

SURVEY OF ADVANCED MICROPROCESSORS

SURVEY OF ADVANCED MICROPROCESSORS

Andrew M. Veronis, Ph. D.

Professor

Department of Computer Science

University of Maryland, Eastern Shore

and

Whiting School of Engineering

The Johns Hopkins University

江苏工业学院图书馆
藏书章



VAN NOSTRAND REINHOLD

New York

Copyright © 1991 by Van Nostrand Reinhold

Library of Congress Catalog Card Number 90-21750

ISBN 0-442-00120-7

All rights reserved. No part of this work covered by the copyright hereon may be reproduced or used in any form or by any means -- graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems -- without written permission of the publisher.

Manufactured in the United States of America

Published by Van Nostrand Reinhold

115 Fifth Avenue

New York, New York 10003

Chapman and Hall

2-6 Boundary Row

London, SE1 8HN

Thomas Nelson Australia

102 Dodds Street

South Melbourne 3205

Victoria, Australia

Nelson Canada

1120 Birchmount Road

Scarborough, Ontario M1K 5G4, Canada

16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

Library of Congress Cataloging-in-Publication Data

Veronis, Andrew.

Survey of advanced microprocessors / Andrew Veronis.

p. cm.

Includes index.

ISBN 0-442-00120-7

1. Microprocessors. 2. Computer architecture. I. Title.

QA76.5.A48 1991

004.22--dc20

90-21750

CIP

Preface

Microprocessors have come a long way since their conception. They have become formidable processing tools, and we encounter them in almost every part of our daily activities, from the kitchen with its microwave oven to the cockpit of a sophisticated aircraft.

The purposes of this book are to “walk through” the current microprocessor technology and briefly to describe some of the most advanced microprocessors available. The book is a survey of advanced microprocessors, aimed particularly at the engineering manager rather than the design engineer.

Chapter One outlines the history of microprocessors and describes some terminology used in computer architecture.

Chapter Two discusses advanced computer concepts, such as data and data types, addressing modes, pipelining, and cache memory.

Chapter Three describes new computer architectures, such as reduced-instruction-set computers (RISCs) and very-long-instruction-word computers. RISC architecture has become very popular among designers.

Chapter Four discusses an architecture, data-flow, which is a departure from the conventional von Neumann architecture. NEC has applied the dataflow architecture on the design of a very sophisticated image processing chip, the NEC PD7281.

Chapters Five and Six are case studies, describing the Am29000 and the Transputer, respectively.

Chapter Seven describes microprocessors specifically designed for digital signal processing.

Chapter Eight discusses micromultiprocessing and describes the various topologies currently used.

Contents

Preface vii

Chapter 1. INTRODUCTION 1

THE HISTORY OF MICROPROCESSORS 1

ADVANCED MICROPROCESSORS 7

TERMINOLOGY 9

Von Neumann Architecture vis-a-vis Harvard Architecture 12

Parallel versus Multiprocessor 14

Multiprogramming 16

Topology 16

Orthogonality 17

Granularity 17

SUMMARY 17

NOTES 18

Chapter 2. ADVANCED COMPUTER CONCEPTS 19

GENERAL COMMENTS 19

DATA AND DATA TYPES 19

Boolean 20

Integer 20

Ordinal 21

Real Operands 21

Literal 22

Pixels 22

Arrays and Records 22

SEMANTIC GAP 23

X contents

INSTRUCTION SETS	23
ADDRESSING MODES	24
Register Direct Modes	27
Memory Address Modes	27
Special Address Modes	31
Program Control Modes	32
Inherent Mode	33
PIPELINING	34
General Description	34
Reduction of Pipeline Latency	37
Depth	39
MEMORY MANAGEMENT	41
General Design Considerations	41
Segmentation	42
Paging	46
Cache Memory	48
Virtual Memory	57
SUMMARY	61
NOTES	61

Chapter 3 . REDUCED-INSTRUCTION-SET, WRITABLE-INSTRUCTION-SET, AND VERY-LONG-INSTRUCTION-WORD COMPUTERS 63

REDUCED-INSTRUCTION-SET COMPUTER (RISC)	63
Philosophy	63
Evolution	64
Characteristics	65
Implementation in the Berkeley RISC	67
Pipelining Caveats	73
Motorola 88100	78
The AMD Processors	93
The MIPS Processor	99
Hewlett-Packard SPECTRUM	102
ACORN ARM (VL86C010)	102

Integrgraph Clipper	105
Gallium Arsenide RISC Processor	109
WRITABLE-INSTRUCTION-SET COMPUTER (WISC)	116
VERY-LONG-INSTRUCTION-WORD COMPUTER (VLIW)	119
Philosophy	119
Intel i860	122
SUMMARY	127
NOTES	127

Chapter 4. DATA-FLOW COMPUTERS 129

CONCEPT BEHIND DATA FLOW	129
EXAMPLES OF DATA-FLOW COMPUTERS	140
Static Data-Flow Computers	141
Dynamic Data-Flow Computers	147
THE NEC PD7281 DATA-FLOW MICROPROCESSOR	152
SUMMARY	155
NOTES	156

Chapter 5. CASE STUDY: THE AM29000 159

INTRODUCTION	159
INTERNAL ORGANIZATION	159
Pipeline Structure	161
Registers	161
Instruction Fetch Unit	164
Branch Target Cache	166
Program Counter Unit	167
EXECUTION UNIT	169
Register File	169
Address Unit	169
Arithmetic-Logic Unit	170

XII contents

MEMORY-MANAGEMENT UNIT	171
DEVELOPMENT TOOLS	171
Hardware	171
Software	174
SUMMARY	177

Chapter 6. CASE STUDY: THE INMOS TRANSPUTER 179

INTRODUCTION	179
HARDWARE DESCRIPTION OF TRANSPUTER	179
LANGUAGE	182
Primitive Processes and Constructs	183
Compiler	185
Hierarchical Decomposition	186
Alternative Construct	187
Parallel Construct	188
Arrays	189
Replicator	190
Time	190
Types and Data Structures	191
Compile-Time Allocation	191
Program Development	192
Configuration	193
Programs	194
Output Guards	197
SUMMARY	198
NOTES	199

Chapter 7. MICROPROCESSORS FOR SIGNAL AND IMAGE PROCESSING 201

INTRODUCTION	201
TEXAS INSTRUMENTS' TMS320 FAMILY	201

TMS32010 and TMS320C10 Digital Signal Processors	203
TMS32011 Digital Signal Processor	203
TMS32010-25 Digital Signal Processor	205
TMS32020 and TMS320C25 Digital Signal Processors	205
WE DSP32 (AT&T) DIGITAL SIGNAL PROCESSOR	207
Data Arithmetic Unit	207
Control Arithmetic Unit	210
Memory	211
Serial and Parallel Input-Output	211
INMOS A100 DIGITAL SIGNAL PROCESSOR	211
IMS A100 Cascadable Signal Processor	212
IMS A110 Image and Signal Processing Subsystem	214
TEXAS INSTRUMENTS' TMS34010 GRAPHICS PROCESSOR	217
SUMMARY	219

Chapter 8 . MICROPROCESSORS AND MULTIPROCESSING 221

GENERAL	221
TOPOLOGIES	222
Link-Oriented Networks	224
Bus-Oriented Topologies	230
CASE STUDY OF MULTIPROCESSING WITH THE TRANSPUTER	232
SUMMARY	242

Index	245
-------	-----

Chapter 1

Introduction

THE HISTORY OF MICROPROCESSORS

The transistor and the integrated circuit are among the fundamental discoveries in electronics. To date, however, the most significant invention may be the microprocessor.

"Announcing a New Era in Integrated Electronics," Intel advertised its first microprocessor in November 1971. The prophecy already has been fulfilled with delivery of the 8008 in 1972, the 8080 in 1974, the 8085 in 1976, the 8086 in 1978, the 80286 in 1983-84, the 80386 in 1986, and, more recently, the i860 and i960. When this book is published, Intel will have introduced the 80486.

During this interval, throughput has improved 100-fold, the price of a microprocessor has declined from \$500 to \$5, and the availability of microprocessors has revolutionized design concepts in countless applications. Microcomputers have entered our homes, offices, and automobiles. A four-transputer board (described later), plugged into a personal computer, will yield the same performance as a much larger VAX 8600.

In the late 1960s, it became clear that the practical use of large-scale integrated circuits depended on defining chips to have:

- a high gate-to-pin ratio;
- a regular cell structure; and
- a large standard-part market.

Intel Corporation was formed, in 1968, to exploit the semiconductor memory market with circuits that fulfilled these criteria. Easy semiconductor read-write memories (RAM), read-only memories

2 survey of advanced microprocessors

(ROM), and shift registers were welcomed wherever small memories were needed, particularly in calculators and CRT terminals.

In 1969, Intel engineers began to study ways by which to integrate and partition the control functions of these systems into LSI chips. At the same time, other companies (notably Texas Instruments) explored ways to reduce the time needed to develop custom integrated circuits.

An alternative approach envisioned a customer's application as a computer system that required a control program, i/o monitoring, and arithmetic routines, rather than as a collection of special-purpose logic chips. Intel drew on its strength in memory to partition systems into RAM, ROM, and single controller chips—i.e., CPU.

Intel embarked on the design of two customer-sponsored microprocessors—the 4004 for a calculator and the 8008 for a CRT terminal. The 4004 replaced six chips usable by only one customer. Since the first microcomputer applications were known and easy to understand, instruction sets and architectures were defined in a matter of weeks. As programmable computers, their uses could be extended indefinitely.

The 4004 and the 8008 are both complete CPU on a chip and have similar characteristics, but, because the 4004 was designed for serial BCD arithmetic and the 8008 for 8-bit character handling, their instruction sets are different.

Computer Terminal Corporation (subsequently known as Datapoint) contracted for Intel to make a pushdown stack chip for a processor to be used in a CRT terminal. Datapoint intended to build a bit-serial processor in TTL logic, using shift register memory. Intel counterproposed that the entire processor could be implemented on one chip. This processor was to become the 8008 and, along with the 4004, was to be fabricated using PMOS, the then current memory-fabrication technology. Due to the long lead time that Intel required, Datapoint marketed its serial processor, and, thus, compatibility constraints were imposed on the 8008.

Datapoint specified most of the instruction set and register organization. Intel modified the instruction set so that the processor would fit on one chip and added instructions to make the device more “general purpose,” for, although Intel was developing the 8008 for a

specific customer, the company wanted the option of selling the chip to others. Since, in those days, Intel was using only 16- and 18-pin packages, the company chose to use 18 pins for the 8008, rather than design a new package for a supposedly low-volume chip.

The 8008 architecture is quite simple compared with those of today's microprocessors. The data-handling facilities provide for byte data only. The memory space is limited to 16 kilobytes, and the stack (small memory area for the temporary storage of data) is on the chip and holds only eight data items. The instruction set is small but symmetrical, with only a few operand addressing modes available. An interrupt mechanism is provided, but there is no way to disable interrupts.

The chip supports eight 8-bit input ports and 24 8-bit output ports. The instruction set can address each port directly. The chip's designers had felt that output ports were more important than input ports, which always could be multiplexed by external hardware under control of additional output ports.

One of the interesting things about the advent of the microprocessor era is that, for the first time, users were given access to the memory bus and could define their own memory structure. No longer were users confined to what vendors offered, as users had been with minicomputers. A user had, for example, the option of putting i/o ports inside a memory address (memory-mapped i/o) instead of a separate i/o space (direct i/o).

The 8008 contains two register files and four 1-bit flags (carry, parity, sign, and zero). The register files are the scratchpad and address stack.

The scratchpad contains an 8-bit accumulator (A) and six additional 3-bit registers (B, C, D, E, H, and L). All of the arithmetic operations use the accumulator as one of the operands and the memory for the result, as is typical in 8-bit microprocessors. All seven registers can be used interchangeably for on-chip temporary storage.

The instruction set of the 8008 consists of scratchpad-register instructions, accumulator-specific instructions, transfer-of-control instructions, i/o instructions, and processor-control instructions.

By 1973, memory-fabrication technology had advanced from PMOS to NMOS. As an engineering exercise, Intel used the 8008

4 survey of advanced microprocessors

layout masks with the NMOS process, to obtain increased operating speeds and, after a short study, determined that a new layout was required. The company decided to enhance the processor at the same time and to utilize a new 40-pin package, made practicable for high-volume calculator chips. The result was the 8080 processor.

The 8080 was the first processor specifically designed for the microprocessor market. The chip was constrained to include all of the 8008 instructions, although not necessarily with the same op-codes. Thus, a user's software was portable, but the ROM chips that contained the firmware required replacement. The main objectives in the design of the 8080 were a 10-to-1 improvement in throughput, elimination of many of the then-apparent shortcomings of the 8008, and provision of processing capabilities not found in the 8008. These capabilities included the handling of 16-bit data types (mainly for address calculations) and BCD arithmetic, enhancement of addressing modes, and improvement of interrupt processing.

The lower price of memory chips had made larger memory capacity more practicable. Thus, another goal was the addressing of more than 16K bytes. Symmetry, however, was not a goal, as the benefits of symmetrical extension could not justify the resultant increase in chip size and opcode space.

The 8080 contains a file of seven 8-bit, general-purpose registers, a 16-bit program counter, a 16-bit stack pointer, and five 1-bit flags.

The registers in the 8080 are the same seven 8-bit registers in the 8008 scratchpad. To handle 16-bit data, however, certain 8080 instructions operate on register pairs BC, DE, and HL.

The seven registers can be used interchangeably for on-chip temporary storage. The three register pairs can be used for address manipulations, but their roles are not interchangeable; an 8080 instruction allows operations on DE but not BC. There are address modes that access memory indirectly through BC and DE but not through HL.

The A register serves as the accumulator. The HL pair is used as a memory pointer.

The 8080 has a total memory-access capacity of 64K. The stack is contained in memory, rather than on the chip, to remove the restriction of only seven levels of nested subroutines. The entries on the

stack are 16 bits wide. The 16-bit stack pointer is used to locate the stack in memory. The execution of a "call" instruction causes the contents of the program counter to be "pushed" onto the stack, and the "return" instruction causes the last stack entry to be "popped" into the program counter. The stack pointer contains addresses that lead toward the lower end of memory (hence the term *pushdown stack*). This action simplifies both indexing into the stack from the user's program (positive indexing) and displaying the contents of the stack from a front panel.

A programmer can access the stack pointer in the 8080 directly, unlike the 8008. Furthermore, the stack itself is directly accessible, and instructions enable a programmer to push and pop 16-bit items onto the stack.

A couple of flags were added to the 8080 processor status code — i.e., the auxiliary carry flag, useful in BCD arithmetic, and the overflow flag.

Expanded input-output instructions allow the contents of any of the 256 8-bit ports to be transferred either to or from the accumulator. Nonetheless, the bottlenecking weakness inherent in uniprocessors, particularly ones that use registers dedicated to specific tasks, is here. Although the 8080 is better than its predecessor the 8008, it still uses the A register for both i/o and arithmetic operations and, thus, precludes concurrency of these operations.

By 1976, technological advances allowed Intel to consider enhancement of the 8080. Among the several weaknesses in the 8080 are the use of a triple power supply and the need for both an oscillator chip and a system controller chip. The company set out to produce a processor that would operate with a single power supply and would require fewer, if any, support chips. The resultant processor was called the 8085 and was constrained to be machine-code compatible with the 8080. Thus, extensions to the instruction set could use only 12 unused opcodes of the 8080.

Closer examination of the so-called enhanced 8085 reveals that its internal organization does not differ substantially from that of the 8080. The chips have the same register sets and the same flags. A couple of new instructions (Set Interrupt Mask and Reset Interrupt Mask) had been added. The only significant enhancements are the

use of an on-chip oscillator and the use of a single power supply. Several new instructions that had been included never were announced because of their ramifications on the 8085 support software and because of the compatibility constraints that the instructions would have imposed on the forthcoming 8086.

The 8086, a processor with a 16-bit-wide data structure, was designed to provide an order-of-magnitude increase over both the 8080 and the 8085. The processor was to be compatible with the 8080 at the assembly language level, so that existing 8080 software could be reassembled and correctly executed on the 8086. Therefore, the 8080 register set and instruction set were to appear as logical subsets of the 8086 registers and instructions. By utilizing a general-register structure, Intel could capitalize on its experience with the 8080 to obtain a processor with the highest degree of sophistication. Strict 8080 compatibility, however, was not attempted, especially in areas where compatibility would have compromised the final design.

The organization of the 8086 processor comprises a memory structure, a register structure, an instruction set, and an external interface. The processor can access up to one megabyte of memory and up to 64K input-output ports. By today's microprocessor standards, these capabilities do not seem significant, but they were when the 8086 was introduced in the mid-1970s.

The 8086 uses "segmented" addressing; that is, memory is divided into 64K byte segments, and referencing of a datum inside a segment is accomplished by the addition of an address and the contents of a segment register. The problems arising from segmented addressing are discussed in Chapter 2.

The 8086 contains a total of 13 16-bit registers and 9 1-bit flags. The registers are divided into three files of four registers each. The 13th register, called the *instruction pointer*, is not directly accessible by a programmer; this register is manipulated with control-transfer instructions. One cannot help but notice the compatibility of the 8086 register set with those of the 8080/8085.

The 8086 processor has seen days of glory since IBM decided to use a subset, the 8088, as well as an enhanced version, the 80286, in the IBM line of personal computers. Many IBM-compatible personal computers, called clones, followed suit.

Intel microprocessors were chosen for description as a tribute to the first manufacturer of these marvels of technology. Other semiconductor manufacturers, however, did not stand idle. Zilog, a company formed by former Intel engineers, produced the Z80, which outsold every other 8-bit microprocessor on the market. The Z80, an enhanced version of the 8080, is upward-compatible with the latter; that is, it can execute code written for the 8080, but it also has its own additional instructions.

Motorola, another semiconductor pioneer, came out with a full line of 8-bit and 16-bit microprocessors. The Motorola 68000 family of microprocessors offers several advantages over the 8086 processor.

The Motorola 68020 is a true 32-bit microprocessor. This processor has 16 32-bit, general-purpose registers, a 32-bit program counter, a 16-bit status register, a 32-bit vector base register, two 32-bit alternate function code registers, and two 32-bit cache-memory-handling registers. The 68020 can manipulate bit, byte, 16-bit, 32-bit, and 64-bit operands. The direct-addressing capacity of this chip is four gigabytes!

Another line of microprocessors is the bit-slice device. Some of these devices offer definite advantages but have not become as popular as the other microprocessors.

Microprocessors also have given us the tools with which to produce multimicroprocessor systems, such as the Hypercube and the Connection Machine. These types of systems can meet the demand for additional computing power to satisfy new requirements and to support complex applications.

The demand for faster and more sophisticated microprocessors has led to continued enhancement of microprocessor internal organizations.

ADVANCED MICROPROCESSORS

As the number of applications with more elaborate computational demands increases, more processing power is needed. This need can

8 survey of advanced microprocessors

be met, at the processor level, by relying on technological improvements to push the microprocessor beyond its current maximum capabilities or, at the system level, by extending the capabilities of a single microprocessor through concurrent execution.

Bottlenecking, which is inherent in von Neumann-type processors, is minimized through the use of pipelining.

The demand for more enhanced microprocessors and the technological improvements of very-large-scale integration have led to introduction of new microprocessors. A new type of organization, called reduced-instruction-set computer (RISC), was designed and prototyped at the University of California at Berkeley.

The RISC architecture is a radical departure from conventional thinking about microprocessor design. This type of design includes four important constraints. The execution time is one instruction per cycle. Instructions are intended to be as simple and fast as microinstructions on computers like the VAX and, to simplify implementation, also have the same size. To simplify system design, system memory is accessed only with LOAD and STORE instructions. This feature is well matched for a microprocessor optimized for keeping operands in internal registers. Finally, the RISC design epitomizes support of high-level languages.

Several microprocessor manufacturers have adopted the RISC architecture.

Other enhancements are the departure from the control-flow, von Neumann-type architecture and the introduction of the so-called data-flow architecture.

Sophisticated microcontrollers and dedicated digital-signal-processing microprocessors are now available. Texas Instruments markets a microprocessor with LISP on the chip. Intel markets the i860, which, according to the company, “delivers supercomputing performance in a single VLSI component”—a supercomputer on a chip. This processor uses a new architecture called very-long-instruction-word (VLIW), which, some claim, will replace RISC. The i860 is described in a later chapter.

The i960, a microprocessor designed for embedded applications, displays a lot of the features of the ill-fated 432.

INMOS markets the Transputer, a 32-bit microprocessor