

大学计算机教育国外著名教材系列 (影印版)

MODERN PROCESSOR DESIGN  
FUNDAMENTALS OF  
SUPERSCALAR PROCESSORS

# 现代处理器设计

John P. Shen Mikko Lipasti 著



清华大学出版社

# MODERN PROCESSOR DESIGN

## FUNDAMENTALS OF SUPERSCALAR PROCESSORS

# 现代处理器设计

## 内容简介

本书是关于处理器设计的最新、最权威教材，主要论述了：(1)处理器的设计方法和原理；(2)流水线技术；(3)主存与I/O系统；(4)超标量组织与技术；(5)PowerPC 620和Intel P6等示例；(6)超标量处理器设计；(7)先进的指令流技术、存储器数据流技术；(8)多线程技术等。

本书适合作为计算机及相关专业的“处理器设计”课程的教材，也是有关专业人员很有价值的参考用书。

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（不包括中国香港、  
澳门特别行政区和中国台湾地区）销售发行。



<http://www.mhhe.com>

ISBN 978-7-302-15357-3



9 787302 153573 >

定价:62.00元

大学计算机教育国外著名教材系列（影印版）

# Modern Processor Design

Fundamentals of Superscalar Processors

## 现代处理器设计

John P. Shen

Mikko Lipasti

清华大学出版社

北 京

John P. Shen and Mikko Lipasti

**Modern Processor Design: Fundamentals of Superscalar Processors**

EISBN: 0-07-057064-7

Copyright © 2005 by The McGraw-Hill Companies, Inc.

Original language published by The McGraw-Hill Companies, Inc. All Rights reserved. No part of this publication may be reproduced or distributed by any means, or stored in a database or retrieval system, without the prior written permission of the publisher.

Authorized English language edition jointly published by McGraw-Hill Education (Asia) Co. and Tsinghua University Press. This edition is authorized for sale only to the educational and training institutions, and within the territory of the People's Republic of China (excluding Hong Kong, Macao SAR and Taiwan). Unauthorized export of this edition is a violation of the Copyright Act. Violation of this Law is subject to Civil and Criminal Penalties.

本书英文影印版由清华大学出版社和美国麦格劳-希尔教育出版(亚洲)公司合作出版。此版本仅限在中华人民共和国境内(不包括中国香港、澳门特别行政区及中国台湾地区)针对教育及培训机构之销售。未经许可之出口,视为违反著作权法,将受法律之制裁。

未经出版者预先书面许可,不得以任何方式复制或抄袭本书的任何部分。

北京市版权局著作权合同登记号 图字: 01-2005-4073

本书封面贴有 McGraw-Hill 公司激光防伪标签,无标签者不得销售。

版权所有,侵权必究。侵权举报电话:010-62782989 13501256678 13801310933

#### 图书在版编目(CIP)数据

现代处理器设计=Modern Processor Design: 英文/(美)谢(Shen, J.P.)等著. —影印本. —北京:清华大学出版社, 2007.8

(大学计算机教育国外著名教材系列)

ISBN 978-7-302-15357-3

I. 现… II. 谢… III. 微处理器—系统设计—高等学校—教材—英文 IV. TP332

中国版本图书馆 CIP 数据核字(2007)第 082348 号

出 版 者: 清华大学出版社

<http://www.tup.com.cn>

[c-service@tup.tsinghua.edu.cn](mailto:c-service@tup.tsinghua.edu.cn)

社 总 机: 010-62770175

投稿咨询: 010-62772015

地 址: 北京清华大学学研大厦

邮 编: 100084

邮购热线: 010-62786544

客户服务: 010-62776969

印 刷 者: 北京市昌平环球印刷厂

装 订 者: 三河市金元印装有限公司

发 行 者: 全国新华书店

开 本: 185×230 印张: 41.75

版 次: 2007 年 8 月第 1 版 2007 年 8 月第 1 次印刷

印 数: 1~3000

定 价: 62.00 元

本书如存在文字不清、漏印、缺页、倒页、脱页等印装质量问题,请与清华大学出版社出版部联系调换。联系电话: 010-62770177 转 3103 产品编号: 018520-01

## 出版说明

进入 21 世纪, 世界各国的经济、科技以及综合国力的竞争将更加激烈。竞争的中心无疑是对人才的竞争。谁拥有大量高素质的人才, 谁就能在竞争中取得优势。高等教育, 作为培养高素质人才的事业, 必然受到高度重视。目前我国高等教育的教材更新较慢, 为了加快教材的更新频率, 教育部正在大力促进我国高校采用国外原版教材。

清华大学出版社从 1996 年开始, 与国外著名出版公司合作, 影印出版了“大学计算机教育丛书(影印版)”等一系列引进图书, 受到国内读者的欢迎和支持。跨入 21 世纪, 我们本着为我国高等教育教材建设服务的初衷, 在已有的基础上, 进一步扩大选题内容, 改变图书开本尺寸, 一如既往地请有关专家挑选适用于我国高校本科及研究生计算机教育的国外经典教材或著名教材, 组成本套“大学计算机教育国外著名教材系列(影印版)”, 以飨读者。深切期盼读者及时将使用本系列教材的效果和意见反馈给我们。更希望国内专家、教授积极向我们推荐国外计算机教育的优秀教材, 以利我们把“大学计算机教育国外著名教材系列(影印版)”做得更好, 更适合高校师生的需要。

清华大学出版社

# John Paul Shen



**John Paul Shen** is the Director of Intel's Microarchitecture Research Lab (MRL), providing leadership to about two-dozen highly skilled researchers located in Santa Clara, CA; Hillsboro, OR; and Austin, TX. MRL is responsible for developing innovative microarchitecture techniques that can potentially be used in future microprocessor products from Intel. MRL researchers collaborate closely with microarchitects from product teams in joint advanced-development efforts. MRL frequently hosts visiting faculty and Ph.D. interns and conducts joint research projects with academic research groups.

Prior to joining Intel in 2000, John was a professor in the electrical and computer engineering department of Carnegie Mellon University, where he headed up the CMU Microarchitecture Research Team (CMuART). He has supervised a total of 16 Ph.D. students during his years at CMU. Seven are currently with Intel, and five have faculty positions in academia. He won multiple teaching awards at CMU. He was an NSF Presidential Young Investigator. He is an IEEE Fellow and has served on the program committees of ISCA, MICRO, HPCA, ASPLOS, PACT, ICCD, ITC, and FTCS.

He has published over 100 research papers in diverse areas, including fault-tolerant computing, built-in self-test, process defect and fault analysis, concurrent error detection, application-specific processors, performance evaluation, compilation for instruction-level parallelism, value locality and prediction, analytical modeling of superscalar processors, systematic microarchitecture test generation, performance simulator validation, precomputation-based prefetching, database workload analysis, and user-level helper threads.

John received his M.S. and Ph.D. degrees from the University of Southern California, and his B.S. degree from the University of Michigan, all in electrical engineering. He attended Kimball High School in Royal Oak, Michigan. He is happily married and has three daughters. His family enjoys camping, road trips, and reading *The Lord of the Rings*.

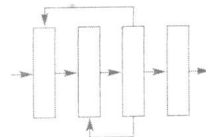
# Mikko Lipasti



**Mikko Lipasti** has been an assistant professor at the University of Wisconsin-Madison since 1999, where he is actively pursuing various research topics in the realms of processor, system, and memory architecture. He has advised a total of 17 graduate students, including two completed Ph.D. theses and numerous M.S. projects, and has published more than 30 papers in top computer architecture conferences and journals. He is most well known for his seminal Ph.D. work in value prediction. His research program has received in excess of \$2 million in support through multiple grants from the National Science Foundation as well as financial support and equipment donations from IBM, Intel, AMD, and Sun Microsystems.

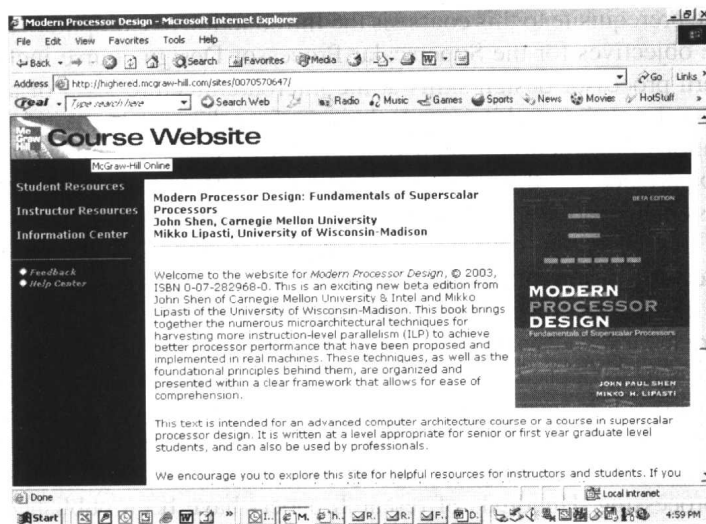
The Eta Kappa Nu Electrical Engineering Honor Society selected Mikko as the country's Outstanding Young Electrical Engineer for 2002. He is also a member of the IEEE and the Tau Beta Pi engineering honor society. He received his B.S. in computer engineering from Valparaiso University in 1991, and M.S. (1992) and Ph.D. (1997) degrees in electrical and computer engineering from Carnegie Mellon University. Prior to beginning his academic career, he worked for IBM Corporation in both software and future processor and system performance analysis and design guidance, as well as operating system kernel implementation. While at IBM he contributed to system and microarchitectural definition of future IBM server computer systems. He has served on numerous conference and workshop program committees and is co-organizer of the annual Workshop on Duplicating, Deconstructing, and Debunking (WDDD). He has filed seven patent applications, six of which are issued U.S. patents; won the Best Paper Award at MICRO-29; and has received IBM Invention Achievement, Patent Issuance, and Technical Recognition Awards.

Mikko has been happily married since 1991 and has a nine-year-old daughter and a six-year old son. In his spare time, he enjoys regular exercise, family bike rides, reading, and volunteering his time at his local church and on campus as an English-language discussion group leader at the International Friendship Center.



# Additional Resources

In addition to the comprehensive coverage within the book, a number of additional resources are available with Shen/Lipasti's MODERN PROCESSOR DESIGN through the book's website at [www.mhhe.com/shen](http://www.mhhe.com/shen).



## Instructor Resources

- **Solutions Manual**—A complete set of solutions for the chapter-ending homework problems are provided.
- **PowerPoint Slides**—Two sets of MS PowerPoint slides, from Carnegie Mellon University and the University of Wisconsin-Madison, can be downloaded to supplement your lecture presentations.
- **Figures**—A complete set of figures from the book are available in eps format. These figures can be used to create your own presentations.
- **Sample Homework Files**—A set of homework assignments with answers from Carnegie Mellon University are provided to supplement your own assignments.
- **Sample Exams**—A set of exams with answers from Carnegie Mellon University are also provided to supplement your own exams.
- **Links to [www.simplescalar.com](http://www.simplescalar.com)**—We provide several links to the SimpleScalar tool set, which are available free for non-commercial academic use.

# Preface

This book emerged from the course *Superscalar Processor Design*, which has been taught at Carnegie Mellon University since 1995. Superscalar Processor Design is a mezzanine course targeting seniors and first-year graduate students. Quite a few of the more aggressive juniors have taken the course in the spring semester of their junior year. The prerequisite to this course is the Introduction to Computer Architecture course. The objectives for the Superscalar Processor Design course include: (1) to teach modern processor design skills at the microarchitecture level of abstraction; (2) to cover current microarchitecture techniques for achieving high performance via the exploitation of instruction-level parallelism (ILP); and (3) to impart insights and hands-on experience for the effective design of contemporary high-performance microprocessors for mobile, desktop, and server markets. In addition to covering the contents of this book, the course contains a project component that involves the microarchitectural design of a future-generation superscalar microprocessor.

During the decade of the 1990s many microarchitectural techniques for increasing clock frequency and harvesting more ILP to achieve better processor performance have been proposed and implemented in real machines. This book is an attempt to codify this large body of knowledge in a systematic way. These techniques include deep pipelining, aggressive branch prediction, dynamic register renaming, multiple instruction dispatching and issuing, out-of-order execution, and speculative load/store processing. Hundreds of research papers have been published since the early 1990s, and many of the research ideas have become reality in commercial superscalar microprocessors. In this book, the numerous techniques are organized and presented within a clear framework that facilitates ease of comprehension. The foundational principles that underlie the plethora of techniques are highlighted.

While the contents of this book would generally be viewed as graduate-level material, the book is intentionally written in a way that would be very accessible to undergraduate students. Significant effort has been spent in making seemingly complex techniques to appear as quite straightforward through appropriate abstraction and hiding of details. The priority is to convey clearly the key concepts and fundamental principles, giving just enough details to ensure understanding of implementation issues without massive dumping of information and quantitative data. The hope is that this body of knowledge can become widely possessed by not just microarchitects and processor designers but by most B.S. and M.S. students with interests in computer systems and microprocessor design.

Here is a brief summary of the chapters.

## Chapter 1: Processor Design

This chapter introduces the art of processor design, the instruction set architecture (ISA) as the specification of the processor, and the microarchitecture as the implementation of the processor. The dynamic/static interface that separates compile-time

software and run-time hardware is defined and discussed. The goal of this chapter is not to revisit in depth the traditional issues regarding ISA design, but to erect the proper framework for understanding modern processor design.

## **Chapter 2: Pipelined Processors**

This chapter focuses on the concept of pipelining, discusses instruction pipeline design, and presents the performance benefits of pipelining. Pipelining is usually introduced in the first computer architecture course. Pipelining provides the foundation for modern superscalar techniques and is presented in this chapter in a fresh and unique way. We intentionally avoid the massive dumping of bar charts and graphs; instead, we focus on distilling the foundational principles of instruction pipelining.

## **Chapter 3: Memory and I/O Systems**

This chapter provides a larger context for the remainder of the book by including a thorough grounding in the principles and mechanisms of modern memory and I/O systems. Topics covered include memory hierarchies, caching, main memory design, virtual memory architecture, common input/output devices, processor-I/O interaction, and bus design and organization.

## **Chapter 4: Superscalar Organization**

This chapter introduces the main concepts and the overall organization of superscalar processors. It provides a “big picture” view for the reader that leads smoothly into the detailed discussions in the next chapters on specific superscalar techniques for achieving performance. This chapter highlights only the key features of superscalar processor organizations. Chapter 7 provides a detailed survey of features found in real machines.

## **Chapter 5: Superscalar Techniques**

This chapter is the heart of this book and presents all the major microarchitecture techniques for designing contemporary superscalar processors for achieving high performance. It classifies and presents specific techniques for enhancing instruction flow, register data flow, and memory data flow. This chapter attempts to organize a plethora of techniques into a systematic framework that facilitates ease of comprehension.

## **Chapter 6: The PowerPC 620**

This chapter presents a detailed analysis of the PowerPC 620 microarchitecture and uses it as a case study to examine many of the issues and design tradeoffs introduced in the previous chapters. This chapter contains extensive performance data of an aggressive out-of-order design.

## **Chapter 7: Intel’s P6 Microarchitecture**

This is a case study chapter on probably the most commercially successful contemporary superscalar microarchitecture. It is written by the Intel P6 design team led by Bob Colwell and presents in depth the P6 microarchitecture that facilitated the implementation of the Pentium Pro, Pentium II, and Pentium III microprocessors. This chapter offers the readers an opportunity to peek into the mindset of a top-notch design team.

## Chapter 8: Survey of Superscalar Processors

This chapter, compiled by Prof. Mark Smotherman of Clemson University, provides a historical chronicle on the development of superscalar machines and a survey of existing superscalar microprocessors. The chapter was first completed in 1998 and has been continuously revised and updated since then. It contains fascinating information that can't be found elsewhere.

## Chapter 9: Advanced Instruction Flow Techniques

This chapter provides a thorough overview of issues related to high-performance instruction fetching. The topics covered include historical, currently used, and proposed advanced future techniques for branch prediction, as well as high-bandwidth and high-frequency fetch architectures like trace caches. Though not all such techniques have yet been adopted in real machines, future designs are likely to incorporate at least some form of them.

## Chapter 10: Advanced Register Data Flow Techniques

This chapter highlights emerging microarchitectural techniques for increasing performance by exploiting the program characteristic of *value locality*. This program characteristic was discovered recently, and techniques ranging from software memoization, instruction reuse, and various forms of value prediction are described in this chapter. Though such techniques have not yet been adopted in real machines, future designs are likely to incorporate at least some form of them.

## Chapter 11: Executing Multiple Threads

This chapter provides an introduction to thread-level parallelism (TLP), and provides a basic introduction to multiprocessing, cache coherence, and high-performance implementations that guarantee either sequential or relaxed memory ordering across multiple processors. It discusses single-chip techniques like multithreading and on-chip multiprocessing that also exploit thread-level parallelism. Finally, it visits two emerging technologies—implicit multithreading and preexecution—that attempt to extract thread-level parallelism automatically from single-threaded programs.

In summary, Chapters 1 through 5 cover fundamental concepts and foundational techniques. Chapters 6 through 8 present case studies and an extensive survey of actual commercial superscalar processors. Chapter 9 provides a thorough overview of advanced instruction flow techniques, including recent developments in advanced branch predictors. Chapters 10 and 11 should be viewed as advanced topics chapters that highlight some emerging techniques and provide an introduction to multiprocessor systems.

This is the first edition of the book. An earlier beta edition was published in 2002 with the intent of collecting feedback to help shape and hone the contents and presentation of this first edition. Through the course of the development of the book, a large set of homework and exam problems have been created. A subset of these problems are included at the end of each chapter. Several problems suggest the use of the

SimpleScalar simulation suite available from the SimpleScalar website at <http://www.simplescalar.com>. A companion website for the book contains additional support material for the instructor, including a complete set of lecture slides ([www.mhhe.com/shen](http://www.mhhe.com/shen)).

## Acknowledgments

Many people have generously contributed their time, energy, and support toward the completion of this book. In particular, we are grateful to Bob Colwell, who is the lead author of Chapter 7, Intel's P6 Microarchitecture. We also acknowledge his coauthors, Dave Papworth, Glenn Hinton, Mike Fetterman, and Andy Glew, who were all key members of the historic P6 team. This chapter helps ground this textbook in practical, real-world considerations. We are also grateful to Professor Mark Smotherman of Clemson University, who meticulously compiled and authored Chapter 8, Survey of Superscalar Processors. This chapter documents the rich and varied history of superscalar processor design over the last 40 years. The guest authors of these two chapters added a certain radiance to this textbook that we could not possibly have produced on our own. The PowerPC 620 case study in Chapter 6 is based on Trung Diep's Ph.D. thesis at Carnegie Mellon University. Finally, the thorough survey of advanced instruction flow techniques in Chapter 9 was authored by Gabriel Loh, largely based on his Ph.D. thesis at Yale University.

In addition, we want to thank the following professors for their detailed, insightful, and thorough review of the original manuscript. The inputs from these reviews have significantly improved the first edition of this book.

- David Andrews, *University of Arkansas*
- Angelos Bilas, *University of Toronto*
- Fred H. Carlin, *University of California at Santa Barbara*
- Yinong Chen, *Arizona State University*
- Lynn Choi, *University of California at Irvine*
- Dan Connors, *University of Colorado*
- Karel Driesen, *McGill University*
- Alan D. George, *University of Florida*
- Arthur Glaser, *New Jersey Institute of Technology*
- Rajiv Gupta, *University of Arizona*
- Vincent Hayward, *McGill University*
- James Hoe, *Carnegie Mellon University*
- Lizy Kurian John, *University of Texas at Austin*
- Peter M. Kogge, *University of Notre Dame*
- Angkul Kongmunvattana, *University of Nevada at Reno*
- Israel Koren, *University of Massachusetts at Amherst*
- Ben Lee, *Oregon State University*
- Francis Leung, *Illinois Institute of Technology*
- Walid Najjar, *University of California Riverside*
- Vojin G. Oklabdzija, *University of California at Davis*
- Soner Onder, *Michigan Technological University*
- Parimal Patel, *University of Texas at San Antonio*
- Jih-Kwon Peir, *University of Florida*
- Gregory D. Peterson, *University of Tennessee*
- Amir Roth, *University of Pennsylvania*
- Kevin Skadron, *University of Virginia*
- Mark Smotherman, *Clemson University*
- Miroslav N. Velez, *Georgia Institute of Technology*
- Bin Wei, *Rutgers University*
- Anthony S. Wojcik, *Michigan State University*
- Ali Zaringhalam, *Stevens Institute of Technology*
- Xiaobo Zhou, *University of Colorado at Colorado Springs*

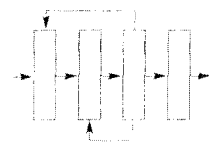
This book grew out of the course *Superscalar Processor Design* at Carnegie Mellon University. This course has been taught at CMU since 1995. Many teaching assistants of this course have left their indelible touch in the contents of this book. They include Bryan Black, Scott Cape, Yuan Chou, Alex Dean, Trung Diep, John Faistl, Andrew Huang, Deepak Limaye, Chris Nelson, Chris Newburn, Derek Noonburg, Kyle Oppenheim, Ryan Rakvic, and Bob Rychlik. Hundreds of students have taken this course at CMU; many of them provided inputs that also helped shape this book. Since 2000, Professor James Hoe at CMU has taken this course even further. We both are indebted to the nurturing we experienced while at CMU, and we hope that this book will help perpetuate CMU's historical reputation of producing some of the best computer architects and processor designers.

A draft version of this textbook has also been used at the University of Wisconsin since 2000. Some of the problems at the end of each chapter were actually contributed by students at the University of Wisconsin. We appreciate their test driving of this book.

John Paul Shen, *Director,*  
*Microarchitecture Research, Intel Labs, Adjunct Professor,*  
*ECE Department, Carnegie Mellon University*

Mikko H. Lipasti, *Assistant Professor,*  
*ECE Department, University of Wisconsin*

June 2004  
*Soli Deo Gloria*



# Table of Contents

|   |            |
|---|------------|
| Table of Contents                                   | v          |
| Additional Resources                                | xi         |
| Preface   | xii        |
| <b>1 Processor Design</b>                           | <b>1</b>   |
| 1.1 The Evolution of Microprocessors                | 2          |
| 1.2 Instruction Set Processor Design                | 4          |
| 1.2.1 Digital Systems Design                        | 4          |
| 1.2.2 Architecture, Implementation, and Realization | 5          |
| 1.2.3 Instruction Set Architecture                  | 6          |
| 1.2.4 Dynamic-Static Interface                      | 8          |
| 1.3 Principles of Processor Performance             | 10         |
| 1.3.1 Processor Performance Equation                | 10         |
| 1.3.2 Processor Performance Optimizations           | 11         |
| 1.3.3 Performance Evaluation Method                 | 13         |
| 1.4 Instruction-Level Parallel Processing           | 16         |
| 1.4.1 From Scalar to Superscalar                    | 16         |
| 1.4.2 Limits of Instruction-Level Parallelism       | 24         |
| 1.4.3 Machines for Instruction-Level Parallelism    | 27         |
| 1.5 Summary   | 32         |
| <b>2 Pipelined Processors</b>                       | <b>39</b>  |
| 2.1 Pipelining Fundamentals                         | 40         |
| 2.1.1 Pipelined Design                              | 40         |
| 2.1.2 Arithmetic Pipeline Example                   | 44         |
| 2.1.3 Pipelining Idealism                           | 48         |
| 2.1.4 Instruction Pipelining                        | 51         |
| 2.2 Pipelined Processor Design                      | 54         |
| 2.2.1 Balancing Pipeline Stages                     | 55         |
| 2.2.2 Unifying Instruction Types                    | 61         |
| 2.2.3 Minimizing Pipeline Stalls                    | 71         |
| 2.2.4 Commercial Pipelined Processors               | 87         |
| 2.3 Deeply Pipelined Processors                     | 94         |
| 2.4 Summary   | 97         |
| <b>3 Memory and I/O Systems</b>                     | <b>105</b> |
| 3.1 Introduction                                    | 105        |
| 3.2 Computer System Overview                        | 106        |
| 3.3 Key Concepts: Latency and Bandwidth             | 107        |

|          |   |            |
|----------|---|------------|
| 3.4      | Memory Hierarchy                                | 110        |
| 3.4.1    | Components of a Modern Memory Hierarchy         | 111        |
| 3.4.2    | Temporal and Spatial Locality                   | 113        |
| 3.4.3    | Caching and Cache Memories                      | 115        |
| 3.4.4    | Main Memory                                     | 127        |
| 3.5      | Virtual Memory Systems                          | 136        |
| 3.5.1    | Demand Paging                                   | 138        |
| 3.5.2    | Memory Protection                               | 141        |
| 3.5.3    | Page Table Architectures                        | 142        |
| 3.6      | Memory Hierarchy Implementation                 | 145        |
| 3.7      | Input/Output Systems                            | 153        |
| 3.7.1    | Types of I/O Devices                            | 154        |
| 3.7.2    | Computer System Busses                          | 161        |
| 3.7.3    | Communication with I/O Devices                  | 165        |
| 3.7.4    | Interaction of I/O Devices and Memory Hierarchy | 168        |
| 3.8      | Summary   | 170        |
| <b>4</b> | <b>Superscalar Organization</b>                 | <b>177</b> |
| 4.1      | Limitations of Scalar Pipelines                 | 178        |
| 4.1.1    | Upper Bound on Scalar Pipeline Throughput       | 178        |
| 4.1.2    | Inefficient Unification into a Single Pipeline  | 179        |
| 4.1.3    | Performance Lost Due to a Rigid Pipeline        | 179        |
| 4.2      | From Scalar to Superscalar Pipelines            | 181        |
| 4.2.1    | Parallel Pipelines                              | 181        |
| 4.2.2    | Diversified Pipelines                           | 184        |
| 4.2.3    | Dynamic Pipelines                               | 186        |
| 4.3      | Superscalar Pipeline Overview                   | 190        |
| 4.3.1    | Instruction Fetching                            | 191        |
| 4.3.2    | Instruction Decoding                            | 195        |
| 4.3.3    | Instruction Dispatching                         | 199        |
| 4.3.4    | Instruction Execution                           | 203        |
| 4.3.5    | Instruction Completion and Retiring             | 206        |
| 4.4      | Summary   | 209        |
| <b>5</b> | <b>Superscalar Techniques</b>                   | <b>217</b> |
| 5.1      | Instruction Flow Techniques                     | 218        |
| 5.1.1    | Program Control Flow and Control Dependences    | 218        |
| 5.1.2    | Performance Degradation Due to Branches         | 219        |
| 5.1.3    | Branch Prediction Techniques                    | 223        |
| 5.1.4    | Branch Misprediction Recovery                   | 228        |
| 5.1.5    | Advanced Branch Prediction Techniques           | 231        |
| 5.1.6    | Other Instruction Flow Techniques               | 236        |
| 5.2      | Register Data Flow Techniques                   | 237        |
| 5.2.1    | Register Reuse and False Data Dependences       | 237        |
| 5.2.2    | Register Renaming Techniques                    | 239        |
| 5.2.3    | True Data Dependences and the Data Flow Limit   | 244        |

|          |   |            |
|----------|---|------------|
| 5.2.4    | The Classic Tomasulo Algorithm          | 246        |
| 5.2.5    | Dynamic Execution Core                  | 254        |
| 5.2.6    | Reservation Stations and Reorder Buffer | 256        |
| 5.2.7    | Dynamic Instruction Scheduler           | 260        |
| 5.2.8    | Other Register Data Flow Techniques     | 261        |
| 5.3      | Memory Data Flow Techniques             | 262        |
| 5.3.1    | Memory Accessing Instructions           | 263        |
| 5.3.2    | Ordering of Memory Accesses             | 266        |
| 5.3.3    | Load Bypassing and Load Forwarding      | 267        |
| 5.3.4    | Other Memory Data Flow Techniques       | 273        |
| 5.4      | Summary                                 | 279        |
| <b>6</b> | <b>The PowerPC 620</b>                  | <b>301</b> |
| 6.1      | Introduction                            | 302        |
| 6.2      | Experimental Framework                  | 305        |
| 6.3      | Instruction Fetching                    | 307        |
| 6.3.1    | Branch Prediction                       | 307        |
| 6.3.2    | Fetching and Speculation                | 309        |
| 6.4      | Instruction Dispatching                 | 311        |
| 6.4.1    | Instruction Buffer                      | 311        |
| 6.4.2    | Dispatch Stalls                         | 311        |
| 6.4.3    | Dispatch Effectiveness                  | 313        |
| 6.5      | Instruction Execution                   | 316        |
| 6.5.1    | Issue Stalls                            | 316        |
| 6.5.2    | Execution Parallelism                   | 317        |
| 6.5.3    | Execution Latency                       | 317        |
| 6.6      | Instruction Completion                  | 318        |
| 6.6.1    | Completion Parallelism                  | 318        |
| 6.6.2    | Cache Effects                           | 318        |
| 6.7      | Conclusions and Observations            | 320        |
| 6.8      | Bridging to the IBM POWER3 and POWER4   | 322        |
| 6.9      | Summary                                 | 324        |
| <b>7</b> | <b>Intel's P6 Microarchitecture</b>     | <b>329</b> |
| 7.1      | Introduction                            | 330        |
| 7.1.1    | Basics of the P6 Microarchitecture      | 332        |
| 7.2      | Pipelining                              | 334        |
| 7.2.1    | In-Order Front-End Pipeline             | 334        |
| 7.2.2    | Out-of-Order Core Pipeline              | 336        |
| 7.2.3    | Retirement Pipeline                     | 337        |
| 7.3      | The In-Order Front End                  | 338        |
| 7.3.1    | Instruction Cache and ITLB              | 338        |
| 7.3.2    | Branch Prediction                       | 341        |
| 7.3.3    | Instruction Decoder                     | 343        |
| 7.3.4    | Register Alias Table                    | 346        |
| 7.3.5    | Allocator                               | 353        |

|          |  |            |
|----------|--|------------|
| 7.4      | The Out-of-Order Core  | 355        |
| 7.4.1    | Reservation Station  | 355        |
| 7.5      | Retirement   | 357        |
| 7.5.1    | The Reorder Buffer   | 357        |
| 7.6      | Memory Subsystem   | 361        |
| 7.6.1    | Memory Access Ordering   | 362        |
| 7.6.2    | Load Memory Operations   | 363        |
| 7.6.3    | Basic Store Memory Operations                                      | 363        |
| 7.6.4    | Deferring Memory Operations  | 363        |
| 7.6.5    | Page Faults  | 364        |
| 7.7      | Summary  | 364        |
| 7.8      | Acknowledgments  | 365        |
| <b>8</b> | <b>Survey of Superscalar Processors</b>                            | <b>369</b> |
| 8.1      | Development of Superscalar Processors                              | 369        |
| 8.1.1    | Early Advances in Uniprocessor Parallelism:<br>The IBM Stretch     | 369        |
| 8.1.2    | First Superscalar Design: The IBM Advanced<br>Computer System      | 372        |
| 8.1.3    | Instruction-Level Parallelism Studies                              | 377        |
| 8.1.4    | By-Products of DAE: The First<br>Multiple-Decoding Implementations | 378        |
| 8.1.5    | IBM Cheetah, Panther, and America                                  | 380        |
| 8.1.6    | Decoupled Microarchitectures                                       | 380        |
| 8.1.7    | Other Efforts in the 1980s   | 382        |
| 8.1.8    | Wide Acceptance of Superscalar                                     | 382        |
| 8.2      | A Classification of Recent Designs                                 | 384        |
| 8.2.1    | RISC and CISC Retrofits  | 384        |
| 8.2.2    | Speed Demons: Emphasis on Clock Cycle Time                         | 386        |
| 8.2.3    | Brainiacs: Emphasis on IPC   | 386        |
| 8.3      | Processor Descriptions   | 387        |
| 8.3.1    | Compaq / DEC Alpha   | 387        |
| 8.3.2    | Hewlett-Packard PA-RISC Version 1.0                                | 392        |
| 8.3.3    | Hewlett-Packard PA-RISC Version 2.0                                | 395        |
| 8.3.4    | IBM POWER  | 397        |
| 8.3.5    | Intel i960   | 402        |
| 8.3.6    | Intel IA32—Native Approaches                                       | 405        |
| 8.3.7    | Intel IA32—Decoupled Approaches                                    | 409        |
| 8.3.8    | x86-64   | 417        |
| 8.3.9    | MIPS   | 417        |
| 8.3.10   | Motorola   | 422        |
| 8.3.11   | PowerPC—32-bit Architecture  | 424        |
| 8.3.12   | PowerPC—64-bit Architecture  | 429        |
| 8.3.13   | PowerPC-AS   | 431        |
| 8.3.14   | SPARC Version 8  | 432        |
| 8.3.15   | SPARC Version 9  | 435        |