

《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

出版社：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以及在线试读，请支持正版图书。

更多资源请访问：www.tushu111.com

《计算机组成与设计》

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

《计算机组成与设计:硬件/软件接口(英文版·第4版·ARM版)》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,《计算机组成与设计:硬件/软件接口(英文版·第4版·ARM版)》还包括一些关于x86架构的介绍。

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。

采用ARMv6(ARM 11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 Ultra SPARC 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《计算机组成与设计:硬件/软件接口(英文版·第4版·ARM版)》的实例。

用SPEC CPU2006组件更新了所有处理器性能实例。

海报:

《计算机组成与设计》

作者简介

John L. Hennessy 斯坦福大学校长，IEEE和DACM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，因在RISC技术方面做出的突出贡献而荣获2001年的Eckert—Mauchly奖。同时他也是2001年SeymourCray计算机工程奖得主，并且和本书另一位作者David A. Paterson分享了2000年的IEEE John von Neumann奖。

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会员，曾因成功的启发式教育方法被IEEE授予James H. M. Aligan, 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。在RAID技术方面的成就为他赢得了1999年IEEE Reyrloid Johnson信息存储奖。2000年他和John L. Hennessy分享了冯·诺伊曼奖。

Contents Preface CHAPTERS

1 Computer Abstractions and Technology 1.1 Introduction 1.2 Below Your Program 1.3 Under the Covers 1.4 Performance 1.5 The Power Wall 1.6 The Sea Change: The Switch from Uniprocessors to Multiprocessors 1.7 Real Stuff: Manufacturing and Benchmarking the AMD Opteron X4 1.8 Fallacies and Pitfalls 1.9 Concluding Remarks 1.10 Historical Perspective and Further Reading 1.11 Exercises

2 Instructions: Language of the Computer 2.1 Introduction 2.2 Operations of the Computer Hardware 2.3 Operands of the Computer Hardware 2.4 Signed and Unsigned Numbers 2.5 Representing Instructions in the Computer 2.6 Logical Operations 2.7 Instructions for Making Decisions 2.8 Supporting Procedures in Computer Hardware 2.9 Communicating with People 2.10 ARM Addressing for 32-Bit immediates and More Complex Addressing Modes 2.11 Parallelism and Instructions: Synchronization 2.12 Translating and Starting a Program 2.13 A C Sort Example to Put It All Together : This icon identifies material on the 2.14 Arrays versus Pointers 2.15 Advanced Material: Compiling C and Interpreting Java 2.16 Real Stuff: MIPS Instructions 2.17 Real Stuff: x86 Instructions 2.18 Fallacies and Pitfalls 2.19 Concluding Remarks 2.20 Historical Perspective and Further Reading 2.21 Exercises

3 Arithmetic for Computers 3.1 Introduction 3.2 Addition and Subtraction 3.3 Multiplication 3.4 Division 3.5 Floating Point 3.6 Parallelism and Computer Arithmetic: Associativity 3.7 Real Stuff: Floating Point in the x86 3.8 Fallacies and Pitfalls 3.9 Concluding Remarks 3.10 Historical Perspective and Further Reading 3.11 Exercises

4 The Processor 4.1 Introduction 4.2 Logic Design Conventions 4.3 Building a Datapath 4.4 A Simple Implementation Scheme 4.5 An Overview of Pipelining 4.6 Pipelined Datapath and Control 4.7 Data Hazards: Forwarding versus Stalling 4.8 Control Hazards 4.9 Exceptions 4.10 Parallelism and Advanced Instruction-Level Parallelism 4.11 Real Stuff: the AMD Opteron X4 (Barcelona) Pipeline 4.12 Advanced Topic: an Introduction to Digital Design Using a Hardware Design Language to Describe and Model a Pipeline and More Pipelining Illustrations 4.13 Fallacies and Pitfalls 4.14 Concluding Remarks 4.15 Historical Perspective and Further Reading 4.16 Exercises

5 Large and Fast: Exploiting Memory Hierarchy 5.1 Introduction 5.2 The Basics of Caches 5.3 Measuring and Improving Cache Performance 5.4 Virtual Memory 5.5 A Common Framework for Memory Hierarchies 5.6 Virtual Machines 5.7 Using a Finite-State Machine to Control a Simple Cache 5.8 Parallelism and Memory Hierarchies: Cache Coherence 5.9 Advanced Material: Implementing Cache Controllers 5.10 Real Stuff: the AMD Opteron X4 (Barcelona) and Intel Nehalem Memory Hierarchies 5.11 Fallacies and Pitfalls 5.12 Concluding Remarks 5.13 Historical Perspective and Further Reading 5.14 Exercises

6 Storage and Other I/O Topics 6.1 Introduction 6.2 Dependability, Reliability, and Availability 6.3 Disk Storage 6.4 Flash Storage 6.5 Connecting Processors, Memory, and I/O Devices 6.6 Interfacing I/O Devices to the Processor, Memory, and Operating System 6.7 I/O Performance Measures: Examples from Disk and File Systems 6.8 Designing an I/O System 6.9 Parallelism and I/O: Redundant Arrays of Inexpensive Disks 6.10 Real Stuff: Sun Fire x4150 Server 6.11 Advanced Topics: Networks 6.12 Fallacies and Pitfalls 6.13 Concluding Remarks 6.14 Historical Perspective and Further Reading 6.15 Exercises

7 Multicores, Multiprocessors, and Clusters 7.1 Introduction 7.2 The Difficulty of Creating Parallel Processing Programs 7.3 Shared Memory Multiprocessors 7.4 Clusters and Other Message-Passing Multiprocessors 7.5 Hardware Multithreading 7.6 SISD, MIMD, SIMD, SPMD, and Vector 7.7 Introduction to Graphics Processing Units 7.8 Introduction to Multiprocessor Network Topologies 7.9 Multiprocessor Benchmarks 7.10 Roofline: A Simple Performance Model 7.11 Real Stuff: Benchmarking Four Multicores Using the Roofline Model 7.12 Fallacies and Pitfalls 7.13 Concluding Remarks 7.14 Historical Perspective and Further Reading 7.15 Exercises

Index I-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。

精彩短评

- 1、概念性的东西
- 2、书本身是本好书，不过还是觉得字太小，能有电子版的该多好
- 3、都是些基础学习的书，感兴趣的可以试试
- 4、对于想理解计算机底层硬件组织的同学来说，是一部不错的入门书
- 5、这个版本最好。1)mips版cd配套网上有，对比后，arm版appendix比mips要多几个~~；2)网上已有mips版的revised。对比后，主要第二章不同，其他都一样~~3)GPU的附录确实没在cd里，华章网站上有提供下载...不知再版后有无加入cd，没得化就应该谴责，确实让人很不快滴；... 阅读更多
- 6、买的书看上去不错的，不过对我来说买了本英语版的的确有点小失误了，毕竟我的英语水平不是很高，不过看原版的还是挺好的了.....
- 7、CD-ROM CONTENT没包含附录A的全部东西！
- 8、总之我买错了。。这是ARM版，原来旧版的是MIPS版，但是如今MIPS版不在亚洲地区卖了，所以大家看清楚！！
- 9、大神写的经典教材
- 10、还不错，顺便练习一下英文。
- 11、计算机组成经典教材
- 12、风格大福给对方给第三方
- 13、打折的时候买的，47.5，就是纸张有点暗，灰色调的感觉不太适应。其他的还好，准备好好学学。
- 14、还行, 讲得清楚, 不过废话稍多了一点
- 15、计算机体系结构的经典教材
- 16、附录只有A11，没有A1~A10，原书900多页，这个只有600页，印刷质量非常差，机械工业出版社真是臭名昭著啊，怒火中烧中！谁买谁后悔！真是焚琴煮鹤啊！
- 17、书写的很好，字小了点，看的很累
- 18、挺好的 还是学到点东西。英文的东西照搬到中文，确实在理解和进一步提高上有了很多障碍，本来很简单的从生活中的单词来的计算机术语，到了中文又冒出那么多的莫名其妙的概念。--是个大问题
- 19、本来想买mips版的
- 20、只有A.11，其他的A.1-A.10都没有，本来就是冲A买的
- 21、书的质量很好，包装也很仔细。
- 22、经典好书的有一本

《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:www.tushu111.com