

<<计算机组成与设计>>

图书基本信息

书名：<<计算机组成与设计>>

13位ISBN编号：9787111302889

10位ISBN编号：7111302885

出版时间：2010-4

出版时间：机械工业出版社

作者：(美)帕特森//亨尼西

页数：689

版权说明：本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问：<http://www.tushu007.com>

前言

We believe that learning in computer science and engineering should reflect the current state of the field, as well as introduce the principles that are shaping computing. We also feel that readers in every specialty of computing need to appreciate the organizational paradigms that determine the capabilities, performance, and, ultimately, the success of computer systems. Modern computer technology requires professionals of every computing specialty to understand both hardware and software. The interaction between hardware and software at a variety of levels also offers a framework for understanding the fundamentals of computing. Whether your primary interest is hardware or software, computer science or electrical engineering, the central ideas in computer organization and design are the same. Thus, our emphasis in this book is to show the relationship between hardware and software and to focus on the concepts that are the basis for current computers. The recent switch from uniprocessor to multicore microprocessors confirmed the soundness of this perspective, given since the first edition. While programmers could once ignore that advice and rely on computer architects, compiler writers, and silicon engineers to make their programs run faster without change, that era is now over. For programs to run faster, they must become parallel. While the goal of many researchers is to make it possible for programmers to be unaware of the underlying parallel nature of the hardware they are programming, it will take many years to realize this vision. Our view is that for at least the next decade, most programmers are going to have to understand the hardware/software interface if they want programs to run efficiently on parallel computers. The audience for this book includes those with little experience in assembly language or logic design who need to understand basic computer organization as well as readers with backgrounds in assembly language and/or logic design who want to learn how to design a computer or understand how a system works and why it performs as it does.

<<计算机组成与设计>>

内容概要

这本最畅销的计算机组成书籍经过全面更新，关注现今发生在计算机体系结构领域的革命性变革：从单处理器发展到多核微处理器。

此外，出版这本书的ARM版是为了强调嵌入式系统对于全亚洲计算行业的重要性，并采用ARM处理器来讨论实际计算机的指令集和算术运算，因为ARM是用于嵌入式设备的最流行的指令集架构，而全世界每年约销售40亿个嵌入式设备。

与前几版一样，本书采用了一个MIPS处理器来展示计算机硬件技术、流水线、存储器层次结构以及I/O等基本功能。

此外，本书还包括一些关于x86架构的介绍。

本书主要特点 · 采用ARMv6 (ARM11系列) 为主要架构来展示指令系统和计算机算术运算的基本功能。

· 覆盖从串行计算到并行计算的革命性变革，新增了关于并行化的一章，并且每章中还有一些强调并行硬件和软件主题的小节。

· 新增一个由NVIDIA的首席科学家和架构主管撰写的附录，介绍了现代GPU的出现和重要性，首次详细描述了针对可视计算进行了优化的高度并行化、多线程、多核的处理器。

· 描述一种度量多核性能的独特方法——“Roofline model”，自带benchmark测试和分析AMD Opteron X4、Intel Xeon 5000、Sun UltraSPARC T2和 IBM Cell的性能。

· 涵盖了一些关于闪存和虚拟机的新内容。

· 提供了大量富有启发性的练习题，内容达200多页。

· 将AMD Opteron X4和Intel Nehalem作为贯穿本书的实例。

· 用SPEC CPU2006组件更新了所有处理器性能实例。

<<计算机组成与设计>>

作者简介

David A.Patterson, 加州大学伯克利分校计算机科学系教授。

美国国家工程研究院院士。

IEEE和ACM会士。

曾因成功的启发式教育方法被IEEE授予James H.Mulligan, Jr教育奖章。

他因为对RISC技术的贡献而荣获1995年IEEE技术成就奖, 而在RAID技术方面的成就为他赢得了1999年IEEE Reynold Johnson信息存储奖。

2000年他和John L.Hennessy分享了John von Neumann奖。

John L.Hennessy, 斯坦福大学校长, IEEE和ACM会士。

美国国家工程研究院院士及美国科学艺术研究院院士。

Hennessy教授因为在RISC技术方面做出了突出贡献而荣获2001年的Eckert-Mauchly奖章.他也是2001年Seymour Cray计算机工程奖得主。

并且和David A.Patterson分享了2000年John von Neumann奖。

<<计算机组成与设计>>

书籍目录

Contents Preface xv CHAPTERS Computer Abstractions and Technology 2 1.1 Introduction 3 1.2 Below Your Program 10 1.3 Under the Covers 13 1.4 Performance 26 1.5 The Power Wall 39 1.6 The Sea Change: The Switch from Uniprocessors to Multiprocessors 41 1.7 Real Stuff: Manufacturing and Benchmarking the AMD Opteron X4 44 1.8 Fallacies and Pitfalls 51 1.9 Concluding Remarks 54 1.10 Historical Perspective and Further Reading 55 1.11 Exercises 56 Instructions: Language of the Computer 74 2.1 Introduction 76 2.2 Operations of the Computer Hardware 77 2.3 Operands of the Computer Hardware 80 2.4 Signed and Unsigned Numbers 86 2.5 Representing Instructions in the Computer 93 2.6 Logical Operations 100 2.7 Instructions for Making Decisions 104 2.8 Supporting Procedures in Computer Hardware 113 2.9 Communicating with People 122 2.10 ARM Addressing for 32-Bit immediates and More Complex Addressing Modes 127 2.11 Parallelism and Instructions: Synchronization 133 2.12 Translating and Starting a Program 135 2.13 A C Sort Example to Put It All Together 143 : This icon identifies material on the CD 2.14 Arrays versus Pointers 152 2.15 Advanced Material: Compiling C and Interpreting Java 156 2.16 Real Stuff: MIPS Instructions 156 2.17 Real Stuff: x86 Instructions 161 2.18 Fallacies and Pitfalls 170 2.19 Concluding Remarks 171 2.20 Historical Perspective and Further Reading 174 2.21 Exercises 174 Arithmetic for Computers 214 3.1 Introduction 216 3.2 Addition and Subtraction 216 3.3 Multiplication 220 3.4 Division 226 3.5 Floating Point 232 3.6 Parallelism and Computer Arithmetic: Associativity 258 3.7 Real Stuff: Floating Point in the x86 259 3.8 Fallacies and Pitfalls 262 3.9 Concluding Remarks 265 3.10 Historical Perspective and Further Reading 268 3.11 Exercises 269 The Processor 284 4.1 Introduction 286 4.2 Logic Design Conventions 289 4.3 Building a Datapath 293 4.4 A Simple Implementation Scheme 302 4.5 An Overview of Pipelining 316 4.6 Pipelined Datapath and Control 330 4.7 Data Hazards: Forwarding versus Stalling 349 4.8 Control Hazards 361 4.9 Exceptions 370 4.10 Parallelism and Advanced Instruction-Level Parallelism 377 4.11 Real Stuff: the AMD Opteron X4 (Barcelona) Pipeline 390 4.12 Advanced Topic: an Introduction to Digital Design Using a Hardware Design Language to Describe and Model a Pipeline and More Pipelining Illustrations 392 4.13 Fallacies and Pitfalls 393 4.14 Concluding Remarks 394 4.15 Historical Perspective and Further Reading 395 4.16 Exercises 395 Large and Fast: Exploiting Memory Hierarchy 436 5.1 Introduction 438 5.2 The Basics of Caches 443 5.3 Measuring and Improving Cache Performance 461 5.4 Virtual Memory 478 5.5 A Common Framework for Memory Hierarchies 504 5.6 Virtual Machines 511 5.7 Using a Finite-State Machine to Control a Simple Cache 515 5.8 Parallelism and Memory Hierarchies: Cache Coherence 520 5.9 Advanced Material: Implementing Cache Controllers 524 5.10 Real Stuff: the AMD Opteron X4 (Barcelona) and Intel Nehalem Memory Hierarchies 525 5.11 Fallacies and Pitfalls 529 5.12 Concluding Remarks 533 5.13 Historical Perspective and Further Reading 534 5.14 Exercises 534 Storage and Other I/O Topics 554 6.1 Introduction 556 6.2 Dependability, Reliability, and Availability 559 6.3 Disk Storage 561 6.4 Flash Storage 566 6.5 Connecting Processors, Memory, and I/O Devices 568 6.6 Interfacing I/O Devices to the Processor, Memory, and Operating System 572 6.7 I/O Performance Measures: Examples from Disk and File Systems 582 6.8 Designing an I/O System 584 6.9 Parallelism and I/O: Redundant Arrays of Inexpensive Disks 585 6.10 Real Stuff: Sun Fire x4150 Server 592 6.11 Advanced Topics: Networks 598 6.12 Fallacies and Pitfalls 599 6.13 Concluding Remarks 603 6.14 Historical Perspective and Further Reading 604 6.15 Exercises 605 Multicores, Multiprocessors, and Clusters 616 7.1 Introduction 618 7.2 The Difficulty of Creating Parallel Processing Programs 620 7.3 Shared Memory Multiprocessors 624 7.4 Clusters and Other Message-Passing Multiprocessors 627 7.5 Hardware Multithreading 631 7.6 SISD, MIMD, SIMD, SPMD, and Vector 634 7.7 Introduction to Graphics Processing Units 640 7.8 Introduction to Multiprocessor Network Topologies 646 7.9 Multiprocessor Benchmarks 650 7.10 Roofline: A Simple Performance Model 653 7.11 Real Stuff: Benchmarking Four Multicores Using the Roofline Model 661 7.12 Fallacies and Pitfalls 670 7.13 Concluding Remarks 672 7.14 Historical Perspective and Further Reading 674 7.15 Exercises 674 Index I-1 CD-ROM CONTENT Graphics and Computing GPUs A-2 A.1 Introduction A-3 A.2 GPU System Architectures A-7 A.3 Scalable Parallelism – Programming GPUs A-12 A.4 Multithreaded Multiprocessor Architecture A-25 A.5 Parallel Memory System G.6 Floating Point A-36 A.6 Floating Point Arithmetic A-41 A.7 Real Stuff: The NVIDIA GeForce 8800 A-46 A.8 Real

<<计算机组成与设计>>

Stuff: Mapping Applications to GPUs A-55 A.9 Fallacies and Pitfalls A-72 A.10 Concluding Remarks A-76 A.11 Historical Perspective and Further Reading A-77 ARM and Thumb Assembler Instructions B1-2 B1.1 Using This Appendix B1-3 B1.2 Syntax B1-4 B1.3 Alphabetical List of ARM and Thumb Instructions B1-8 B1.4 ARM Assembler Quick Reference B1-49 B1.5 GNU Assembler Quick Reference B1-60 ARM and Thumb Instruction Encodings B2-2 B2.1 ARM Instruction Set Encodings B2-3 B2.2 Thumb Instruction Set Encodings B2-9 B2.3 Program Status Registers B2-11 Instruction Cycle Timings B3-2 B3.1 Using the Instruction Set Cycle Timing Tables B3-3 B3.2 ARM7TDMI Instruction Cycle Timings B3-5 B3.3 ARM9TDMI Instruction Cycle Timings B3-6 B3.4 StrongARM1 Instruction Cycle Timings B3-8 B3.5 ARM9E Instruction Cycle Timings B3-9 B3.6 ARM10E Instruction Cycle Timings B3-11 B3.7 Intel XScale Instruction Cycle Timings B3-12 B3.8 ARM11 Cycle Timings B3-14 C The Basics of Logic Design C-2 C.1 Introduction C-3 C.2 Gates, Truth Tables, and Logic Equations C-4 C.3 Combinational Logic C-9 C.4 Using a Hardware Description Language C-20 C.5 Constructing a Basic Arithmetic Logic Unit C-26 C.6 Faster Addition: Carry Lookahead C-38 C.7 Clocks C-48 C.8 Memory Elements: Flip-Flops, Latches, and Registers C-50 C.9 Memory Elements: SRAMs and DRAMs C-58 C.10 Finite-State Machines C-67 C.11 Timing Methodologies C-72 C.12 Field Programmable Devices C-78 C.13 Concluding Remarks C-79 C.14 Exercises C-80 D Mapping Control to Hardware D-2 D.1 Introduction D-3 D.2 Implementing Combinational Control Units D-4 D.3 Implementing Finite-State Machine Control D-8 D.4 Implementing the Next-State Function with a Sequencer D-22 D.5 Translating a Microprogram to Hardware D-28 D.6 Concluding Remarks D-32 D.7 Exercises D-33 ADVANCED CONTENT Section 2.15 Compiling C and Interpreting Java Section 4.12 An Introduction to Digital Design Using a Hardware Design Language to Describe and Model a Pipeline and More Pipelining Illustrations Section 5.9 Implementing Cache Controllers Section 6.11 Networks HISTORICAL PERSPECTIVES & FURTHER READING Chapter 1 Computer Abstractions and Technology: Section 1.10 Chapter 2 Instructions: Language of the Computer: Section 2.20 Chapter 3 Arithmetic for Computers: Section 3.10 Chapter 4 The Processor: Section 4.15 Chapter 5 Large and Fast: Exploiting Memory Hierarchy: Section 5.13 Chapter 6 Storage and Other I/O Topics: Section 6.14 Chapter 7 Multicores, Multiprocessors, and Clusters: Section 7.14 Appendix A Graphics and Computing GPUs: Section A.11 TUTORIALS VHDL Verilog SOFTWARE Xilinx FPGA Design, Simulation and Synthesis Software QEMU <http://www.nongnu.org/qemu/about.html> Glossary G-1 Index I-1 Further Reading FR-1

章节摘录

插图：Diameters of hard disks vary by more than a factor of 3 today, from 1 inch to 35 inches, and have been shrunk over the years to fit into new products; workstation servers, personal computers, laptops, palmtops, and digital cameras have all inspired new disk form factors. Traditionally, the widest disks have the highest performance and the smallest disks have the lowest unit cost. The best cost per gigabyte varies. Although most hard drives appear inside computers, as in Figure 17, hard drives can also be attached using external interfaces such as universal serial bus (USB). The use of mechanical components means that access times for magnetic disks are much slower than for DRAMs: disks typically take 5-20 milliseconds, while DRAMs take 50-70 nanoseconds—making DRAMs about 100,000 times faster. Yet disks have much lower costs than DRAM for the same storage capacity, because the production costs for a given amount of disk storage are lower than for the same amount of integrated circuit. In 2008, the cost per gigabyte of disk is 30 to 100 times less expensive than DRAM. Thus, there are three primary differences between magnetic disks and main memory: disks are nonvolatile because they are magnetic; they have a slower access time because they are mechanical devices; and they are cheaper per gigabyte because they have very high storage capacity at a modest cost. Many have tried to invent a technology cheaper than DRAM but faster than disk to fill that gap, but many have failed. Challengers have never had a product to market at the right time. By the time a new product would ship, DRAMs and disks had continued to make rapid advances, costs had dropped accordingly, and the challenging product was immediately obsolete. Flash memory, however, is a serious challenger. This semiconductor memory is nonvolatile like disks and has about the same bandwidth, but latency is 100 to 1000 times faster than disk. Flash is popular in cameras and portable music players because it comes in much smaller capacities, it is more rugged, and it is more power efficient than disks, despite the cost per gigabyte in 2008 being about 6 to 10 times higher than disk. Unlike disks and DRAM, flash memory bits wear out after 100,000 to 1,000,000 writes. Thus, file systems must keep track of the number of writes and have a strategy to avoid wearing out storage, such as by moving popular data. Chapter 6 describes flash in more detail. Although hard drives are not removable, there are several storage technologies in use that include the following: Optical disks, including both compact disks (CDs) and digital video disks (DVDs), constitute the most common form of removable storage. The Blu-Ray (BD) optical disk standard is the heir-apparent to DVD. Flash-based removable memory cards typically attach to a USB connection and are often used to transfer files. Magnetic tape provides only slow serial access and has been used to back up disks, a role now often replaced by duplicate hard drives.

<<计算机组成与设计>>

媒体关注与评论

“本版特别之处在于采用ARM取代了早先使用MIPS作为核心处理器来讲述计算机设计的根本原则，约本书增加了另一个层面的内涵。

ARM作为嵌入式领域的量流处理器，在嵌入式计算领域具有非常重要的意义。

本书弥补了现有数学体系中的空白，即有针对性地向学习嵌入式系统的学生讲授计算机组成的基本原理。

同以往版本一样，本书仍主要介绍计算机硬件/软件接口，弄巧妙地将其与嵌入式系统设计的基本知识相联系。

” ——Ranjani Parthasarathi Anna大学，印度钦奈

<<计算机组成与设计>>

编辑推荐

《计算机组成与设计:硬件/软件接口(英文版·第4版·ARM版)》为经典原版书库。

<<计算机组成与设计>>

版权说明

本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问:<http://www.tushu007.com>