a formal design method. It wasn't handed down from the high fastnesses of theoretical computer science as a way to produce theoretically perfect software. Nor is it that perennial executive's mirage, some way to magically extract innovative but reliable software on too short a deadline from unmotivated, badly managed, and underpaid programmers.

The Unix philosophy (like successful folk traditions in other engineering disciplines) is bottom-up, not top-down. It is pragmatic and grounded in experience. It is not to be found in official methods and standards, but rather in the implicit half-reflexive knowledge, the *expertise* that the Unix culture transmits. It encourages a sense of proportion and skepticism—and shows both by having a sense of (often subversive) humor.

Doug McIlroy, the inventor of Unix pipes and one of the founders of the Unix tradition, had this to say at the time [McIlroy78]:

- (i) Make each program do one thing well. To do a new job, build afresh rather than complicate old programs by adding new features.
- (ii) Expect the output of every program to become the input to another, as yet unknown, program. Don't clutter output with extraneous information. Avoid stringently columnar or binary input formats. Don't insist on interactive input.
- (iii) Design and build software, even operating systems, to be tried early, ideally within weeks. Don't hesitate to throw away the clumsy parts and rebuild them.
- (iv) Use tools in preference to unskilled help to lighten a programming task, even if you have to detour to build the tools and expect to throw some of them out after you've finished using them.

He later summarized it this way (quoted in *A Quarter Century of Unix* [Salus]):

This is the Unix philosophy: Write programs that do one thing and do it well. Write programs to work together. Write programs to handle text streams, because that is a universal interface.

Rob Pike, who became one of the great masters of C, offers a slightly different angle in *Notes on C Programming* [Pike]:

Rule 1. You can't tell where a program is going to spend its time. Bottlenecks occur in surprising places, so don't try to second guess and put in a speed hack until you've proven that's where the bottleneck is.

Rule 2. Measure. Don't tune for speed until you've measured, and even then don't unless one part of the code overwhelms the rest.

Rule 3. Fancy algorithms are slow when n is small, and n is usually small. Fancy algorithms have big constants. Until you know that n is frequently going to be big, don't get fancy. (Even if n does get big, use Rule 2 first.)

Rule 4. Fancy algorithms are buggier than simple ones, and they're much harder to implement. Use simple algorithms as well as simple data structures.

Rule 5. Data dominates. If you've chosen the right data structures and organized things well, the algorithms will almost always be self-evident. Data structures, not algorithms, are central to programming.⁶

6. Pike's original adds "(See Brooks p. 102.)" here. The reference is to an early edition of *The Mythical Man-Month* [Brooks]; the quote is "Show me your flow charts and conceal your tables and I shall continue to be mystified, show me your tables and I won't usually need your flow charts; they'll be obvious".

Rule 6. There is no Rule 6.

Ken Thompson, the man who designed and implemented the first Unix, reinforced Pike's rule 4 with a gnomic maxim worthy of a Zen patriarch:

When in doubt, use brute force.

More of the Unix philosophy was implied not by what these elders said but by what they did and the example Unix itself set. Looking at the whole, we can abstract the following ideas:

1. Rule of Modularity: Write simple parts connected by clean interfaces.
2. Rule of Clarity: Clarity is better than cleverness.
3. Rule of Composition: Design programs to be connected to other programs.
4. Rule of Separation: Separate policy from mechanism; separate interfaces from engines.
5. Rule of Simplicity: Design for simplicity; add complexity only where you must.
6. Rule of Parsimony: Write a big program only when it is clear by demonstration that nothing else will do.
7. Rule of Transparency: Design for visibility to make inspection and debugging easier.
8. Rule of Robustness: Robustness is the child of transparency and simplicity.
9. Rule of Representation: Fold knowledge into data so program logic can be stupid and robust.
10. Rule of Least Surprise: In interface design, always do the least surprising thing.
11. Rule of Silence: When a program has nothing surprising to say, it should say nothing.
12. Rule of Repair: When you must fail, fail noisily and as soon as possible.
13. Rule of Economy: Programmer time is expensive; conserve it in preference to machine time.
14. Rule of Generation: Avoid hand-hacking; write programs to write programs when you can.

15. Rule of Optimization: Prototype before polishing. Get it working before you optimize it.
16. Rule of Diversity: Distrust all claims for “one true way”.
17. Rule of Extensibility: Design for the future, because it will be here sooner than you think.

If you're new to Unix, these principles are worth some meditation. Software-engineering texts recommend most of them; but most other operating systems lack the right tools and traditions to turn them into practice, so most programmers can't apply them with any consistency. They come to accept blunt tools, bad designs, overwork, and bloated code as normal—and then wonder what Unix fans are so annoyed about.

1.6.1 Rule of Modularity: Write simple parts connected by clean interfaces.

As Brian Kernighan once observed, “Controlling complexity is the essence of computer programming” [Kernighan-Plauger]. Debugging dominates development time, and getting a working system out the door is usually less a result of brilliant design than it is of managing not to trip over your own feet too many times.

Assemblers, compilers, flowcharting, procedural programming, structured programming, “artificial intelligence”, fourth-generation languages, object orientation, and software-development methodologies without number have been touted and sold as a cure for this problem. All have failed as cures, if only because they ‘succeeded’ by escalating the normal level of program complexity to the point where (once again) human brains could barely cope. As Fred Brooks famously observed [Brooks], there is no silver bullet.

The only way to write complex software that won't fall on its face is to hold its global complexity down—to build it out of simple parts connected by well-defined interfaces, so that most problems are local and you can have some hope of upgrading a part without breaking the whole.

1.6.2 Rule of Clarity: Clarity is better than cleverness.

Because maintenance is so important and so expensive, write programs as if the most important communication they do is not to the computer that executes them but to the human beings who will read and maintain the source code in the future (including yourself).

In the Unix tradition, the implications of this advice go beyond just commenting your code. Good Unix practice also embraces choosing your algorithms and imple-

mentations for future maintainability. Buying a small increase in performance with a large increase in the complexity and obscurity of your technique is a bad trade—not merely because complex code is more likely to harbor bugs, but also because complex code will be harder to read for future maintainers.

Code that is graceful and clear, on the other hand, is less likely to break—and more likely to be instantly comprehended by the next person to have to change it. This is important, especially when that next person might be yourself some years down the road.

Never struggle to decipher subtle code three times. Once might be a one-shot fluke, but if you find yourself having to figure it out a second time—because the first was too long ago and you’ve forgotten details—it is time to comment the code so that the third time will be relatively painless.

—Henry Spencer

1.6.3 Rule of Composition: Design programs to be connected with other programs.

It’s hard to avoid programming overcomplicated monoliths if none of your programs can talk to each other.

Unix tradition strongly encourages writing programs that read and write simple, textual, stream-oriented, device-independent formats. Under classic Unix, as many programs as possible are written as simple *filters*, which take a simple text stream on input and process it into another simple text stream on output.

Despite popular mythology, this practice is favored not because Unix programmers hate graphical user interfaces. It’s because if you don’t write programs that accept and emit simple text streams, it’s much more difficult to hook the programs together.

Text streams are to Unix tools as messages are to objects in an object-oriented setting. The simplicity of the text-stream interface enforces the encapsulation of the tools. More elaborate forms of inter-process communication, such as remote procedure calls, show a tendency to involve programs with each others’ internals too much.

To make programs composable, make them independent. A program on one end of a text stream should care as little as possible about the program on the other end. It should be made easy to replace one end with a completely different implementation without disturbing the other.

GUIs can be a very good thing. Complex binary data formats are sometimes unavoidable by any reasonable means. But before writing a GUI, it’s wise to ask if the tricky interactive parts of your program can be segregated into one piece and the workhorse algorithms into another, with a simple command stream or application protocol connecting the two. Before devising a tricky binary format to pass data

around, it's worth experimenting to see if you can make a simple textual format work and accept a little parsing overhead in return for being able to hack the data stream with general-purpose tools.

When a serialized, protocol-like interface is not natural for the application, proper Unix design is to at least organize as many of the application primitives as possible into a library with a well-defined API. This opens up the possibility that the application can be called by linkage, or that multiple interfaces can be glued on it for different tasks.

(We discuss these issues in detail in Chapter 7.)

1.6.4 Rule of Separation: Separate policy from mechanism; separate interfaces from engines.

In our discussion of what Unix gets wrong, we observed that the designers of X made a basic decision to implement “mechanism, not policy”—to make X a generic graphics engine and leave decisions about user-interface style to toolkits and other levels of the system. We justified this by pointing out that policy and mechanism tend to mutate on different timescales, with policy changing much faster than mechanism. Fashions in the look and feel of GUI toolkits may come and go, but raster operations and compositing are forever.

Thus, hardwiring policy and mechanism together has two bad effects: It makes policy rigid and harder to change in response to user requirements, and it means that trying to change policy has a strong tendency to destabilize the mechanisms.

On the other hand, by separating the two we make it possible to experiment with new policy without breaking mechanisms. We also make it much easier to write good tests for the mechanism (policy, because it ages so quickly, often does not justify the investment).

This design rule has wide application outside the GUI context. In general, it implies that we should look for ways to separate interfaces from engines.

One way to effect that separation is, for example, to write your application as a library of C service routines that are driven by an embedded scripting language, with the application flow of control written in the scripting language rather than C. A classic example of this pattern is the *Emacs* editor, which uses an embedded Lisp interpreter to control editing primitives written in C. We discuss this style of design in Chapter 11.

Another way is to separate your application into cooperating front-end and back-end processes communicating through a specialized application protocol over sockets; we discuss this kind of design in Chapter 5 and Chapter 7. The front end implements policy; the back end, mechanism. The global complexity of the pair will often be far

lower than that of a single-process monolith implementing the same functions, reducing your vulnerability to bugs and lowering life-cycle costs.

1.6.5 Rule of Simplicity: Design for simplicity; add complexity only where you must.

Many pressures tend to make programs more complicated (and therefore more expensive and buggy). One such pressure is technical machismo. Programmers are bright people who are (often justly) proud of their ability to handle complexity and juggle abstractions. Often they compete with their peers to see who can build the most intricate and beautiful complexities. Just as often, their ability to design outstrips their ability to implement and debug, and the result is expensive failure.

The notion of “intricate and beautiful complexities” is almost an oxymoron. Unix programmers vie with each other for “simple and beautiful” honors — a point that’s implicit in these rules, but is well worth making overt.

—Doug McIlroy

Even more often (at least in the commercial software world) excessive complexity comes from project requirements that are based on the marketing fad of the month rather than the reality of what customers want or software can actually deliver. Many a good design has been smothered under marketing’s pile of “checklist features” — features that, often, no customer will ever use. And a vicious circle operates; the competition thinks it has to compete with chrome by adding more chrome. Pretty soon, massive bloat is the industry standard and everyone is using huge, buggy programs not even their developers can love.

Either way, everybody loses in the end.

The only way to avoid these traps is to encourage a software culture that knows that small is beautiful, that actively resists bloat and complexity: an engineering tradition that puts a high value on simple solutions, that looks for ways to break program systems up into small cooperating pieces, and that reflexively fights attempts to gussy up programs with a lot of chrome (or, even worse, to design programs *around* the chrome).

That would be a culture a lot like Unix’s.

1.6.6 Rule of Parsimony: Write a big program only when it is clear by demonstration that nothing else will do.

‘Big’ here has the sense both of large in volume of code and of internal complexity. Allowing programs to get large hurts maintainability. Because people are reluctant to throw away the visible product of lots of work, large programs invite overinvestment in approaches that are failed or suboptimal.

(We’ll examine the issue of the right size of software in more detail in Chapter 13.)

1.6.7 Rule of Transparency: Design for visibility to make inspection and debugging easier.

Because debugging often occupies three-quarters or more of development time, work done early to ease debugging can be a very good investment. A particularly effective way to ease debugging is to design for *transparency* and *discoverability*.

A software system is *transparent* when you can look at it and immediately understand what it is doing and how. It is *discoverable* when it has facilities for monitoring and display of internal state so that your program not only functions well but can be *seen* to function well.

Designing for these qualities will have implications throughout a project. At minimum, it implies that debugging options should not be minimal afterthoughts. Rather, they should be designed in from the beginning—from the point of view that the program should be able to both demonstrate its own correctness and communicate to future developers the original developer’s mental model of the problem it solves.

For a program to demonstrate its own correctness, it needs to be using input and output formats sufficiently simple so that the proper relationship between valid input and correct output is easy to check.

The objective of designing for transparency and discoverability should also encourage simple interfaces that can easily be manipulated by other programs—in particular, test and monitoring harnesses and debugging scripts.

1.6.8 Rule of Robustness: Robustness is the child of transparency and simplicity.

Software is said to be *robust* when it performs well under unexpected conditions which stress the designer’s assumptions, as well as under normal conditions.

Most software is fragile and buggy because most programs are too complicated for a human brain to understand all at once. When you can’t reason correctly about the guts of a program, you can’t be sure it’s correct, and you can’t fix it if it’s broken.

It follows that the way to make robust programs is to make their internals easy for human beings to reason about. There are two main ways to do that: transparency and simplicity.

For robustness, designing in tolerance for unusual or extremely bulky inputs is also important. Bearing in mind the Rule of Composition helps; input generated by other programs is notorious for stress-testing software (e.g., the original Unix C compiler reportedly needed small upgrades to cope well with Yacc output). The forms involved often seem useless to humans. For example, accepting empty lists/strings/etc., even in places where a human would seldom or never supply an empty string, avoids having to special-case such situations when generating the input mechanically.

—Henry Spencer

One very important tactic for being robust under odd inputs is to avoid having special cases in your code. Bugs often lurk in the code for handling special cases, and in the interactions among parts of the code intended to handle different special cases.

We observed above that software is *transparent* when you can look at it and immediately see what is going on. It is *simple* when what is going on is uncomplicated enough for a human brain to reason about all the potential cases without strain. The more your programs have both of these qualities, the more robust they will be.

Modularity (simple parts, clean interfaces) is a way to organize programs to make them simpler. There are other ways to fight for simplicity. Here's another one.

1.6.9 Rule of Representation: Fold knowledge into data, so program logic can be stupid and robust.

Even the simplest procedural logic is hard for humans to verify, but quite complex data structures are fairly easy to model and reason about. To see this, compare the expressiveness and explanatory power of a diagram of (say) a fifty-node pointer tree with a flowchart of a fifty-line program. Or, compare an array initializer expressing a conversion table with an equivalent switch statement. The difference in transparency and clarity is dramatic. See Rob Pike's Rule 5.

Data is more tractable than program logic. It follows that where you see a choice between complexity in data structures and complexity in code, choose the former. More: in evolving a design, you should actively seek ways to shift complexity from code to data.

The Unix community did not originate this insight, but a lot of Unix code displays its influence. The C language's facility at manipulating pointers, in particular, has encouraged the use of dynamically-modified reference structures at all levels of coding from the kernel upward. Simple pointer chases in such structures frequently do duties

that implementations in other languages would instead have to embody in more elaborate procedures.

(We also cover these techniques in Chapter 9.)

1.6.10 Rule of Least Surprise: In interface design, always do the least surprising thing.

(This is also widely known as the Principle of Least Astonishment.)

The easiest programs to use are those that demand the least new learning from the user—or, to put it another way, the easiest programs to use are those that most effectively connect to the user’s pre-existing knowledge.

Therefore, avoid gratuitous novelty and excessive cleverness in interface design. If you’re writing a calculator program, ‘+’ should always mean addition! When designing an interface, model it on the interfaces of functionally similar or analogous programs with which your users are likely to be familiar.

Pay attention to your expected audience. They may be end users, they may be other programmers, or they may be system administrators. What is least surprising can differ among these groups.

Pay attention to tradition. The Unix world has rather well-developed conventions about things like the format of configuration and run-control files, command-line switches, and the like. These traditions exist for a good reason: to tame the learning curve. Learn and use them.

(We’ll cover many of these traditions in Chapter 5 and Chapter 10.)

The flip side of the Rule of Least Surprise is to avoid making things superficially similar but really a little bit different. This is extremely treacherous because the seeming familiarity raises false expectations. It’s often better to make things distinctly different than to make them *almost* the same.

—Henry Spencer

1.6.11 Rule of Silence: When a program has nothing surprising to say, it should say nothing.

One of Unix’s oldest and most persistent design rules is that when a program has nothing interesting or surprising to say, it should *shut up*. Well-behaved Unix programs do their jobs unobtrusively, with a minimum of fuss and bother. Silence is golden.

This “silence is golden” rule evolved originally because Unix predates video displays. On the slow printing terminals of 1969, each line of unnecessary output was a serious drain on the user’s time. That constraint is gone, but excellent reasons for terseness remain.

I think that the terseness of Unix programs is a central feature of the style. When your program's output becomes another's input, it should be easy to pick out the needed bits. And for people it is a human-factors necessity—important information should not be mixed in with verbosity about internal program behavior. If all displayed information is important, important information is easy to find.

—Ken Arnold

Well-designed programs treat the user's attention and concentration as a precious and limited resource, only to be claimed when necessary.

(We'll discuss the Rule of Silence and the reasons for it in more detail at the end of Chapter 11.)

1.6.12 Rule of Repair: Repair what you can—but when you must fail, fail noisily and as soon as possible.

Software should be transparent in the way that it fails, as well as in normal operation. It's best when software can cope with unexpected conditions by adapting to them, but the worst kinds of bugs are those in which the repair doesn't succeed and the problem quietly causes corruption that doesn't show up until much later.

Therefore, write your software to cope with incorrect inputs and its own execution errors as gracefully as possible. But when it cannot, make it fail in a way that makes diagnosis of the problem as easy as possible.

Consider also Postel's Prescription:⁷ “Be liberal in what you accept, and conservative in what you send”. Postel was speaking of network service programs, but the underlying idea is more general. Well-designed programs cooperate with other programs by making as much sense as they can from ill-formed inputs; they either fail noisily or pass strictly clean and correct data to the next program in the chain.

However, heed also this warning:

The original HTML documents recommended “be generous in what you accept”, and it has bedeviled us ever since because each browser accepts a different superset of the specifications. It is the *specifications* that should be generous, not their interpretation.

—Doug McIlroy

7. Jonathan Postel was the first editor of the Internet RFC series of standards, and one of the principal architects of the Internet. A tribute page <<http://www.postel.org/postel.html>> is maintained by the Postel Center for Experimental Networking.

McIlroy adjures us to *design* for generosity rather than compensating for inadequate standards with permissive implementations. Otherwise, as he rightly points out, it's all too easy to end up in tag soup.

1.6.13 Rule of Economy: Programmer time is expensive; conserve it in preference to machine time.

In the early minicomputer days of Unix, this was still a fairly radical idea (machines were a great deal slower and more expensive then). Nowadays, with every development shop and most users (apart from the few modeling nuclear explosions or doing 3D movie animation) awash in cheap machine cycles, it may seem too obvious to need saying.

Somehow, though, practice doesn't seem to have quite caught up with reality. If we took this maxim really seriously throughout software development, most applications would be written in higher-level languages like Perl, Tcl, Python, Java, Lisp and even shell—languages that ease the programmer's burden by doing their own memory management (see [Ravenbrook]).

And indeed this is happening within the Unix world, though outside it most applications shops still seem stuck with the old-school Unix strategy of coding in C (or C++). Later in this book we'll discuss this strategy and its tradeoffs in detail.

One other obvious way to conserve programmer time is to teach machines how to do more of the low-level work of programming. This leads to...

1.6.14 Rule of Generation: Avoid hand-hacking; write programs to write programs when you can.

Human beings are notoriously bad at sweating the details. Accordingly, any kind of hand-hacking of programs is a rich source of delays and errors. The simpler and more abstracted your program specification can be, the more likely it is that the human designer will have gotten it right. Generated code (at *every* level) is almost always cheaper and more reliable than hand-hacked.

We all know this is true (it's why we have compilers and interpreters, after all) but we often don't think about the implications. High-level-language code that's repetitive and mind-numbing for humans to write is just as productive a target for a code generator as machine code. It pays to use code generators when they can raise the level of abstraction—that is, when the specification language for the generator is simpler than the generated code, and the code doesn't have to be hand-hacked afterwards.

In the Unix tradition, code generators are heavily used to automate error-prone detail work. Parser/lexer generators are the classic examples; makefile generators and GUI interface builders are newer ones.

(We cover these techniques in Chapter 9.)

1.6.15 Rule of Optimization: Prototype before polishing. Get it working before you optimize it.

The most basic argument for prototyping first is Kernighan & Plauger's; "90% of the functionality delivered now is better than 100% of it delivered never". Prototyping first may help keep you from investing far too much time for marginal gains.

For slightly different reasons, Donald Knuth (author of *The Art of Computer Programming*, one of the field's few true classics) popularized the observation that "Premature optimization is the root of all evil".⁸ And he was right.

Rushing to optimize before the bottlenecks are known may be the only error to have ruined more designs than feature creep. From tortured code to incomprehensible data layouts, the results of obsessing about speed or memory or disk usage at the expense of transparency and simplicity are everywhere. They spawn innumerable bugs and cost millions of man-hours—often, just to get marginal gains in the use of some resource much less expensive than debugging time.

Disturbingly often, premature local optimization actually hinders global optimization (and hence reduces overall performance). A prematurely optimized portion of a design frequently interferes with changes that would have much higher payoffs across the whole design, so you end up with both inferior performance and excessively complex code.

In the Unix world there is a long-established and very explicit tradition (exemplified by Rob Pike's comments above and Ken Thompson's maxim about brute force) that says: *Prototype, then polish. Get it working before you optimize it.* Or: Make it work first, then make it work fast. 'Extreme programming' guru Kent Beck, operating in a different culture, has usefully amplified this to: "Make it run, then make it right, then make it fast".

The thrust of all these quotes is the same: get your design right with an un-optimized, slow, memory-intensive implementation before you try to tune. Then, tune systematically, looking for the places where you can buy big performance wins with the smallest possible increases in local complexity.

8. In full: "We should forget about small efficiencies, say about 97% of the time: premature optimization is the root of all evil". Knuth himself attributes the remark to C. A. R. Hoare.

Prototyping is important for system design as well as optimization—it is much easier to judge whether a prototype does what you want than it is to read a long specification. I remember one development manager at Bellcore who fought against the “requirements” culture years before anybody talked about “rapid prototyping” or “agile development”. He wouldn’t issue long specifications; he’d lash together some combination of shell scripts and awk code that did roughly what was needed, tell the customers to send him some clerks for a few days, and then have the customers come in and look at their clerks using the prototype and tell him whether or not they liked it. If they did, he would say “you can have it industrial strength so-many-months from now at such-and-such cost”. His estimates tended to be accurate, but he lost out in the culture to managers who believed that requirements writers should be in control of everything.

—Mike Lesk

Using prototyping to learn which features you don’t have to implement helps optimization for performance; you don’t have to optimize what you don’t write. The most powerful optimization tool in existence may be the delete key.

One of my most productive days was throwing away 1000 lines of code.

—Ken Thompson

(We’ll go into a bit more depth about related ideas in Chapter 12.)

1.6.16 Rule of Diversity: Distrust all claims for “one true way”.

Even the best software tools tend to be limited by the imaginations of their designers. Nobody is smart enough to optimize for everything, nor to anticipate all the uses to which their software might be put. Designing rigid, closed software that won’t talk to the rest of the world is an unhealthy form of arrogance.

Therefore, the Unix tradition includes a healthy mistrust of “one true way” approaches to software design or implementation. It embraces multiple languages, open extensible systems, and customization hooks everywhere.

1.6.17 Rule of Extensibility: Design for the future, because it will be here sooner than you think.

If it is unwise to trust other people’s claims for “one true way”, it’s even more foolish to believe them about your own designs. Never assume you have the final answer.

Therefore, leave room for your data formats and code to grow; otherwise, you will often find that you are locked into unwise early choices because you cannot change them while maintaining backward compatibility.

When you design protocols or file formats, make them sufficiently self-describing to be extensible. Always, *always* either include a version number, or compose the format from self-contained, self-describing clauses in such a way that new clauses can be readily added and old ones dropped without confusing format-reading code. Unix experience tells us that the marginal extra overhead of making data layouts self-describing is paid back a thousandfold by the ability to evolve them forward without breaking things.

When you design code, organize it so future developers will be able to plug new functions into the architecture without having to scrap and rebuild the architecture. This rule is not a license to add features you don't yet need; it's advice to write your code so that adding features later when you *do* need them is easy. Make the joints flexible, and put "If you ever need to..." comments in your code. You owe this grace to people who will use and maintain your code after you.

You'll be there in the future too, maintaining code you may have half forgotten under the press of more recent projects. When you design for the future, the sanity you save may be your own.

1.7 The Unix Philosophy in One Lesson

All the philosophy really boils down to one iron law, the hallowed 'KISS principle' of master engineers everywhere:



Unix gives you an excellent base for applying the KISS principle. The remainder of this book will help you learn how.

1.8 Applying the Unix Philosophy

These philosophical principles aren't just vague generalities. In the Unix world they come straight from experience and lead to specific prescriptions, some of which we've already developed above. Here's a by no means exhaustive list:

- Everything that can be a source- and destination-independent filter *should* be one.
- Data streams should if at all possible be textual (so they can be viewed and filtered with standard tools).
- Database layouts and application protocols should if at all possible be textual (human-readable and human-editable).
- Complex front ends (user interfaces) should be cleanly separated from complex back ends.
- Whenever possible, prototype in an interpreted language before coding C.
- Mixing languages is better than writing everything in one, if and only if using only that one is likely to overcomplicate the program.
- Be generous in what you accept, rigorous in what you emit.
- When filtering, never throw away information you don't need to.
- Small is beautiful. Write programs that do as little as is consistent with getting the job done.

We'll see the Unix design rules, and the prescriptions that derive from them, applied over and over again in the remainder of this book. Unsurprisingly, they tend to converge with the very best practices from software engineering in other traditions.⁹

1.9 Attitude Matters Too

When you see the right thing, do it—this may look like more work in the short term, but it's the path of least effort in the long run. If you don't know what the right thing

9. One notable example is Butler Lampson's *Hints for Computer System Design* [Lampson], which I discovered late in the preparation of this book. It not only expresses a number of Unix dicta in forms that were clearly discovered independently, but uses many of the same tag lines to illustrate them.

is, do the minimum necessary to get the job done, at least until you figure out what the right thing is.

To do the Unix philosophy right, you have to be loyal to excellence. You have to believe that software design is a craft worth all the intelligence, creativity, and passion you can muster. Otherwise you won't look past the easy, stereotyped ways of approaching design and implementation; you'll rush into coding when you should be thinking. You'll carelessly complicate when you should be relentlessly simplifying—and then you'll wonder why your code bloats and debugging is so hard.

To do the Unix philosophy right, you have to value your own time enough never to waste it. If someone has already solved a problem once, don't let pride or politics suck you into solving it a second time rather than re-using. And never work harder than you have to; work smarter instead, and save the extra effort for when you need it. Lean on your tools and automate everything you can.

Software design and implementation should be a joyous art, a kind of high-level play. If this attitude seems preposterous or vaguely embarrassing to you, stop and think; ask yourself what you've forgotten. Why do you design software instead of doing something else to make money or pass the time? You must have thought software was worthy of your passion once....

To do the Unix philosophy right, you need to have (or recover) that attitude. You need to *care*. You need to *play*. You need to be willing to *explore*.

We hope you'll bring this attitude to the rest of this book. Or, at least, that this book will help you rediscover it.

Index

- a option, 243
- Abstraction, 102–103
- Access Control Lists (ACLs), 471
- Accidental complexity, 299–300
- Ada, 328
- Adams, Scott, 53
- Ad-hoc code generation, 225–229
 - for *ascii* displays, 225–227
 - generating HTML code for a tabular list, 227–229
- Adhocity trap, 297
- ADOS, 65
- Advanced shell programming, 325
- Aegis, 368–369
- AF_INET socket family, 173
- Agile programming, 477–478
- AIX, 77
- Algol, 84, 324
- Allen, Paul, 36
- AmigaOs, 67
- ANSI C, 11
- Apache, 9, 48, 49, 108, 131, 172, 175, 336, 426, 440, 446
- Applets, 339
- Application programming interfaces (APIs), 8, 16, 58, 70, 85, 150, 402
- Application protocol, 173
- Application protocol design, 123–127
 - IMAP, 123, 126–127
 - POP3, 123, 124–126
 - SMTP, 123, 124
- Application protocol metaformats, 127–132
 - CDDB/freedb.org database, 129
 - classical Internet application metaprotocol, 127–128
 - HTTP, 128–129
- Arnold, Ken, 21, 34, 88, 115, 171, 173
- ARPANET, 9, 31, 35–36, 44–45
 - fusion with Unix culture, 45–47
- Art of Computer Programming, The* (Knuth), xxxii, 23
- Artificial intelligence, 14
- Artistic License, 388, 458
- ascii*, 90, 217–218
- ash*, 244
- ASR-33 teletype, 30
- Assemblers, 14
- at*, 275
- AT&T, 32–33, 35, 37, 40–42, 61, 176. *See also*
 - System V Unix
 - Bell Labs, 30–36
 - divestiture, impact of, 6
 - and XENIX, 35
- atd*, 275
- Attitude, 26–27
- audacity*, 135–136, 258, 265
- autoconf*, 363–364, 447–448
- autoheader*, 447
- automake*, 364, 447
- Automated version control, 366
- awk*, 75, 89, 112, 183, 200–202, 244, 330, 354
- b option, 243
- Backup scripts, 166
- bash shell, 71
- bc*, 165, 183, 187, 203–205, 262–264
- BCPL, 395
- Bechtolsheim, Andreas, 36
- Beck, Kent, 23, 360
- BEEP (Blocks Extensible Exchange Protocol), 130–131
- Bell Labs, 33–36, 367, 464
 - netlib code, 196
- Bellovin, Steven M., 330
- BeOS, 71–72, 79
- Berkeley Unix, 35–36
- Berners-Lee, Tim, 299
- Bernstein chaining, 167–168
- bib*, 195
- Big-endian, 98
- BIND, 183

- binhex*, 245
- bison*, 353, 441
- Blaauw, Gerrit, 98
- Blivet trap, 297
- blq*, 333–334
- BNF (Backus-Naur Form), 355
- bogofilter*, 218
- bootpd*, 172
- Bostic, Keith, 41
- Brooks, Fredrick P., 14, 87, 98, 215
- Brooks’s Law, 86
- Browser DOM, 206
- BROWSER environment variable, 240
- BSD Class License, 457–458
- BSD sockets, 39, 89, 90, 176, 182, 400, 463, 464
- BSD Unix, 4, 41, 43, 49, 90, 171, 175–176, 388, 402. *See also* Berkeley Unix
 - open-source variants, 61, 175
 - and reliable signals, 173
 - version 4.2, 35
- bsh*, 244
- Bug tracking, 365
- Byte stream, 58, 106, 172
- Byte-stuffed payload, 127–128
- bzip*, 245, 248
- c option, 244
- C language, 5, 11, 16, 19, 22, 26, 30, 32, 36–37, 45, 70, 84, 88–89, 103, 106, 209, 323–325, 334, 338, 409–410
 - early history of, 395–396
 - evolution of, 394–398
 - fetchmail, 327
 - portability of, 409–410
 - standards, 396–398
 - as “thin glue”, 98–99
- C programming, 12
- C++ language, 22, 37, 56, 89, 98–99, 106, 191, 209, 238, 322–325, 327–300, 334, 337–349, 353, 356–357, 370, 384, 397, 410
 - portability of, 410
 - Qt toolkit, 329–330
- Caldera, 43
- Cantrip pattern, 268
- Case studies, xxx–xxxii
- Casual programming:
 - for MVS, 74
 - and Unix, 60–61
- cat*, 256, 467
- cc*, 269
- cdda2wav*, 276
- CDDB/freedb.org database, 129
- cdrecord*, 276
- cdrtools*, xxx
- CGI (Common Gateway Interface), 228, 282–284
- Change tracking, 365
- Checklist features, 17
- checkpassword*, 167
- Christensen, Clayton, 52
- ci*, 246
- Clarity, 14–15
 - rule of, 13, 102, 154–155
- Clark, Alexander, 231
- Classical Internet application metaprotocol, 127–128
- clear*, 268
- CLI server, 123
 - pattern, 182, 278–279
- Client, 59
- clone*, 466
- CMS Pipelines, 75
- co*, 246
- Code, 11
 - clarity of, 14–15
- Code generation, 215–229
 - ad-hoc code generation, 225–229
 - for *ascii* displays, 225–227
 - generating HTML code for a tabular list, 227–229
 - data-driven programming, 216–225
 - ascii*, 217–218
 - defined, 216–217
 - fetchmailconf*, metaclass hacking in, 219–225
 - statistical spam filtering, 218–219
- Codebase size, 297
- Coherent, 37
- COLUMNS environment variable, 239
- Comer, Douglas, 33
- Command-line interfaces (CLIs), 58–59
- Command-line options, 242–248
 - a option, 243
 - b option, 243
 - c option, 244
 - d option, 244
 - D option, 244
 - e option, 244
 - f option, 244–245
 - h option, 245
 - i option, 245
 - I option, 245
 - k option, 245
 - l option, 245–246
 - m option, 246
 - n option, 246
 - o option, 246
 - p option, 246
 - portability to other operating systems, 248
 - q option, 246

- R option, 246–247
- r option, 246–247
- s option, 247
- t option, 247
- u option, 247
- v option, 247
- V option, 247
- w option, 247
- x option, 247–248
- y option, 248
- z option, 248
- “Communications of the ACM”, paper, 33
- Compactness, 84, 87–89, 202, 225
 - and the strong single center, 91, 92–94
- Compatible Time-Sharing System (CTSS), 29, 234
- Compiler pattern, 269
- Compilers, 14
- Complexity, 295–317
 - accidental, 299–300
 - compromise, 312–314
 - ed.*, 304–305
 - Emacs, 307–308, 314–315
 - essential, 299–300
 - interaction with helper subprocesses, 303
 - interface and implementation complexity,
 - tradeoffs between, 298–299
 - mapping, 300–302
 - optional, 299–300
 - output parsing, 303
 - plain-text editing, 302
 - problems, identifying, 309–312
 - rich-text editing, 302–303
 - Sam editor, 306–307
 - software, right size of, 316–317
 - sources of, 296–298
 - syntax awareness, 303
 - vi* editor, 305–306
 - Wily* editor, 308–309
- Complexity level, choosing, 207–208
- Composition, rule of, 13, 15–16
- Comprehensive Perl Archive Network, 332
- Concision, of interface designs, 257–258
- Concurrent Version System (CVS), 367–368
- condredirect*, 167
- Configuration, 231–252
 - command-line options, 242–248
 - control information in startup-time
 - environment, 233–234
 - environment variables, 238–242
 - portability to other operating systems, 242
 - system, 238–239
 - user, 240
 - when to use, 240–242
- methods:
 - choosing among, 248–252
 - fetchmail*, 249–250
 - XFree86 Server, 251–252
- run-control files, 234–238
 - .netrc* file, 236–237
- portability to other operating systems, 238
- Conventions, used in book, xxix–xxx
- Cookie-jar format, 115–116
- Cooperating processes, 55–56
- Cooperative multitasking, 54–55
- Costello, Joseph, 417
- cp*, 246, 465
- CPAN archive, 384, 456
- cpio*, 245
- crontab*, 244, 247
- csh*, 92
- csplit*, 247
- CSSC (Compatibility Stupid Source Control), 367
- ctags*, 104
- Culture, 3–4
- CUP, 355
- curses* libraries, 144, 292
- Custom grammar, writing, 210
- cut*, 112, 113
- cvs*, 247, 248
- Cygwin, 70
- CTSS. *See* Compatible Time-Sharing System
- d option, 244
- D option, 244
- Da Vinci, Leonardo, 253
- Daemons, 170
 - fetchmail*, 172
- Data file metaformats, 112–123
 - cookie-jar format, 115–116
 - defined, 112
 - DSV style, 113–114
 - record-jar format, 116–117
 - RFC 822 format, 114–115
 - Unix textual file format conventions, 120–122
 - Windows INI format, 119–120
 - XML, 117–119
- Data vs. program logic, 215–216
- Data-driven programming, 216–225
 - ascii*, 217–218
 - defined, 216–217
 - fetchmailconf*, metaclass hacking in, 219–225
 - statistical spam filtering, 218–219
- date*, 476
- dbx*, 371

- dc*, 165, 183, 187, 203–205, 262–264
- Dead code, 155
- debugfs*, 152
- Debugging, 14
- DEC (Digital Equipment Corporation), 37, 40, 61, 75
- Defense Advanced Research Projects Agency (DARPA), 35
- Delimiter-Separated Values (DSV), 113–114
- Dependency derivation, 362
- Designing for the future, 24–25
- Detachment, value of, 94
- df*, 243
- diff*, 38, 93, 441
- Discoverability, 18, 133–134, 258–259
 - coding for, 150–151
 - designing for, 148–149
- Disruptive technology, 52
- ditroff*, 107
- Diversity, rule of, 14, 24, 102
- DLL hell problem, 70
- DocBook, 426, 427–434
 - document type definitions, 427–429
 - editing tools, 432–433
 - migration tools, 431–432
 - related standards/practices, 433
 - SGML, 433
 - toolchain, 429–431
 - writing Unix documentation, best practices for, 434–435
 - XML-DocBook references, 433–434
- DocBook: The Definitive Guide* (Walsh), 434
- Documentation, 417–435
 - concepts, 418–420
 - cultural style, 421–422
 - formats, 422–426
 - DocBook, 426, 427–434
 - Documenter’s Workbench tools, 422–424
 - HTML, 426
 - POD (Plain Old Documentation), 425
 - TEX, 424–425
 - Texinfo*, 425
 - troff*, 422–424
 - large-document bias, 420–421
- Documenter’s Workbench (DWB) tools, 195–199, 204, 422–424
- Domain-specific languages, 183–184
- Don’t Repeat Yourself rule. *See* SPOT rule
- Dormant code, 155
- DSV style, 113–114
- DTD (Document Type Definition), 429
- DTSS, 78
- du*, 243
- DVI (TEX Device-Independent format), 331, 430
- Dynamically-linked libraries (DLLs), 100
 - e option, 244
 - Ease, of interface designs, 257–258
 - Economy, rule of, 13, 22, 375
 - ed*, 188, 304–305, 309, 476
 - ed* pattern, 270
 - EDITOR environment variable, 240
 - egrep*, 188, 244
 - Einstein, Albert, 295
 - Elegant code, 134
 - Elements of Programming Style, The* (Kernighan/Plauger), 104
 - ELIZA program, 205
 - elm*, 256, 272
 - elvis*, 351
 - emacs*, 245, 247
 - Emacs, xxix, 16, 47, 106, 205, 235, 255, 272, 279, 307–308, 311, 314–315, 351–352, 437
 - combining tools with, 370–373
 - and *make*, 371
 - and profiling, 372
 - and runtime debugging, 371
 - and version control, 371–372
 - Emacs Lisp, 88–89, 187, 205, 342–343, 344, 369, 412
 - portability, 343, 412
 - Emacs Lisp libraries, 134
 - Embedded scripts, 334
 - EMX, 67
 - Encapsulation, 85–87, 107
 - Environment variables, 238–242
 - portability to other operating systems, 242
 - system, 238–239
 - user, 240
 - when to use, 240–242
 - EPSF, 203
 - eqn*, 183, 195–196
 - EROS, 472
 - Essential complexity, 299–300
 - etags*, 104
 - Ethernet, 4
 - ex*, 247
 - Expressiveness, of interface designs, 257–258
 - Extending/embedding languages, 209
 - Extensibility, 308, 311, 336
 - rule of, 14, 24–25
 - f option, 244–245
 - f77*, 245
 - faces*, 247
 - fcntl*, 464, 469, 471
 - Feldman, Stuart, 358–360
 - Fellowship theme, 31

- fetchmail*, xxx, 119, 155, 165–166, 172, 186, 200, 219–221, 229, 237, 243, 245–247, 249–250, 274, 327, 358, 359, 361
 - as a pipeline, 165–166
 - run-control syntax, 186, 199–200, 356
 - v option, 136–139, 151
- fetchmailconf*, 219–225, 274, 338–339
 - metaclass hacking in, 219–225
- File compression, 122–123
- File servers, 464
- file* (tool), 122–123, 338
- File Transfer Protocol (FTP), 277
- Filter pattern, 266–267
- Filtering:
 - and discarded information, 267
 - and noise, 267
- Filters, 15
- find*, 476
- First Guide to Postscript (Web site), 202
- flex*, 243, 247, 353, 354–355
- Flex, 441
- Flowcharting, 14
- FO (Formatting Objects), 331, 430
- FOP, 430–431
- fork*, 158
- FORTH, 203
- FORTRAN, 324, 328, 344
- Fourth-generation languages, 14
- Free Software Foundation, 46, 195, 326, 382, 432
- Free Software Movement, 45–47, 49–51
- Free Standards Group, 402–403
- FreeBSD, 76, 426, 466
- Freeciv, 174–175
 - data files, 146–148
- FreeNet, 342
- Freshmeat, 344–345, 384, 454
- fsck*, 248
- fsdb*, 152
- FSF. *See* Free Software Foundation
- ftp*, 270, 277, 465
- ftpd*, 277
- Function-call overhead, 289
- fuser*, 247

- Gabriel, Richard, 298–299
- Galloway, Tom, 105
- Gancarz, Mike, xxviii
- gated*, 172
- Gates, Bill, 36
- gcc*, 245, 269
- GCC (GNU C compiler), 47, 139–140, 382
- GCOS, 78
- gdb*, 270, 371
- Geeks, 44–45
- Gelernter, David, 133, 208
- GEM, 67
- General Public License (GPL), 46, 458–459
- Generated vs. hand-hacked code, 22
- Generation, rule of, 13, 22–23
- get*, 244
- getenv*, 238
- gettext*, GNU, 413–414
- getty*, 244
- Gettys, Jim, 7, 108, 181, 291, 407
- GhostScript Project, 202
- ghostview*, 276
- gif2png*, 269
- GIMP, xxx, 261, 279, 280, 281, 346, 383
 - plugins, 100–101
- Glade*, 185, 191–193, 353, 356–357
- glob* expressions, 188
- glob* facility, 93
- Glue layers, 97–98
- GNOME project, 66, 178, 315, 346, 383, 426, 476
- GNU, 37, 39–40, 42, 48, 248, 270, 326
- GNU Compiler Collection, 39, 326, 329, 384, 427, 439
- GNU Emacs, 352
- GNU Emacs Lisp Reference Manual*, 343
- GNU project, 39–40, 46–47, 248, 356, 422, 425, 440
- GNU Texinfo, 432
- GPL. *See* General Public License
- gprof*, 289, 370
- grap*, 183, 195–196
- Graphical user interfaces (GUIs), xxx, 15–16, 58, 79, 96, 97, 197, 256, 279, 281, 468, 476
- grep*, 93, 112, 119, 249, 256, 266, 354
- grn*, 195
- groff*, 164, 195, 199, 422–423
- grops*, 247
- gs*, 276
- GTK, 346–347, 409
- Guile*, 209
- gunzip*, 123, 269
- gv*, 276
- gzip*, 122, 269

- h option, 245
- Hackers, 5, 9, 10–11, 33–34, 37, 40–41, 43–51, 76, 84, 123, 155, 202, 272, 298, 305, 328, 332, 337, 344–345, 350, 404, 420, 440, 475, 478
 - origins/history of, 43–49

- Hamming, Dick, 7
- Hand-hacking, avoiding, 22–23
- Harness program, 278
- Hatton, Les, 183
- head*, 246
- Heuristics, 93–94
- High-level-language code, 22
- Hoare, C. A. R., 83, 287
- HOME environment variable, 239
- HP-UX, 77
- HTML, 96–97, 205, 229, 281, 284, 331, 426
- HTTP, 97, 128–130

- i option, 245
- I option, 245
- ibiblio*, 384
- IBM, 4, 10, 33, 36–37, 39–41, 43, 65–67, 72–76, 271, 395, 478, 495
 - alliance with Microsoft, 65
- identify*, 338
- Idris*, 37
- IEEE-X, 400
- IETF. *See* Internet Engineering Task Force
- ifconfig*, 246
- imake*, 363
- IMAP, 123, 126–127, 293
- Imgsizer, 338
- inetd*, 123, 172, 278
- infocmp*, 145, 152
- InfoZip, 383
- Institute of Electrical and Electronic Engineers (IEEE), 399–400
- Intel, 43
- Interfaces, 10, 253–285
 - calculator program, writing, 262–264
 - CLI and visual interfaces, tradeoffs between, 259–262
 - configurability, 264–266
 - expressiveness, 264–266
 - interface design, evaluating, 257–259
 - Rule of Least Surprise, applying, 254–255
 - Rule of Silence, 284–285
 - transparency, 264–266
 - Unix interface design, history of, 256–257
 - Unix interface design patterns, 266–280
 - applying, 280–281
 - cantrip pattern, 268
 - CLI server pattern, 278–279
 - compiler pattern, 269
 - ed* pattern, 270
 - filter pattern, 266–267
 - language-based patterns, 279–280
 - polyvalent-program pattern, 281
 - roguelike pattern, 270–271
 - “separated engine and interface” pattern, 273–278
 - sink pattern, 269
 - source pattern, 268
 - Web browser as universal front end, 281–284
- Internal boundaries, 57
- Internationalization, 413–414
- Internet, 8–9, 476
- Internet Engineering Task Force (IETF), 44, 49
 - Printer Working Group, 129
 - and RFC standards process, 403–405
- Internet Printing Protocol (IPP), 129–130
- Internet Standards Process (Revision 3)*, 404
- Internet technical culture, 10
- Interprocess communication (IPC), 55, 158
- Introspection, 222
- iocctl*, 464, 471
- ITS, 44, 78

- Jabber, 131
- Jack, 355
- Jargon File, 44, 45
- Java, 22, 50–51, 66–67, 70, 88–89, 96, 106, 131, 187, 207, 209, 222, 282–283, 311, 324, 337, 339–342, 369, 411–412
 - FreeNet, 342
 - Native Method Interface, 339
 - portability, 411–412
- Java Applets page, 385
- JavaScript, 96, 187, 205–207, 281
- JCL (Job Control Language), 73
- Johnson, Steve, 10, 84, 208, 232–233, 276, 290, 354, 395, 396, 397, 468
- Jolitz, William, 41
- Joy, Bill, 35–36
- Jython, 337

- k option, 245
- Kay, Alan, 461
- KDE project, 179, 315, 383, 426, 476
- Kernighan, Brian, 5, 14, 259, 335, 395
- Khosla, Vinod, 36
- KISS principle, 25, 119, 219
- kmail*, 140–144, 154, 265
- Knuth, Donald, 23, 370, 430
- Korn, David, 180, 201, 496
- Korn shell, 209, 330–331, 496
- ksh*, 244

- l option, 245–246
- Lampson, Burt, 26

- Language vs. application protocol, 212–213
- Language-based interface patterns, 279–280
- Lasermoon, 401
- L^AT_EX, 211, 424–425, 432, 433
- Law of Software Envelopment, 313
- Least Surprise, rule of, 13, 20, 134, 197, 210, 253, 255, 285
 - applying, 254–255
- Lesk, Mike, 24, 99, 262, 354, 419
- less*, 163, 247
- lex*, 93, 183, 185, 210, 352, 353–356
- LGPL, 388
- Libraries, 99–100
- LINES environment variable, 239
- Linux operating system, 4, 6, 7–8, 41–43, 48–51, 61, 65–67, 72, 76–79, 100, 117, 119, 134, 152, 155, 158, 172, 175–176, 178, 207, 241, 248, 251, 256, 269, 280, 331, 402, 440
 - archives, 384
 - as a cultural phenomenon, 41–43
 - kernel, 383
 - Red Hat, 429–430
 - Sorcerer GNU/Linux, 332
- Linux Standards Base (LSB), 402
- Lion, John, 34
- Lisp, 16, 22, 30, 46, 88, 194, 222, 337, 343
- Little-endian, 98
- Live code, 155
- Lock file, 171
- LOGNAME environment variable, 238
- lpd*, 275
- lpr*, 269, 275
- ls*, 162, 245, 249, 268, 465
- lynx*, 144, 359
 - m option, 246
- m4*, 185, 193–194, 210
- Mac Interface Guidelines, 64
- Macintosh, 475–477
 - culture, convergence with Unix culture, 477
- MacOS, 4, 50, 62, 64–65, 176, 242, 248, 329, 333, 335, 338, 341, 346, 347, 351, 409–410, 471, 476–477
- Macro expansion, 194, 211
- Macros, 210–212
- Mail transport agent (MTA), 124
- MAILER environment variable, 240
- Maintainability, designing for, 154–155
- make*, 89, 183, 197, 246, 357–364
 - basic theory of, 357–358
 - for document-file translation, 359
 - generating makefiles, 362–364
 - autoconf*, 363–364
 - automake*, 364
 - imake*, 363
 - makedepend*, 362–363
 - in non-C/C++ development, 359
 - utility productions, 359–362
- makedepend*, 362–363
- man* macros, 88–89
- man* pages, 432
- Manularity trap, 297
- Mapping complexity, 300–302
- Marshalling, 106
- MATLAB, 290
- McCarthy, John, 30
- McIlroy, Doug, 5, 7, 11–12, 17, 21–22, 30, 32, 49, 55, 56, 77, 90, 92–93, 161–162, 177, 212, 232, 250, 267, 314, 322, 450
- McKusick, Marshall Kirk, 41, 61
- mdoc*, 422
- Memory management unit (MMU), 57, 63, 69, 71
- Metaclass hacking, 222–223
- Metcalf, Robert, 4
- MICAH!, 232
- Microsoft, 36, 40, 47, 65–66, 67, 70, 72
 - Windows, 10, 50, 75, 78–79, 146, 176
 - Windows NT, 68–71
 - Word, macro viruses, 78, 207
- MIME, 114, 131, 140, 480
- Minilanguages, 183–213
 - applying, 187–199
 - awk*, 200–202
 - bc* and *dc*, 203–205
 - designing, 184, 206–213
 - complexity level, choosing, 207–208
 - custom grammar, writing, 210
 - extending/embedding languages, 209
 - language vs. application protocol, 212–213
 - macros, 210–212
- Documenter's Workbench (DWB) tools, 195–199
- Emacs Lisp, 205
- fetchmail* run-control syntax, 199–200
- Glade*, 191–193
- JavaScript, 205–206
- m4*, 193–194
- PostScript, 202–203
- regular expressions, 188–191
 - basic, 188
 - examples, 189
 - extended, 188
 - glob* expressions, 188
 - operations, 190
 - Perl, 189
 - sng*, 187–188

- spectrum of, 186
- taxonomy of languages, 185–187
- XSLT, 194–195
- MIT Artificial Intelligence Lab, 44
- MIT or X Consortium License, 388, 457
- mkisofs*, 276
- mm*, 422
- mmap*, 175, 449, 462
- Moded interface, 305–306
- Model/view/controller pattern, 273–274
- Modularity, 14, 83–104
 - coding for, 103–104
 - compactness, 87–89
 - and the strong single center, 92–94
 - encapsulation, 85–87
 - libraries, 99–100
 - optimal module size, 85–87
 - orthogonality, 89–91
 - rule of, 13, 84, 85, 159
 - SPOT rule, 91–92
- Moodss system, 336
- Moore's Law, 287
- more*, 162, 163
- mountd*, 172
- Mozilla, 9, 49, 255, 440
- Mozilla Public License (MPL), 388, 459
- MPE, 78
- MPL. *See* Mozilla Public License
- ms*, 422
- MTS, 78
- Multics, 29–32, 44, 61, 74–75, 78, 438
- Multics Emacs: The History, Design, and Implementation* (Greenberg), 308
- Multiprogramming, 157–182
 - bc* and *dc*, 165
 - Bernstein chaining, 167–168
 - complexity control, separating from
 - performance tuning, 159–160
 - defined, 157
 - fetchmail*, as a pipeline, 165–166
 - filters, 162–166
 - mutt* mail user agent, 161
 - peer-to-peer interprocess communication,
 - 169–176
 - fetchmail*'s use of signals, 172
 - shared memory, 175–176
 - signals, 170–171
 - sockets, 172–175
 - system daemons and conventional signals,
 - 172
 - tempfiles, 169–170
 - pic2graph*, 164–165
 - pipes, 161–166
 - piping to a pager, 163
 - problems/methods to avoid, 176–181
 - obsolete Unix IPC methods, 176–178
 - remote procedure calls (RPCs), 178–180
 - threads, 180–181
 - process partitioning at the design level,
 - 181–182
 - process spawning, 158
 - redirection, 162–166
 - slave processes, 168–169
 - scp* and *ssh*, 169
 - specialist programs, handing off tasks to,
 - 160–161
 - Unix IPC methods, taxonomy of, 160–161
 - word lists, making, 164
 - wrappers, 166
 - backup scripts, 166
 - security wrappers, 167–168
- Multitasking, 54–55
 - cooperative, 54–55
 - preemptive, 55
- Multitrack editing, 136
- mutt*, xxx, 144, 161, 256, 272
- mv*, 465
- MVS (Multiple Virtual Storage), 72–74
 - n option, 246
 - named*, 172
 - Named pipes, 163
 - NetBSD, 76
 - Netscape, 49–50, 67, 255
 - NFS (Network File System), 178
 - nfsd*, 172
 - No Silver Bullet* (Brooks), 300
 - Noncompact designs, 89
 - Nonlocality, 290–291
 - Novell, 43
 - nroff*, 32, 246, 422
 - nsgmls*, 247
 - NUMA (non-uniform memory access), 290
 - o option, 246
 - O notation, 288
 - Object orientation, 14, 216
 - Object-oriented programming (OO), 101–103
 - Obsolete Unix IPC methods, 176–178
 - STREAMS, 177–178, 462
 - System V IPC, 177
 - od*, 247
 - Olsen, Ken, 40
 - “On Holy Wars and a Plea for Peace”, paper, 98
 - On the Design of Application Protocols* (Rose), 127

- Open Software Foundation (OSF), 40
- Open source, 39, 41–43, 47–52, 61, 65, 67, 70–72, 76–77, 79, 117, 130–131, 146, 155, 174–175, 178, 191, 194–196, 201–203, 205–206, 209, 251, 256, 280, 284, 290, 298, 322, 324, 326, 328, 329, 332, 340–341, 344–346, 350, 355, 357, 359, 362, 367–369, 372–373, 380–390, 399, 401–402, 408–410, 412–415, 424, 426–427, 437–459
 - license selection, 456
 - licensing, 457–459
 - Artistic License, 458
 - BSD Class License, 457–458
 - General Public License, 46, 458–459
 - MIT or X Consortium License, 457
 - Mozilla Public License, 459
 - Open Source Definition (OSD), 7, 386–387, 457
 - seeking legal assistance, 390
 - software, issues in using, 385
 - standard licenses, 388–390
 - using, 457
 - and Unix, 438–440
 - working with developers, 440–456
 - communication practice, 454–456
 - development practice, 447–450
 - distribution-making practice, 450–454
 - patching practice, 440–444
 - project- and archive-naming practice, 444–447
- OpenBSD, 76
- Open-source community, 9
- Open-source movement, 49–51
- Operating systems:
 - binary file formats, 58
 - comparisons, 61–78
 - BeOS, 71–72, 79
 - Linux, 76–78
 - MacOS, 64–65
 - MVS (Multiple Virtual Storage), 72–74
 - OS/2, 65–67
 - VM/CMS, 74–76
 - VMS, 61–64
 - Windows NT, 68–71
 - elements of style, 53–61
 - entry barriers to development, 60–61
 - file attributes, 57–58
 - intended audience, 59–60
 - internal boundaries, 57
 - inter-process communication (IPC), 55
 - multitasking, 54–55
 - preferred user interface style, 58–59
 - record structures, 57–58
- Optimal module size, 85–87
- Optimization, 287–294
 - choosing to do nothing, 287–288
 - measuring before, 288–290
 - nonlocality, 290–291
 - rule of, 14, 23–24
 - throughput vs. latency, 291–294
 - batching operations, 292
 - caching operation results, 293–294
 - overlapping operations, 293
- Optional complexity, 299–300
- Orthogonality, 84, 87, 89–91, 93–94, 99, 225, 301, 305
- OS/2, 65–67
- Ossana, Joseph, 417
- Ousterhout, John, 411
- Output parsing, 303
- p option, 246
- Packard, Keith, 47, 117, 407–408
- PAGER environment variable, 240
- Palladium, 475
- Parser/lexer generators, 23
- Parsimony, rule of, 13, 18
- Pascal, 324
- PassiveTeX, 430
- passwd*, 245
- paste*, 113
- patch*, 38, 39, 42, 93, 437, 441
- Paterson, Tim, 36
- PATH environment variable, 239
- Pattern Language, A* (Alexander), xxxii
- PDF. *See* Portable Document Format
- PDP computer Web FAQ, 30
- PDP computers, 123
 - PDP-1, 44
 - PDP-7, 30–32
 - PDP-10, 37, 44–45. *See also* TOPS-10 and TOPS-20
 - PDP-11, 32, 34, 40, 84, 98, 395, 481
- Peer-to-peer interprocess communication, 169–176
 - fetchmail*'s use of signals, 172
 - shared memory, 175–176
 - signals, 170–171
 - sockets, 172–175
 - system daemons and conventional signals, 172
 - tempfiles, 169–170
- perl*, 244
- Perl, xxvi, 22, 39, 49, 75, 88–89, 103, 187, 201–202, 209, 228, 311, 332–334, 338, 369, 410
 - blq*, 333–334

- keeper, 334
- portability, 410
- Perl XS, 339
- Perlis, Alan, 157, 200
- pic*, 164–165, 183, 195–199, 210
- pic2plot*, 199
- pidfile*, 171
- Pike, Rob, 12, 19, 23, 154, 288, 466
- pine*, 272
- Pipes, 11, 36, 45, 55, 56, 58, 75, 107, 158, 161–166, 176, 197, 242, 267, 278, 466
- Piping to a pager, 163
- PL/1, 84, 324
- Plain-text editing, 302
- Plan 9 from Bell Labs*, 306, 308, 312, 464–466, 471–472
- Plugins, 100
 - GIMP, 100–101
- PLY, 355
- PNG (Portable Network Graphics) file format, 111–112, 130, 142–144, 152, 154, 185, 187, 269, 339, 414
- POD (Plain Old Documentation), 425, 426, 432
- poll*, 182
- Polyvalent-program pattern, 281
- POP3, 123, 124–126, 293
- popen*, 168
- popfile*, 218
- Portability, 362, 393–417
 - C language:
 - early history of, 395–396
 - evolution of, 394–398
 - standards, 396–398
 - and choice of language, 409–412
 - C, 409–410
 - C++, 410
 - Emacs Lisp, 412
 - Java, 411–412
 - Perl, 410
 - Python, 410–411
 - shell, 410
 - Tcl, 411
 - internationalization, 413–414
 - Internet Engineering Task Force (IETF), and RFC standards process, 403–405
 - and open standards/open source, 414–415
 - programming for, 408–413
 - system dependencies, avoiding, 412–413
 - tools for, 413
 - Unix standards, 398–403
 - in the open-source world, 402–403
 - and Unix wars, 398–401
- Portability layer, 449
- Portable Document Format (PDF), 420, 426, 430
- Portable Operating System Standard (IEEE), 8
- POSIX, 8, 39–40, 67, 71, 170, 175–176, 181, 188–189, 326, 398, 400–401, 414, 469
- Postel, Jonathan, 21
- Postel’s Prescription, 21, 267
- PostgreSQL, 174
- postmaster*, 277
- PostScript, 164, 186, 202–203, 331, 420, 426, 431
- pppd*, 172
- pr*, 247
- Preemptive multitasking, 55
- Presentation-level markup, 419
- Principle of Least Astonishment, 20
- Procedural programming, 14
- Process spawning, 158
- Profiling, 370
- Programmer time, conserving, 22
- Protocol family, 173
- Prototyping, 23–24
- ps*, 245, 247, 268
- PS/2 port, 39
- psql*, 152, 277
- PyChecker, 448
- Python, xxvi, 22, 75, 88–89, 106, 187, 209, 228–229, 244, 311, 327, 333, 336–339, 369, 383, 410–411
 - fetchmailconf*, 338–339
 - home page, 384
 - imgsizer*, 338
 - portability, 410–411
 - Python Imaging Library (PIL), 339
 - Python interpreter, 293–294
 - Python Software Activity, 456
- q option, 246
- QDOS, 36
- qed editor, 304
- qmail*, 275
- qmail* POP3 server, 167
- qmail-popup*, 167–168
- qmail-smtpd*, 168
- QNX, 37
- Qt toolkit, 329–330, 409
- r option, 246–247
- R option, 246–247
- Race condition, 171
- Rationale for the Structure of the Model and Protocol for the Internet Printing Protocol*, 129–130
- rblsmtpd*, 167–168
- read*, 177
- Record-jar format, 116–117

- Redirection, 162–166
- Refactoring, 90–91
- refer*, 195
- Registry creep, 69
- Regular expressions, 188–191
 - basic, 188
 - examples, 189
 - extended, 188
 - glob* expressions, 188
 - operations, 190
 - Perl, 189
- rehash** directive, in *cs*, 92
- Remote procedure call (RPC), 68, 178–180
- Repair, rule of, 13, 21–22, 147, 154, 252
- Representation, rule of, 13, 19–20, 102
- Requests for Comment (RFCs), 44
- Reuse, 375–390
 - moving to open source, 380–381
 - new programmers, 376–379
 - and transparency, 379–380
- Revision Control System (RCS), 367
- Rexx scripting language, 66, 75
- RFC 822 format, 114–115
- r*fork, 466
- Rich-text editing, 302–303
- Ritchie, Dennis, 5, 30–33, 34, 36, 49, 98–99, 161, 177, 395
- rm*, 268
- Robustness, rule of, 13, 18–19, 154
- rogue*, 256, 270
- Roguelike pattern, 270–271
- Rootless Root*, 499–509
- Rule of Clarity, 13, 102, 154–155
- Rule of Composition, 13, 15–16
- Rule of Diversity, 14, 24, 102
- Rule of Economy, 13, 22, 375
- Rule of Extensibility, 14, 24–25
- Rule of Generation, 13, 22–23
- Rule of Least Surprise, 13, 20, 134, 197, 210, 253, 255, 285
 - applying, 254–255
- Rule of Minimality, 316
- Rule of Modularity, 13, 84, 85, 159
- Rule of Optimization, 14, 23–24
- Rule of Parsimony, 13, 18
- Rule of Repair, 13, 21–22, 147, 154, 252
- Rule of Representation, 13, 19–20, 102
- Rule of Robustness, 13, 18–19, 154
- Rule of Separation, 13, 16–17
- Rule of Silence, 13, 20–21, 284–285
- Rule of Simplicity, 13, 17, 102, 178
- Rule of Transparency, 13, 18, 19, 102, 151, 264
- Run-control files, 234–238
 - .netrc* file, 236–237
 - portability to other operating systems, 238
- Russell, Bertrand, 183
 - s option, 247
 - SAIL, 44
 - Saint-Exupéry, Antoine de, 99
 - Samba, 426
 - Sandboxed languages, 207
 - Santa Cruz Operation (SCO), 35, 43
 - Santayana, George, 29
 - Saxon, 429
 - Scheme, 279
 - scp*, 169
 - screen*, 470
 - Scriptability, of interface designs, 257–259
 - Scripting languages, xxx, 112, 122, 166, 183, 187–188, 201–202, 204, 209, 226, 321–324, 337–338, 344, 359, 410
 - ScrollKeeper, 433, 435
 - sdb*, 371
 - Security wrappers, 167–168
 - sed*, 89, 112, 183, 188, 330
 - select*, 182
 - Semi-compact designs, 88
 - sendmail*, 183–184, 193, 275, 277
 - “Separated engine and interface” pattern, 273–278
 - client/server pair, 277–278
 - configurator/actor pair, 274
 - driver/engine pair, 275–277
 - spooler/daemon pair, 275
 - Separation, rule of, 13, 16–17
 - Server, 59
 - setuid* bit, 159
 - SGML, 205, 426, 432–434
 - sh*, 244, 247, 270
 - shar*, 247
 - Shared culture, xxv
 - Shared memory, 175–176
 - shell*, 75, 88–89, 187, 410
 - SHELL environment variable, 239
 - Shell:
 - languages, 22
 - portability, 410
 - programs, 330–332
 - Sorcerer GNU/Linux, 332
 - xmlto*, 331–332
 - Shelling out, 160
 - Signal handler, 170
 - Silence, rule of, 13, 20–21, 284–285
 - Simplicity, rule of, 13, 17, 102, 178
 - Single Point of Truth. *See* SPOT rule

- Single Unix Specification, 175, 401
- Sink pattern, 269
- Slave processes, 168–169
 - scp* and *ssh*, 169
- slrn*, 144, 272
- smbclient*, 152, 237
- SMTP, 123, 124
- SNA (Systems Network Architecture), 73–74
- sng*, 152, 187–188
- SNG (Scriptable Network Graphics format), 142–144, 154, 185
- snmpnetstat*, 246
- SOAP (Simple Object Access Protocol), 131, 179
- Sockets, 163, 172–175
 - Freeciv, 174–175
 - PostgreSQL, 174
- Software engineering, 3–4
- Software-development methodologies, 14
- Solaris, 4, 77
- Sorcerer GNU/Linux, 332
- sort*, 249, 266, 330
- Source Code Control System (SCCS), 367
- Source pattern, 268
- SourceForge, 344–345, 384
- Space Travel game, 30–31, 44
- SPACEWAR, 44
- spambayes*, 218
- Spencer, Henry, 3, 15, 19, 20, 100, 107, 145, 153, 204, 216, 250, 264, 354, 365, 376, 397, 415, 439, 469
- Sponge, 269
- SPOT rule, 91–92, 220, 224, 293, 301
- ssh*, 169
- Stallman, Richard (RMS), 37, 39, 45–47, 48, 51, 314, 400
- startx*, 268
- stat*, 449
- Stephenson, Neal, xxv
- STL (Standard Template Library), 328
- STREAMS, 177–178, 462
- Structured programming, 14
- Stylesheet, 419, 428
- Subversion, 368–369
- Sun Community Source License (SCSL), 340
- Sun Microsystems, 36–38, 40, 42, 50, 178, 339, 345, 399
 - Network Window System (NeWS), 256
- Syntactic sugar, 200
- Syntax awareness, 303
- system*, 160
- System III Unix, 90, 171, 367, 396, 398
- System V Unix, 37–38, 171, 176–178, 238, 299, 395, 398–400, 462. *See also* AT&T
 - t option, 247
- tail*, 246
- Tales of Zen Master Greg*, 500
- Tao of Programming, The*, 499
- Tao Te Ching, 375
- tar*, 166, 243–244, 248
- tbl*, 183, 195–197
- Tcl, 22, 131, 186, 187, 209, 212, 228, 279, 281, 322, 327, 333, 334–339, 342–347, 369, 371, 394, 410, 411
 - expect* package, 212
 - Moodss system, 336
 - portability of, 411
 - TkMan, 336
- tcpdump*, 246
- TCP/IP, 8–9, 35, 36, 66, 73–74, 174, 176, 178, 239, 278, 462, 463
- TECO, 184
- TERM environment variable, 239
- termcap* capability format, 144
- terminfo*, 145–146
- Terminfo* database, 144–146
- T_EX, 211, 261, 424–425, 430
- Texinfo, 425
- Text streams, 15
- Textuality, 105–131
 - application protocol design, 123–127
 - IMAP, 123, 126–127
 - POP3, 123, 124–126
 - SMTP, 123, 124
 - application protocol metaformats, 127–132
 - CDDB/freedb.org database, 129
 - classical Internet application metaprotocol, 127–128
 - HTTP, 128–129
 - data file metaformats, 112–123
 - cookie-jar format, 115–116
 - defined, 112
 - DSV style, 113–114
 - record-jar format, 116–117
 - RFC 822 format, 114–115
 - Unix textual file format conventions, 120–122
 - Windows INI format, 119–120
 - XML, 117–119
 - file compression, 122–123
 - importance of, 107–109
- Internet Printing Protocol (IPP), 129–130
 - BEEP (Blocks Extensible Exchange Protocol), 130–131
 - Jabber, 131
 - SOAP (Simple Object Access Protocol), 131, 179
 - XML-RPC, 131, 158

- marshalling, 106
 - .newsrc format, 110–111
 - PNG (Portable Network Graphics) file format, 111–112, 185
 - Unix password file format, 109–110
 - unmarshalling, 106
- Textualizers, 152
- Thin-glue principle, 97–98
- Thompson, Ken, 4, 5, 16, 23, 24, 30–32, 34–35, 37, 60, 84, 98–99, 104, 139, 155, 161, 203, 304, 395, 466, 472
- Threads, 180–181
- Throughput vs. latency, 291–294
 - batching operations, 292
 - caching operation results, 293–294
 - overlapping operations, 293
- Throwing an exception, 470
- Timeless Way of Building* (Alexander), xxxii
- Timesharing operating systems, 61–62
- tin*, 272
- Tk, 409
- Tk toolkit, 334
- TkMan, 336
- Toolkits, 256–257
- Tools, 349–373
 - developer-friendly operating system, 349–373
 - editors:
 - choosing, 350–353
 - Emacs, 351–352
 - using both *vi* and Emacs, 352
 - vi*, 351
 - Emacs:
 - combining tools with, 370–373
 - and *make*, 371
 - and profiling, 372
 - and runtime debugging, 371
 - and version control, 371–372
 - make*, 357–364
 - autoconf*, 363–364
 - automake*, 364
 - basic theory of, 357–358
 - for document-file translation, 359
 - generating makefiles, 362–364
 - imake*, 363
 - makedepend*, 362–363
 - in non-C/C++ development, 359
 - utility productions, 359–362
 - profiling, 370
 - runtime debugging, 369
 - special-purpose code generators, 352–357
 - fetchmailrc* grammar, 356
 - Glade*, 356–357
 - lex*, 353–356
 - yacc*, 353–356
- version control:
 - Aegis, 368–369
 - automated, 366
 - Concurrent Version System (CVS), 367–368
 - by hand, 365
 - purpose of, 364–365
 - Revision Control System (RCS), 367
 - Source Code Control System (SCCS), 367
 - Subversion, 368–369
 - systems, 364–369
 - Unix tools for, 367–369
- Top-down vs. bottom-up design, 95–97
- TOPS-10, 78
- TOPS-20, 44, 78, 468
- Torvalds, Linus, 41, 47, 48–49, 406, 437
- touch*, 268
- tr*, 112, 266
- Transparency, 133–155
 - case studies:
 - audacity*, 135–136
 - fetchmail*, *-v* option, 136–139
 - FreeCiv data files, 146–148
 - GCC (GNU C compiler), 139–140
 - kmail*, 140–144
 - Terminfo* database, 144–146
 - of code, 150–151
 - of designs, 148–149
 - and discoverability, 133–134, 148–151
 - editable representations, 152–153
 - fault diagnosis, 153–154
 - fault recovery, 153–154
 - of interface designs, 257–258
 - maintainability, designing for, 154–155
 - overprotectiveness, avoiding, 151–152
 - rule of, 13, 18, 19, 102, 151, 264
 - Zen of, 149–150
- troff*, 89, 107, 183, 186, 197, 199, 246, 261, 417, 421, 422–424
 - markup, 89
- Trusted Computing Platform Alliance, 474–475
- TSO, 73
- Turing-complete languages, 187, 201, 205
- u* option, 247
- ucspi-tcp*, 168
- uNETix, 37
- Unix:
 - APIs, 90
 - application programming interface (API), 8
 - basic philosophy, 11–14
 - binary file formats, 58
 - and casual programming, 60–61
 - code, 11

- code generation, 215–229
- complexity, 295–317
- configuration, 231–252
- cross-platform portability, 8
- culture:
 - case against learning, 5–6
 - expertise transmitted by, 11–12
 - problems in, 475–477
- data file metaformats, 112–123
- documentation, 417–435
 - concepts, 418–420
 - cultural style, 421–422
 - formats, 422–426
 - large-document bias, 420–421
- durability/adaptability of, 4–5
- environment, problems in, 473–475
- files, 6, 466–467
 - attributes, 57–58
 - deletion, 6
- flexibility of, 9–10
- fusion with ARPANET culture, 45–47
- future of, 461–478
- and hackers, 10–11, 33, 43
- heritage of, 6
- history of, 29–43
 - lessons learned from, 51–52
- interfaces, 253–285
 - styles, 7
- internal boundaries, 57
- inter-process communication (IPC), 55
- job control, 6
- koans of Master Foo, 499–504
- language evaluations, 325–343
 - C language, 326–327
 - C++, 327–330
 - Emacs Lisp, 342–343
 - Java, 339–342
 - Perl, 332–334
 - Python, 336–339
 - shell programs, 330–332
 - Tcl, 334–336
- languages, 321–347
 - C language, 323–325
 - cornucopia of, 321–322
 - future trends, 344–345
 - interpreted languages, 325
- minilanguages, 183–213
- modularity, 83–104
- multiprogramming, 157–182
- multitasking, 54–55
- objections to, 6–7
- and object-oriented languages, 101–103
- open source, 437–459
- open standards, 8
- open-source applications, 9
 - as open-source software, 7–8
- operating system, elements of style, 53–61
- optimization, 287–294
- options, 7
- origin of, 29–30
- philosophy, 3–27, 369
 - applying, 26
 - compared to others, 53–80
- portability, 393–417
- preferred user interface style, 58–59
- record structures, 57–58
- reuse, 375–390
- specification level, 215–229
- textuality, 105–131
- tools, 349–373
- transparency, 133–155
 - and universality, 6
 - use of term, xxix
 - Version 7 release, 34
 - wars, 35–41, 398–401
 - X toolkit, choosing, 346–347
- Unix interface design, history of, 256–257
- Unix interface design patterns, 266–280
 - applying, 280–281
 - cantrip pattern, 268
 - CLI server pattern, 182, 278–279
 - compiler pattern, 269
 - ed* pattern, 270
 - filter pattern, 266–267
 - language-based patterns, 279–280
 - polyvalent-program pattern, 281
 - roguelike pattern, 270–271
 - “separated engine and interface” pattern, 273–278
 - client/server pair, 277–278
 - configurator/actor pair, 274
 - driver/engine pair, 275–277
 - spooler/daemon pair, 275
 - sink pattern, 269
 - source pattern, 268
- Unix System V. *See*
- Unix textual file format conventions, 120–122
- Unmarshalling, 106
- unzip*, 245
- URI (Universal Resource Indicator), 128
- URL (Uniform Resource Locator), 9, 128
- Usenet, 35, 42, 110, 114, 205, 272, 305, 343, 383, 440, 451, 455, 495, 497
- USER environment variable, 238
- User-interface design, 140, 475
- Utility productions, 359–362
- UUCP, 35
- uupc*, 247

- v option, 247
- V option, 247
- Validating parser, 428
- VAX/VMS operating system, 10, 35, 367
- Venix, 37
- Version control:
 - Aegis, 368–369
 - automated, 366
 - by hand, 365
 - purpose of, 364–365
 - Subversion, 368–369
 - systems, 364–369
 - Unix tools for, 367–369
 - Concurrent Version System (CVS), 367–368
 - Revision Control System (RCS), 367
 - Source Code Control System (SCCS), 367
- vi*, 35, 247, 272, 311, 351, 352
- vile*, 351
- vim*, 351
- VM/CMS, 4, 74–76
- VMS, 40, 61–64, 468

- w option, 247
- Wall, Larry, 39, 458
- Waterfall model, 406
- wc*, 162
- White Book, 94, 395
- who*, 268
- Wily* editor, 308–309
- Window manager, 257
- Windows operating system, 10, 40, 50, 66, 67, 75, 77–79, 100, 113, 135, 146–147, 152–153, 171, 173, 175, 176, 179, 181, 193, 207, 238, 242, 248, 252, 256–257, 265, 324, 329, 340, 349–352, 367, 409–410, 412, 450, 453
 - Windows Help markup, 426
 - Windows INI format, 119–120
 - Windows NT, 68–71
- Wittgenstein, Ludwig, 321
- Workplace Shell (WPS), 66
- World Wide Web (WWW), 8–9, 42, 281, 476
- “Worse Is Better”, paper (Gabriel), 298, 438
- Wrappers, 166, 330
 - backup scripts, 166
 - security wrappers, 167–168
- write*, 177
- wxWindows, 346–347, 409
- WYSIWYG, 196, 258, 418–419

- x option, 247–248
- X clients, 251
- X servers, 251

- X toolkit, 191
 - choosing, 346–347
- X windowing system, 7, 47, 49, 63, 108, 183, 463
 - and XFree86, 42,
 - Xalan, 429
 - xcalc*, 262–263
 - xcdroast*, xxxi, 276
 - xdb*, 371
 - XEmacs, 352
 - XENIX, 35
 - Xerox PARC, 256, 273, 461, 463, 467–468, 480
 - XFree86Config file, 251–252
 - XFree86, 42, 251–252, 383
 - XHTML, 426
 - XML, 117–119, 196, 205, 284, 331, 463
 - XML Cover Pages, 434
 - XML-DocBook, 426
 - xmllint*, 147
 - XML-RPC, 131, 158
 - xmlroff*, 431
 - xmltk* toolkit, 118–119
 - xmto*, xxx-xxxi, 331–332, 425–426, 429, 435
 - X/Open, 401
 - XSLT, 186, 194–195, 331–332
 - xsltproc*, 429–430
 - xterm*, 144, 244, 403
 - xvi*, 351

- y option, 248
- yacc*, 10, 93, 183, 185, 210, 352, 353–356
- Yacc/M, 355
- ypbind*, 172

- z option, 248
- Zawinski’s Law, 313, 316
- zcat*, 248
- Zen, xxix, 13, 149–150, 225, 287, 317, 477
 - and value of detachment, 94
- zip*, 122, 248