
Preface

Although Mac OS X is a relatively new operating system, its lineage is rather colorful, and the history of most of its components is much older. Mac OS X, and Apple in general, has attracted a lot of attention in recent years. Given Apple’s “cult-like” status, and the unique cultural and technical composition of Mac OS X, it is not surprising that people with vastly different backgrounds and interests—both users and non-users of Mac OS X—are curious about the system.

After working on, programming on, and extending several operating systems over the years, I was introduced to Mac OS X on April 1, 2003.¹ I was soon curious about the structure of the system. Despite the availability of several good books on Mac OS X, I was dismayed to find that I could not learn the details of Mac OS X internals from a book—*no such book existed*. There are books describing how to perform various tasks on Mac OS X; how to configure, customize, and tweak the system; and how certain user-visible aspects differ between Mac OS X and Windows. There are also books covering specific Mac OS X programming topics, such as the Cocoa and Carbon APIs. Other books make life easier for a Unix² user migrating to Mac OS X—such books typically discuss the command-line interface to the operating system. Although these books play important roles in making Mac OS X accessible, the core architecture and implementation of Mac OS X and its components have remained mysteries. To make matters worse, besides the lack of information, it is common to find misinformation on the composition of

1. The date is interesting because coincidentally, Apple was founded on April 1, 1976.

2. I use the term “Unix” to represent one of a UNIX system, a UNIX-derived system, or a UNIX-like system.

Mac OS X. Consequently, the system is often misunderstood because of long-standing myths and stereotypes, or is perceived as a black box.

The purpose of this book is to deconstruct and demystify Mac OS X by describing it from a system-design perspective. It takes an implementation-oriented approach to understanding the system. Consider the example of interprocess communication (IPC). Mac OS X has numerous user-level and kernel-level IPC mechanisms, some of which are well known and documented. Rather than just showing how to use these mechanisms, the book explains the design and implementation of the most fundamental mechanism (Mach IPC), and then discusses how the other mechanisms are layered on top of one another. My goal is not to teach you how to do something specific—my goal is to provide you with enough knowledge and examples so that after reading the book, depending on your interests and background, you can build upon your newly gained knowledge and do what you choose.

Along with text, the book uses detailed illustrations, function call graphs, annotated code excerpts, and programming examples to present a detailed examination of Mac OS X. To keep the subject matter interesting and accessible—even to the casual reader—the book contains relevant trivia, digressions, and other tidbits.³

Who This Book Is For

I hope the book will be useful to anyone curious about the composition and working of Mac OS X.

Application programmers can gain a deeper understanding of how their applications interact with the system. System programmers can use the book as a reference and to construct a better picture of how the core system works. In my experience as a programmer, a solid understanding of system internals is immensely useful in design, development, and debugging. For example, you know what the system is capable of, what is feasible and what is not, what the “best” option is in a given situation, and what the plausible reasons are for certain program behavior. This book’s primary goal is to build a strong foundation for anyone who programs on Mac OS X.

Mac OS X users can read the book to better understand how the system is designed and how it comes together. System administrators and technical support staff will also find value in the book.

3. And a healthy dose of footnotes, too!

Besides those who use Mac OS X, the intended audience includes members of other technical communities, such as the BSD, Linux, and Windows communities. Given that many internal aspects of Mac OS X are radically different from these systems (for example, how the Mach kernel is used), the book will help such readers broaden their knowledge, and will assist them in comparing and contrasting Mac OS X with other operating systems.

The book will also be useful in an advanced operating systems course, particularly if you wish to do a case study on Mac OS X. The book is not suitable, however, as an introductory text. Most of the material is not presented at an introductory level, although I introduce many of the advanced topics with at least some background information.

The Book's Structure

Modern operating systems have become so large and complex that it is impossible to reasonably describe an entire system in a book. This book is somewhat ambitious in that it attempts to cover Mac OS X in substantial breadth *and* depth. The most important contributors to the book's depth are the carefully selected programming examples. The book is organized into twelve chapters. Although much of the book's content is rather technical, each chapter has sections that should be accessible to non-programmers.

Chapter 1, "Origins of Mac OS X," describes the technical history of Mac OS X and the systems it derives from. An unabridged version of Chapter 1, which covers all of Apple's past and present operating systems, is available on this book's accompanying web site.

Chapter 2, "An Overview of Mac OS X," is an eclectic tour of Mac OS X and its important features. It contains brief overviews of various layers that constitute the system.

Chapter 3, "Inside an Apple," describes the PowerPC architecture, using the PowerPC 970 ("G5") processor line as a specific example. It also discusses the PowerPC assembly language and calling conventions.

Chapter 4, "The Firmware and the Bootloader," describes both Open Firmware and the Extensible Firmware Interface (EFI), along with their respective bootloaders. It discusses the roles the firmware and the bootloader play in the system's operation, usage scenarios, and events that occur during early bootstrapping.

Chapter 5, "Kernel and User-Level Startup," describes the sequence of events—including initializations of kernel subsystems—from where the kernel

starts executing to the point where the first user-space program (launchd) is run by the kernel. The discussion includes launchd's function and implementation.

Chapter 6, "The xnu Kernel," describes the core kernel architecture of Mac OS X. The discussion includes system call families and their implementation, low-level tracing and debugging mechanisms, and special features such as the virtual machine monitor in the PowerPC version of the kernel.

Chapter 7, "Processes," describes abstractions such as tasks, threads, and processes, the various forms in which they exist in Mac OS X subsystems, and processor scheduling. The discussion includes using various kernel-level and user-level interfaces for manipulating the aforementioned abstractions.

Chapter 8, "Memory," describes the Mac OS X memory subsystem's architecture, including discussions of the Mach virtual memory architecture, paging, the unified buffer cache, the working-set detection mechanism, kernel-level and user-level memory allocators, and support for 64-bit addressing.

Chapter 9, "Interprocess Communication," describes various IPC and synchronization mechanisms available in Mac OS X. In particular, it discusses the implementation and usage of Mach IPC.

Chapter 10, "Extending the Kernel," describes the I/O Kit, the object-oriented driver subsystem in Mac OS X.

Chapter 11, "File Systems," describes the overall file system layer in Mac OS X, including brief discussions of each file system type. The discussion also covers partitioning schemes, disk management, and the Spotlight search technology.

Chapter 12, "The HFS Plus File System," describes the internals of the HFS Plus file system. The discussion is aided by the use of a custom file system debugger written for this chapter.

Appendix A, "Mac OS X on x86-Based Macintosh Computers," highlights the key differences between the x86-based and PowerPC-based versions of Mac OS X. Besides this appendix, the book covers the details of several key x86-specific topics, such as EFI, GUID-based partitioning, and Universal Binaries. *Most of Mac OS X is architecture-independent, and consequently, the majority of the book is architecture-independent.*

Given the book's length, I chose to exclude several topics that are well covered in other texts. The TCP/IP stack is an example—there is no "networking" chapter in the book since the Mac OS X TCP/IP stack is largely a derivative of the FreeBSD stack, which is already well documented. In general, information that is generic across Unix variants and can be found in standard texts is not included in this book.

How to Read This Book

Since the first two chapters provide the background and overall picture of Mac OS X, respectively, I recommend that you read these chapters first. The subsequent chapters are best read sequentially, although, depending on your interests and familiarity with the topics, you can skip certain sections (and perhaps even chapters) and still gain value from the book.

It will help if you have some familiarity with operating system concepts and have used a Unix operating system.

Given that the book has a large number of C programs and program excerpts, you should have some programming experience, and in particular, knowledge of the C programming language. I sometimes use code not only to demonstrate the working of a concept, but also to *describe* the concept. I realize that it is usually considered “difficult” to “read” code, and authors often expect that many readers will simply skip the code. My belief is that reading the code (as opposed to only running it) in this book will be particularly helpful to programmers.

Despite the book’s technical nature, several parts of the book can be read casually by both programmers and non-programmers.

I hope that as a reference on Mac OS X internals, the book and its examples will be useful to its readers for a long time to come.

How to Use the Examples

I have included a number of self-contained examples in the book. Many of these examples are non-trivial in that they do something that is both useful and interesting. I hope that these examples will serve as food for thought and building blocks for other projects. Almost all of the examples in the book are shown along with the command lines used to compile and run them.

The examples were tested on both PowerPC-based and x86-based Macintosh computers where applicable. It is interesting to note that in the cases where the code is PowerPC-only, say, in a PowerPC assembly language example, it can usually be both compiled and run on an x86-based Macintosh—such code will run under the Rosetta binary translation software. However, a small number of examples in the book require a PowerPC Macintosh—they will not run under Rosetta.

Related Material

Technology moves so fast these days that it is almost impossible to publish a fully up-to-date book. Thankfully, Internet access allows the author and publisher to

make various materials available to readers after the book is published. The most useful resource for this book is its accompanying web site, www.osxbook.com, which includes the following:

- Errata and updates
- Source code from the book
- The book's blog, with news and announcements about the availability of new material
- A set of discussion forums where topics related to the book (and to Mac OS X in general) can be discussed
- A bonus content section, which contains additional articles, presentations, binaries, and source code relevant to the book
- Sample content from the book, including a detailed table of contents