

COMPUTER ARCHITECTURE AND IMPLEMENTATION

This textbook provides a clear and concise introduction to computer architecture and implementation. Two important themes are interwoven throughout the book. The first is an overview of the major concepts and design philosophies of computer architecture and organization. The second is the early introduction and use of analytic modeling of computer performance.

The author begins by describing the classic von Neumann architecture and then presents in detail a number of performance models and evaluation techniques. He goes on to cover user instruction set design, including RISC architecture. A unique feature of the book is its memory-centric approach – memory systems are discussed before processor implementations. The author also deals with pipelined processors, input/output techniques, queuing modes, and extended instruction set architectures. Each topic is illustrated with reference to actual IBM and Intel architectures.

The book contains many worked examples and over 130 homework exercises. It is an ideal textbook for a one-semester undergraduate course in computer architecture and implementation.

Harvey Cragon holds the Ernest J. Cockrell, Jr., Centennial Chair Emeritus in Engineering at the University of Texas at Austin. He is a member of the U.S. National Academy of Engineering, a Life Fellow of the IEEE, and a Fellow of the Association for Computing Machinery. He has over 30 years' industrial experience, including 25 years with Texas Instruments. He has written over 50 technical papers and holds 9 patents in the field of computer engineering.

Cambridge University Press
0521657059 - Computer Architecture and Implementation
Harvey G. Cragon
Frontmatter
[More information](#)

COMPUTER ARCHITECTURE AND IMPLEMENTATION

HARVEY G. CRAGON

University of Texas at Austin



Cambridge University Press
0521657059 - Computer Architecture and Implementation
Harvey G. Cragon
Frontmatter
[More information](#)

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 2RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521651684

© Cambridge University Press 2000

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2000
This digitally printed first paperback version 2005

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data

Cragon, Harvey G.

Computer architecture and implementation / Harvey G. Cragon.

p. cm.

ISBN 0-521-65168-9 (hard.)

1. Computer architecture. I. Title.

QA76.9.A73C74 2000

004.2'2-dc21

99-16243

CIP

ISBN-13 978-0-521-65168-4 hardback

ISBN-10 0-521-65168-9 hardback

ISBN-13 978-0-521-65705-1 paperback

ISBN-10 0-521-65705-9 paperback

CONTENTS

<i>Preface</i>	<i>page viii</i>
1 COMPUTER OVERVIEW	1
1.0 Introduction	1
1.1 von Neumann Model	2
1.2 The von Neumann Architecture	5
1.2.1 The von Neumann Instruction Set Architecture	6
1.2.2 Instruction Interpretation Cycle	10
1.2.3 Limitations of the von Neumann Instruction Set Architecture	13
1.3 Historical Notes	13
1.3.1 Precursors to the von Neumann Instruction Set Architecture	14
REFERENCES	23
EXERCISES	24
2 PERFORMANCE MODELS AND EVALUATION	26
2.0 Introduction	26
2.1 Performance Models	26
2.2 Amdahl's Law	37
2.3 Moore's Law	39
2.4 Learning Curve	42
2.5 Grosch's Law	44
2.6 Steady-State Performance	45
2.7 Transient Performance	46
REFERENCES	47
EXERCISES	47
3 USER INSTRUCTION SET DESIGN	50
3.0 Introduction	50
3.1 Lexical Level	51

vi Contents

3.2	Instruction Set Architecture	51
3.2.1	Addresses	53
3.2.2	Operations	62
3.2.3	Data Types	74
3.2.3.1	User Data Types	75
3.2.3.2	Program Sequencing Data Types	88
3.2.4	Instruction Composition	93
3.3	CISC and RISC Architectural Styles	98
3.4	Static and Dynamic Instruction Statistics	103
3.5	Arithmetic	109
3.5.1	Addition/Subtraction	111
3.5.2	Multiplication and Division	114
3.5.3	Floating-Point Arithmetic	118
3.5.3.1	Precision Treatment	121
	REFERENCES	127
	EXERCISES	128
4	MEMORY SYSTEMS	131
4.0	Introduction	131
4.1	Hierarchical Memory	132
4.2	Paged Virtual Memory	133
4.3	Caches	145
4.4	Interleaved Real Memory	163
4.5	Virtual- and Real-Address Caches	171
4.6	Segmented Virtual Memory	173
4.7	Disk Memory	176
4.8	System Operation	179
	REFERENCES	181
	EXERCISES	181
5	PROCESSOR CONTROL DESIGN	183
5.0	Introduction	183
5.1	Hardwired Control	185
5.2	Microprogrammed Control	193
	REFERENCES	201
	EXERCISES	202
6	PIPELINED PROCESSORS	204
6.0	Introduction	204
6.1	Performance Models	205
6.2	Pipeline Partitioning	207
6.3	Pipeline Delays	211
6.3.1	Branch Delays	212
6.3.2	Structural Hazards	225
6.3.3	Data Dependencies	227

6.4	Interrupts	231
6.5	Superscalar Pipelines	234
6.6	Pipelined Processor Memory Demand	237
	REFERENCES	239
	EXERCISES	239
7	INPUT/OUTPUT	241
7.0	Introduction	241
7.1	I/O System Architecture	241
7.2	I/O Device Requirements	249
7.3	Buses and Controllers	252
7.3.1	Bus Design Examples	258
7.4	Serial Communication	258
7.5	Clock Synchronization	263
7.6	Queuing Theory	265
7.6.1	Open-System Queuing Model	265
7.6.2	Closed-System Queuing Model	271
	REFERENCES	273
	EXERCISES	274
8	EXTENDED INSTRUCTION SET ARCHITECTURES	275
8.0	Introduction	275
8.1	Operating System Kernel Support	276
8.2	Virtual-Memory Support	283
8.3	Interrupt Support	290
8.4	Input/Output Support	293
8.5	Cache Support	295
8.6	Multiprocessor Support	298
8.7	Other User ISA Enhancements, MMX	302
8.7.1	Other Intel ISA Support	307
	REFERENCES	308
	EXERCISES	308
	<i>Index</i>	310
	<i>Trademarks</i>	318

PREFACE

This textbook is intended for use in a one-semester, upper-division undergraduate course. It is expected that the students will have had courses in logic design with state machines and assembly language programming. A course in data structures as well as basic operating systems would be helpful.

There are two major themes running through this book. The first theme is an overview of the major ideas and design philosophies of computer architecture and implementation with some insight into their historical development. We are inclined to forget the work of our predecessors in computer design; the few pages devoted to this topic may help to rectify this neglect.

The second theme is the introduction and use of analytic modeling of computer performance. I believe that engineering students should take an engineering approach to the study of computer architecture and implementation. Various design trade-off issues will be examined with analytical models. What is an engineering approach? Consider the following two quotations.

“What distinguished the engineer from the technician is largely the ability to formulate and carry out the detailed calculations of forces and deflections, concentrations and flows, voltages and currents, that are required to test a proposed design on paper with regard to failure criteria. The ability to calculate is the ability to predict the performance of a design before it is built and tested.” (Petroske 1996).

“... it is much easier to have some vague notion about any subject, no matter what, than to arrive at the real truth about a single question, however simple that may be.” (Descartes 1629).

In addition to the two major themes, this book will introduce students to a new vocabulary. Whenever a new topic is approached, it is inevitable that one learns a new vocabulary. Therefore, whenever a new word or term is introduced, this book attempts to provide a definition along with common synonyms. We recommend that students obtain a copy of a computer dictionary to support their studies. Reference can also be made to an on-line dictionary of computer terms at (<http://wombat.doc.ic.ac.uk/foldoc/index.html>).

One of the problems in the computer field is the lack of a standard set of terms. For example, the Intel x86 architecture uses the term linear address whereas I prefer the

term virtual address, as used with the IBM S370 and S390. The reason for the preference is that the IBM S370 predates the virtual-memory version of the Intel x86 and established precedence. Thus, in the text, wherever possible, the term with precedence is used.

The von Neumann architecture is used in Chapter 1 to introduce the basic principles of computer architecture. Students should know what von Neumann proposed with this architecture and what its limitations were. These limitations were overcome in subsequent architectures while the underlying design remained unchanged.

Two architecture families will be extensively used for illustrations: The IBM S360/370/390 and the Intel x86 up to the Pentium Pro with MMX extensions. These two families represent the two most widely used processors today. For mainframe applications, the IBM family is preeminent in the world. The Intel family is found in approximately 90% of all personal computers. These two families are the most likely to be encountered by students after they finish their studies.

The sequence of presentation in this book differs from that found in other texts. It starts with a computer overview in Chapter 1 that covers the von Neumann architecture and some of its precursors. Chapter 2 introduces analytical performance modeling, learning curves, and other related evaluation techniques. The modeling technique will be used throughout the balance of the book.

Chapter 3 describes the issues in user instruction set design. The instruction set architecture taxonomy of addresses, data types, and operations is introduced. This chapter also provides a basic treatment of computer arithmetic and the IEEE floating-point standard.

Because a memory-centric view of a computer is taken, the text turns to memory implementation in Chapter 4. I believe that students should be exposed to memory systems before processors, as the process state is held memory and is the focus of how problems are solved. Further, I approach the memory system by first looking at virtual memory that provides the functionality required of large programs and large data sets.

Performance follows functionality, from a designer's point of view; thus caches are covered after virtual memory. The gap between the cache and the disk is then covered by describing interleaved memory and DRAM organizations. Performance models for hierarchical memories are used throughout the chapter to illustrate design trade-offs.

Following the chapter on memory, the issues of processor implementation are addressed. First, hardwired and microprogrammed implementations are covered in Chapter 5. Then Chapter 6 addresses the ideas and issues of pipelining. The techniques of performance models are interwoven with the functional descriptions of processor implementation.

Chapter 7 covers input/output. In this chapter we look at requirements for various workloads and examine how these requirements are satisfied in contemporary systems. Bus design and serial communications are discussed along with clock recovery and clock synchronization. A brief introduction to queuing theory is included in this chapter.

The concluding chapter, Chapter 8, covers extensions to the user instruction set architecture that support operating systems, virtual memory, caches, and multiprocessors.

x Preface

This chapter concludes with a description of the MMX extensions to the Intel Pentium Pro.

A final goal of this book is to enable students to read the current literature on computers with understanding and appreciation of the ideas presented. This is accomplished by exposing students to the basic principles of architecture and implementation, and by careful use of the nonstandard terminology of this field.

REFERENCES

- Petroski, H. (1996). *Invention by Design*, Harvard U. Press, Cambridge, MA, pp. 89–90.
- Descartes, R. (1692). *Rules for the Direction of the Mind*.