

国外优秀信息科学与技术系列教学用书

计算机 系统结构基础

(影印版)

ESSENTIALS OF COMPUTER
ARCHITECTURE

■ Douglas E. Comer



高等教育出版社
Higher Education Press

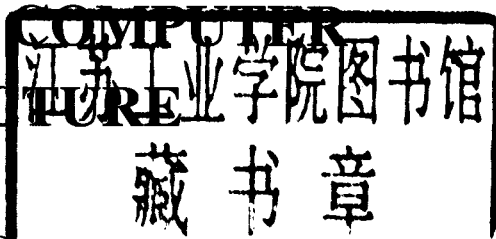
国外优秀信息科学与技术系列教学用书

计算机 系统结构基础

(影印版)

ESSENTIALS OF COMPUTER
ARCHITECTURE

Douglas E. Comer



高等教育出版社

图字:01 - 2005 - 0964 号

Essentials of Computer Architecture

Douglas E. Comer

本书封面贴有 Pearson Education(培生教育出版集团)激光防伪标签。无标签者不得销售。

English reprint edition copyright © 2005 by **PEARSON EDUCATION ASIA LIMITED** and **HIGHER EDUCATION PRESS**. (Essentials of Computer Architecture from Pearson Education's edition of the Work)

Essentials of Computer Architecture, le by *Douglas E. Comer*, Copyright © 2005.

All Rights Reserved.

Published by arrangement with the original publisher, Pearson Education, Inc., publishing as Pearson Education, Inc..

This edition is authorized for sale only in the People's Republic of China (excluding the Special Administrative Regions of Hong Kong and Macau).

原版 ISBN:0 - 13 - 149179 - 2

For sale and distribution in the People's Republic of China exclusively (except Taiwan, Hong Kong SAR and Macao SAR).

仅限于中华人民共和国境内(但不允许在中国香港、澳门特别行政区和中国台湾地区)销售发行。

图书在版编目(CIP)数据

计算机系统结构基础 = Essentials of Computer Architecture / (美)科默(Comer, D. E.). —影印本.

北京:高等教育出版社, 2005. 12

ISBN 7 - 04 - 017816 - 8

I. 计... II. 科... III. 计算机体系结构 - 高等学校 - 教材 - 英文 IV. TP303

中国版本图书馆 CIP 数据核字(2005)第 132966 号

出版发行	高等教育出版社	购书热线	010 - 58581118
社 址	北京市西城区德外大街 4 号	免费咨询	800 - 810 - 0598
邮政编码	100011	网 址	http://www.hep.edu.cn
总 机	010 - 58581000		http://www.hep.com.cn
经 销	北京蓝色畅想图书发行有限公司	网上订购	http://www.landaco.com
印 刷	北京外文印刷厂		http://www.landaco.com.cn
开 本	787 × 1092 1/16	版 次	2005 年 12 月第 1 版
印 张	25	印 次	2005 年 12 月第 1 次印刷
字 数	480 000	定 价	38.00 元

本书如有缺页、倒页、脱页等质量问题,请到所购图书销售部门联系调换。

版权所有 侵权必究

物料号 17816-00

出版说明

20 世纪末,以计算机和通信技术为代表的信息科学和技术对世界经济、科技、军事、教育和文化等产生了深刻影响。信息科学技术的迅速普及和应用,带动了世界范围信息产业的蓬勃发展,为许多国家带来了丰厚的回报。

进入 21 世纪,尤其随着我国加入 WTO,信息产业的国际竞争将更加激烈。我国信息产业虽然在 20 世纪末取得了迅猛发展,但与发达国家相比,甚至与印度、爱尔兰等国家相比,还有很大差距。国家信息化的发展速度和信息产业的国际竞争能力,最终都将取决于信息科学技术人才的质量和数量。引进国外信息科学与技术优秀教材,在有条件的学校推动开展英语授课或双语教学,是教育部为加快培养大批高质量的信息技术人才采取的一项重要举措。

为此,教育部要求由高等教育出版社首先开展信息科学与技术教材的引进试点工作。同时提出了两点要求,一是要高水平,二是要低价格。在高等教育出版社和信息科学技术引进教材专家组的努力下,经过比较短的时间,第一批由教育部高等教育司推荐的 20 多种引进教材已经陆续出版。这套教材出版后受到了广泛的好评,其中有不少是世界信息科学技术领域著名专家、教授的经典之作和反映信息科学技术最新进展的优秀作品,代表了目前世界信息科学技术教育的一流水平,而且价格也是最优惠的,与国内同类自编教材相当。这套教材基本覆盖了计算机科学与技术专业的课程体系,体现了权威性、系统性、先进性和经济性等特点。

在引进教材的同时,我们还应做好消化吸收,注意学习国外先进的教学思想和教学方法,提高自编教材的水平,使我们的教学和教材在内容体系上,在理论与实践的结合上,在培养学生的动手能力上能有较大的突破和创新。

希望这些教学用书的引进出版,对于提高我国高等学校信息科学技术的教学水平,缩小与国际先进水平的差距,加快培养一大批具有国际竞争力的高质量信息技术人才,起到积极的推动作用。同时也欢迎广大教师和专家们对教材引进工作提出宝贵的意见和建议。联系方式: hep.cs@263.net。

高等教育出版社
二〇〇四年十月

前言

当我被派去帮忙拯救本科生的计算机系统结构这门课程时，便开始酝酿本书的编写了。计算机系统结构课程遭受了多年的忽视：虽然很多教授和访问学者都讲授过这门课程，但他们往往对数字硬件缺乏兴趣，或缺乏相关的知识背景，从而使这门课程已经退化成各种硬件体系结构相关知识的大杂烩。在一些学期里，学生们将一门课程的所有时间都花在学习布尔代数上，却毫不涉及实际的硬件；而在另外一些学期里，学生们只学习某一种汇编语言晦涩难懂的细节，却完全不了解其他的语言。

那么，计算机系统结构这门课程还值得保留吗？绝对值得保留！在众多计算机科学的教学计划中，计算机系统结构这门课程是学生们了解有关计算机结构的基本概念的唯一机会。对硬件的理解使学生能够编写效率更高、错误更少的程序。广义上讲，体系结构的基本知识有助于理解编程选择的因果关系，从而帮助程序员提高程序效率。此外，了解硬件的工作原理使程序员能够迅速找出程序出错的根源，从而改进编程过程。最后，毕业生们必须理解体系结构的基本概念，才能通过像 Intel 和 Microsoft 这样的公司的招聘测试。

拯救系统结构这门课程的步骤之一就是寻找一本合适的教材。我们发现，现有的教材大致可以分为两类：一类面向将来会从事硬件设计的工程类低年级学生；另一类则面向计算机科学系的学生，并试图囊括编译器、操作系统和网络协议（至少一种）等诸多内容。这些教材都不适合作为单独的入门级计算机系统结构课程的教材。我们需要这样一本教材：（1）重点介绍概念而非工程细节（因为我们的学生并不会去深究硬件设计）；（2）从程序员的角度来阐述系统结构这一主题，并强调其对程序员的重要意义；（3）不应涵盖多门课程的内容。当我们无法找到这样的教材时，唯一的办法就是编写一本。

本书主要分为五个部分。第 I 部分讲述数字逻辑、逻辑门和数据表示的基础知识，重点是数据表示这一章，这是因为二进制补码运算和整数值域的概念是编程的基础。第 II、第 III 和第 IV 部分分别介绍体系结构的三个基本组件：处理器、存储器和 I/O 系统。各章都为学生提供了充足的背景知识，以便理解系统操作机制及其对于程序员的意义。最后，第 V 部分讨论了一些高级论题，例如并行处理、流水线技术以及性能。

附录主要介绍此课程的一个重要组成部分：实验，学生们可以通过实验进行学习。虽然大多数实验问题都要求通过编程解决，但在开始的几个星期内，学生们应该在实验室里将一些门电路焊接在电路板上。这些器材并不贵（平均每个学生花费不到 15 美元，用于购买固定设备；每个学生花费不到 20 美元就可以买到一套属于自己的芯片）。

我们已经建立了本书的辅助网站：

<http://www.eca.cs.purdue.edu>

Rajesh Subraman 已经同意维护该网站，其中包括了一系列由作者与 Rajesh 分别编写的课堂教学材料。我们还邀请了其他教师贡献出自己的教学材料。

本书和实验练习已经在普度大学采用，学生们给予了极高的评价。我们收到了许多关于本书及课程的感谢留言。许多学生在实验室第一次接触到了硬件，并且表现出极大的热情。

我要感谢许多对此书的出版做出贡献的人。Bernd Wolfinger 对此书进行了广泛的审阅，并对本书的主题和方向提出了诸多重要的建议。Dan Ardelean、James Cernak 和 Tim Korb 对许多章节进行了详细的评论。Dave Capka 审阅了前面几章的内容。Rajesh Subraman 采用本书进行教学并对书中内容提出了自己的建议。普度大学计算机科学系 250 班的下列学生分别发现了手稿中的一个或多个错误：Nitin Alreja、Alex Cox、David Ehrmann、Roger Maurice Elion、Andrew Lee、Stan Luban、Andrew L. Soderstrom 和 Brandon Wuest。

最后，我要感谢我的妻子 Chris，她耐心而仔细的编辑工作以及颇有见地的建议，为本书增色不少。

Douglas E. Comer

2004 年 6 月

关于作者

Douglas Comer 博士在计算机系统领域有着广泛的知识背景，而且一直从事着软件和硬件的教学与研究。Comer 在软件方面的研究工作横跨计算机系统的各个领域，包括编译器和操作系统。他构建过一个完整的操作系统，包括进程管理器、内存管理器以及串行接口和并行接口的设备驱动程序。Comer 也为常规计算机和网络处理器编写过网络协议软件和网络设备驱动程序。他开发的操作系统 Xinu 和 TCP/IP 协议栈已经被应用于商业产品中。

Comer 的硬件工作经验包括设计离散元件、逻辑门电路以及基本的硅工艺经验。他已经编写了几本畅销的关于网络处理器体系结构的教材。在贝尔实验室，Comer 研究过 VLSI 设计并设计了一种 VLSI 芯片。

Comer 是普度大学计算机科学系的著名教授，他致力于计算机系统结构、操作系统、网络和 Internet 方面的开发、教学和研究。Comer 建立了几个创新实验室，使得学生可以建立和测试各种系统，例如操作系统和 IP 路由器；Comer 的所有课程都包括实验。他还不断地在世界各地的大学、工业界和会议上进行咨询和演讲。

除了编写一系列享誉世界的计算机操作系统、网络、TCP/IP 和计算机技术等方面的书籍外，Comer 还兼任 *Software: Practice and Experience* 杂志的主编。他是 ACM 成员、普度大学教学学会成员，获得过无数嘉奖，其中包括 USENIX 终身成就奖（USENIX Lifetime Achievement Award）。

关于其本人的更多信息可查阅网站：

www.cs.purdue.edu/people/comer

关于 Comer 著作的信息可查阅网站：

www.comerbooks.com

Preface

This book began when I was assigned to help salvage an undergraduate computer organization course. The course had suffered years of neglect: it had been taught by a series of professors, mostly visitors, who had little or no interest or background in digital hardware, and the curriculum had deteriorated to a potpourri of topics that were only loosely related to hardware architectures. In some semesters, students spent the entire class studying Boolean Algebra, without even the slightest connection to actual hardware. In others, students learned the arcane details of one particular assembly language, without a notion of alternatives.

Is a computer organization course worth saving? Absolutely! In many Computer Science programs, the computer organization course is the only time students are exposed to fundamental concepts that explain the structure of the computer they are programming. Understanding the hardware makes it possible to construct programs that are more efficient and less prone to errors. In a broad sense, a basic knowledge of architecture helps programmers improve program efficiency by understanding the consequences of programming choices. Knowing how the hardware works can also improve the programming process by allowing programmers to pinpoint the source of bugs quickly. Finally, graduates need to understand basic architectural concepts to pass job application tests given by firms like Intel and Microsoft.

One of the steps in salvaging our architecture course consisted in looking at textbooks. We discovered the texts could be divided into roughly two types: texts aimed at beginning engineering students who would go on to design hardware, and texts written for CS students that attempt to include topics from compilers, operating systems, and (in at least one case) a complete explanation of how Internet protocols operate. Neither approach seemed appropriate for a single, introductory course on the subject. We wanted a book that (1) focused on the concepts rather than engineering details (because our students are not focused on hardware design); (2) explained the subject from a programmer's point of view, and emphasized consequences for programmers; and (3) did not try to cover several courses' worth of material. When no text was found, it seemed that the only solution was to create one.

The text is divided into five parts. Part 1 covers the basics of digital logic, gates, and data representation. We emphasize the representation chapter because notions of two's-complement arithmetic and ranges of integer values are essential in programming. Parts 2, 3, and 4 cover the three essential areas of architecture: processors, memories, and I/O systems. In each case, the chapters give students enough background to under-

stand how the mechanisms operate and the consequences for programmers. Finally, Part 5 covers advanced topics like parallelism, pipelining, and performance.

An Appendix describes an important aspect of the course: a hands-on lab where students can learn by doing. Although most lab problems focus on programming, students should spend the first few weeks in lab wiring a few gates on a breadboard. The equipment is inexpensive (we spent less than fifteen dollars per student on permanent equipment; students purchase their own set of chips for under twenty dollars).

We have set up a web site to accompany the book at:

<http://www.eca.cs.purdue.edu>

Rajesh Subraman has agreed to manage the site, which contains a set of class presentation materials created by the author as well as a set created by Rajesh. We invite other instructors to contribute their materials.

The text and lab exercises have been used at Purdue; students have been extremely positive about both. We received notes of thanks for the text and course. For many students, the lab is their first experience with hardware, and they are enthusiastic.

My thanks to the many individuals who contributed to the book. Bernd Wolfinger provided extensive reviews and made several important suggestions about topics and direction. Dan Ardelean, James Cernak, and Tim Korb gave detailed comments on many chapters. Dave Capka reviewed early chapters. Rajesh Subraman taught from the book and provided his thoughts about the content. In the CS 250 class at Purdue, the following students each identified one or more typos in the manuscript: Nitin Alreja, Alex Cox, David Ehrmann, Roger Maurice Elion, Andrew Lee, Stan Luban, Andrew L. Soderstrom, and Brandon Wuest.

Finally, I thank my wife, Chris, for her patient and careful editing and valuable suggestions that improve and polish each book.

Douglas E. Comer

June, 2004

About The Author

Dr. Douglas Comer has an extensive background in computer systems, and has worked with both hardware and software. Comer's work on software spans most aspects of systems, including compilers and operating systems. He created a complete operating system, including a process manager, a memory manager, and device drivers for both serial and parallel interfaces. Comer has also implemented network protocol software and network device drivers for conventional computers and network processors. Both his operating system, Xinu, and TCP/IP protocol stack have been used in commercial products.

Comer's experience with hardware includes work with discrete components, building circuits from logic gates, and experience with basic silicon technology. He has written popular textbooks on network processor architectures, and at Bell Laboratories, Comer studied VLSI design and fabricated a VLSI chip.

Comer is a Distinguished Professor of Computer Science at Purdue University, where he develops and teaches courses and does research on computer organization, operating systems, networks, and Internets. Comer has created innovative laboratories in which students can build and measure systems such as operating systems and IP routers; all of Comer's courses include hands-on lab work. He continues to consult and lecture at universities, industries, and conferences around the world.

In addition to writing a series of internationally acclaimed technical books on computer operating systems, networks, TCP/IP, and computer technologies, Comer serves as the editor-in-chief of the journal *Software — Practice and Experience*. He is a Fellow of the ACM, a Fellow of the Purdue Teaching Academy, and a recipient of numerous awards, including a Usenix Lifetime Achievement award.

Additional information can be found at:

www.cs.purdue.edu/people/comer

and information about Comer's books can be found at:

www.comerbooks.com

Contents

Preface	xxi
----------------	------------

Chapter 1 Introduction And Overview	1
--	----------

- 1.1 The Importance Of Architecture* 1
- 1.2 Learning The Essentials* 1
- 1.3 Organization Of The Text* 2
- 1.4 What We Will Omit* 3
- 1.5 Terminology: Architecture And Design* 3
- 1.6 Summary* 3

PART I Basics

Chapter 2 Fundamentals Of Digital Logic	7
--	----------

- 2.1 Introduction* 7
- 2.2 Electrical Terminology: Voltage And Current* 7
- 2.3 The Transistor* 8
- 2.4 Logic Gates* 9
- 2.5 Symbols Used For Gates* 10
- 2.6 Construction Of Gates From Transistors* 11
- 2.7 Example Interconnection Of Gates* 12
- 2.8 Multiple Gates Per Integrated Circuit* 14
- 2.9 The Need For More Than Combinatorial Circuits* 15
- 2.10 Circuits That Maintain State* 15
- 2.11 Transition Diagrams* 16
- 2.12 Binary Counters* 17
- 2.13 Clocks And Sequences* 18
- 2.14 The Important Concept Of Feedback* 20
- 2.15 Starting A Sequence* 22
- 2.16 Iteration In Software Vs. Replication In Hardware* 22
- 2.17 Gate And Chip Minimization* 23
- 2.18 Using Spare Gates* 24

2.19	<i>Power Distribution And Heat Dissipation</i>	24
2.20	<i>Timing</i>	25
2.21	<i>Physical Size And Process Technologies</i>	26
2.22	<i>Circuit Boards And Layers</i>	27
2.23	<i>Levels Of Abstraction</i>	27
2.24	<i>Summary</i>	28

Chapter 3 Data And Program Representation

29

3.1	<i>Introduction</i>	29
3.2	<i>Digital Logic And Abstraction</i>	29
3.3	<i>Bits And Bytes</i>	30
3.4	<i>Byte Size And Possible Values</i>	30
3.5	<i>Binary Arithmetic</i>	31
3.6	<i>Hexadecimal Notation</i>	32
3.7	<i>Notation For Hexadecimal And Binary Constants</i>	33
3.8	<i>Character Sets</i>	34
3.9	<i>Unicode</i>	35
3.10	<i>Unsigned Integers, Overflow, And Underflow</i>	35
3.11	<i>Numbering Bits And Bytes</i>	36
3.12	<i>Signed Integers</i>	37
3.13	<i>An Example Of Two's Complement Numbers</i>	38
3.14	<i>Sign Extension</i>	39
3.15	<i>Floating Point</i>	40
3.16	<i>Special Values</i>	42
3.17	<i>Range Of IEEE Floating Point Values</i>	42
3.18	<i>Data Aggregates</i>	42
3.19	<i>Program Representation</i>	43
3.20	<i>Summary</i>	43

PART II Processors

Chapter 4 The Variety Of Processors And Computational Engines

47

4.1	<i>Introduction</i>	47
4.2	<i>Von Neumann Architecture</i>	47
4.3	<i>Definition Of A Processor</i>	48
4.4	<i>The Range Of Processors</i>	48
4.5	<i>Hierarchical Structure And Computational Engines</i>	49
4.6	<i>Structure Of A Conventional Processor</i>	51
4.7	<i>Definition Of An Arithmetic Logic Unit (ALU)</i>	52

4.8	<i>Processor Categories And Roles</i>	52
4.9	<i>Processor Technologies</i>	54
4.10	<i>Stored Programs</i>	54
4.11	<i>The Fetch-Execute Cycle</i>	55
4.12	<i>Clock Rate And Instruction Rate</i>	56
4.13	<i>Control: Getting Started And Stopping</i>	57
4.14	<i>Starting The Fetch-Execute Cycle</i>	57
4.15	<i>Summary</i>	58

Chapter 5 Processor Types And Instruction Sets

61

5.1	<i>Introduction</i>	61
5.2	<i>Mathematical Power, Convenience, And Cost</i>	61
5.3	<i>Instruction Set And Representation</i>	62
5.4	<i>Opcodes, Operands, And Results</i>	63
5.5	<i>Typical Instruction Format</i>	63
5.6	<i>Variable-Length Vs. Fixed-Length Instructions</i>	63
5.7	<i>General-Purpose Registers</i>	64
5.8	<i>Floating Point Registers And Register Identification</i>	65
5.9	<i>Programming With Registers</i>	65
5.10	<i>Register Banks</i>	66
5.11	<i>Complex And Reduced Instruction Sets</i>	67
5.12	<i>RISC Design And The Execution Pipeline</i>	68
5.13	<i>Pipelines And Instruction Stalls</i>	69
5.14	<i>Other Causes Of Pipeline Stalls</i>	71
5.15	<i>Consequences For Programmers</i>	71
5.16	<i>Programming, Stalls, And No-Op Instructions</i>	72
5.17	<i>Forwarding</i>	72
5.18	<i>Types Of Operations</i>	73
5.19	<i>Program Counter, Fetch-Execute, And Branching</i>	73
5.20	<i>Subroutine Calls, Arguments, And Register Windows</i>	75
5.21	<i>An Example Instruction Set</i>	76
5.22	<i>Minimalistic Instruction Set</i>	78
5.23	<i>The Principle Of Orthogonality</i>	79
5.24	<i>Condition Codes And Conditional Branching</i>	80
5.25	<i>Summary</i>	80

Chapter 6 Operand Addressing And Instruction Representation

83

6.1	<i>Introduction</i>	83
6.2	<i>Zero, One, Two, Or Three Address Designs</i>	83
6.3	<i>Zero Operands Per Instruction</i>	84

6.4	<i>One Operand Per Instruction</i>	85
6.5	<i>Two Operands Per Instruction</i>	85
6.6	<i>Three Operands Per Instruction</i>	86
6.7	<i>Operand Sources And Immediate Values</i>	86
6.8	<i>The Von Neumann Bottleneck</i>	87
6.9	<i>Explicit And Implicit Operand Encoding</i>	88
6.10	<i>Operands That Combine Multiple Values</i>	89
6.11	<i>Tradeoffs In The Choice Of Operands</i>	90
6.12	<i>Values In Memory And Indirect Reference</i>	91
6.13	<i>Operand Addressing Modes</i>	92
6.14	<i>Summary</i>	93

Chapter 7 CPUs: Microcode, Protection, And Processor Modes

95

7.1	<i>Introduction</i>	95
7.2	<i>A Central Processor</i>	95
7.3	<i>CPU Complexity</i>	96
7.4	<i>Modes Of Execution</i>	97
7.5	<i>Backward Compatibility</i>	97
7.6	<i>Changing Modes</i>	98
7.7	<i>Privilege And Protection</i>	99
7.8	<i>Multiple Levels Of Protection</i>	99
7.9	<i>Microcoded Instructions</i>	100
7.10	<i>Microcode Variations</i>	102
7.11	<i>The Advantage Of Microcode</i>	102
7.12	<i>Making Microcode Visible To Programmers</i>	103
7.13	<i>Vertical Microcode</i>	103
7.14	<i>Horizontal Microcode</i>	104
7.15	<i>Example Horizontal Microcode</i>	105
7.16	<i>A Horizontal Microcode Example</i>	107
7.17	<i>Operations That Require Multiple Cycles</i>	108
7.18	<i>Horizontal Microcode And Parallel Execution</i>	109
7.19	<i>Look-Ahead And High Performance Execution</i>	110
7.20	<i>Parallelism And Execution Order</i>	111
7.21	<i>Out-Of-Order Instruction Execution</i>	111
7.22	<i>Conditional Branches And Branch Prediction</i>	112
7.23	<i>Consequences For Programmers</i>	113
7.24	<i>Summary</i>	113

Chapter 8 Assembly Languages And Programming Paradigm**115**

8.1	<i>Introduction</i>	115
8.2	<i>Characteristics Of A High-level Programming Language</i>	115
8.3	<i>Characteristics Of A Low-Level Programming Language</i>	116
8.4	<i>Assembly Language</i>	117
8.5	<i>Assembly Language Syntax And Opcodes</i>	118
8.6	<i>Operand Order</i>	120
8.7	<i>Register Names</i>	121
8.8	<i>Operand Types</i>	122
8.9	<i>Assembly Language Programming Paradigm And Idioms</i>	122
8.10	<i>Assembly Code For Conditional Execution</i>	123
8.11	<i>Assembly Code For A Conditional Alternative</i>	124
8.12	<i>Assembly Code For Definite Iteration</i>	124
8.13	<i>Assembly Code For Indefinite Iteration</i>	125
8.14	<i>Assembly Code For Procedure Invocation</i>	125
8.15	<i>Assembly Code For Parameterized Procedure Invocation</i>	126
8.16	<i>Consequence For Programmers</i>	127
8.17	<i>Assembly Code For Function Invocation</i>	128
8.18	<i>Interaction Between Assembly And High-Level Languages</i>	128
8.19	<i>Assembly Code For Variables And Storage</i>	129
8.20	<i>Two-Pass Assembler</i>	130
8.21	<i>Assembly Language Macros</i>	131
8.22	<i>Summary</i>	134

PART III Memories**Chapter 9 Memory And Storage****137**

9.1	<i>Introduction</i>	137
9.2	<i>Definition</i>	137
9.3	<i>The Key Aspects Of Memory</i>	138
9.4	<i>Characteristics Of Memory Technologies</i>	138
9.5	<i>The Important Concept Of A Memory Hierarchy</i>	140
9.6	<i>Instruction And Data Store</i>	140
9.7	<i>The Fetch-Store Paradigm</i>	141
9.8	<i>Summary</i>	141

Chapter 10 Physical Memory And Physical Addressing**143**

- 10.1 Introduction* 143
- 10.2 Characteristics Of Computer Memory* 143
- 10.3 Static And Dynamic RAM Technologies* 144
- 10.4 Measures Of Memory Technology* 145
- 10.5 Density* 146
- 10.6 Separation Of Read And Write Performance* 146
- 10.7 Latency And Memory Controllers* 146
- 10.8 Synchronized Memory Technologies* 147
- 10.9 Multiple Data Rate Memory Technologies* 148
- 10.10 Examples Of Memory Technologies* 148
- 10.11 Memory Organization* 148
- 10.12 Memory Access And Memory Bus* 149
- 10.13 Memory Transfer Size* 150
- 10.14 Physical Addresses And Words* 150
- 10.15 Physical Memory Operations* 150
- 10.16 Word Size And Other Data Types* 151
- 10.17 An Extreme Case: Byte Addressing* 151
- 10.18 Byte Addressing With Word Transfers* 152
- 10.19 Using Powers Of Two* 153
- 10.20 Byte Alignment And Programming* 154
- 10.21 Memory Size And Address Space* 154
- 10.22 Programming With Word Addressing* 155
- 10.23 Measures Of Memory Size* 155
- 10.24 Pointers And Data Structures* 156
- 10.25 A Memory Dump* 156
- 10.26 Indirection And Indirect Operands* 158
- 10.27 Memory Banks And Interleaving* 158
- 10.28 Content Addressable Memory* 159
- 10.29 Ternary CAM* 160
- 10.30 Summary* 160

Chapter 11 Virtual Memory Technologies And Virtual Addressing**163**

- 11.1 Introduction* 163
- 11.2 Definition* 163
- 11.3 A Virtual Example: Byte Addressing* 164
- 11.4 Virtual Memory Terminology* 164
- 11.5 An Interface To Multiple Physical Memory Systems* 164
- 11.6 Address Translation Or Address Mapping* 166
- 11.7 Avoiding Arithmetic Calculation* 167
- 11.8 Discontiguous Address Spaces* 168

11.9	Other Memory Organizations	169
11.10	Motivation For Virtual Memory	169
11.11	Multiple Virtual Spaces And Multiprogramming	170
11.12	Multiple Levels Of Virtualization	171
11.13	Creating Virtual Spaces Dynamically	171
11.14	Base-Bound Registers	172
11.15	Changing The Virtual Space	172
11.16	Virtual Memory, Base-Bound, And Protection	173
11.17	Segmentation	174
11.18	Demand Paging	175
11.19	Hardware And Software For Demand Paging	175
11.20	Page Replacement	176
11.21	Paging Terminology And Data Structures	176
11.22	Address Translation In A Paging System	177
11.23	Using Powers Of Two	178
11.24	Presence, Use, And Modified Bits	179
11.25	Page Table Storage	180
11.26	Paging Efficiency And A Translation Lookaside Buffer	181
11.27	Consequences For Programmers	182
11.28	Summary	183

Chapter 12 Caches And Caching

185

12.1	Introduction	185
12.2	Definition	185
12.3	Characteristics Of A Cache	186
12.4	The Importance Of Caching	187
12.5	Examples Of Caching	188
12.6	Cache Terminology	188
12.7	Best And Worst Case Cache Performance	189
12.8	Cache Performance On A Typical Sequence	190
12.9	Cache Replacement Policy	190
12.10	LRU Replacement	191
12.11	Multi-level Cache Hierarchy	191
12.12	Preloading Caches	192
12.13	Caches Used With Memory	192
12.14	TLB As A Cache	193
12.15	Demand Paging As A Form Of Caching	193
12.16	Physical Memory Cache	194
12.17	Write Through And Write Back	194
12.18	Cache Coherence	195
12.19	L1, L2, and L3 Caches	196
12.20	Sizes Of L1, L2, And L3 Caches	197