

“The ‘must-have’ PC architecture reference set.”

—PC Magazine’s “Read Only” column

ISA SYSTEM ARCHITECTURE

THIRD EDITION



MINDSHARE, INC.

Tom Shanley / Don Anderson

**PC SYSTEM
ARCHITECTURE
S E R I E S**



world-class technical training

Are your company's technical training needs being addressed in the most effective manner?

MindShare has over 25 years experience in conducting technical training on cutting-edge technologies. We understand the challenges companies have when searching for quality, effective training which reduces the students' time away from work and provides cost-effective alternatives. MindShare offers many flexible solutions to meet those needs. Our courses are taught by highly-skilled, enthusiastic, knowledgeable and experienced instructors. We bring life to knowledge through a wide variety of learning methods and delivery options.

training that fits your needs

MindShare recognizes and addresses your company's technical training issues with:

- Scalable cost training
- Customizable training options
- Reducing time away from work
- Just-in-time training
- Overview and advanced topic courses
- Training delivered effectively globally
- Training in a classroom, at your cubicle or home office
- Concurrently delivered multiple-site training

MindShare training courses expand your technical skillset

- ☞ PCI Express 2.0®
- ☞ Intel Core 2 Processor Architecture
- ☞ AMD Opteron Processor Architecture
- ☞ Intel 64 and IA-32 Software Architecture
- ☞ Intel PC and Chipset Architecture
- ☞ PC Virtualization
- ☞ USB 2.0
- ☞ Wireless USB
- ☞ Serial ATA (SATA)
- ☞ Serial Attached SCSI (SAS)
- ☞ DDR2/DDR3 DRAM Technology
- ☞ PC BIOS Firmware
- ☞ High-Speed Design
- ☞ Windows Internals and Drivers
- ☞ Linux Fundamentals
- ... and many more.

All courses can be customized to meet your group's needs. Detailed course outlines can be found at www.mindshare.com

bringing life
to knowledge.

real-world tech training put into practice worldwide



MindShare Classroom



In-House Training



Public Training

Classroom Training

Invite MindShare to train you in-house, or sign-up to attend one of our many public classes held throughout the year and around the world. No more boring classes, the 'MindShare Experience' is sure to keep you engaged.

MindShare Virtual Classroom



Virtual In-House Training



Virtual Public Training

Virtual Classroom Training

The majority of our courses live over the web in an interactive environment with WebEx and a phone bridge. We deliver training cost-effectively across multiple sites and time zones. Imagine being trained in your cubicle or home office and avoiding the hassle of travel. Contact us to attend one of our public virtual classes.

MindShare eLearning



Intro eLearning Modules



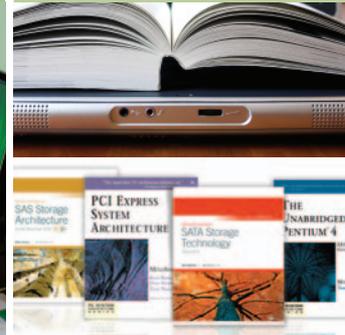
Comprehensive eLearning Modules

eLearning Module Training

MindShare is also an eLearning company. Our growing list of interactive eLearning modules include:

- **Intro to Virtualization Technology**
- **Intro to IO Virtualization**
- **Intro to PCI Express 2.0 Updates**
- **PCI Express 2.0**
- **USB 2.0**
- **AMD Opteron Processor Architecture**
- **Virtualization Technology**
- **...and more**

MindShare Press



Books



eBooks

MindShare Press

Purchase our books and eBooks or publish your own content through us. MindShare has authored over 25 books and the list is growing. Let us help make your book project a successful one.

Engage MindShare

Have knowledge that you want to bring to life? MindShare will work with you to "Bring Your Knowledge to Life." Engage us to transform your knowledge and design courses that can be delivered in classroom or virtual classroom settings, create online eLearning modules, or publish a book that you author.

We are proud to be the preferred training provider at an extensive list of clients that include:

ADAPTEC • AMD • AGILENT TECHNOLOGIES • APPLE • BROADCOM • CADENCE • CRAY • CISCO • DELL • FREESCALE
 GENERAL DYNAMICS • HP • IBM • KODAK • LSI LOGIC • MOTOROLA • MICROSOFT • NASA • NATIONAL SEMICONDUCTOR
 NETAPP • NOKIA • NVIDIA • PLX TECHNOLOGY • QLOGIC • SIEMENS • SUN MICROSYSTEMS • SYNOPSYS • TI • UNISYS

ISA System Architecture

Third Edition

MINDSHARE, INC.

*TOM SHANLEY
AND
DON ANDERSON*

*EDITED AND REVISED BY
JOHN SWINDLE*



Addison-Wesley Publishing Company

Reading, Massachusetts • Menlo Park, California • New York

Don Mills, Ontario • Wokingham, England • Amsterdam

Bonn • Sydney • Singapore • Tokyo • Madrid • San Juan

Paris • Seoul • Milan • Mexico City • Taipei

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and Addison-Wesley was aware of a trademark claim, the designations have been printed in initial capital letters or all capital letters.

The authors and publishers have taken care in preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

Library of Congress Cataloging-in-Publication Data

ISBN: 0-201-40996-8

Copyright © 1995 by MindShare, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. Printed in the United States of America. Published simultaneously in Canada.

Sponsoring Editor: Keith Wollman
Project Manager: Eleanor McCarthy
Production Coordinator: Deborah McKenna
Cover design: Barbara T. Atkinson
Set in 10 point Palatino by MindShare, Inc.

1 2 3 4 5 6 7 8 9 -MA- 9998979695

First printing, February 1995

Addison-Wesley books are available for bulk purchases by corporations, institutions, and other organizations. For more information please contact the Corporate, Government, and Special Sales Department at (800) 238-9682.

Contents

Foreword.....	xxiii
Acknowledgments.....	xxv

About This Book

The MindShare Architecture Series.....	1
Organization of This Book.....	2
Who This Book Is For.....	2
Prerequisite Knowledge.....	2
Documentation Conventions.....	3
Hex Notation.....	3
Binary Notation.....	3
Decimal Notation.....	3
Signal Name Representation.....	3
Identification of Bit Fields (logical groups of bits or signals)	4
We Want Your Feedback	4

Overview

System Kernel	6
Memory Subsystems.....	6
ISA Subsystem.....	7
Origins of ISA	7
The IBM PC	7
The IBM PC/AT.....	8
The ISA Concept.....	8

Part 1: The System Kernel

Chapter 1: Intro to Microprocessor Communications

Instruction Fetch and Execution	11
General.....	11
In-Line Code Fetching.....	13
Reading and Writing.....	15
Type of Information Read from Memory.....	16
Type of Information Written to Memory.....	16
The Buses	16
The Address Bus.....	17
Control Bus – Transaction Type and Synchronization	19
The Data Bus — Data Transfer Path	19

ISA System Architecture

Chapter 2: Introduction to the Bus Cycle

Introduction	21
Automatic Dishwasher – Classic State Machine Example.....	21
The System Clock – a Metronome	22
Microprocessor's Bus Cycle State Machine	23
Address Time	24
Data Time.....	25
The Wait State	26

Chapter 3: Addressing I/O and Memory

Evolution of Memory and I/O Address Space.....	29
Intel 8080 Microprocessor Address Space.....	29
8086 and 8088 Microprocessor Address Space	33
286 and 386SX Address Space	34
386DX, 486 and Pentium Processor Address Space.....	35
Memory Mapped I/O	36
The I/O Device	36

Chapter 4: The Address Decode Logic

The Address Decoder Concept.....	39
Data Bus Contention (Address Conflicts).....	41
How Address Decoders Work	41
Example 1– PC and PC/XT ROM Address Decoder	42
Background	42
The PC/XT ROM Address Decode Logic	44
Example 2 – System Board I/O Address Decoder	47

Chapter 5: The 80286 Microprocessor

The 80286 Functional Units.....	53
The Instruction Unit	55
The Execution Unit	55
General Registers	56
The Status and Control Registers	59
The Address Unit.....	63
The Segment Registers	63
Segment Register Usage in Real Mode	63
Code Segment (CS) & Instruction Pointer (IP) Registers	67
The Data Segment (DS) Register	69
The Extra Segment (ES) Register	70
Stack Segment (SS) & Stack Pointer (SP) Registers	71
Little-Endian Byte-Ordering Rule	74
Definition of Extended Memory.....	75
Accessing Extended Memory in Real Mode	76

The Bus Unit	81
Address Latches and Drivers	81
Instruction Prefetcher and 6-byte Prefetch Queue	81
Processor Extension Interface	82
Bus Control Logic	83
Data Transceivers	83
80286 Hardware Interface to External Devices	83
The Address Bus	83
How 80286 Addresses External Locations	85
The Data Bus	87
The Cardinal Rules	88
Cardinal Rule Number One	88
Cardinal Rule Number Two	89
Cardinal Rule Number Three	89
The Control Bus	89
Bus Cycle Definition Lines	90
Bus Mastering Lines	90
Protecting Access To Shared Resource	92
Ready Line	94
Interrupt Lines	95
Processor Extension Interface Lines	96
The Clock Line	97
The Reset Line	97
Protected Mode	98
Intro to Protected Mode and Multitasking Operating Systems	99
Segment Register Usage in Protected Mode	101

Chapter 6: The Reset Logic

The Power Supply Reset	107
Reset Button	108
Shutdown Detect	108
Hot Reset	109
Alternate (Fast) Hot Reset	111
Ctrl-Alt-Del Soft Reset	111

Chapter 7: The Power-Up Sequence

The Power Supply – Primary Reset Source	113
How RESET Affects the Microprocessor	115
Processor Reaction When Output Voltages Stabilize	115
The First Bus Cycle	116

Chapter 8: The 80286 System Kernel: the Engine

The Bus Control Logic	119
-----------------------------	-----

ISA System Architecture

The Address Latch	121
Address Pipelining	123
The Data Bus Transceivers	124
Data Bus Steering Logic	126
Scenario One – Read Even-Addressed Location in 8-Bit Device.....	128
Scenario Two – 8-Bit Read from Odd-Addressed Location in 8-Bit Device	129
Scenario Three – 8-Bit Write to Odd-Addressed Location in 8-Bit Device	131
Scenario Four – 16-Bit Write to 8-Bit Device	132
Scenario Five – 16-Bit Read from 8-Bit Device.....	133
Scenario Six – 8-Bit Read from 16-Bit Device	134
Scenario Seven – 16-Bit Read from 16-Bit Device.....	134
Scenario Eight – 8-Bit Write to 16-Bit Device	134
Scenario Nine – 16-Bit Write to 16-Bit Device.....	134
The Ready Logic	134
Access Time	134
Stretching the Transfer Time.....	135
The Default Ready Timer	135
Custom Ready Timers.....	136
Extending the Default Timing	137
Shortening the Default Timing	138

Chapter 9: Detailed View of the 80286 Bus Cycle

Address and Data Time Revisited	139
The Read Bus Cycle	141
Bus Cycle A	144
Bus Cycle B	145
The Write Bus Cycle	146
The Halt or Shutdown Bus Cycle	150
Halt	150
Shutdown.....	151

Chapter 10: The 80386 DX and SX Microprocessors

Introduction	153
The 80386 Functional Units	154
General	154
Code Prefetch Unit	155
Instruction Decode Unit.....	155
Execution Unit.....	156
General.....	156
The Registers	156
General Registers.....	156
Status, MSW and Instruction Registers	157
Debug Registers.....	159

Test Registers	160
Segmentation Unit	160
Paging Unit	161
Bus Unit	162
Protected Mode	162
Page Translation	164
Virtual Paging	164
Translation Lookaside Buffer	170
Virtual-8086 Mode	172
Automatic Self-Test	174
80386DX External Interface	175
The Address Bus	175
The Data Bus	183
The Control Bus	184
Bus Cycle Definition Outputs	185
Processor Extension Lines	185
Address Status Output	185
Pipelining Control Input	186
Dynamic Bus Sizing (BS16#)	186
80386SX External Interface	187
Interface Signal Differences	187
A0 or BLE#	188
Addressing Scheme, Data Bus Width Ramifications	188
Throughput and Compatibility Considerations	189

Chapter 11: The 80386 System Kernel

Introduction	191
The 80386SX System Kernel	192
The 80386DX System Kernel	194
80386DX System Kernel with Dynamic Bus Sizing	194
Introduction	194
Reading from an 8-Bit Device	196
One-Byte Read from an 8-Bit Device	196
Two-Byte Read from an 8-Bit Device	197
Four-Byte Read from an 8-Bit Device	198
Writing to an 8-Bit Device	200
One-Byte Write to an 8-Bit Device	200
Two-Byte Write to an 8-Bit Device	201
Four-Byte Write to an 8-Bit Device	203
Reading from a 16-Bit Device	205
Two-Byte Read from a 16-Bit Device	205
Four-Byte Read from a 16-Bit Device	206
Writing to a 16-Bit Device	207

ISA System Architecture

Two-Byte Write to a 16-Bit Device	207
Four-Byte Write to a 16-Bit Device.....	208
Reading from a 32-Bit Device	210
Writing to a 32-Bit Device	211
80386DX System Kernel without Dynamic Bus Sizing.....	212
Introduction.....	212
Reading from a 16-Bit Device	214
Two-Byte Read from a 16-Bit Device.....	214
Four-Byte Read from a 16-Bit Device	215
Writing to a 16-Bit Device	216
Two-Byte Write to a 16-Bit Device.....	216
Four-Byte Write to a 16-Bit Device.....	217

Chapter 12: Detailed View of the 80386 Bus Cycles

The Bus Cycle Types.....	221
Address Pipelining Overview	222
Memory or I/O Read Bus Cycle.....	222
Memory or I/O Write Bus Cycle.....	224
Address Pipelining Example	226
Interrupt Acknowledge Bus Cycle	230
Halt or Shutdown Bus Cycle	230

Part 2: Memory Subsystems

Chapter 13: RAM Memory: Theory of Operation

Dynamic RAM (DRAM) Memory	235
DRAM Addressing Sequence	236
Row and Column Address Source	238
DRAM Addressing Logic	239
Detailed Description of DRAM Addressing Sequence.....	242
How Data is Stored in DRAM.....	244
DRAM Refresh	244
Refresh Logic and RAS-only Refresh.....	245
CAS-before-RAS Refresh.....	248
Hidden Refresh	248
Self-Refresh.....	248
Destructive Read: Pre-Charge Delay and Cycle Time.....	250
DRAM Bank	251
DRAM Bank Width	253
DRAM Error Detection and Correction.....	256
DRAM Parity.....	256
Error-Checking-and-Correcting Memory	259
Page-Mode DRAM and Its Variations	259

Page Mode DRAM	259
Enhanced Page Mode DRAM	264
Burst and Nibble Mode DRAM	265
Static Column RAM (SCRAM)	267
Synchronous DRAM	268
Interleaved Memory Architecture	269
Static RAM (SRAM)	271

Chapter 14: Cache Memory Concepts

The Problem	273
The Solution	274
Principles of Locality	277
Temporal Locality	277
Spatial Locality	277
Cache Performance	278
Overall System Performance	278
Cache Consistency	279
Components of a Cache Subsystem	279
Cache Memory	279
Cache Management Logic	281
Cache Memory Directory	281
Intro to Cache Architecture, Coherency, Write Policies and Organization	282
Cache Architectures	283
Look-Through Cache	283
Look-Aside Cache	287
First- and Second-Level Caches	289
Combined (Unified) and Split (Dedicated) Caches	290
Cache Consistency (Coherency)	291
Causes of Cache Consistency Problems	292
Write Policy	292
Write-Through Cache Designs	293
Buffered Write-Through Designs	293
Write-Back Cache Designs	294
Bus Master/Cache Interaction	295
Bus Snooping/Snarfing	295
Coherency via Cache Flushing	297
Software-Enforced Coherency	298
Cache Organization and Size	299
Fully-Associative Cache	299
Direct-Mapped Cache (One-Way Set-Associative)	301
Two-Way Set-Associative Cache	304
Four-Way Set-Associative Cache	307
Least-Recently Used (LRU) Algorithm	307

ISA System Architecture

Cache Line Size	308
Cache Size	310
First-Level Cache Size	310
Second-Level Cache Size	310
Cache Addressing	311
I/O Information Not Cached	311
Non-Cacheable Memory	312
Testing Memory	313

Chapter 15: ROM Memory

ROM Memory — Theory of Operation	315
Introduction	315
Fusible-Link PROM	317
Masked ROM (MROM)	317
Eraseable Programmable Read-Only Memory (EPROM)	318
Electrically Eraseable Programmable Read-Only Memory (EEPROM)	319
Flash EEPROM	320
ROM's Interface to System	320
System Board ROM Memory	322
Testing	322
Shadow RAM	324
Shadow RAM and ROM Occupying Different Address Spaces	324
Shadow RAM and ROM Occupy Same Address Space	325
Double Mapping ROM and Shadow RAM Address Space	325
Recovering Unused ROM Address Space	326
32KB System Board ROM Configuration	327
64KB System Board ROM Configuration	329
ROMs on ISA Cards (Device ROMs)	330

Part 3: The Industry Standard Architecture

Chapter 16: ISA Bus Structure

Introduction	335
Address Bus	338
Data Bus	339
Bus Cycle Definition	339
Bus Cycle Timing	341
Device Size	343
Reset	344
DMA	345
Interrupts	347
Error Reporting Signal	349
Miscellaneous Signals	350

Chapter 17: Types of ISA Bus Cycles

Introduction	351
Transfers with 8-bit Devices	352
Transfers with 16-bit Devices	355
Standard 16-bit Memory Device ISA Bus Cycle	355
Standard 16-bit I/O Device ISA Bus Cycle	359
0-Wait State Access to 16-bit Memory Device	362

Chapter 18: The Interrupt Subsystem

What Is an Interrupt?	365
Microprocessor Response to Interrupt Request	366
Interrupt Acknowledge Bus Cycles.....	369
Saving Pointer to Interrupted Program	373
Clearing the Interrupt Enable Flag	378
Jumping to the ISR.....	378
Resuming the Interrupted Program	379
When One 8259 Interrupt Controller Isn't Enough	383
Servicing Requests to Slave Interrupt Controller	387
Interrupt Table Entry Assignments	388
IRQ2 Redirect	389
Shareable Interrupts in ISA Machines	391
Generating the Interrupt Request.....	391
Interrupt Table Initialization – Add-in Devices	392
Shared Interrupt Procedure	393
Phantom Interrupts	394
Programming the 8259	396
Introduction.....	396
Programming the Registers	396
Non-Maskable Interrupt Requests (NMI)	399
Software Interrupts	402
Software Exceptions	402
Software Interrupt Instruction	403
Protected Mode Interrupts	406

Chapter 19: Direct Memory Access (DMA)

DMA Concept	409
DMA Example	410
DMA Controller (DMAC)	414
DMA Transfer Types.....	414
DMA Transfer Modes	414
Single Transfer Mode.....	415
Block Transfer Mode	416

ISA System Architecture

Demand Transfer Mode.....	416
Cascade Mode.....	417
DMAC Priority Logic.....	417
DMA Bus Cycle.....	420
Byte or Word Transfers.....	423
DMAC Addressing Capability.....	423
Addressing ISA Memory.....	427
Addressing Local Bus Memory.....	428
Address Translation.....	430
Data Bus Steering.....	430
DMA Transfer Rate.....	431
DMAC Initialization During POST.....	431

Chapter 20: ISA Bus Masters

ISA Bus Master Capability.....	433
ISA Bus Masters in 80386DX and Higher Systems.....	435
Address Translation.....	438
Data Bus Steering.....	438
Bus Masters and DRAM Refresh.....	439

Chapter 21: RTC and Configuration RAM

Introduction.....	441
Accessing the RAM Locations.....	442
Address Decoder and RTC's Hardware Interface.....	442
Real-Time Clock Function.....	443
Using BIOS to Control Real-Time Clock.....	445
Configuration RAM Usage.....	446

Chapter 22: Keyboard/Mouse Interface

Keyboard/Mouse Interface.....	449
Keyboard.....	450
Mouse.....	450
8042 Local I/O Ports.....	450
Hot Reset.....	451
A20 Gate.....	451
Local Port Definition.....	452
System Interface.....	453
Command/Status Port.....	454
Data Port.....	456
BIOS Routine.....	457

Chapter 23: Numeric Coprocessor

Introduction.....	459
Setup.....	460

Contents

With Numeric Coprocessor Installed.....	460
Numeric Coprocessor Reset.....	463
Without Numeric Coprocessor Installed (Emulation).....	465
Weitek Numeric Coprocessor.....	467
Introduction.....	467
Systems Incorporating Just the Weitek.....	467
General.....	467
Weitek Handling of Intel Coprocessor Instructions.....	467
Systems Incorporating Both Coprocessor Types.....	468

Chapter 24: ISA Timers

ISA Oscillator, OSC.....	469
System Timer (Timer 0).....	469
Refresh Timer (Timer 1).....	471
Speaker Timer (Timer 2).....	471
Watchdog Timer.....	472
Slowdown Timer.....	474

Appendices

I/O Address Map.....	479
Glossary of Terms.....	491
Index.....	507

Chapter 1

This Chapter

This chapter defines the microprocessor's role in the system, the usage of memory, and defines the role of the address, control and data buses.

The Next Chapter

The next chapter, “Introduction to the Bus Cycle,” introduces the concept of the bus cycle and defines the basic sequence of events when the microprocessor uses a bus cycle to communicate with memory or an I/O device.

Instruction Fetch and Execution

General

The microprocessor is an engine with only one task in life — to continually read instructions from memory and execute them (perform the operations specified by the instructions).

An instruction tells the microprocessor to perform one of three basic types of operations:

- Read data from an external device.
- Write data to an external device.
- Perform an internal operation that doesn't involve reading from or writing to the outside world (such as math functions).

Note — As used in this book, the term external device refers to a device external to the microprocessor chip itself. Refer to figure 1-1.

The instructions fetched (read) from memory tell the microprocessor what to do. When a microprocessor-based system is initially powered up, the microprocessor knows what address in memory to fetch (read) its first instruction from. This location is known as the power-on restart address. A group of in-

ISA System Architecture

structions that cause the microprocessor to perform a particular task is referred to as a program. After fetching the first instruction from the restart address, the microprocessor is totally dependent on the program to tell it what to do.

In a properly functioning system, the microprocessor is never idle. At any given moment in time, the microprocessor is reading data from an external device, writing data to an external device, or executing an instruction that doesn't require that a read or write take place (for example, an instruction to add the contents of two of the microprocessor's internal registers together).

The microprocessor communicates with all external devices by reading data from them or writing data to them. The terms read and write are extremely important in any discussion of microprocessors. Always think of reading and writing from the point of view of the microprocessor rather than from that of the device the microprocessor is communicating with. The microprocessor does the reading and writing. Devices are read from and written to by the microprocessor. If you think of the terms read and write from the device's point of view, much of the information in this book will not make sense.

Chapter 1: Intro to Microprocessor Communications

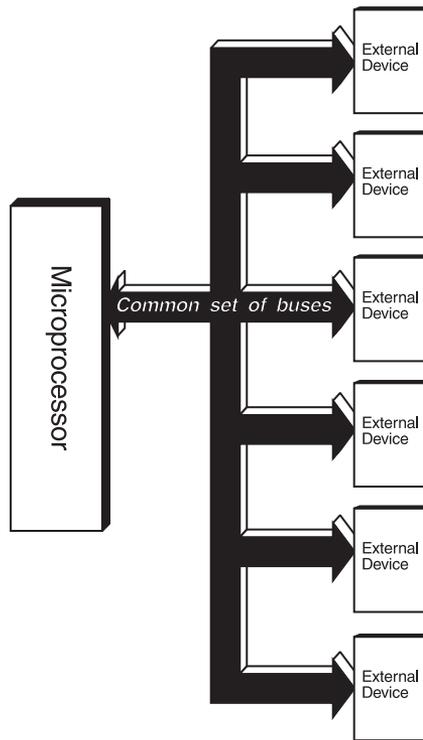


Figure 1-1. Microprocessor's Relationship to External Devices

In-Line Code Fetching

When a microprocessor fetches an instruction from a memory location and executes it, the address of the next instruction may or may not be specified as part of the currently executing instruction.

Case 1: The currently executing instruction doesn't specify the memory address of the next instruction. The microprocessor automatically assumes the next instruction is to be fetched from the next sequential memory location.

or

Case 2: The currently executing instruction specifies the memory address of the next instruction. Execution of the instruction, called a jump instruction, causes the microprocessor to alter its program flow (which is usually sequential). Rather than fetching its next instruction from the

ISA System Architecture

next sequential memory location, the microprocessor fetches it from the memory location specified by the instruction.

After executing a JUMP instruction, the microprocessor resumes fetching instructions from sequential memory locations until another JUMP instruction is executed. Most well-written programs do not contain an excessive number of JUMPs. Statistically then, the microprocessor is executing non-jump instructions the majority of the time. This means that the microprocessor is performing what is commonly referred to as in-line code fetches most of the time.

As an example, refer to figure 1-2. Assume that the microprocessor has just fetched the ADD instruction from location 00000h in memory. The flow of instructions fetched and executed proceeds as follows:

1. The microprocessor executes the ADD instruction fetched from location 00000h. Since it's not a JUMP instruction, the microprocessor fetches its next instruction from memory location 00001h.
2. The microprocessor executes the SUBTRACT instruction fetched from location 1. Since it's not a JUMP, the microprocessor fetches its next instruction from memory location 00002h.
3. The microprocessor executes the MOVE instruction fetched from location 2. Since it's not a JUMP, the microprocessor fetches its next instruction from memory location 00003h.
4. The microprocessor executes the JUMP 10000 instruction fetched from location 3. The JUMP instruction alters program flow. Rather than fetch its next instruction from the next sequential memory location (00004h), the microprocessor fetches it from memory location 10000h.
5. The microprocessor executes the MOVE instruction fetched from location 10000h. Since it's not a JUMP, the microprocessor fetches its next instruction from memory location 10001h.
6. And so on.

Chapter 1: Intro to Microprocessor Communications

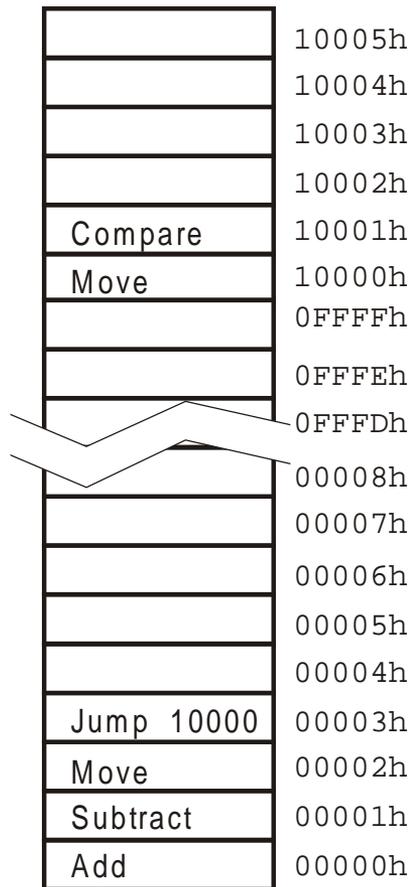


Figure 1-2. Sample Instruction Sequence

Reading and Writing

The microprocessor communicates with external devices under the following circumstances:

- To fetch (read) the next instruction from memory.
- When the currently executing instruction directs the microprocessor to read data from an external device.
- When the currently executing instruction directs the microprocessor to write data to an external device.

Chapter 2

The Previous Chapter

In “Introduction to Microprocessor Communications,” the microprocessor’s role in the system, usage of memory, and the role of the address, control and data buses were defined.

This Chapter

When the microprocessor must communicate with an external device, it uses its buses. The sequence of events necessary when using the buses to perform a read or write transaction is referred to as a bus cycle. This chapter introduces the microprocessor’s bus unit, the concept of a state machine, and defines address time, data time and the wait state.

The Next Chapter

The next chapter, “Addressing I/O and Memory,” provides a definition of an I/O device. It defines the method used by the x86 processors to distinguish memory and I/O addresses and the range of memory and I/O addresses available to the 8088, 8086, 80286, 80386, 80486 and Pentium microprocessors.

Introduction

The microprocessor’s internal bus unit is a state machine that performs the required bus cycle when the microprocessor must communicate with another device. The concept of the state machine is introduced, as well as the concept of the clock, or timebase, used by the state machine to define the duration of each state.

Automatic Dishwasher – Classic State Machine Example

Any task is easier to perform when divided into logical steps. A state machine is a device (either mechanical, electronic, or software-based) designed to per-

ISA System Architecture

form a task that can be divided into steps. Prior to starting the task, the state machine is said to be idle. When commanded to perform the task, the state machine leaves the idle state and moves through a series of steps, or states. Predefined portions of the overall task are performed during each state. The duration of each state is defined by a clock, or timebase.

In the case of the automatic dishwasher, the dishwasher is idle until you start it. A timer then begins to run, defining the duration of each state the dishwasher must pass through in order to accomplish the overall task of washing dishes. These states are:

- During the first state, the state machine wets down the dishes.
- During the second state, the dishes are soaped.
- During the third state, the dishes are rinsed.
- During the fourth and final state, the dishes are dried.
- The state machine then returns to the idle state.

Most dishwashers also have one or more switches the user can manipulate to alter the sequence of states or possibly to cause the machine to execute a particular state more than once. A perfect example would be the “Pot Scrubber” button. When pressed, this might cause the state machine to execute the “rinse” state two times instead of one.

The System Clock – a Metronome

Every x86 microprocessor has an input called CLOCK (or a similar name). This signal is produced when a voltage is applied to a crystal oscillator. It then begins to generate an electrical signal of a specific frequency. In essence, the oscillator acts as a highly accurate electronic tuning fork. The signal looks like Figure 2-1.



Figure 2-1. Crystal Oscillator Output

All microprocessors perform their operations in a highly organized, predefined fashion. Portions of each task are always performed during predefined time slots. The duration of each time slot is defined by the output of the crystal oscillator.

Chapter 2: Introduction to the Bus Cycle

Indirectly, the 80286 and 80386 microprocessors use the CLOCK input to define the length of a time slot. They divide the frequency of the CLOCK input by two to yield an internal timebase referred to as the Processor Clock (PCLK). The CLOCK input is referred to as a double-frequency clock because its frequency is double that of the required PCLK. The 486DX microprocessor, on the other hand, uses the CLOCK input as the PCLK without dividing it. Clock-doubling or clock-multiplying processors, such as the DX4, use phase-locked loops or similar technology to multiply (instead of divide) the CLOCK signal to produce a PCLK that is faster than CLOCK.

PCLK is the real metronome that defines the duration of the time slots during which the microprocessor performs a task or a pre-defined portion of a task. Refer to Figure 2-2.

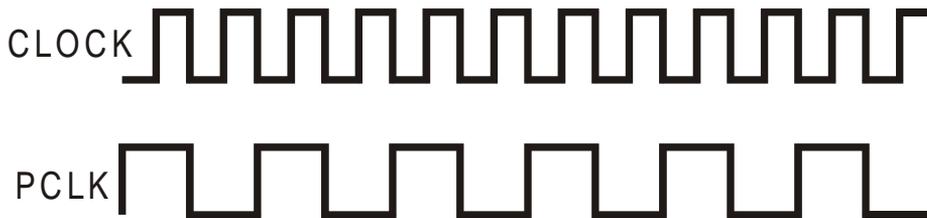


Figure 2-2. Relationship of CLOCK and PCLK

As illustrated in figure 2-2, PCLK is half the frequency of CLOCK. When a PC manufacturer refers to the operating speed of their computer as 8MHz (mega-hertz, or million cycles-per-second), 10MHz, etc., they are referring to the microprocessor's PCLK frequency, not that of CLOCK.

Microprocessor's Bus Cycle State Machine

When the microprocessor performs a read or a write operation, it initiates a sequence of events called a bus cycle.

During a bus cycle, the microprocessor places the address on the address bus, sets the control bus lines to indicate the type of transaction (such as a memory read or I/O write bus cycle), and transfers the data between the target location and itself. This happens in a very orderly fashion, with each step occurring at the proper point during the appropriate time slot.

x86 microprocessor chips include a subsystem called a bus unit, tasked with the job of running bus cycles when required. The bus unit is a state machine

ISA System Architecture

that is stepped through its various states by the PCLK signal. The duration of each state is one cycle of PCLK.

Refer to figure 2-3. As an example, if the CLOCK input frequency is 40MHz (40 million cycles per second), the PCLK frequency is half that frequency, or 20MHz. In order to determine the duration of one cycle of PCLK, just divide 20 million cycles-per-second into one second. In this example, a PCLK cycle is 50ns in duration (50 nanoseconds; 0.000000050 seconds or 50 billionths of a second). Since the bus unit state machine's states are each one PCLK in duration, this means that each bus unit state is 50ns in duration.

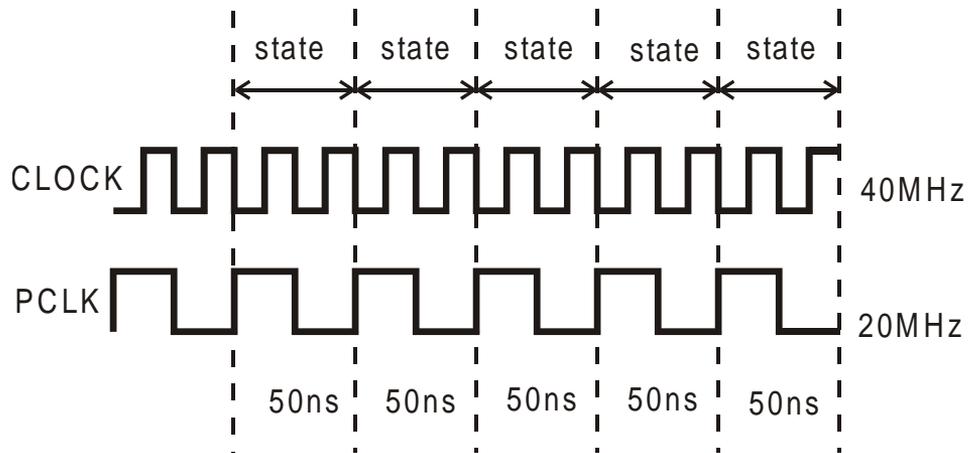


Figure 2-3. Each Cycle of PCLK Defines the Duration of a State

If not currently engaged in a bus cycle (a read or a write operation), the bus unit state machine is said to be in the idle state. It remains in the idle state until the microprocessor must perform a bus cycle.

Address Time

When the microprocessor must perform a read or a write, its bus unit initiates a bus cycle. The bus unit leaves the idle state and enters a state that will be referred to in this text as address time. Although Intel uses a different name for this state, address time is the name that will be used in this text for clarity's sake. During address time, one PCLK cycle in duration, the microprocessor places the address on the address bus and the bus cycle definition (type of bus cycle) on the control bus lines.

Chapter 2: Introduction to the Bus Cycle

Data Time

After performing all actions required during address time, the bus unit state machine immediately enters the state we will refer to as data time. As with address time, Intel uses a different name for this state, but, for clarity's sake, data time is the name that will be used in this text. During data time, one PCLK cycle in duration, the microprocessor expects the data to be transferred between itself and the currently addressed device. At the end, or trailing edge, of data time (refer to reference point 1 in figure 2-4), the microprocessor samples (tests the state of) its READY# input to see if the currently addressed device is ready to complete the bus cycle. If the READY# input is asserted (low, as indicated by the pound sign), the bus unit state machine terminates the bus cycle.

If a read were in progress, the bus unit interprets an asserted level on READY# to mean that the currently addressed device has placed the requested data onto the data bus for the microprocessor. The microprocessor reads the data from the data bus and terminates the bus cycle.

If a write were in progress, the bus unit interprets an asserted level on READY# to mean that the currently addressed device has accepted the data being written to it. The microprocessor terminates the bus cycle. The chapter entitled "The 80286 System Kernel: the Engine" provides a detailed operational description of the Ready logic. Figure 2-4 illustrates a bus cycle consisting of address time and data time.

The 0-wait-state bus cycle is the fastest type of bus cycle the 80286, 80386, 80486 and Pentium microprocessors are capable of performing. In other words, the fastest x86 bus cycle takes two "ticks" (cycles) of PCLK. If this were an 80386 running at 20MHz (PCLK speed), a 0-wait-state bus cycle would take 100ns (50ns each for address time and data time).

Chapter 3

The Previous Chapter

The previous chapter, "Introduction To the Bus Cycle," provided a basic description of the microprocessor's communications scheme. Using the three buses, the microprocessor performs bus cycles to read information from or write information to memory or I/O devices.

This Chapter

This chapter explores the addressing scheme used by the x86 processors. It defines the addressing range provided by the 8086, 8088, 286, 386, 486 and Pentium microprocessors. A basic definition of the term I/O device is also provided.

The Next Chapter

Having covered the addressing scheme used by the x86 processors in this chapter, the next chapter, "The Address Decode Logic," introduces the concept of the address decoder. Several examples of address decoders are examined in detail.

Evolution of Memory and I/O Address Space

The Intel 8080 microprocessor is the ancestor of the entire x86 microprocessor family. Many of the family's characteristics stem from this common ancestor.

Intel 8080 Microprocessor Address Space

The 8080 microprocessor's address bus consists of 16 address lines, A[15:0]. With 16 address lines, the microprocessor can place any address pattern on the address bus from all zeros (0000h) to all ones (FFFFh). In other words, the 8080 can place any one of 65536 addresses on the address bus and can therefore address any one of 65536 locations located in external devices. In the computer

ISA System Architecture

industry, 65536 is referred to as 64K (K stands for Kilo, Greek for 1,000, but in the computer industry, kilo is 2^{10} or 1,024). Refer to figure 3-1. This means that the designer of an 8080-based system can include memory devices with a total of no more than 64K locations in the system.

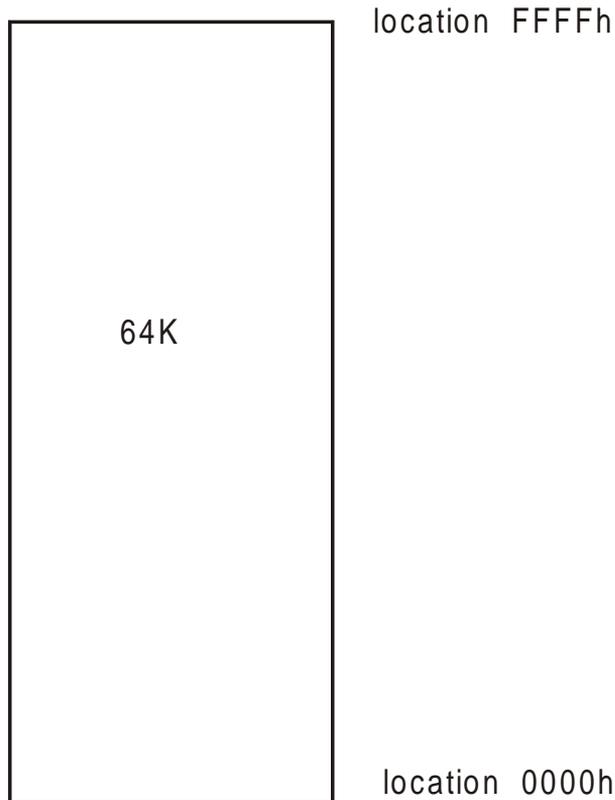


Figure 3-1. Single 64K Address Space

In addition to memory devices used for program and data storage, the designer must also use a number of locations for input/output ports (I/O ports) used to communicate with I/O devices. In simple terms, any device that the microprocessor can read data from or write data to that isn't a memory device (RAM or ROM) is an I/O device. Implementing these ports certainly is necessary, but depletes the number of addressable locations available to be assigned to memory devices.

Chapter 3: Addressing I/O and Memory

Rather than use up already scarce memory locations to implement I/O devices, Intel provided two new instructions to the 8080 microprocessor: IN and OUT, and added a way for the 8080 to indicate to the system that the 8080 wishes to communicate with an I/O device instead of a memory device. The system typically included an Intel 8228 controller chip which detected the 8080's intention to communicate with an I/O device.

Some of the functionality of the 8228 was incorporated in the 8085 microprocessor, a successor of the 8080. The 8085 has 16 address lines, as does the 8080, so the 8085 can address 64K locations.

The 8085 has a new output pin called IO/M#. In Intel signal names, the “/” character should always be read as the English word “or.” IO/M# therefore stands for “IO or Memory.” Since there is no pound sign after the “IO”, a high on this signal line means an I/O address is being output by the microprocessor. The pound sign after the “M” portion of the signal name indicates that a memory address is present on the address bus when IO/M# is low.

The addition of the IO/M# pin actually created two separate address spaces (or maps), one for Memory and the other for I/O. In essence, the IO/M# pin “points” to the appropriate address space when the microprocessor places an address on the address bus.

Refer to figure 3-2. As previously mentioned, the 8085 provides 64K of memory address space. However, due to the limited number of I/O devices typically found in systems, Intel only implemented 256 I/O locations.

ISA System Architecture

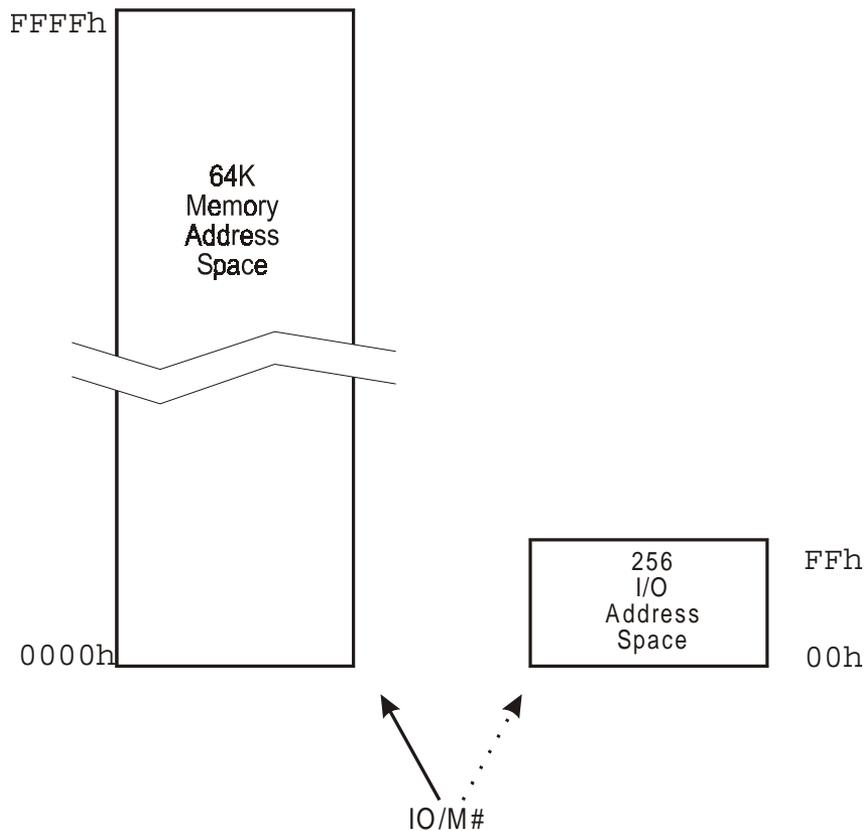


Figure 3-2. The 8085's Memory and I/O Address Spaces

As an example, if the microprocessor places address 0010h on the address bus and places a high on IO/M#, it is addressing I/O location 0010h, not memory location 0010h. On the other hand, if the microprocessor places address 0010h on the address bus and places a low on IO/M#, it is addressing memory location 0010h, not I/O location 0010h.

In addition to the IO/M# pin, Intel also added two I/O instructions, IN and OUT, to the 8085's instruction set. When executed, the IN instruction causes the microprocessor to perform a read from the I/O address specified by the programmer. This address is placed on the address bus and the IO/M# pin is set high, thereby indicating the presence of an I/O address on the address bus to external logic. In the same way, the OUT instruction also causes the microprocessor to place an I/O address on the address bus with IO/M# set high to indicate the I/O address it wishes to write data to.

Chapter 3: Addressing I/O and Memory

A separate instruction, *MOV*, is used by the programmer to specify a memory address to move data to or from. When executed, the *MOV* instruction causes the specified memory address to be placed on the address bus and the *IO/M#* pin to be set low, indicating the presence of a memory address.

These changes allowed the designer to implement I/O ports without encroaching on memory space.

8086 and 8088 Microprocessor Address Space

When Intel designed the 8086 and 8088 microprocessors, they increased the number of address lines from 16 to 20. This means that the 8086 and 8088 microprocessors can place any address from 00000h to FFFFFh on the address bus, giving them an address range of 1,048,576 or 1M locations. “M” stands for the Greek prefix Mega, meaning large or great. In physics, M is 1,000,000, but in the computer industry, M is 2^{20} or 1,048,576. Some computer companies use M to mean 1,000,000 when referring to disk storage while still using M to mean 2^{20} when referring to memory size, so beware!

As with the 8085 microprocessor, the 8086 and 8088 microprocessors have a pin to indicate memory or I/O addresses. Note that the pin name and function is slightly different (or completely different, depending on your point of view). It’s named *M/IO#*, meaning that the 8086 and 8088 set the *M/IO#* pin low when executing an *IN* or *OUT* instruction. The 8086 and 8088 set the *M/IO#* pin high when executing a *MOV* instruction to access memory. This is just the opposite of the way the 8085 did it.

For the 8086 and 8088 microprocessors, Intel also increased the size of the I/O address space to 64K. The programmer may specify any memory address within the 1M memory address range, but is restricted to I/O addresses in the first 64K locations of address space (0000h to FFFFh). There were two reasons for this decision:

1. The 8086 and 8088 microprocessors remain backward-compatible with the I/O instructions in the 8080 instruction set, although these instructions will not take full advantage of the processors address capability.
2. Virtually no designer requires more than 64K I/O ports (locations) to implement a full complement of I/O devices, no matter how complex they might be.

Chapter 4

The Previous Chapter

The previous chapter, “Addressing I/O and Memory,” explored the addressing scheme used by the x86 processors. It defined the addressing range provided by the 8086, 8088, 80286, 80386, 80486 and Pentium microprocessors, and provided a basic definition of the term “I/O device.”

This Chapter

This chapter describes the function performed by address decoders, discusses the address decoder design process, defines data bus contention, and provides a detailed analysis of two example address decoders found in ISA systems.

The Next Chapter

The next chapter, “The 80286 Microprocessor,” provides a detailed description of the internal hardware architecture of the 80286 microprocessor, describes its register set, and supplies definitions of real and protected modes.

The Address Decoder Concept

Refer to figure 4-1. When the microprocessor must read data from or write data to an external device, it places the address on the address bus. Every device that can be read from or written to has an associated piece of logic called an address decoder. The address decoder selects its associated device when an address within its defined range is detected.

In most cases, the selected device contains multiple internal locations. When selected, the addressed device examines the address to identify the exact internal location with which the microprocessor wishes to perform a data transfer.

ISA System Architecture

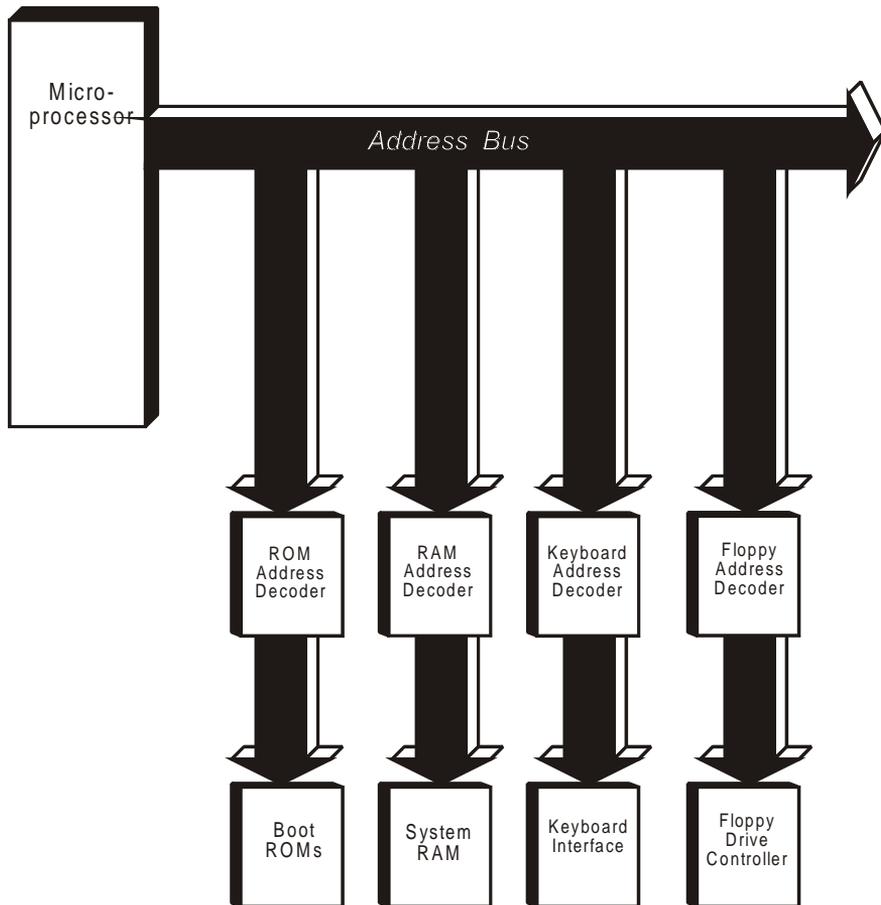


Figure 4-1. Relationship of Address Bus, Address Decoders and Devices

A memory or I/O chip does not respond to any pre-assigned address range. Rather, the device's address decode logic informs it that the address currently on the address bus is assigned to one of its internal locations by asserting the device's chip-select input pin.

The overall range of memory and I/O locations addressable by a particular microprocessor are commonly referred to as address maps. The design of its associated address decoder, not that of the memory or I/O device itself, defines where a device lives within the overall memory or I/O address map of a microprocessor.

Chapter 4: The Address Decode Logic

Data Bus Contention (Address Conflicts)

Under no circumstances should the address decoders for two devices installed in a system be designed in such a fashion that they both detect the same address ranges. This would cause data bus contention.

Every expansion card installed in an ISA system has its own address decode logic that defines the address range the card responds to. Address conflicts are a common problem in the PC world. Consider the problem that arises if the address decoders on two expansion cards are both designed to detect the same address range.

As an example, assume that two cards respond to the I/O address range 03F0h to 03F7h. If the microprocessor executes an instruction that causes a read from I/O location 03F1h, both devices are chip-selected by their respective address decoders. Each then places the contents of its respective location 03F1h onto the common data bus simultaneously. One card may be driving a particular data bus signal line high while the other card is driving the same data bus signal line low. In this situation, we have two separate current sources trying to drive five volts and zero volts onto the same piece of wire.

At best, this situation will cause garbled data and, at worst, hardware damage. The proper term for this condition is data bus contention, a problem quite common in the PC world. A user populates a PC with many cards purchased from various manufacturers. The prospect of an address conflict is very real.

In order to permit resolution of address conflicts, a card designer should allow the user to easily alter the address range the board responds to. In ISA systems, this is usually accomplished with configuration switches (or jumpers) on the board. By changing the switch setting, the user causes the card's address decoder to detect a different address range, thereby resolving the conflict.

How Address Decoders Work

The address decoder function is similar to that performed by the post office. Because it more or less identifies a neighborhood, the zip code is analogous to the high-order part of the address (upper address bits). The low-order portion of the address identifies the exact location within the neighborhood.

In much the same way, the address decoder associated with a device examines the high-order part of the address being output by the microprocessor. It de-

ISA System Architecture

termines if the address is within the range assigned to its device. If it is, the address decoder chip-selects the device. The chip-selected device then examines the lower part of the address to determine the exact location within it the microprocessor is addressing.

It should be noted that an address decoder doesn't have to look at the entire address to determine that the current address is within the address range assigned to its respective device. If you were traveling in a car and were told to keep an eye out for the two thousand block, you would only need to look at the thousands digit of the address to determine if it was within the desired range. The lower digits would be insignificant to you.

The following two sections describe example address decoders and how they function. The first example is that of the ROM address decoder in an IBM PC/XT. The second example is that of the address decoder that detects addresses assigned to the I/O devices found on a typical ISA system board.

Example 1– PC and PC/XT ROM Address Decoder

Background

The original IBM PC was designed around the 8088 microprocessor. With an address bus consisting of 20 lines, designated A[19:0], the 8088 could address any one of 1,048,576 memory locations. The designers chose to populate this memory map as illustrated in Figure 4-2.

The memory address range from F0000h to FFFFFh was assigned to the system board ROM. This ROM is frequently referred to as the boot or POST/BIOS ROM. The 64K addresses from F0000h to FFFFFh are set aside for the boot ROM. It should therefore be selected whenever the uppermost digit on the address bus is Fh and the M/IO# pin is high (indicating the presence of a memory address on the address bus). Table 4-1 illustrates the state of the address lines when any memory address in this range is accessed.

Chapter 4: The Address Decode Logic

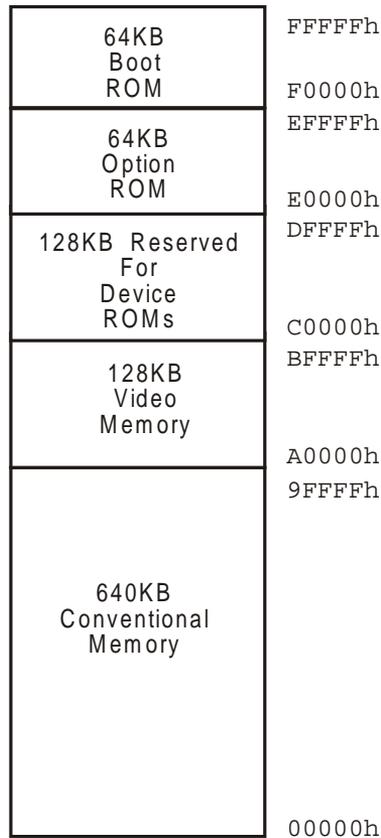


Figure 4-2. PC Memory Address Map

Table 4-1. State of the Address Lines When Any Memory Address in Range F0000h to FFFFFh Is Addressed

19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	1	1	1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
F				X				X				X				X			

In addition, memory address range E0000h to EFFFFh was assigned to the optional ROM that could be installed in the empty ROM socket provided on the system board. This capability was provided to allow the installation of a ROM containing the BIOS routines and interrupt service routines for optional devices that could be added to the machine. This add-in ROM is commonly referred to as the option ROM. The 64K addresses from E0000h to EFFFFh are set aside for the option ROM. It should therefore be selected whenever the uppermost digit on the address bus is Eh and the M/IO# pin is high (indicating the presence of

Chapter 5

The Previous Chapter

The basics of microprocessor communication have been introduced. This includes concept of the bus cycle and the types of devices that the microprocessor communicates with (memory and I/O), along with how these devices are addressed.

This Chapter

This chapter provides a description of the 80286 microprocessor. Emphasis is placed on the hardware interface between the 80286 microprocessor and the remainder of the system. Memory address formation in real mode is covered in detail and an introduction to protected mode is provided.

The Next Chapter

The next chapter provides a description of the six possible sources for various reset signals found in ISA systems and the effect each of them has on the system. It also explains the actions taken when Ctrl/Alt/Delete are depressed.

The 80286 Functional Units

Any microprocessor is really a system consisting of a number of functional units. Each unit has a specific job to do and all of the units working together comprise the microprocessor. Figure 5-1 illustrates the units which make up the 80286 microprocessor.

ISA System Architecture

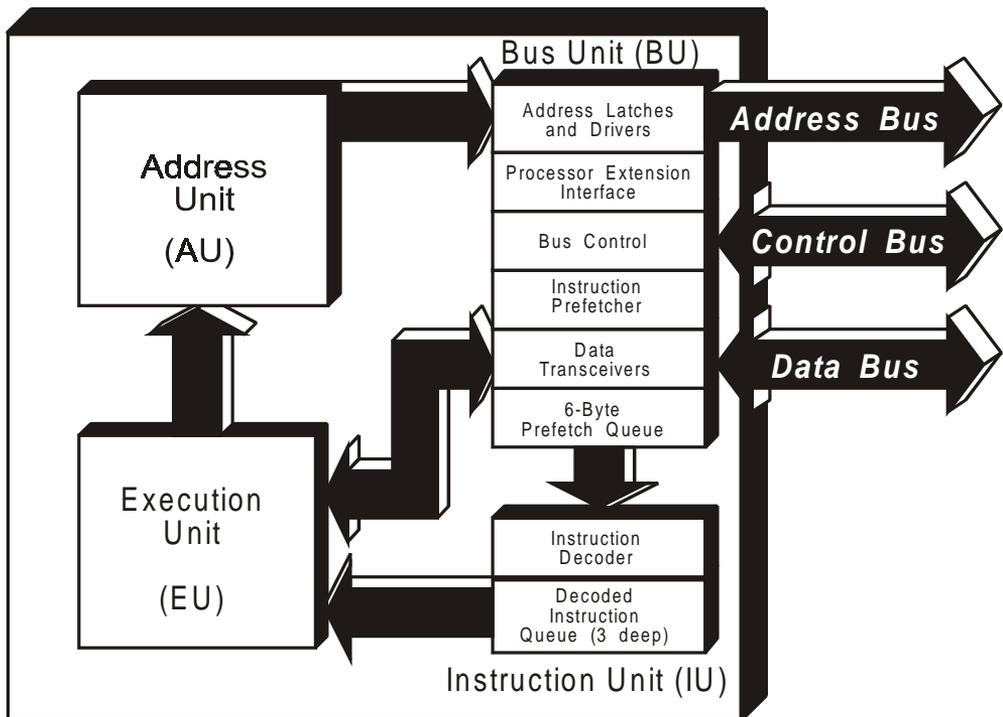


Figure 5-1. 80286 Microprocessor Block Diagram

Refer to 5-1. When the microprocessor starts operation, it will automatically begin fetching (reading) instructions from memory. It is the bus unit's task to perform the bus cycles required to fetch all instructions from memory. Instructions are read from memory over the data bus and placed into the instruction Prefetch Queue. From there they go to the instruction unit where they are decoded and sent to the execution unit one at a time for processing.

Execution of some instructions requires that data be read from or written to memory or I/O devices. If a memory device is specified, the execution unit directs the address unit to form the memory address of the location specified by the instruction. I/O addresses specified by an instruction do not require special address formation as do memory address. Instead they are sent directly from the execution unit to the bus unit to run the bus cycle specified by the instruction.

Chapter 5: The 80286 Microprocessor

In summary, two types of information may be read into the microprocessor; instructions (code) and data. Instructions are read from memory over the data bus, go through the data transceivers, and are placed into the prefetch queue. Data to be processed is read from either memory or I/O devices over the data bus and is transferred via the data transceivers to the execution unit.

An instruction may also specify that data be written to a memory or I/O location, in which case data is sent from the execution unit to the data transceivers for transfer over the data bus.

In short, the 80286 microprocessor consists of four functional units:

1. **Instruction unit.** Decodes the instruction prior to passing it to the execution unit for execution.
2. **Execution unit.** Handles the actual execution of the instruction.
3. **Address unit.** When the microprocessor must address a memory location, the address unit forms the memory address that is driven out onto the address bus by the bus unit during the bus cycle.
4. **Bus unit.** Handles communication with the world outside the microprocessor chip (by performing bus cycles).

The following sections provide a more detailed description of each of these units.

The Instruction Unit

One at a time, the 80286 instruction unit pops (reads) instructions from the prefetch queue, decodes them into their component parts and places them into the decoded instruction queue to be forwarded to the execution unit.

The Execution Unit

Instructions are executed by the 80286 execution unit. In general, instructions cause either internal processing of data already stored within the execution unit registers, or reading data from or writing data to devices external to the microprocessor.

When the currently executing instruction requires writing or reading a memory or I/O location, the execution unit issues a bus cycle request to the bus unit. The execution unit then stalls until the bus unit completes the requested data transfer.

ISA System Architecture

Two types of information may be read into the microprocessor: instructions (code) or data. Instructions are always read from memory and are placed into the instruction prefetch queue. From there, they are routed to the instruction unit where they are decoded and sent to the execution unit for processing. Data to be processed goes directly to registers inside the execution unit. These registers have several functions as described in the following paragraphs.

A register can be thought of as a storage location within the microprocessor. Some registers can only be read from, and some can only be written to, while others are both readable and writable. Registers within the execution unit are used by the programmer or the microprocessor itself for the following purposes:

- to temporarily hold values to be used in calculations
- to receive data being read from an external memory or I/O location
- to provide the data to be written to an external memory or I/O location
- to keep track of the microprocessor's current status
- to control certain aspects of the microprocessor's operation

Although they are storage locations, the programmer refers to the microprocessor's registers by names (rather than by the hex addresses used to address external memory and I/O locations). Most of the 80286 registers are 100% backward-compatible with the 8086 and 8088 microprocessor's registers. The registers inside the execution unit can be logically divided into two categories:

- General Registers
- Status and Control Registers

The following sections describe each of the registers in the General and the Status and Control register sets.

General Registers

This section offers a description of each of the general registers and examples of their usage. Figure 5-2 details the registers in the general register set.

Chapter 5: The 80286 Microprocessor

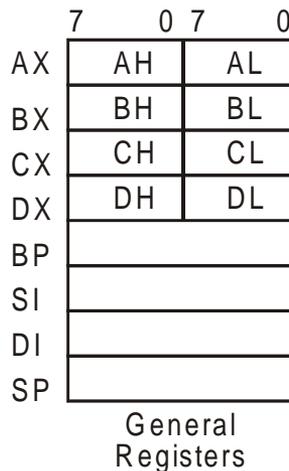


Figure 5-2. The 80286 General Registers

The AX, BX, CX and DX Registers—Each of these four registers can hold 16 bits (two bytes) of information. As an example, the following instruction will move the 16-bit value 1234h into the AX register:

```
MOV AX,1234 ;move the value 1234h into AX
```

In addition, the programmer can individually refer to each of the 8-bit registers that make up the upper and lower half of each of these 16-bit registers. As an example, the AX register actually consists of two 8-bit registers: the AL (Lower half of AX register) and AH (upper half of the AX register) registers. The same is true of the BX, CX and DX registers.

If the programmer wants to move only one byte of information, he or she would specify one of the 8-bit registers rather than a 16-bit register. In the following example, the programmer wants to write the value 02h into I/O location 60h:

```
MOV AL,02 ;move value 02h into AL  
OUT 60,AL ;write AL contents to I/O port 60h
```

In the example above, the programmer first moves the value 02h into the 8-bit AL register and then performs an I/O write instruction to write the contents of the AL register to I/O location 60h. This will cause an I/O write bus cycle with address 000060h on the address bus and 02h on the lower part of the data bus

Chapter 6

The Previous Chapter

The first five chapters introduced the concepts of microprocessor communications and provide a detailed understanding of the 80286 microprocessor's interface signals.

This Chapter

ISA products usually have six separate source for reset, two of which are generated only by hardware and four that can be initiated by software commands. This chapter details the reset sources within typical ISA systems.

The Next Chapter

The next chapter explains the sequence of events that occur immediately after an ISA system is powered on.

The Power Supply Reset

The power supply provides the operating voltages necessary for system operation. When the power switch is first placed in the on position, it takes some time for the power supply's output voltages to reach their proper operating levels. If the system components were allowed to begin operating before the voltages stabilized, it would result in erratic operation.

Every PC power supply produces an output signal commonly called POWERGOOD. On the system board, the POWERGOOD signal is used to produce the RESET signal. During the period required for stabilization of the output voltages, the POWERGOOD signal is kept deasserted by the power supply. While POWERGOOD is deasserted, the RESET signal is kept asserted.

While RESET is asserted, it has two effects on the microprocessor and other system components:

ISA System Architecture

1. Keeps any activity from occurring until power has stabilized.
2. Presets the microprocessor and other system devices to a known state prior to letting them begin to do their job. This ensures that the machine will always start up the same way.

When POWERGOOD becomes asserted, the RESET signal is deasserted and the microprocessor can begin to fetch and execute instructions. The first instruction is always fetched from the power-on restart address which is always located 16 locations from the very top of the memory address space. For the 8088 microprocessor, this would be FFFF0h, FFFFF0h for the 80286 and 80386SX, and FFFFFFF0h for the 80386DX, 80486 and Pentium microprocessors. This location contains the first instruction of the POST.

Reset Button

Some systems provide a momentary-contact switch which the user may press to force the RESET signal to be asserted to the microprocessor and to the rest of the system. The effect on the system is the same as deasserting POWERGOOD from the power supply, but the power supply is not actually turned off.

Shutdown Detect

When the microprocessor detects an exception while attempting to execute the handler for a prior exception, the two exceptions are generally handled serially. Certain combinations of exceptions cannot be handled serially, however. These exceptions are caused by protection violations in protected mode. For example, in protected mode, if an application attempts to access data that is outside its data segment, the microprocessor will start running the general protection fault exception handler. If, say, a stack overflow were to occur while the microprocessor ran the general protection fault exception handler, the microprocessor would not be able to handle the two faults serially. The microprocessor would invoke the double-fault exception instead.

If the microprocessor detects another exception while attempting to service the double-fault, it goes into a shutdown condition. This is sometimes referred to as a triple fault. The microprocessor indicates the shutdown condition to external logic by running the halt or shutdown bus cycle. Since the microprocessor has abnormally ceased to fetch and execute instructions, it can't run a program to inform the user that a severe error has occurred. Therefore, in an ISA machine, the shutdown detect logic outside the microprocessor detects the shut-

Chapter 6: The Reset Logic

down and toggles the microprocessor's RESET line, causing the system to reboot.

See the chapters entitled “Detailed View of the 80286 Bus Cycle” and “Detailed View of the 80386 Bus Cycle” for additional information on the halt or shut-down bus cycle.

Hot Reset

MS-DOS was written specifically for the Intel 8088 microprocessor using 8088-specific instructions. Since the 8088 only has 20 address lines, it is incapable of generating a memory address greater than FFFFh, or 1MB. Furthermore, protected mode wasn't introduced until the advent of the 80286 microprocessor. Consequently, MS-DOS can only address the lower 1MB and doesn't understand protected mode.

When an 80286 is powered up, it operates in real mode. In other words, it operates as if it were an 8088. This means that, although the 80286 has twenty-four address lines and can physically address up to 16MB of memory address space, it is limited to the lower 1MB when in real mode.

Most current applications programs are written to run under MS-DOS. Many of these programs require access to more memory space than allowed under MS-DOS. As an example, many Lotus 1-2-3 spreadsheets require very large amounts of memory space, well in excess of that allowed under MS-DOS and the 8088 microprocessor.

If a system is based on the 80286 microprocessor, the problem can be solved by switching the microprocessor into protected mode. This allows the microprocessor to access up to 16MB of memory space, but MS-DOS presents a problem. It doesn't understand protected mode and is therefore incapable of generating addresses above 1MB.

Special software that understands protected mode operation prepares the segment descriptor tables in memory prior to switching into protected mode. It also must save the address of the next instruction to be executed when the microprocessor returns to real mode so MS-DOS can continue to run at the point where it went to protected mode. The 80286 is then switched into protected mode and the extended memory above 1MB accessed (for example, to store spreadsheet data).

Once the extended memory access has been completed, the microprocessor must be switched back into real mode so the applications program can continue

ISA System Architecture

to run under MS-DOS. Here's where the problem comes in. In order to switch the 80286 microprocessor from protected to real mode, the microprocessor must be reset.

To do this, the following actions must be taken:

1. The programmer stores an address pointer that points to the address of the real mode instructions that are responsible for restoring the system to its previous operating condition. This pointer is stored in locations 0040:0067 to 0040:006A. The contents of these locations will be used later when the processor returns to real mode to find the address in memory where the real mode instructions reside.
2. A special value (05h or 0Ah) is stored in the configuration CMOS RAM location (0Fh) to indicate the reason for the reset. Refer to table 6-1 for definitions of the reset byte in CMOS RAM.
3. The system has now been prepared to return to real mode. The Hot Reset command must now be issued to the keyboard/mouse interface.* This is accomplished by writing a FEh to the keyboard/mouse interface's command port at I/O address 0064h.
4. In response, the keyboard/mouse interface pulses its Hot Reset output one time, causing the hardware to generate a reset to the 80286.
5. When reset becomes deasserted, the microprocessor begins to fetch and execute instructions at the power-on restart address exactly as if a power-up had just occurred.
6. At the beginning of the POST, the programmer reads the value stored in the configuration RAM location to ascertain the reset's cause. In this case, the value (05h or 0Ah) indicates that it was caused by Hot Reset to get back to real mode and continue program execution.
7. POST then retrieves the previously stored real mode address pointer found at location 0040:0067 and jumps to the indicated address and picks up where it left off in real mode.

* Some manufacturers employ a Hot Reset "intercept" where external circuitry recognizes an FEh written to the keyboard controller and immediately generates a Fast Hot Reset. This provides a faster reset signal compared to the slower generation via a command to the keyboard controller. Later versions of Intel's microcontrollers (typically used as keyboard controllers) provide an internal intercept that also results in a fast Hot Reset.

Another method used to produce a fast Hot Reset is to cause the microprocessor to go into the shutdown condition by intentionally causing faults using software. This causes the shutdown detect logic to reset the microprocessor.

Alternate (Fast) Hot Reset

The Alternate or Fast Hot Reset command is found in some ISA systems that implement system control port A at I/O address 0092h, sometimes called the PS/2 compatibility port. Alternate Hot Reset performs the same function as the Hot Reset command. However, the microprocessor is reset more quickly using this method than when using Hot Reset. The Hot Reset command must be interpreted by the ROM-based program inside of the keyboard/mouse interface, while the Alternate Hot Reset command is executed by the hardware much more quickly.

Alternate Hot Reset is generated by setting bit 0 in system control port A to a one. This generates a pulse on the Alternate Hot Reset signal line which, in turn, causes a pulse on the Hot Reset signal. This resets the microprocessor. For a more extensive explanation of Hot Reset, see the previous section.

Ctrl-Alt-Del Soft Reset

When the control and alternate keys are depressed and held, followed by the delete key being depressed, three “make” codes for these keys are stored in the keyboard buffer in memory. When the sequence of “make” codes are detected by the keyboard interrupt service routine, the soft reset routine is invoked. The soft reset routine causes the system to:

1. Place the value 1234h into the reset flag location (0040:0072) in main memory. This value tells the system to skip a portion of the POST when rebooting.
2. Store a special value (00h) in configuration RAM to indicate that a Control/Alternate/Delete is causing the reset (also used for power-on reset).
3. Depending on the computer model and manufacturer one of two methods is typically used to reset the processor.
 - Write a FEh to port 64h, causing a CPU reset. When the reset is removed, the microprocessor begins to fetch and execute code from location FFFF0h, the beginning of POST.
 - Jump to location FFFF0h and begin program execution.
4. Near the beginning of the POST, the programmer reads the value stored in the configuration RAM location to ascertain the reset's cause. In this case, the value indicates that it was caused by Control/Alternate/Delete or power-on reset.

Chapter 7

The Previous Chapter

The previous chapter, “The Reset Logic,” provided a detailed description of the possible sources for RESET.

This Chapter

This chapter provides a description of the ISA machine power-up sequence from the moment power is applied until the microprocessor begins to fetch and execute instructions.

The Next Chapter

The next chapter, “The 80286 System Kernel: the Engine,” describes the support logic necessary to allow the 80286 microprocessor to communicate with 8- and 16-bit devices.

The Power Supply – Primary Reset Source

The power supply provides the operating voltages necessary for system operation. When the power switch is first placed in the on position, a period of time is required for the power supply's output voltages to reach their proper operating levels. If the system components are allowed to begin operation before the supply voltages have stabilized, erratic operation would result.

Figure 7-1 illustrates the power supply-related logic that produces the microprocessor's RESET signal. Every ISA PC power supply produces an output signal commonly called POWERGOOD. On the system board, the RESET signal is derived from the POWERGOOD signal. In order to guarantee proper operation of the microprocessor, the RESET input must change state in synchronism with the CLK2 signal (the microprocessor's double-frequency clock input). On each rising edge of CLK2, the state of the POWERGOOD input is latched by the flip-flop, inverted to its opposite state, and presented on the flip-flop's output, RESET. During the period required for stabilization of the output voltages, the

ISA System Architecture

POWERGOOD signal is held deasserted by the power supply. While POWERGOOD is deasserted, the RESET signal is held asserted. The flip-flop in the figure provides a RESET output that is synchronized with CLK2 and is the invert of the POWERGOOD signal from the power supply.

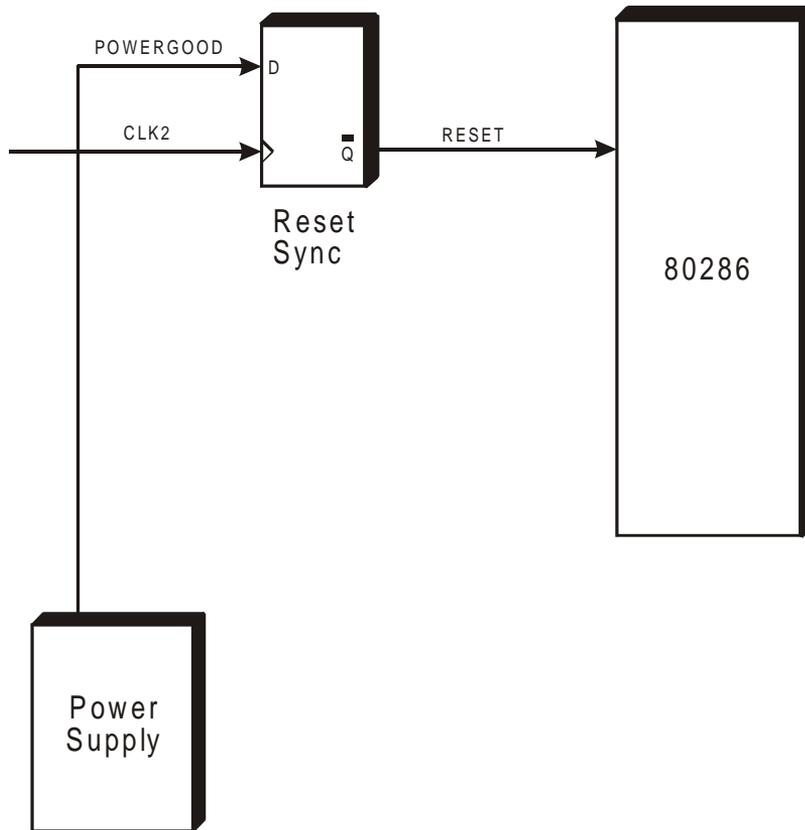


Figure 7-1. RESET Is Derived from POWERGOOD

While RESET is asserted, it has two effects on the microprocessor and other system components:

1. Keeps any activity from occurring until power has stabilized.
2. Presets the microprocessor and other system devices to a known state prior to letting them begin normal operation. This ensures that the machine will always start up the same way.

Chapter 7: The Power-Up Sequence

How RESET Affects the Microprocessor

Table 7-1 defines the values forced into the 8086 and 8088 microprocessor's registers when RESET is asserted.

Table 7-1. Values Preset Into 8086 and 8088 Microprocessor Registers by RESET

Register	Contents
FLAGS	0002h
MSW	FFF0h
IP	FFF0h
CS	F000h
DS	0000h
ES	0000h
SS	0000h

While RESET is asserted, the microprocessor cannot fetch and execute instructions and its outputs are set as illustrated in tables 7-2 and 7-3.

Table 7-2. State of 80386DX Outputs While RESET is Asserted

Pin Name	Pin State
LOCK#, D/C#, ADS#, A[31:2]	high
W/R#, M/IO#, HLDA, BE#[3:0]	low
D[31:0]	tri-state (floating)

Table 7-3. State of 80286 Outputs While RESET Is Asserted

Pin Name	Pin State
A[23:0], S0#, S1#, PEACK#, BHE#, LOCK#	high
M/IO#, COD/INTA#, HLDA	low
D[15:0]	tri-state (floating)

Processor Reaction When Output Voltages Stabilize

When the power supply output voltages have stabilized, the POWERGOOD signal is asserted and the logic on the system board responds by deasserting RESET. The deasserted level on RESET allows the microprocessor to begin functioning. It should be noted that the 386, 486 and Pentium microprocessors can perform a self-test when RESET becomes deasserted. Details of the processor self-test are discussed in the chapter entitled "The 80386DX and SX Micro-

ISA System Architecture

processors” and in the MindShare Architecture Series books *80486 System Architecture* and *Pentium Processor System Architecture* published by Addison-Wesley. The microprocessor initiates the fetching and executing of instructions from memory, using the code segment (CS) and instruction pointer (IP) register contents to point to the memory location containing the first instruction.

While RESET was asserted, the machine status word (MSW) register had FFF0h forced into it. Since the protect enable bit, bit 0, is therefore 0, the microprocessor always begins operation in real mode. In real mode, the microprocessor always appends an extra 0h on the end of the CS register contents, F000h, resulting in a code segment start address of F0000h. It then adds the offset portion of the address, FFF0h, contained in the IP register, to the segment start address:

$$\begin{array}{rcl} \text{CS} & & \text{F0000h} \\ \text{IP} & + & \underline{\text{FFF0h}} \\ & & \text{FFFF0h} = \text{physical memory address} \end{array}$$

The resultant memory address, FFFF0h, is referred to as the power-on restart address. Since the address is formed exactly the same way every time the system is powered up, the 8086 or 8088 microprocessor always fetches its first instruction from the power-on restart address.

The First Bus Cycle

The microprocessor then initiates a memory read bus cycle to fetch the first instruction from the power-on restart address in memory. The power-on restart address is always located in the boot ROMs. This is because the first instruction must be located in *non-volatile memory*. If the first instruction were fetched from RAM memory, the information returned would be junk. This is because RAM memory is *volatile*.

The first instruction fetched from the power-on restart address is always the first instruction of the power-on self-test, or POST. The POST program is contained in the boot ROMs and is always the first program to be run.

Although they always begin operation in real mode and should therefore function as a fast 8086, the 286, 386, 486 and Pentium microprocessors have more than twenty address lines. When the 286 or higher microprocessors place the power-on restart address on A[19:0] of the address bus, the address lines above A19 are automatically set high. This results in address FFFFF0h being placed

Chapter 7: The Power-Up Sequence

on the 286 and 386SX address bus, rather than 0FFFF0h. The 386DX, 486 and Pentium processors place FFFFFFF0h on the address bus.

The processor keeps these upper address lines high until the program performs a jump to another code segment; in other words, until a new value is loaded into the CS register. A jump instruction that reloads both the CS and the IP registers is referred to as an “inter-segment jump” or a “far jump.” The upper address lines are then set low and remain low unless the processor is switched into Protected Mode by the program. In practice, the first instruction executed from system ROM is always a jump instruction, thus a new CS value is immediately loaded when the first instruction is executed.

After the upper bits are set low, they will not be set high again while the processor remains in real mode. This means that the processor cannot access extended memory (memory above the lower megabyte) while in real mode. There is one exception to this rule. It is discussed in the chapter entitled “The 80286 Microprocessor” under the heading “Accessing Extended Memory in Real Mode.”

Chapter 8

The Previous Chapter

The preceding chapters have explained how the 80286 microprocessor communicates with memory and I/O devices and how the 80286 interfaces with external devices. Sources of system reset and the power-up sequence have also been covered.

This Chapter

This chapter introduces and explains the external logic required by the 80286 microprocessor to communicate with 8- and 16-bit devices. Additional logic needed to synchronize memory and I/O devices of varying speeds is also covered. A thorough understanding of this chapter is crucial to understanding the 80386 kernel.

The Next Chapter

The next chapter provides a detailed look at the bus cycles that an 80286 can run. The chapter includes timing diagrams and explanations for read, write, and the halt or shutdown bus cycles.

The Bus Control Logic

In order to perform all of the actions required to complete a bus cycle, the 80286 microprocessor requires the aid of external logic. One of the major pieces of logic involved in this process is referred to as the bus control logic. Refer to figure 8-1.

Logic external to the microprocessor can detect the beginning of a bus cycle by looking at the 80286's bus cycle definition outputs. When either S0# or S1# is detected going low, this indicates the start of the bus cycle. This triggers the bus control logic's state machine that works with the microprocessor's Bus Unit to accomplish a data transfer during a bus cycle. The bus control logic uses CLK2, the double frequency clock, to define time slots for its state machine. At the

ISA System Architecture

proper points during a bus cycle, the bus control logic performs the appropriate actions to accomplish the bus cycle.

The functions performed by the bus control logic are described in this chapter.

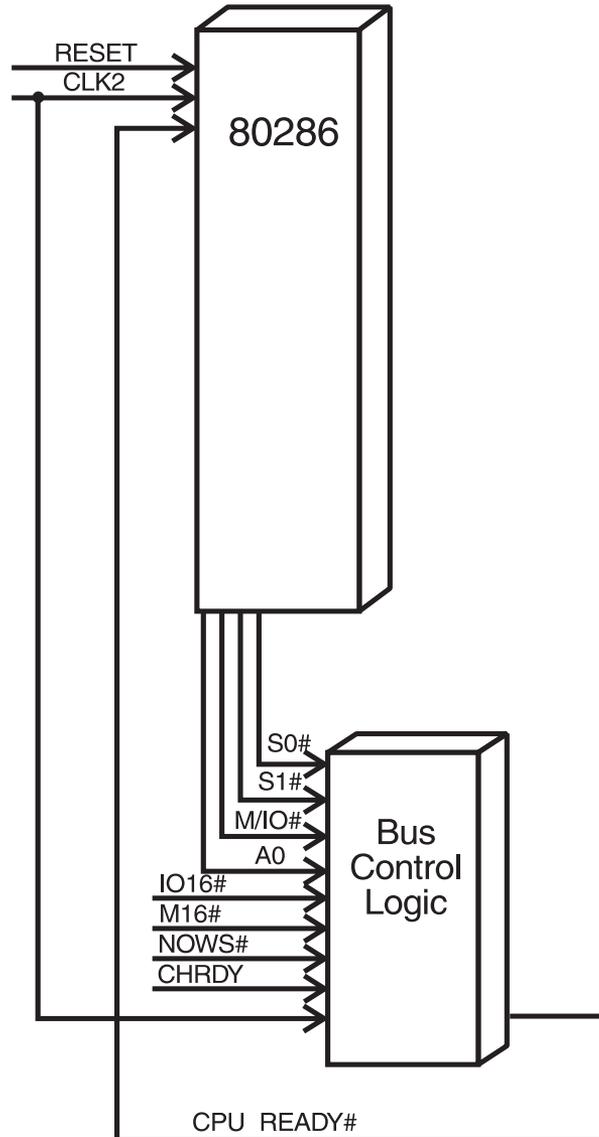


Figure 8-1. The Bus Control Logic

Chapter 8: The 80286 System Kernel: The Engine

The Address Latch

Refer to figure 8-2. Intel dictates that every 80286-based system must incorporate an external device known as an address latch. When the microprocessor outputs the address onto its local address bus during a bus cycle, the bus control logic signal ALE (address latch enable) commands the address latch to hold (latch) the address and remember it. Once latched, the address latch outputs the address to the system on the system address (SA) bus, SA[19:1]. Address decoders throughout the system can then examine the latched address to see if the microprocessor is attempting to communicate with their respective device.

ISA System Architecture

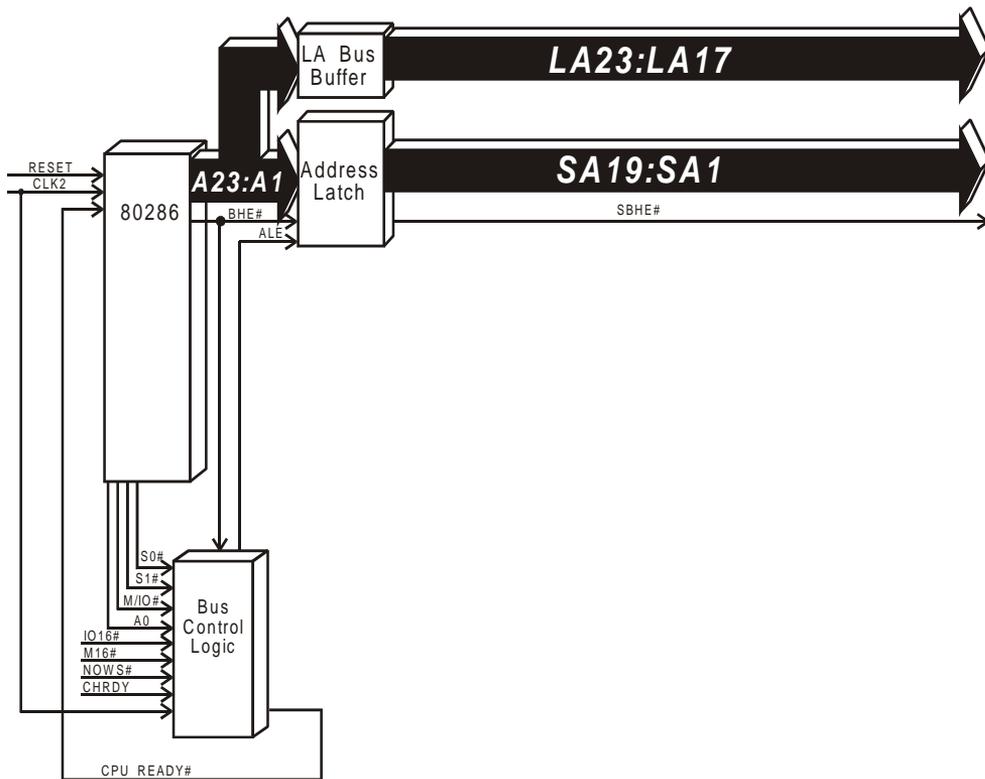


Figure 8-2. The Address Latch

The address latch and LA lines (latchable address lines) must be included to support the 80286's address pipelining capability. Notice that address lines A[23:20] are not latched. These address lines go directly to the LA bus through a buffer, along with address lines A[19:17]. From the buffer, the LA bus is connected directly to the 16-bit slots on the ISA bus. The purpose of the LA bus will be described later.

Notice that the address latch only latches A[19:1], but not A0 (address bit A0). The reason for this is clarified later. Rather than going to the address latch, A0 goes to the bus control logic where it is allowed to pass through onto SA0 at the same time as A[19:1] are latched and presented on SA[19:1].

Chapter 8: The 80286 System Kernel: The Engine

Address Pipelining

When the 80286 microprocessor must perform back-to-back memory transfers, it uses address pipelining. This involves placing the address for the next cycle on the microprocessor's local address bus during the current bus cycle. This is possible because the current address will have been latched on the system address bus (SA Bus) and held for the duration of the current bus cycle. As a result, the address for the next bus cycle can be output by the microprocessor since the address for the current bus cycle has already been latched.

A portion of the next bus cycle's address, LA[23:17], goes directly to the ISA bus via an address buffer. LA[23:17] (the latchable address bus) are simply a buffered version of the uppermost bits of the microprocessor's address bus. 16-bit memory expansion boards are designed to decode the LA address lines. This allows 16-bit memory expansion boards to chip select their memory devices much earlier than would be possible if they waited for the latched address (SA lines).

At the end of address time of the next bus cycle, the bus control logic again pulses the address latch enable (ALE) line, commanding the address latch to latch address lines SA[19:0] and SBHE#. Since the high order bits (LA[23:17]) have already been used to generate the chip select, address decoding can be completed very quickly. Additional information regarding address pipelining can be found in the chapter entitled "Detailed View of the Bus Cycle."

Although the 80286 microprocessor pipelines out its entire 24-bit address, the ISA bus only provides seven of those lines, LA[23:17]. It would have been more expensive to include all 24 of the pipelined address lines and system designers are always looking for ways to reduce cost while maintaining or improving performance. If the designers had provided all 24 lines, the cost would have gone up but the performance would not have improved. There are two reasons for this: 1) the vast majority of bus cycles are memory reads since the processor is constantly fetching instructions and 2) the 4164 memory chip was popular when the PC/AT was designed. The 4164 has 64Kbits of information in it, but it's only one bit wide. Sixteen 4164's are needed to make a bank of memory for the 80286 microprocessor. Sixteen times 64Kbits yields 128KB, which is the smallest block of addresses that can be produced using LA[23:17]. With LA[23:17], the memory subsystem has enough information to select the proper bank of memory chips. Additional information about memory banks can be found in the chapter entitled "RAM Memory: Theory of Operation."

Chapter 9

The Previous Chapter

Previous chapters have covered how the microprocessor communicates with memory and I/O devices and have explained the support logic needed to allow the microprocessor to communicate with both 8- and 16-bit devices.

This Chapter

The exact timing of read and write bus cycles for both memory and I/O are covered in this chapter. Detailed explanations of the halt and shutdown bus cycles, along with the causes of each are also covered.

The Next Chapter

The next several chapters detail the 80386 microprocessor and the support circuitry required for the 80386 kernel.

Address and Data Time Revisited

The concept of the bus cycle was introduced in the chapter entitled "Introduction to the Bus Cycle." The microprocessor performs a bus cycle when it has to transfer information between itself and a memory or I/O location. The microprocessor's bus unit uses the address, data and control buses to address a device, indicate the type of transaction in progress and to transfer the data between the microprocessor and the currently addressed location.

The microprocessor uses the bus cycle definition lines to indicate the type of transaction. The types of bus cycles are indicated in table 9-1.

With the exception of the interrupt acknowledge bus cycle, this chapter provides a detailed description of each bus cycle type. The interrupt acknowledge bus cycle is described in the chapter entitled "The Interrupt Subsystem."

ISA System Architecture

Table 9-1. 80286 Bus Cycle Definition

M/IO#	S1#	S0#	Bus Cycle Type
0	0	0	Interrupt Acknowledge
0	0	1	I/O Read
0	1	0	I/O Write
1	0	0	Halt or Shutdown
1	0	1	Memory Read
1	1	0	Memory Write

When it must perform a bus cycle, the microprocessor's bus unit leaves the idle state and enters address time. During this "tick" of PCLK, the microprocessor places the address and bus cycle definition on the buses. Address decoders throughout the system start decoding the address during this time slot.

Address time is always followed by data time. During this state, the microprocessor expects the data to be transferred between itself and the currently addressed device.

In reality, these aren't the state names Intel uses for the 80286 and 80386 bus unit states. Table 9-2 lists the state names used by Intel:

Table 9-2. 80286 and 80386 Bus Cycle State Names

Nickname	80286 Name	80386 Name
Address Time	Ts	T1
Data Time	Tc	T2

These states were renamed for clarity's sake because the names Ts and Tc provide little information regarding the actions that occur during these states. Intel calls address time Ts, meaning "Send Status" time, because one of the things that occurs during Ts is the output of the bus cycle definition on M/IO#, S0#, and S1#. Intel refers to S0# and S1# as the microprocessor's status outputs, hence the name "Send Status" time. Data time is called Tc, "Perform Command" time, because it's time to perform the command, or data transfer, that the microprocessor initiated.

The various actions that must take place during a bus cycle have been described in previous chapters. This chapter provides an in-depth look at the sequence and exact timing of these actions in relation to each other.

Chapter 9: Detailed View of the 80286 Bus Cycle

The Read Bus Cycle

With the exception of the M/IO# output, I/O and memory read bus cycles are identical. This section describes the exact timing of the actions performed during a read bus cycle. Figure 9-1 is a timing diagram illustrating a series of typical 80286 read bus cycles.

PCLK has been placed across the top of the timing diagram as a point of reference. Remember that each cycle of PCLK defines the duration of one bus unit state. Vertical dotted lines have been superimposed on the diagram to define each cycle of PCLK (and, therefore, each state). In addition, the name of each state is written across the top of the diagram.

A general idea of the information presented can be derived from the state names across the top. In this example, the sequence of states is as follows:

Tc	Data Time
Ts	Address Time
Tc	Data Time
Ts	Address Time
Tc	Data Time
Tc	Data Time
Ts	Address Time

Every bus cycle consists of at least an address time and data time pair. Additional data times (wait states) may be inserted in the bus cycle if the microprocessor samples its READY# input deasserted at the end of data time. Based on these rules, the states illustrated on the diagram can be interpreted as follows:

Tc	Data Time	end of previous bus cycle
Ts	Address Time	start of bus cycle A
Tc	Data Time	1st data time of bus cycle A
Ts	Address Time	start of bus cycle B
Tc	Data Time	1st data time of bus cycle B
Tc	Data Time	wait state inserted in bus cycle B
Ts	Address Time	start of next bus cycle

In other words, the diagram illustrates the end of a bus cycle followed by a 0-wait-state bus cycle (bus cycle A), a 1-wait-state bus cycle (bus cycle B) and the start of another bus cycle.

ISA System Architecture

When reading a timing diagram, one scans from left-to-right (from earlier in time to later in time) looking for the signals that change first, second, and so on. Since PCLK should always be pulsing (changing), it's not necessary to note the fact that it changes.

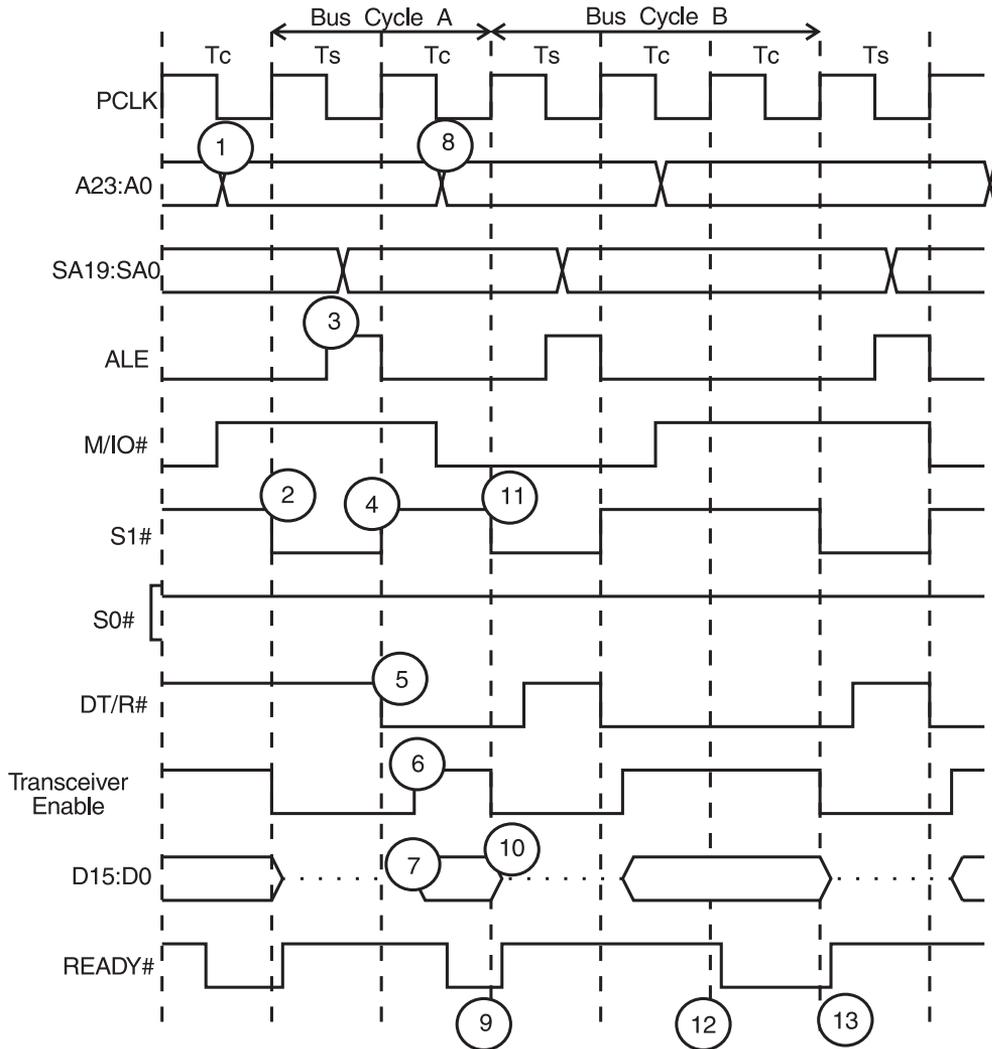


Figure 9-1. The Read Bus Cycle

Chapter 9: Detailed View of the 80286 Bus Cycle

Figure 9-2 shows the 80286 system engine. Referencing both the timing diagram and the functional block diagram will help tie the timing and functions together for a more complete understanding.

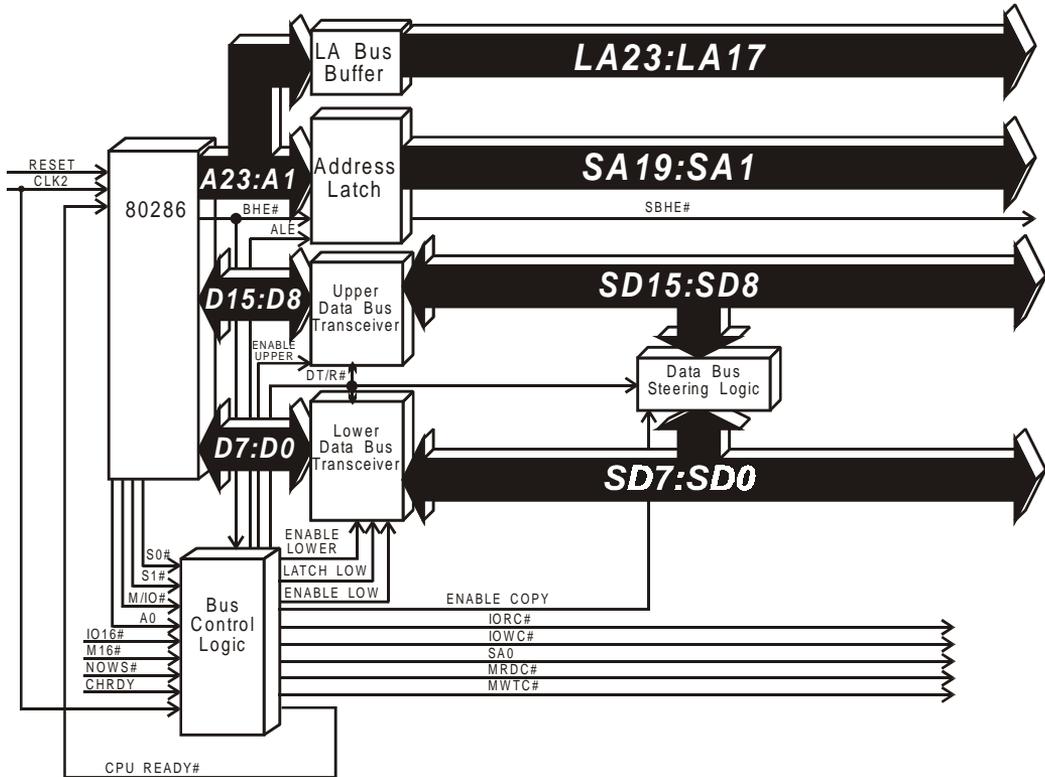


Figure 9-2. The 80286 System Engine

Chapter 10

The Previous Chapter

A detailed analysis of the 80286 microprocessor, the kernel logic necessary to interface it to external devices, and its bus cycle types were covered in prior chapters.

This Chapter

This chapter provides a detailed description of the 80386DX and SX microprocessors. It covers the functional units, processor self-test, enhancements to protected mode, virtual paging, virtual 8086 mode, and the processor's hardware interface to external devices.

The Next Chapter

The next chapter, "The 80386 System Kernel," defines the support logic necessary to interface the 80386DX and SX microprocessors to external devices.

Introduction

This chapter describes the 80386DX and SX microprocessors. The 80386DX and SX microprocessors are identical internally and only differ in their interface to external devices. This chapter is therefore divided into three parts:

- a discussion of the internal structure and operation of the 80386 microprocessor.
- a discussion of the 80386DX interface to external devices.
- a discussion of the 80386SX interface to external devices.

ISA System Architecture

The 80386 Functional Units

General

Figure 10-1 illustrates the units that make up the 80386 microprocessor.

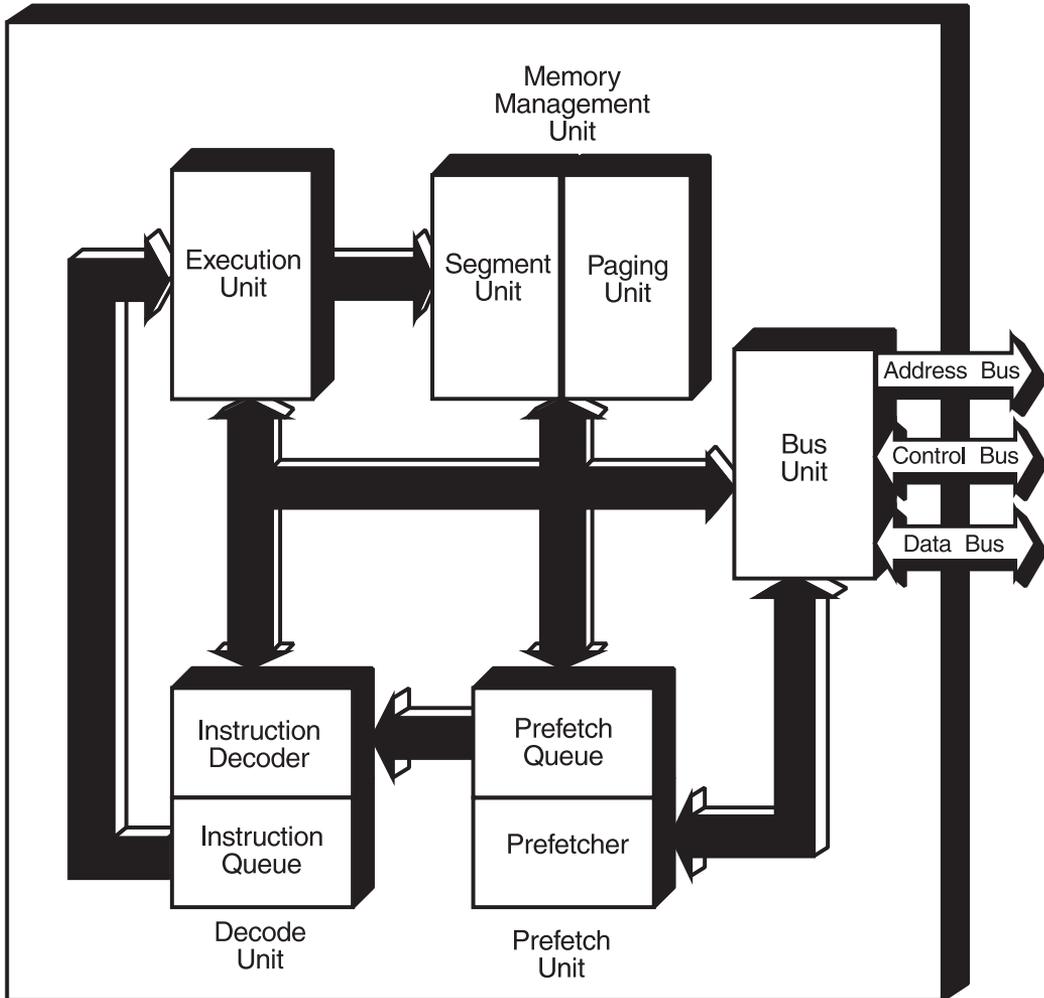


Figure 10-1. 80386 Microprocessor Block Diagram

Chapter 10: The 80386DX and SX Microprocessors

The 80386 microprocessor consists of five functional units:

- **Bus Unit.** Handles communication with devices external to the microprocessor chip.
- **Code Prefetch Unit.** Fetches instructions from memory before the microprocessor actually requests them.
- **Instruction Decode Unit.** Decodes the instruction prior to passing it to the execution unit for execution.
- **Execution Unit.** Handles the actual execution of the instruction.
- **Memory Management Unit (MMU).** When the microprocessor must address a memory location, the MMU forms the physical memory address that is driven out onto the address bus by the bus unit during a bus cycle.

Code Prefetch Unit

The code prefetch unit capitalizes on the predictability of program execution. Statistically, program instructions are fetched from memory sequentially most of the time. Only jump instructions alter program flow, causing it to jump to another area of memory. Once the flow alteration has occurred, however, the program returns to sequential instruction processing until another jump instruction is encountered.

When the bus unit isn't performing bus cycles for the execution unit, the code prefetch unit uses the bus unit to fetch the next sequential instruction from memory and stores it in the 16-byte prefetch queue. The code prefetch unit always attempts to maximize the bus unit's throughput by fetching an entire doubleword (four bytes) at a time. The prefetched instructions are placed in the prefetch queue to await processing by the instruction decode unit.

Code prefetch requests are given a lower priority by the bus unit than requests from the execution unit. This ensures that the execution unit will not be stalled while waiting for a memory or I/O data transfer to complete. For all practical purposes, instruction prefetching reduces the time the execution unit spends waiting for the next instruction to zero.

Instruction Decode Unit

The instruction decode unit takes an instruction from the 16-byte prefetch queue, converts it into microcode and stores it in a three-deep decoded instruction queue for use by the execution unit. Any immediate data and offset associ-

ISA System Architecture

ated with the instruction is also taken from the prefetch queue and stored in the decoded instruction queue.

Execution Unit

General

The execution unit executes each instruction received from the decoded instruction queue. It communicates with the other microprocessor functional units on an as needed basis in order to execute each instruction. The execution unit is comprised of three sub-units:

- The control unit's microcode and special parallel hardware speeds multiplies, divides and effective address calculation.
- The data unit contains the arithmetic logic unit, or ALU, eight general-purpose registers and a 64-bit barrel shifter. The data unit performs data operations requested by the control unit.
- The protection test unit checks for segmentation violations.

The Registers

The 80386 execution unit incorporates the following registers:

- The general registers
- The status and control registers
- The debug registers
- The test registers

The following sections describe each of these register groups.

General Registers

The general registers implemented in the 80386 microprocessor are a superset of those found in the earlier implementations of the X86 processor family. Figure 10-2 illustrates the general registers.

Chapter 10: The 80386DX and SX Microprocessors

	31	23	15	8	7	0
EAX				AH	AX	AL
EBX				BH	BX	BL
ECX				CH	CX	CL
EDX				DH	DX	DL
EBP				BP		
ESI				SI		
EDI				DI		
ESP				SP		

Figure 10-2. The 80386 General Registers

Extended versions of the AX, BX, CX, DX, BP, SI, DI and SP registers are now available to the programmer. The names of each of the extended registers starts with the letter "E". Referring to the EAX, EBX, ECX or EDX registers within a MOV instruction allows the programmer to specify a doubleword, or 32-bit, object to be moved. Extending the width of BP, SI, DI and SP registers to 32 bits permits the programmer to specify any memory address within the 4GB addressing range of the microprocessor without changing the contents of the segment registers.

Status, MSW and Instruction Registers

Figure 10-3 illustrates the 80386's status and instruction registers.

Appendix B

Glossary

- access time** - The time span between a device being addressed and when it delivers or accepts valid data.
- active** - see asserted
- address bus** - The group of signal lines that carry an address from a bus master.
- address decoder** - Every device that can be read from or written to has an associated piece of logic called an address decoder. The address decoder selects its associated device when an address within its defined range is detected.
- address latch** - A device connected to the microprocessor's local address bus. This device drives the address to all devices throughout the system and latches (holds) the address for the duration of the bus cycle. Once latched, the Address Latch outputs the address to the system on the system address (SA) bus, SA19:SA1.
- address pipelining** - A process in which the microprocessor places the address for the next cycle on its local address bus during the current bus cycle. This can reduce access time to system memory if designed to take advantage of address pipelining.
- address translation** - The process of converting one type of address to another. For example: translating the address from an ISA bus master (SA16:SA0, LA23:LA17 and SBHE#) to a 32-bit memory address (A31:A2 and BE3#:BE0#) required by 32-bit memory.
- AEN** - An ISA bus signal or a pin on the 8237 DMA controllers. Note that the same name is used for two differently-acting signals. When the CPU asserts the HLDA signal, the ISA bus AEN signal is asserted, disabling all I/O address decoders to prevent I/O devices from mistaking the address as theirs during DMA transfers. When the DMA controller sees its HLDA input asserted, it asserts its AEN pin to enable the outputs of its address latch and page register.
- ALE** - The Address Latch Enable signal that controls the operation of the Address Latch.
- arbiter** - An arbiter is a device that evaluates the pending requests for access to the bus and grants the bus to a bus master based on a system-specific algorithm.
- asserted** - A signal is asserted when it is at its logic true state. A signal such as HOLD is asserted when it is high (positive logic). A signal such as MRDC# is asserted when it is low (negative logic). A signal such as W/R# is always asserted since both its high and low states represent true logic states. The opposite of asserted is deasserted. The term "asserted" is used in this book instead of "active" because "active" is ambiguous. The term "active" can mean either that a high impedance buffer's output is on (driven), or that a signal is at a logic true state. For example, a tri-state buffer may be actively driving CHRDY low (false), in which case the signal is active (by one of the ambiguous definitions) but inactive (by the other definition). In this case, CHRDY is deasserted.
- BALE** - An ISA bus signal that is used by expansion devices to notify them that a valid address is on the ISA bus. BALE is a buffered version of ALE on 8088- and

ISA System Architecture

80286-based PC's, but it is asserted at different times than ALE on 80386 and higher PC's.

- BCLK** - An ISA bus signal (bus clock) that provides the timing reference for all bus transactions.
- big-endian byte-ordering** - A bus master using big-endian byte-ordering stores the MSB from the specified source register into the start memory address and the less-significant bytes from the register into the ascending memory locations immediately following the start address.
- BIOS** - Basic Input/Output System routines usually contained in system ROM and device ROMs that provide a low level interface to devices.
- bit cell** - A one bit storage location in a typical DRAM consisting mainly of a capacitor.
- block** - Same as cache line in X86 processors. The cache stores information in its cache in blocks of a fixed length. The cache block length is processor design-dependent. As an example, an 80486 cache block is 16 bytes long.
- block transfer mode** - A DMA transfer mode in which an entire block of data is transferred prior to the DMA Controller giving up ownership of the buses.
- bridge** - The device that provides the interface between two independent buses. An example is the Bus Control Logic and associated transceivers between the 80386 host processor's bus and the ISA bus.
- buffered write-through** - A variation of the write-through policy where a look-through cache controller stores an entire write operation in a buffer so that it can complete the memory write to main memory later.
- bus concurrency** - Separate transfers occurring simultaneously on two or more separate buses. An example would be an ISA bus master transferring data to or from another ISA device while the host processor is transferring data to or from system memory. Bus concurrency allows a processor to gain access to memory from its memory cache, while another bus master accesses main memory over the system bus at the same time.
- bus control logic** - The logic that assists the microprocessor in running bus cycles over the ISA bus.
- bus master, ISA** - An ISA card that can gain ownership of the ISA buses and run bus cycles.
- bus snooping** - Bus snooping is the method used by cache subsystems to monitor main memory accesses made by another bus master.
- byte enables** - 80386 and higher X86 processor address bus signals that indicate which data paths will be used during a transfer. The byte enables indicate which bytes within an addressed doubleword are being accessed.
- cache** - A relatively small amount of high-speed static RAM (SRAM) that is used to keep copies of information recently read from system DRAM memory. The cache controller maintains a directory that tracks the information currently resident within the cache. If the host processor should request any of the in-

formation currently resident in the cache, it will be returned to the processor quickly (due to the fast access time of the SRAM).

cache coherency - If the information resident in a cache accurately reflects the information in DRAM memory, the cache is said to be coherent or consistent.

cache consistency - see **cache coherency**

cache controller - A cache memory controller manages cache memory which stores copies of frequently accessed information read from DRAM memory.

cache directory - Memory inside of a cache controller that keeps track of information stored in cache memory. Sometimes called the tag RAM.

cache FLUSH# - A cache controller input that causes the cache controller to clear the valid bit in every directory entry to 0, thereby invalidating all cache entries.

cache line - A line is the smallest unit of data that a cache can keep track of.

cache line fill - When a processor's internal cache, or its external second level cache has a miss on a read attempt by the processor, it will read a fixed amount (referred to as a line) of information from the external cache or system DRAM memory and record it in the cache. This is referred to as a cache line fill. The size of a line of information is cache controller design dependent.

cache memory - Cache memory is a relatively small amount of high cost, fast access SRAM designed to improve access to system memory.

cache page - The size of the cache data RAM in direct-mapped caches and the size of each way in a set-associative cache. The cache views main memory as divided into pages that are the same size as the cache pages.

cache read hit - The process in which the cache memory controller sees the microprocessor initiate a memory read bus cycle, checks to determine if it has a copy of the requested information in cache memory, and if a copy is present, immediately reads the information from the cache and sends it back to the microprocessor at zero wait-states.

cache read miss - The process in which the cache memory controller sees the microprocessor initiate a memory read bus cycle, checks to determine if it has a copy of the requested information in cache memory, and find no copy is present. The cache memory controller then passes the read bus cycle on to slow main memory.

CAS# - A memory enable called Column Address Strobe used to latch the column portion of the address inside a DRAM chip.

CAS before RAS refresh - Some DRAMs incorporate their own row counters to be used for DRAM refresh. The external DRAM refresh logic has only to assert the DRAM's CAS line and then assert its RAS line. The DRAM will automatically increment its internal row counter and refresh (recharge) the next row of storage.

CBR refresh - see **CAS before RAS refresh**

channel check - This signal is asserted by an ISA expansion board to signal a catastrophic error condition to the microprocessor.

ISA System Architecture

channel ready - By using the CHRDY signal (Channel Ready) on the ISA bus, designers can override the default timer and stretch out the bus cycle by the required number of wait states to match the access time of the device.

CHCHK# - see **channel check**

CHRDY - see **channel ready**

CMOS RAM - see **configuration RAM**

coherency - see **cache coherency**

concurrent bus operation - see **bus concurrency**

configuration RAM - A small amount of low power battery-backed memory used to store configuration information used during the boot-up process.

consistency - see **cache coherency**

control bus - A group of signal lines used by the microprocessor to control a variety of miscellaneous functions.

cycle time - The amount of time needed between accesses to DRAM memory, consisting of access time plus precharge delay.

DAKn# - A group of signals on the ISA bus (named DACKn# or DAKn#) used to notify DMA I/O devices that a DMA cycle is about to be performed to transfer data for them.

data bus - The group of signal lines used to transfer data between devices.

data bus contention - An address decoding problem that arises if the address decoders for two devices are both designed to detect the same address range. In such cases both devices would deliver data to the data bus simultaneously during a read operation, causing contention.

data bus steering logic - Logic used to match transfers between devices with different sized data busses. Consists of transceivers connected between the data paths.

deasserted - A signal is deasserted when it is at its logic false state. See **asserted**.

debug registers - Six programmer-accessible debug registers provide on-chip support for debugging and are present in the 80386. Debug registers DR0 through DR3 specify the four linear breakpoint addresses. The Breakpoint Control Register, DR7, is used to set the breakpoints, define them as data or code breakpoints, and whether to break on an attempted read or write. The Breakpoint Status Register, DR6, reflects the current state of the breakpoints.

demand transfer mode - The transfer mode where the DMAC runs DMA bus cycles back-to-back as long as the I/O device can continue to supply data.

device ROM - ROM on an ISA expansion device. Also known as expansion ROM. This ROM typically contains the initialization code and possibly the BIOS and interrupt service routines for its associated device. Device ROMs should reside in memory address space between C0000h and DFFFFh.

direct-mapped cache - A cache organization in which only one directory entry needs to be checked to determine whether the requested memory location is in the cache. A direct-mapped cache is the same size as a (cache) page in memory. When a line of information is fetched from a position within a page of DRAM

1

1.19318MHz clock, 469, 471
14.31818MHz clock, 350, 469
16-bit read from 16-bit device, 134
16-bit read from 8-bit device, 133
16-bit write to 16-bit device, 134
16-bit ISA bus master, 345
16-bit write to 8-bit device, 132

3

32KHz crystal oscillator, 441

8

80286 and 80386 bus cycle state names, 140
80286 bus cycle definition, 140
80286 microprocessor address space, 34
80287, 459
80386DX outputs while RESET is asserted, 115
80386SX microprocessor address space, 34
80387, 459
8042 local I/O ports, 450
8042 microcomputer, 449
8042 sample keyboard/mouse interface command list, 455
8042 test port, 452
8080 microprocessor, 29
8086, 33
8087 numeric coprocessor, 459
8088, 33
8237 DMA controller, 410
8259 programmable interrupt controller, 366
8-bit read from 16-bit device, 134
8-bit write to 16-bit device, 134
8-bit write to odd addressed location in 8-bit device, 131
8-Bit read from odd addressed location in 8-bit device, 129

A

A2 output, 465
A20 GATE, 79, 450
A20 MASK#, 79
access time, 134
address pipelining, 123
address bus, 17
address decoder, 18, 39
address enable, 345
address latch, 121
address translation, 438
AEN, 345, 424
alarm output of real-time clock and configuration RAM chip, 386
ALE (address latch enable), 121
alternate gate A20, 400
alternate hot reset, 400
atomic access, 94

B

BALE, 342
basic input/output system (BIOS) routines, 405
BASIC language, 322
battery, 441
buffered write-through cache, 293
bus control logic, 119
bus snooping, 295
BCLK, 341
BHE#, 83
big-endian byte ordering rule, 74
BIOS, 322
BIOS ROM, 42
BIOS routine, 43, 404
bit cell, 244, 315, 318
BLE#, 188
boot program, 322
boot ROM, 42, 116
BOUND range exceeded exception, 403
breakpoint, 403
buffered address latch enable, 342
bus clock, 341

ISA System Architecture

bus concurrency, 278
bus high enable (BHE#), 83
bus low enable (BLE#), 188
bus master capability, ISA, 433
BUSY#, 461, 464

C

cache coherency, 292
cache FLUSH#, 313
cache handling of non-cacheable memory, 312
cache line, 283
cache memory, 279
cache memory controller, 274
cache read hit, 274
cache read miss, 275
cache ways, 305
cache, advantage of, 276
CAS# (column address strobe), 236, 242
CAS-before-RAS refresh, 248
CBR refresh, 248
cache and testing memory, 313
cache consistency, 279
cache directory, 281
cache handling of I/O information, 311
cache line size, 308
cache management logic, 281
cache performance, 278
cache size, 310
column address decoder, 237
column address latch, 236
column address strobe (CAS#), 242
channel or I/O check, 349, 400
channel ready, 342
CHCHK#, 349
checksum error, 331
checksumming, 322
CHRDY, 137, 342
clear interrupt enable instruction, 368
CLI, 368
CMD0, 465
CMD0#, 465
CMD1, 465
CMOS address decoder, 442
CMOS address port, 442

CMOS century part of time and date function, 446
CMOS checksum value, 446
CMOS data port, 442
CMOS diagnostic status byte, 446
CMOS diskette drive type byte, 446
CMOS equipment installed byte, 446
CMOS fixed disk drive type byte, 446
CMOS indicator for memory above 1MB, 446
CMOS low and high base memory bytes, 446
CMOS RAM memory, 441
CMOS RAM valid, 445
CMOS reset code byte, 446, 447
CMOS system information byte, 446
CMOS, accessing, 442
code cache, 291
column address, 236, 238, 239, 242
combined cache, 290
command outputs, 145
command port, 36
configuration RAM, 460
configuration RAM usage, 446
control bus, 19
control/alternate/delete soft reset, 111, 447
coprocessor chip-select, 465
coprocessor emulation, 460, 465
coprocessor error, 403
coprocessor error conditions, clearing, 463
coprocessor RESET, 463
coprocessor status register, 463
coprocessor's data register, 460
coprocessor's opcode, or instruction, register, 460
coprocessor, systems using Intel and Weitek, 468
coprocessor/processor communication signals, 464
copy protection scheme, 474

D

DAK0#:DAK3#, 345
DAK2#, 412
DAK4#, 417

- DAK5#:DAK7#, 345
- Dallas Semiconductor real-time clock and configuration RAM chip, 441
- data bus, 19
- data bus contention (address conflicts), 41
- data bus transceivers, 124
- data cache, 291
- DATA ENABLE, 126
- data port, 37
- date and time, 441, 470
- data bus steering logic, 126
- default ready timer, 135
- direct-mapped cache, 301
- dynamic RAM (DRAM) memory, 235
- dedicated cache, 290
- descriptor, 101, 162
- descriptor privilege level (DPL), 105
- device ROMs, 330
- divide-by-zero condition, 402
- DMA acknowledge, 412
- DMA acknowledge for channels 0 through 3 and 5 through 7, 345
- DMA address latch, 424, 431
- DMA ADSTB signal, 424
- DMA AEN signal, 424
- DMA block transfer mode, 416
- DMA bus cycle, 420
- DMA byte or word transfers, 423
- DMA channels, 410
- DMA clock, 420
- DMA compressed timing, 421
- DMA controller (DMAC), 409, 414
- DMA controllers cascaded, 419
- DMA demand transfer mode, 416
- DMA direction of transfer with reference to memory, 410
- DMA extended write option, 422
- DMA master and slave controllers, 417
- DMA memory address logic, 425, 426
- DMA page register, 423, 424
- DMA read transfer, 414
- DMA request, 412
- DMA request for channels 0 through 3 and 5 through 7, 345
- DMA single transfer mode, 415
- DMA start memory address, 410
- DMA transfer count, 410
- DMA transfer count exhausted, 413
- DMA transfer modes, 414
- DMA transfer rate, 431
- DMA transfer types, 414
- DMA verify transfer, 414
- DMA write transfer, 414
- DMA, EOP (end-of-process), 413
- DMA, fly-by transfer, 409
- DMA, TC (terminal count reached), 413
- DMAC, 409
- DMAC control register, 431
- DMAC initialization during POST, 431
- DMAC mode register, 431
- DMAC states, 420
- don't care bits, 46
- double exception detected, 403
- DPL, 105
- DRAM, 235
- DRAM addressing sequence, 236, 242
- DRAM access time, 250
- DRAM address decoder, 240
- DRAM addressing logic, 238
- DRAM addressing multiplexer, or MUX, 240
- DRAM bank, 251
- DRAM bank width, 253
- DRAM memory chip block diagram, 237
- DRAM parity, 256
- DRAM refresh logic, 471
- DRAM refresh, 244
- DRAM, burst mode, 265
- DRAM, cycle time, 250
- DRAM, destructive read, 250
- DRAM, enhanced page-mode, 264
- DRAM, interleaved memory architecture, 269
- DRAM, nibble mode, 265
- DRAM, page-mode, 259
- DRAM, pre-charge delay, 250
- DRAM, static column RAM (SCRAM), 267
- DRQ0:DRQ3, 345
- DRQ2, 412
- DRQ4, 417
- DRQ5:DRQ7, 345
- DT/R#, data transmit or receive, 126

ISA System Architecture

E

ECC memory, 259
EEPROM, 319
EM bit, 460, 466
ENABLE LOWER, 126
ENABLE UPPER, 126
EOI (end-of-interrupt) command, 379
EOI (end-of-interrupt) command, for slave, 388
EOP (end-of-process), 413
EPROM, 318
ERROR# output of the numeric coprocessor, 386, 463, 464
error-checking-and-correcting memory, 259
exception 7 interrupt request, 466
exception 7 interrupt service routine, 466
exception interrupts, 402
extended memory block-move return, 447

F

first level cache, 289
four-way set-associative cache, 307
fully associative cache, 299
fixed disk activity LED, 400
flag register, 374
floating gate, 318
floating-point, 459
floating-point calculations, 466
floating-point processor, 459
floppy disk controller, 387
FLUSH#, 313
FORCE A20, 79, 451

G

GDT, 101
GDTR, 102
general protection software exception, 403
global descriptor table (GDT), 101
global descriptor table register (GDTR), 102

H

halt, 150, 230
halt or shutdown bus cycle, 150
hard disk controllers, 387
Hi/Lo byte copier, 130
hidden refresh, 248
high memory area, 79, 327
HLDA, 412
HMA, 79, 327
HOLD, 412, 475
hold time, 149
hot reset, 450

I

I/O address decoder, system board, 47
I/O address range assigned to system board devices, 52
I/O bus, 335
I/O check, 349
I/O ports, 30, 36
I/O read command, 340
I/O size 16, 344
I/O write command, 340
ICW1, 397
ICW2, 371, 397
ICW3, 397
ICW4, 397
IDT, 101
IDTR, 102
interleaved memory architecture, 269
IMR, 397
IN instruction, 32
in-line code fetch, 14
in-service register, 371
instruction cache, 291
INT instruction, 403
INTA, 369
integer math, 465
interrupt acknowledge, 369, 374
interrupt cascade bus, 387
interrupt controller, 366
interrupt controller cascading, 383
interrupt controller, slave, 387

- interrupt descriptor table (IDT), 101
- interrupt descriptor table register (IDTR), 102
- interrupt ID, 371
- interrupt instruction, 403
- interrupt mask register, 397
- interrupt on overflow, 403
- interrupt priority scheme, 368, 384
- interrupt request, 365
- interrupt request (IRQ) lines, 347, 366
- interrupt request level 0, 384
- interrupt request level 7, 384
- interrupt request register (IRR), 371, 399
- interrupt return instruction, 379
- interrupt service routines, 43, 322, 366, 369
- interrupt servicing flowchart, 367
- interrupt servicing, clearing the interrupt enable flag, 378
- interrupt servicing, jumping to the ISR, 378
- interrupt servicing, saving instruction pointer, 374
- interrupt table, 369
- interrupt table entry assignments, 388
- interrupt table entry number, 371
- interrupt table initialization, 371
- interrupt type code, 371
- interrupt vector, 371
- interrupts, dedicated IRQ lines, 386
- interrupts, edge-triggered mode, 397
- interrupts, end-of-interrupt, or EOI, command, 397, 398
- interrupts, exception, 402
- interrupts, initialization command words, or ICWs, 396
- interrupts, initialization sequence for the master interrupt controller, 397
- interrupts, initialization sequence for the slave interrupt controller, 398
- interrupts, in-service register, 398
- interrupts, master's interrupt ID, 397
- interrupts, operation command words, or OCWs, 396
- interrupts, programming the 8259, 396
- interrupts, software interrupts, 402
- interrupts, special fully-nested mode, 397
- interrupts, special mask mode, 399
- INTR, 368
- invalid opcode exception, 403
- invalid task state segment, 403
- IO16#, 128
- IORC#, 340
- IOWC#, 340
- IRET (interrupt return) instruction, 379, 404
- IRQ lines, 347
- IRQ0, 386, 472
- IRQ0 frequency, 470
- IRQ0 interrupt service routine, 470
- IRQ0 latch, 469, 473
- IRQ0 through IRQ7, 386
- IRQ0, resetting, 470
- IRQ1, 372, 386, 449
- IRQ10, 387
- IRQ11, 387
- IRQ12, 387, 450
- IRQ13, 386, 463, 464
- IRQ13 interrupt service routine, 463
- IRQ14, 387
- IRQ15, 387
- IRQ2, 383, 386, 387
- IRQ2 redirect, 387, 389
- IRQ3, 386
- IRQ4, 386
- IRQ5, 386
- IRQ6, 387
- IRQ7, 387
- , 391, 395
- IRQ8, 386, 444
- IRQ8 through IRQ15, 386
- IRQ9, 387
- IRR, 371
- ISA address bus, 338
- ISA bus cycle definition, 339
- ISA bus cycle timing, 341
- ISA bus master capability, 433
- ISA command lines, 340
- ISA connector, 337
- ISA connector, 16-bit, 336
- ISA connector, 8-bit, 336
- ISA device size lines, 343
- ISA DMA transfer rates, 431
- ISA DMA-related signals, 345
- ISA error reporting signal, 349

ISA System Architecture

ISA interrupt request (IRQ) lines, 347
ISA interrupt-related signals, 347
ISA miscellaneous signals, 350
ISA reset signal, 344
ISR, 366
ISRs, 369

J

jump instruction, 13

K

key matrix, 450
keyboard, 449, 450
keyboard BIOS routine, 457
keyboard break scan code, 450
keyboard command port,, 451
keyboard data port, 450
keyboard interrupt request, 450
keyboard interrupt request line, 386
keyboard interrupt service routine, 372, 450
keyboard make scan code, 450
keyboard scan, 450
keyboard scan code, 450
keyboard/mouse command/status port, 454
keyboard/mouse data port, 456
keyboard/mouse interface, 449

L

LA bus, 122
LA23:LA17, 338
latchable address bus, 122, 338
LDT, 101
LDTR, 102
line, 283, 308
little-endian byte ordering rule, 74
least recently used (LRU) algorithm, 307
local data bus, 126
look-aside cache, 287
local descriptor table (LDT), 101
local descriptor table register (LDTR), 102
look-through cache, 283

M

M/IO#, 33
M16#, 128, 344
machine status word (MSW) register, 116
maskable interrupt request line (INTR), 368
master interrupt controller, 383
MASTER16#, 345, 435
MC146818 real-time clock and configuration RAM chip, 441
memory address bus, 236, 241
memory parity error, 399
memory read command, 340
memory size 16, 344
memory write command, 340
memory-based semaphore, 93
miscellaneous ISA signals, 350
MODE command, 475
Motorola MC146818 real-time clock and configuration RAM chip, 441
mouse, 449, 450
mouse BIOS routine, 457
mouse interface, 387
mouse interrupt request, 450
MOV instruction, 33
MP bit, 460, 466
MRDC#, 340
MROM, 317
MS-DOS function calls, 404
MSW register, 460
multitasking operating system, 472
MUX, 240
MWTC#, 340

N

NCA#, 312
NMI, 349, 399, 442, 473
no wait state, 343
non-specific EOI, 379
NOWS#, 138, 343
NPRD#, 465
NPS1#, 464
NPS2, 465
NPWR#, 465

numeric coprocessor, 459
numeric coprocessor control bits in the MSW register, 460
numeric coprocessor installed, description of operation, 460
numeric coprocessor, Weitek, 467
numeric processor read, 465
numeric processor select 1, 464
numeric processor select 2, 465
numeric processor write, 465

O

OCW1, 397
option ROM, 43
OSC, oscillator signal, 350
OUT instruction, 32
overrun condition, 369

P

page fault, 403
parallel port 1, 387
parallel port 2, 386
parity, 256
parity generator/checker, 256
PASS A20, 451
PEACK#, 461, 464
PEREQ, 461, 464
PIC, 366
POST, 322
POST/BIOS ROM, 42
POWERGOOD, 113
power-on reset, 447
power-on restart address, 12, 116
power-on self-test, or POST, 116, 322
page-mode DRAM, 259
principle of locality, 277
prefetcher, 54
processor extension acknowledge (PEACK#), 461, 464
processor extension not available exception, 403
processor extension request (PEREQ), 461, 464

processor extension segment overrun exception, 403
program, 12
PROM, 317
protect enable, 62
protected mode, 62, 98, 162
protected mode to real mode, 447

R

RAM parity check, 400
RAM refresh logic, 433
RAM, dual-ported, 312
RAS#, 236, 242
RAS/CAS delay, 236, 242
RAS/CAS generation logic, 240
RAS-only refresh, 245
read, 12
read bus cycle, 141
read-only memory, 315
ready line, 141
ready timer, 135
READY#, 141
READYO#, ready output, 464
real mode, 369
real-time clock bytes, 443
real-time clock chip, 470
real-time clock function, 443
refresh, 474
refresh and ISA bus masters, 439
refresh interval, 249
refresh logic, 245, 433
refresh request, 400
refresh request signal, 471
refresh row counter, 245
refresh timer, 471
refresh timer, reprogramming, 471
REFRESH#, 247, 350
requestor privilege level (RPL), 102, 105
RESDRV (reset drive) signal, 344
reset code byte, 447
reset register values, 115
RESET signal, 113
ROM address decoder, 42
ROM checksum, 322
ROM chip-enable, 321

ISA System Architecture

ROM chip-select, 321
ROM memory\ : theory of operation, 315
ROM output enable, 321
ROM programmer, 316
ROM scan, 330
ROM, device ROM signature, 331
ROM, device ROMs, 330
ROM, electrically erasable programmable read-only memory (EEPROM), 319
ROM, erasable programmable read-only memory (EPROM), 318
ROM, erasing contents of EPROM, 319
ROM, flash EEPROM, 320
ROM, fusible-link PROM, 317
ROM, masked ROM (MROM), 317
ROM, system board, 322
row address, 236, 238
RPL, 102, 105
read from even addressed location in 8-bit device, 128
refresh timer, 245
row address decoder, 237
row address strobe (RAS#), 242
row latch, 236
RTC 24-hour mode, 444
RTC alarm interrupt enable, 444
RTC alarm interrupt flag, 445
RTC BIOS control, 445
RTC chip, 441
RTC chip's status registers A through D, 444
RTC date mode, 444
RTC daylight-savings enable, 444
RTC interrupt request flag, 445
RTC periodic interrupt enable, 444
RTC periodic interrupt flag, 445
RTC rate selection bits, 444
RTC set bit, 444
RTC square-wave enable, 444
RTC status register A, 444
RTC status register B, 444
RTC status register C, 445
RTC update in progress, 444
RTC update-ended interrupt enable, 444
RTC update-ended interrupt flag, 445

S

SA19:SA0, 338
SBHE#, 338
scan code, 450
SCRAM, 267
SD15:SD8, 339
SD7:SD0, 339
security lock latch, 400
segment descriptor, 101
segment not present, 403
segment wrap-around, 452
segment wraparound, 79
self refresh, 248
self-modifying code, 16
semaphore, 93
serial port 1, 386
serial port 2, 386
set interrupt enable instruction, 368
setup time, 149
shadow RAM, 324
shared resource, 92
shutdown, 108, 151, 185, 230
single step, or trap, interrupt, 403
slave ID, 387, 398
slave ID register, 383, 387
slave interrupt controller, 383, 387
slowdown timer, 474
SMRDC#, 340
SMWTC#, 340
snoop hit, 295
software interrupt instruction, 403
software interrupts, 402
software-enforced coherency, 298
speaker data timer, 471
split cache, 290
SRAM, 271
second level cache, 289
snarfing, 295
snooping, 295
spatial locality, 277
static RAM (SRAM), 271
system address (SA) bus, 121
stack, 374

stack segment overrun or not present exception, 403
stale data, 279
status enable, 464
status port, 37
STEN, 464
STI (Set Interrupt Enable) instruction, 368
synchronous DRAM, 268
system address (SA) bus, 338
system bus high enable, 338
system configuration information, 446, 450
system control port A, 400, 473
system control port B, 400, 471, 473
system data (SD) bus, 126
system data bus, lower path, 339
system data bus, upper path, 339
system memory read command, 340
system memory write command, 340
system timer, 386, 469
system timer frequency, 469
system timer interrupt, 472

T

T1, 140
T2, 140
Tc, 140
TC (DMA terminal count), 345, 413
time delay device, 239
timer 0, 469
timer 0 frequency, 469

timer 1, 471
timer 2, 471
timer 2 (speaker timer) output, 400
transfer complete, 345
trap interrupt, 403
Ts, 140
tag RAM, 281
temporal locality, 277
two-way set-associative cache, 304

U

UMB, 327
unified cache, 290
upper memory blocks, 327

W

watchdog timer, 472
watchdog timer status, 400
Weitek interpretation of the memory address, 468
Weitek numeric coprocessor, 467
Weitek-specific driver routine, 467
wraparound, 79
write, 12
write bus cycle, 146
write policy, 292
write-back cache, 294
write-through cache, 293