

T H I R D E D I T I O N

Computer Organization and Design

THE HARDWARE/SOFTWARE INTERFACE

David A. Patterson

University of California, Berkeley

John L. Hennessy

Stanford University

With a contribution by

Peter J. Ashenden

Ashenden Designs Pty Ltd

James R. Larus

Microsoft Research

Daniel J. Sorin

Duke University



ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Morgan Kaufmann is an Imprint of Elsevier




MORGAN KAUFMANN PUBLISHERS

Contents

Preface ix


CHAPTERS

1 Computer Abstractions and Technology 2

- 1.1 Introduction 3
- 1.2 Below Your Program 11
- 1.3 Under the Covers 15
- 1.4 Real Stuff: Manufacturing Pentium 4 Chips 28
- 1.5 Fallacies and Pitfalls 33
- 1.6 Concluding Remarks 35
-  1.7 Historical Perspective and Further Reading 36
- 1.8 Exercises 36

COMPUTERS IN THE REAL WORLD **Information Technology for the 4 Billion without IT 44**

2 Instructions: Language of the Computer 46

- 2.1 Introduction 48
- 2.2 Operations of the Computer Hardware 49
- 2.3 Operands of the Computer Hardware 52
- 2.4 Representing Instructions in the Computer 60
- 2.5 Logical Operations 68
- 2.6 Instructions for Making Decisions 72
- 2.7 Supporting Procedures in Computer Hardware 79
- 2.8 Communicating with People 90
- 2.9 MIPS Addressing for 32-Bit Immediates and Addresses 95
- 2.10 Translating and Starting a Program 106
- 2.11 How Compilers Optimize 116
-  2.12 How Compilers Work: An Introduction 121

- 2.13 A C Sort Example to Put It All Together 121
- 2.14 Implementing an Object-Oriented Language 130
- 2.15 Arrays versus Pointers 130
- 2.16 Real Stuff: IA-32 Instructions 134
- 2.17 Fallacies and Pitfalls 143
- 2.18 Concluding Remarks 145
- 2.19 Historical Perspective and Further Reading 147
- 2.20 Exercises 147

COMPUTERS IN THE REAL WORLD
Helping Save Our Environment with Data 156

3

Arithmetic for Computers 158

- 3.1 Introduction 160
- 3.2 Signed and Unsigned Numbers 160
- 3.3 Addition and Subtraction 170
- 3.4 Multiplication 176
- 3.5 Division 183
- 3.6 Floating Point 189
- 3.7 Real Stuff: Floating Point in the IA-32 217
- 3.8 Fallacies and Pitfalls 220
- 3.9 Concluding Remarks 225
- 3.10 Historical Perspective and Further Reading 229
- 3.11 Exercises 229

COMPUTERS IN THE REAL WORLD
Reconstructing the Ancient World 236

4

Assessing and Understanding Performance 238

- 4.1 Introduction 240
- 4.2 CPU Performance and Its Factors 246
- 4.3 Evaluating Performance 254
- 4.4 Real Stuff: Two SPEC Benchmarks and the Performance of Recent Intel Processors 259
- 4.5 Fallacies and Pitfalls 266
- 4.6 Concluding Remarks 270
- 4.7 Historical Perspective and Further Reading 272
- 4.8 Exercises 272

COMPUTERS IN THE REAL WORLD
Moving People Faster and More Safely 280

5**The Processor: Datapath and Control 282**


- 5.1 Introduction 284
- 5.2 Logic Design Conventions 289
- 5.3 Building a Datapath 292
- 5.4 A Simple Implementation Scheme 300
- 5.5 A Multicycle Implementation 318
- 5.6 Exceptions 340
- 5.7 Microprogramming: Simplifying Control Design 346
- 5.8 An Introduction to Digital Design Using a Hardware Design Language 346
- 5.9 Real Stuff: The Organization of Recent Pentium Implementations 347
- 5.10 Fallacies and Pitfalls 350
- 5.11 Concluding Remarks 352
- 5.12 Historical Perspective and Further Reading 353
- 5.13 Exercises 354

COMPUTERS IN THE REAL WORLD
Empowering the Disabled 366**6****Enhancing Performance with Pipelining 368**



- 6.1 An Overview of Pipelining 370
- 6.2 A Pipelined Datapath 384
- 6.3 Pipelined Control 399
- 6.4 Data Hazards and Forwarding 402
- 6.5 Data Hazards and Stalls 413
- 6.6 Branch Hazards 416
- 6.7 Using a Hardware Description Language to Describe and Model a Pipeline 426
- 6.8 Exceptions 427
- 6.9 Advanced Pipelining: Extracting More Performance 432
- 6.10 Real Stuff: The Pentium 4 Pipeline 448
- 6.11 Fallacies and Pitfalls 451
- 6.12 Concluding Remarks 452
- 6.13 Historical Perspective and Further Reading 454
- 6.14 Exercises 454

COMPUTERS IN THE REAL WORLD
Mass Communication without Gatekeepers 464

7**Large and Fast: Exploiting Memory Hierarchy 466**

- 7.1 Introduction 468
- 7.2 The Basics of Caches 473
- 7.3 Measuring and Improving Cache Performance 492
- 7.4 Virtual Memory 511
- 7.5 A Common Framework for Memory Hierarchies 538
- 7.6 Real Stuff: The Pentium P4 and the AMD Opteron Memory Hierarchies 546
- 7.7 Fallacies and Pitfalls 550
- 7.8 Concluding Remarks 552
-  7.9 Historical Perspective and Further Reading 555
- 7.10 Exercises 555

COMPUTERS IN THE REAL WORLD
Saving the World's Art Treasures 562**8****Storage, Networks, and Other Peripherals 564**

- 8.1 Introduction 566
- 8.2 Disk Storage and Dependability 569
-  8.3 Networks 580
- 8.4 Buses and Other Connections between Processors, Memory, and I/O Devices 581
- 8.5 Interfacing I/O Devices to the Processor, Memory, and Operating System 588
- 8.6 I/O Performance Measures: Examples from Disk and File Systems 597
- 8.7 Designing an I/O System 600
- 8.8 Real Stuff: A Digital Camera 603
- 8.9 Fallacies and Pitfalls 606
- 8.10 Concluding Remarks 609
-  8.11 Historical Perspective and Further Reading 611
- 8.12 Exercises 611

COMPUTERS IN THE REAL WORLD
Saving Lives through Better Diagnosis 622**9****Multiprocessors and Clusters 9-2**

- 9.1 Introduction 9-4
- 9.2 Programming Multiprocessors 9-8
- 9.3 Multiprocessors Connected by a Single Bus 9-11

- 9.4 Multiprocessors Connected by a Network 9-20
- 9.5 Clusters 9-25
- 9.6 Network Topologies 9-27
- 9.7 Multiprocessors Inside a Chip and Multithreading 9-30
- 9.8 Real Stuff: The Google Cluster of PCs 9-34
- 9.9 Fallacies and Pitfalls 9-39
- 9.10 Concluding Remarks 9-42
- 9.11 Historical Perspective and Further Reading 9-47
- 9.12 Exercises 9-55

A P P E N D I C E S



Assemblers, Linkers, and the SPIM Simulator A-2

- A.1 Introduction A-3
- A.2 Assemblers A-10
- A.3 Linkers A-18
- A.4 Loading A-19
- A.5 Memory Usage A-20
- A.6 Procedure Call Convention A-22
- A.7 Exceptions and Interrupts A-33
- A.8 Input and Output A-38
- A.9 SPIM A-40
- A.10 MIPS R2000 Assembly Language A-45
- A.11 Concluding Remarks A-81
- A.12 Exercises A-82



The Basics of Logic Design B-2

- B.1 Introduction B-3
- B.2 Gates, Truth Tables, and Logic Equations B-4
- B.3 Combinational Logic B-8
- B.4 Using a Hardware Description Language B-20
- B.5 Constructing a Basic Arithmetic Logic Unit B-26
- B.6 Faster Addition: Carry Lookahead B-38
- B.7 Clocks B-47
- B.8 Memory Elements: Flip-flops, Latches, and Registers B-49
- B.9 Memory Elements: SRAMs and DRAMs B-57
- B.10 Finite State Machines B-67
- B.11 Timing Methodologies B-72

- B.12 Field Programmable Devices B-77
- B.13 Concluding Remarks B-78
- B.14 Exercises B-79



Mapping Control to Hardware C-2



- C.1 Introduction C-3
- C.2 Implementing Combinational Control Units C-4
- C.3 Implementing Finite State Machine Control C-8
- C.4 Implementing the Next-State Function with a Sequencer C-21
- C.5 Translating a Microprogram to Hardware C-27
- C.6 Concluding Remarks C-31
- C.7 Exercises C-32



A Survey of RISC Architectures for Desktop, Server, and Embedded Computers D-2

- D.1 Introduction D-3
- D.2 Addressing Modes and Instruction Formats D-5
- D.3 Instructions: The MIPS Core Subset D-9
- D.4 Instructions: Multimedia Extensions of the Desktop/Server RISCs D-16
- D.5 Instructions: Digital Signal-Processing Extensions of the Embedded RISCs D-19
- D.6 Instructions: Common Extensions to MIPS Core D-20
- D.7 Instructions Unique to MIPS64 D-25
- D.8 Instructions Unique to Alpha D-27
- D.9 Instructions Unique to SPARC v.9 D-29
- D.10 Instructions Unique to PowerPC D-32
- D.11 Instructions Unique to PA-RISC 2.0 D-34
- D.12 Instructions Unique to ARM D-36
- D.13 Instructions Unique to Thumb D-38
- D.14 Instructions Unique to SuperH D-39
- D.15 Instructions Unique to M32R D-40
- D.16 Instructions Unique to MIPS16 D-41
- D.17 Concluding Remarks D-43
- D.18 Acknowledgments D-46
- D.19 References D-47

Index I-1

-  Glossary G-1
-  Further Reading FR-1