

INTRODUCTION TO COMPUTER ARCHITECTURE AND ORGANIZATION

HAROLD LORIN

INTRODUCTION TO COMPUTER ARCHITECTURE AND ORGANIZATION

HAROLD LORIN

IBM Systems Research Institute, New York



A Wiley-Interscience Publication

JOHN WILEY & SONS

New York . Chichester . Brisbane . Toronto . Singapore

Copyright © 1982 by John Wiley & Sons, Inc.

All rights reserved. Published simultaneously in Canada.

Reproduction or translation of any part of this work beyond that permitted by Section 107 or 108 of the 1976 United States Copyright Act without the permission of the copyright owner is unlawful. Requests for permission or further information should be addressed to the Permissions Department, John Wiley & Sons, Inc.

Library of Congress Cataloging in Publication Data.

Lorin, Harold.

Introduction to computer architecture and organization.

"A Wiley-Interscience publication."

Includes bibliographies and index.

1. Computer architecture. I. Title.

QA76.9.A73L67 1982 621.3819'52 82-8640

ISBN 0-471-86679-2 AACR2

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

PREFACE

This book is an introduction to computers for people who know something about computing. It is aimed at the population of computer professionals and paraprofessionals, casual users, and all others who have some view of computing but know nothing about the logical and organizational nature of a computing device.

As the computer field has matured, it, like all advancing areas of knowledge, has fragmented into a number of specialties. There are a number of very competent computer professionals who know many aspects of software specification, data base design, large application development technique, and so on, but do not know much about the instrument on which their miracles are performed. There is also a large population of computing professionals who know about the technology of computer manufacture or packaging digital technology but have never formed a complete image of the nature of the device they manufacture or whose subassemblies they design. Finally, there is the growing population of casual users whose curiosity has been piqued by their contact with computers and who would like to know more about them.

This book is for all these people. It assumes that the reader has some knowledge of languages like BASIC, FORTRAN, and COBOL, and, starting with known concepts, guides the reader through various stages of computer hardware architecture and organization. It is my opinion, however, that the slightest familiarity with the general ideas of higher-level languages is a sufficient starting point.

This book was developed over a period of years as a result of my interaction with both graduate and professional students. A number of these students reviewed the manuscript, and the extent of coverage and many specific details are the result of their comments and class interaction.

What I try to do here is discuss ideas of architecture and organization generally to provide an overview of several immortal concepts. While specific architectures and machine organizations are mentioned in the book, for example, the IBM S/370, IBM Series/1, IBM 8100, Sperry Rand UNIVAC 1100, Cray Research CRAY-1, Digital Equipment Corporation VAX-11, Data General NOVA, INTEL 8080, and INTEL 8086, the book is not intended to describe specific machines but to use them as examples of architectural and organizational concepts. The danger in using specific examples is that the concepts they embody seem to age with the machines though the concepts are in fact ageless.

There are two great challenges in undertaking a book of this type. One, of course, is content. Naturally I hope that the content is well chosen, the omissions not disqualifying, and the simplifications appropriate. Aside from content, the greatest challenge is sequence. Topics are so *interrelated* that it is difficult to discuss many points without anticipating later ones. Where necessary, brief characterizing introductions are offered in one area to support discussion in another. This leads to some repetition, but the idea is to give more and more detail each time a topic is discussed. The sequence of chapters was discovered dynamically in class; it has been changed many times, and I am now confident that the conceptual flow of the book minimizes the temptation for readers to look ahead.

HAROLD LORIN

New York
September 1982

CONTENTS

PART ONE ARCHITECTURE

1. Architecture, Organization, and Implementation	3
1. Architecture, 3	
2. Elements of an Architecture, 5	
2.1. Mode of Data Representation, 5	
2.2. Size of the Basic Data Structure, 5	
2.3. Addressing Conventions, 6	
2.4. Register Model, 6	
2.5. Instruction Set, 7	
2.6. Interrupt Mechanism, 9	
2.7. Control States, 9	
2.8. Input/Output, 10	
3. Organization, 10	
4. Implementation, 11	
5. Architecture versus Organization, 11	
6. Languages and Language Processing, 12	
7. Hardware/Software Interfaces, 13	
2. Compilation	17
1. Introduction, 17	
2. Initial Statement Processing, 17	
3. Field Name Recognition, 18	
4. Symbol Table and Operand Stack, 18	
5. Statement Type, 19	
6. Right-Hand Processing, 20	
7. Intermediate Form, 20	
8. Next Statement, 21	
9. Code Generation, 21	
10. Post-Compilation Processes, 25	
11. Final Comment, 25	

3. Data Coding and Reference	27
1. Data Declaration, 27	
2. Data Modes and Representation, 27	
2.1. Pure Binary Characteristics, 29	
2.2. Binary-Coding Schemes, 30	
2.3. Floating Point, 30	
3. Addressable Units, 31	
3.1. Basic Word Size, 32	
3.2. Address Definition, 33	
3.3. Word Sizes and Processor Type, 34	
4. Nonscalar Representation, 34	
5. Self-Describing Data, 35	
4. Register Organization	37
1. Registers, 37	
2. Register and Register Model Characteristics, 38	
3. Register Operand Referencing, 39	
4. Operations on Registers, 41	
5. Multiple Register Sets, 42	
6. Operand Stacks, 43	
7. Stacks and Compilation, 49	
8. Some Example Operand Register Models, 52	
8.1. CRAY-1, 52	
8.2. UNIVAC 1100, 53	
8.3. VAX-11, 54	
9. Software Register Conventions, 54	
10. Basic Control Registers, 54	
11. Placement of Control Information, 55	
5. Memory Addressing Conventions	57
1. Addressing Conventions in Instructions, 57	
2. Fundamental Issue, 57	
3. Address Size Compression, 58	
4. Fewer Addresses, 60	
5. Mixed Addressing Forms, 61	
6. Registers in Address Formation, 62	
7. Indexing, 63	
8. Page, Base, and Segment Addressing, 66	
9. Segmentation, 68	
10. Indirect Addressing, 69	
11. Memory Protection, 70	

11.1.	Protection Keys, 71	
11.2.	Protection Rings, 71	
6.	Instruction Sets	74
1.	Instruction Set Organization, 74	
2.	Informal Groupings of Instructions, 74	
3.	Other Instruction Groupings, 79	
4.	Symmetric Instruction Sets, 79	
5.	Example Instruction Forms, 81	
5.1.	UNIVAC 1100, 81	
5.2.	Data General NOVA, 82	
5.3.	Other Formats, 83	
7.	Changes in Program Sequencing	86
1.	Sequence Alteration, 86	
2.	Branching, 86	
3.	Condition Code Logic, 89	
4.	Skip Logic, 92	
8.	Subroutine Linkage	95
1.	Concept of a Subroutine, 95	
2.	Subroutine Linkage, 96	
3.	Basic Linkage Instruction, 96	
4.	Basic Stack Linkage Control, 99	
5.	Parameter Passing, 101	
6.	Activation Record, 102	
7.	Stacks, Parameters, and Activation Records, 103	
8.	Addressing and Linkage, 105	
9.	Interrupt Mechanisms and Control States	107
1.	Interrupt Mechanisms, 107	
1.1.	Introduction to Interrupt Response, 107	
1.2.	Interrupt Inhibition, 109	
2.	Variations in Interrupt Architecture, 109	
3.	General Description of S/370 Interrupt Architecture, 110	
3.1.	Interrupt Classes, 110	
3.2.	The S/370 Program Status Word, 111	
3.3.	Organization of Lower Memory, 113	
3.4.	Interrupt Response, 114	
4.	Other Interrupt Schemes, 115	
5.	Control States, 116	

10. Virtual Memory	119
1. Concepts of Virtual Memory, 119	
2. Linear Virtual Memory, 121	
2.1. Basic Mechanisms, 122	
2.2. Memory Space Management, 123	
2.3. Address Formation in Paging Systems, 124	
2.4. Page Faults, 125	
2.5. Linear Virtual Memory and Protection, 127	
2.6. Segmentation in Linear Virtual Memories, 127	
2.7. Associative Memory Assist, 127	
3. Alternative Architectural Concepts, 129	
11. Input/Output	132
1. Input/Output, 132	
2. Disk Characteristics, 134	
3. Input/Output Architecture Features, 136	
4. Basic Processor I/O Architecture, 137	
5. Operating Systems Services, 140	
6. Input/Output Supervisor Flow, 142	
7. Control Languages, 143	
8. Higher Levels of Input/Output Support, 145	
Bibliography, 146	
12. The Power of an Architecture	147
1. Concept of Architectural Power, 147	
2. The Army-Navy Study, 148	
2.1. Absolute Criteria, 149	
2.2. Quantitative Criteria, 151	
3. Measures of an Architecture, 154	
4. Functional Instructions, 155	
5. Architectural Comparisons, 156	
6. Ideal Architecture, 158	
6.1. Addressing and Memory Referencing, 159	
6.2. Characteristics of Ideal Architecture, 159	
7. Architecture, Compilability, and Design Complexity, 161	
8. The Great Debate, 163	
9. Level of an Architecture, 163	
9.1. Support of an Operating System, 165	
9.2. Extension Toward Higher-Level Languages, 166	
Bibliography, 168	

PART TWO ORGANIZATION AND IMPLEMENTATION

13. Concepts in Organization and Implementation	173
1. Introduction, 173	
2. Organization and Implementation Decisions, 174	
2.1. Width of Data Paths, 174	
2.2. Degree of Circuit Sharing, 174	
2.3. Definition of Specialized Units, 175	
2.4. Parallelism of Functional Units, 175	
2.5. Buffering and Queuing, 175	
2.6. Prediction, 175	
2.7. Underlying Technology, 176	
2.8. Functional Implementation, 176	
3. Price/Performance Goals, 176	
4. Architecture, Organization, and Price/Performance, 177	
14. Basic Concepts of Instruction Execution	180
1. Basic Cycles, 180	
2. Execution of an Instruction—Instruction Stages, 182	
3. Instruction Stages and Machine Cycles, 186	
15. Organization for Increased Performance	189
1. Organizational Approaches to Faster Machines, 189	
2. Instruction Times, 190	
3. Cycles Reduction by Stage Speedup, 190	
4. Stage Redefinition, 191	
5. Resequencing Decode and Address Formation, 192	
6. Concluding Remarks, 195	
16. Extended Lookahead	197
1. Inter-Instruction Overlap, 197	
2. Storing Values into Memory, 202	
3. Memory Read and Write Contention, 204	
4. Memory Reference Delays, 204	
5. Increasing Lookahead, 206	
6. Details of I/E Function and Relationship, 207	
7. Buffering and Extended Lookahead, 210	
17. Parallel Instruction Execution	214
1. Multiple Instruction Execution, 214	
2. Considerations in Multiple E-Box Design, 215	
3. Populations of E-Boxes, 216	

4.	Delivery of Instructions to E-Boxes, 217	
5.	Delivery of Operands to E-Boxes, 218	
6.	Relationship between E-Boxes and Registers, 219	
7.	Instruction Sequencing, 220	
	7.1. Sequencing Techniques, 221	
	7.2. Intersecting Memory References, 226	
8.	Code Equivalence and Rearrangement, 227	
9.	Characteristics of E-Unit Designs, 228	
	9.1. Single-Level Units, 228	
	9.2. Multilevel Pipeline Units, 229	
10.	Vector Manipulation, 231	
11.	Multiple I-Stream Machines, 233	
18.	Memory Organization	238
1.	Basic Goals of Memory Design, 238	
2.	Enough Memory, 238	
3.	Memory Organization, 239	
	3.1. Banking, 240	
	3.2. Interleaving to Speed Memory Response, 241	
4.	Memory Lookahead, 242	
5.	Memory Partitioning Techniques and Processor Lookahead, 243	
6.	Limits on Memory Partitioning, 243	
7.	Memory Times, 244	
8.	A Memory/Processor Interconnect Organization, 245	
9.	Alternative Interconnection Organization, 248	
19.	Memory Hierarchies	251
1.	Notion of Hierarchy, 251	
2.	Instruction Buffers, 253	
	2.1. Basic Buffer Fill Technique, 256	
	2.2. Flexible Priority Driven Buffer Fill, 257	
	2.3. Short Loop Mode, 257	
3.	Branch Instructions and Instruction Buffering, 257	
4.	Operand Buffers, 259	
5.	Cache, 260	
	5.1. Cache Mapping and Loading, 262	
	5.2. Cache Contents Management, 265	
	5.3. Stores and Changes to Cache, 266	
	5.4. Some Concluding Comments, 266	

6.	Other Hierarchical Notions, 267	
7.	Instruction and Control Memories, 268	
8.	Final Comments on Memory, 269	
20.	Input/Output Design	271
1.	Basic Concepts, 271	
2.	Programmed Input/Output, 272	
3.	Processor Overlap, 276	
4.	More Elaborate Input/Output Design, 278	
5.	Buffered Input/Output, 281	
6.	Direct Memory Access Organizations, 281	
7.	Multiple Units, 282	
8.	General-Purpose Channels, 284	
9.	Input/Output Processor Concepts, 286	
10.	Processor Support of Input/Output Functions, 288	
21.	Overview of Implementation	290
1.	Introductory Observations on Technology, 290	
1.1.	Physical Change, 290	
1.2.	Performance and Price/Performance, 292	
1.3.	Storage, 292	
1.4.	Summary Remarks, 294	
2.	Computer Building Blocks, 294	
3.	Architecture, Design, and Implementation, 296	
4.	Microprocessors and Microcomputers, 297	
5.	Basics of Implementing One Architecture with Another, 299	
6.	Considerations in Implementing One Architecture with Another, 301	
7.	Microcode Types, 303	
8.	Extending an Architecture, 304	
	Bibliography, 305	
	INDEX	307

PART ONE

ARCHITECTURE

Architecture, Organization, and Implementation

1. ARCHITECTURE

The exact meaning of the word *architecture* in the context of computers is a little uncertain. In general, architecture refers to the visible characteristics of a system as seen by a person or a program creating code capable of running on the machine. It is common, however, to use the word to mean the view of a machine shown by its *assembly language*. An assembly language is a low-level programming language in which the basic characteristics of a computer system are more directly represented than in a language like COBOL or FORTRAN. The assembly language programmer is aware of memory locations used in the machine, the actual instructions of a machine, and possibly the general speed of instructions. He* is not necessarily aware of the underlying organization of the machine in terms of operational logical units or of the hardware technology used in the construction of the machine.

The word architecture is frequently used to mean the visible characteristics of only the element actually performing instructions, that is, the *processor* of a computer system. Since a programmer may view other elements of a computer system through the processor, architecture also more generally means the characteristics of all the component elements of a system that might concern a programmer working at machine level. These components include processor(s), *memory*, and *input/output subsystems*.

A processor is a unit that interprets instructions and changes data in conformity with the instructions of a program. A *computer* may have one or more processors either dedicated to the execution of specialized instructions or capable of performing all instructions. Memory is a device that holds instructions to be executed and data to be operated on. Input/output subsystems are collections of units that connect processors and memories with the devices (1) that interface with the outside world (printers, terminals, sensors, etc.) or (2) on

* Throughout this book "he" is used to mean "he" or "she." This avoids the cumbersome "he/she."

which large amounts of additional data or instructions are stored and from or to which data is moved to or from memory.

The word architecture is also used to suggest the relationships among various building block elements of hardware and software systems. A system may be thought of as having several *architectural levels*. Figure 1 shows a representation of architectural levels. Each rectangle represents a set of functions and the horizontal lines are the interfaces between the functions. The first part of this book discusses the architectural level represented by the line labeled *machine-architecture interface*. The second part of the book addresses the levels that support this interface.

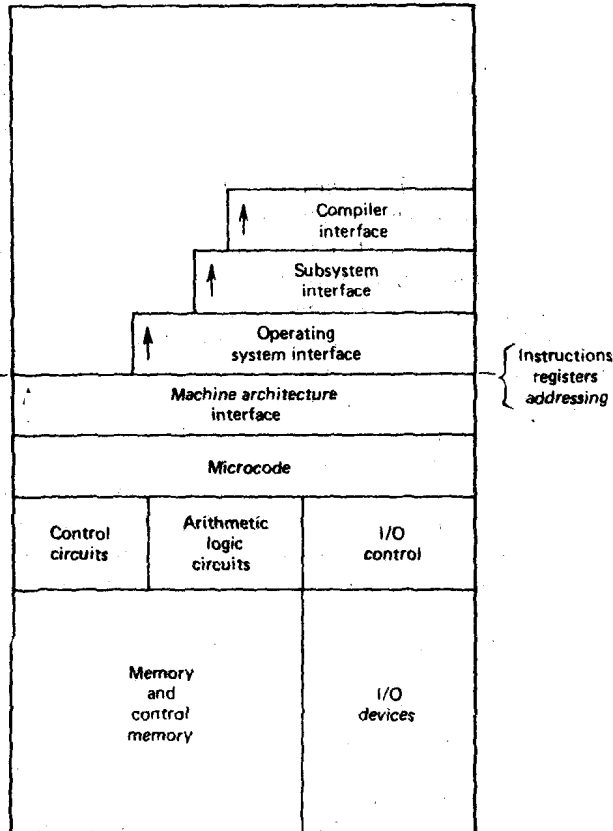


Figure 1. Architectural levels.

2. ELEMENTS OF AN ARCHITECTURE

From the point of view of a programmer looking at a computer system from an architectural level, a system has a set of characteristics that define how data will be represented and referenced, the operations that can be performed on data, and other features that determine how sequences of operations must be executed to achieve a computational result.

2.1. Mode of Data Representation

This refers to the representation of values in the processor and memory and the manner in which strings of binary digits should be interpreted at different points in the architecture. The representation of information may be in either pure base 2 (binary) or encoded form generically called the *binary-coded decimal*. There is also a form of scientific notation, called *floating point*, where values and exponents may be represented separately in either pure binary or coded form.

When a value is to be manipulated by a processor, the circuits of the processor will be designed to treat the bits of that value as either pure binary, binary-coded decimal, or *floating point*. For example, a unit designed to add binary numbers will produce an erroneous result for the logic of a program if the data presented to it is in binary-coded decimal form. As another example, a program must know the mode of representation for various values printed on a system printer. Devices like printers always require data to be in some coded decimal form, which may not be the form in which the data is arithmetically manipulated in the processor. The program must issue instructions to put the data in proper form for printing.

2.2. Size of the Basic Data Structure

This refers to the organization of values in the memory of the system and in the processor. The basic architectural data structures are *bits*, *bytes*, *characters*, and *words*.

A computer system's memory is organized into *locations*, which may be visualized as elements in a one-dimensional array. Each element has a name, or an *address*. Thus we speak of a memory having 4096 locations; each location can be directly referenced according to its position on the list. For example, there are memory locations 0000, 0001, 0002, . . . , 4094, 4095. When one of these numbers appears in control circuits connected to memory either the contents (the value that exists in the location) will be brought from memory to the circuits of the processing unit or information in the processing unit will be stored in the memory, depending on the work associated with the address.

Each location in a memory has a fixed size that represents the number of *bits* (binary digits) held in that location. Thus we speak of machines with 8, 16, 24, and 32-bit memories. This describes the number of bits that will be transferred