July 15, 1997 3:24

Addendums and other updates for this manual can be obtained from
Cyrix Web site: www.cyrix.com.

©1997 Copyright Cyrix Corporation. All rights reserved.
Printed in the United States of America

Trademark Acknowledgments:

Cyrix is a registered trademark of Cyrix Corporation.
6x86 and 6x86MX are trademarks of Cyrix Corporation. MMX is a trademark of Intel Corporation.
All other brand or product names are trademarks of their respective companies.

Order Number: 94329-00
Cyrix Corporation
2703 North Central Expressway
Richardson, Texas 75080-2010
United States of America

Cyrix Corporation (Cyrix) reserves the right to make changes in the devices or specifications described herein without notice. Before design-in or order placement, customers are advised to verify that the information is current on which orders or design activities are based. Cyrix warrants its products to conform to current specifications in accordance with Cyrix' standard warranty. Testing is performed to the extent necessary as determined by Cyrix to support this warranty. Unless explicitly specified by customer order requirements, and agreed to in writing by Cyrix, not all device characteristics are necessarily tested. Cyrix assumes no liability, unless specifically agreed to in writing, for customers' product design or infringement of patents or copyrights of third parties arising from the use of Cyrix devices. No license, either express or implied, to Cyrix patents, copyrights, or other intellectual property rights pertaining to any machine or combination of Cyrix devices is hereby granted. Cyrix products are not intended for use in any medical, life saving, or life sustaining system. Information in this document is subject to change without notice.

6x86MX™ PROCESSOR

Enhanced Sixth-Generation CPU
Compatible with MMX™ Technology

Introduction

♦ Enhanced Sixth-Generation Architecture

- Performance Rating: PR166, PR200, PR233, PR266 and higher
- 64K 4-Way Unified Write-Back Cache
- 2 Level TLB (16 Entry L1, 384 Entry L2)
- Branch Prediction with a 512-entry BTB
- Enhanced Memory Management Unit
- Scratchpad RAM in Unified Cache
- Optimized for both 16- and 32-Bit Code
- High Performance 80-Bit FPU

♦ X86 Instruction Set Includes MMX™ Instructions

- Compatible with MMX™ Technology
- Runs Windows® 95, Windows 3.x, Windows NT, DOS, UNIX®, OS/2®, Solaris®, and others

♦ Other Features

- Socket 7 Pinout Compatible
- 2.9 V Core, 3.3 V I/O
- Flexible Core/Bus Clock Ratios (2x, 2.5x, 3x, 3.5x)
- Leverages Existing Socket Infrastructure

The Cyrix 6x86MX™ processor offers significant enhancements over the 6x86 CPU. The 6x86MX design quadruples the cache size, triples the TLB size, increases the frequency scalability to 200 MHz and beyond, and is compatible with MMX™ technology. The 6x86MX CPU contains a scratchpad RAM feature, supports performance monitoring, and allows caching of both SMI code and SMI data. It delivers high 16- and 32-bit performance while running Windows 95, Windows NT, OS/2, DOS, UNIX, and other operating systems.

The 6x86MX processor achieves top performance through the use of two optimized superpipelined integer units, an on-chip floating point unit, and a 64 KByte unified write-back cache. The superpipelined architecture reduces timing constraints and increase frequency scalability. Advanced architectural techniques include register renaming, out-of-order completion, data dependency removal, branch prediction and speculative execution. Many data dependencies and resource conflicts have been eliminated, allowing higher performance for both 16- and 32-bit software.

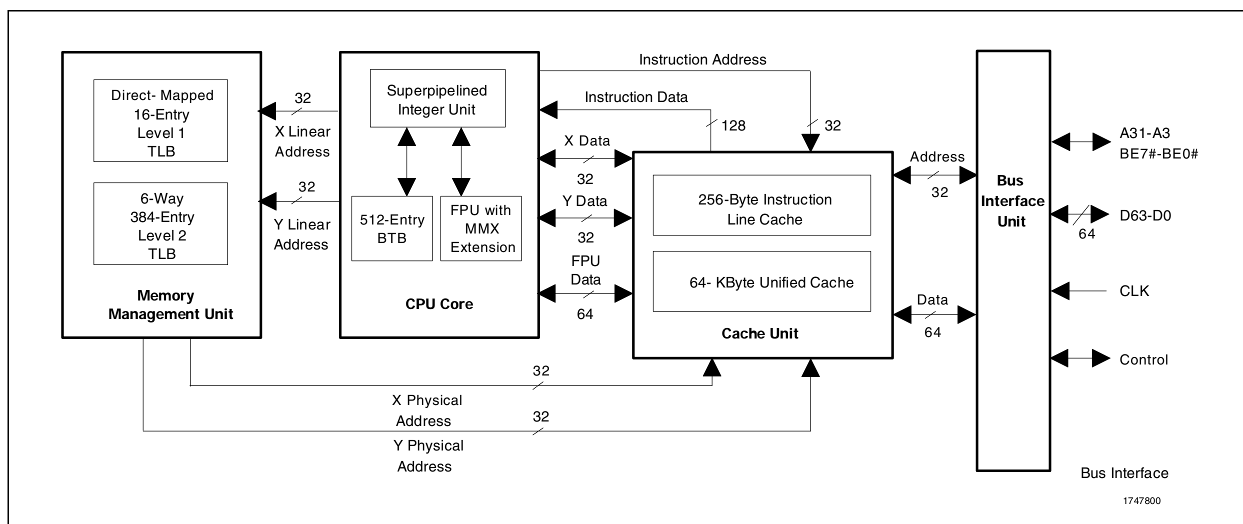

6x86MX™ PROCESSOR

Enhanced Sixth-Generation CPU
Compatible with MMX™ Technology

Cyrix® Advancing the Standards

TABLE OF CONTENTS

1. ARCHITECTURE OVERVIEW
1.1 Major Differences Between the 6x86MX and 6x86 Processors 1-2
1.2 Major Functional Blocks 1-3
1.3 Integer Unit 1-4
1.4 Cache Units 1-14
1.5 Memory Management Unit 1-16
1.6 Floating Point Unit 1-17
1.7 Bus Interface Unit 1-17

2. PROGRAMMING INTERFACE
2.1 Processor Initialization 2-1
2.2 Instruction Set Overview 2-3
2.3 Register Sets 2-4
2.4 System Register Set 2-11
2.5 Model Specific Registers 2-38
2.6 Time Stamp Counter 2-38
2.7 Performance Monitoring 2-38
2.8 Performance Monitoring Counters 1 and 2 2-39
2.9 Debug Registers 2-44
2.10 Test Registers 2-46
2.11 Address Space 2-47
2.12 Memory Addressing Methods 2-48
2.13 Memory Caches 2-57
2.14 Interrupt and Exceptions 2-62
2.15 System Management Mode 2-70
2.16 Shutdown and Halt 2-80
2.17 Protection 2-82
2.18 Virtual 8086 Mode 2-85
2.19 Floating Point Unit Operations 2-86
2.20 MMX Operations 2-89

3. BUS INTERFACE
3.1 Signal Description Table 3-2
3.2 Signal Descriptions 3-7
3.3 Functional Timing 3-23

4. ELECTRICAL SPECIFICATIONS
4.1 Electrical Connections 4-1
4.2 Absolute Maximum Ratings 4-2
4.3 Recommended Operating Conditions 4-3
4.4 DC Characteristics 4-4
4.5 AC Characteristics 4-6

5. MECHANICAL SPECIFICATIONS
5.1 296-Pin SPGA Package 5-1
5.2 Thermal Characteristics 5-7

6. INSTRUCTION SET
6.1 Instruction Set Summary 6-1
6.2 General Instruction Fields 6-2
6.3 CPUID Instruction 6-11
6.4 Instruction Set Tables 6-12
6.5 FPU Instruction Clock Counts 6-30
6.6 6x86MX Processor MMX Instruction Clock Counts 6-37

Appendix, Index and Distributors

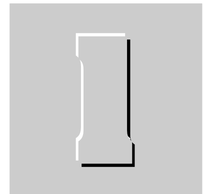
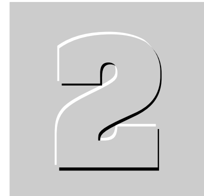

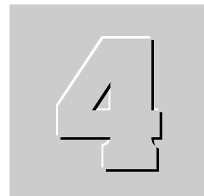

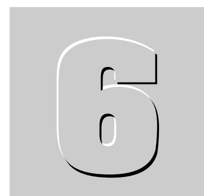

LIST OF FIGURES

Figure Name **Page Number**

Figure 1-1. Integer Unit . . . 1-4
Figure 1-2. Cache Unit Operations . . . 1-15
Figure 1-3. Paging Mechanism within the Memory Management Unit . . . 1-16
Figure 2-1. Application Register Set . . . 2-5
Figure 2-2. General Purpose Registers . . . 2-6
Figure 2-3. Segment Selector in Protected Mode . . . 2-7
Figure 2-4. EFLAGS Register . . . 2-9
Figure 2-5. System Register Set . . . 2-12
Figure 2-6. Control Registers . . . 2-13
Figure 2-7. Descriptor Table Registers . . . 2-16
Figure 2-8. Application and System Segment Descriptors . . . 2-17
Figure 2-9. Gate Descriptor . . . 2-20
Figure 2-10. Task Register. . . . 2-21
Figure 2-11. 32-Bit Task State Segment (TSS) Table . . . 2-22
Figure 2-12. 16-Bit Task State Segment (TSS) Table . . . 2-23
Figure 2-13. 6x86MX Configuration Control Register 0 (CCR0) . . . 2-26
Figure 2-14. 6x86MX Configuration Control Register 1 (CCR1) . . . 2-27
Figure 2-15. 6x86MX Configuration Control Register 2 (CCR2) . . . 2-28
Figure 2-16. 6x86MX Configuration Control Register 3 (CCR3) . . . 2-29
Figure 2-17. 6x86MX Configuration Control Register 4 (CCR4) . . . 2-30
Figure 2-18. 6x86MX Configuration Control Register 5 (CCR5) . . . 2-31
Figure 2-19. 6x86MX Configuration Control Register 6 (CCR6) . . . 2-32
Figure 2-20 Address Region Registers (ARR0 - ARR7) . . . 2-33
Figure 2-21. Region Control Registers (RCR0 -RCR7) . . . 2-36
Figure 2-22. Counter Event Control Register . . . 2-40
Figure 2-23. Debug Registers . . . 2-44
Figure 2-24. Memory and I/O Address Spaces . . . 2-47
Figure 2-25. Offset Address Calculation. . . . 2-49
Figure 2-26. Real Mode Address Calculation . . . 2-50
Figure 2-27. Protected Mode Address Calculation . . . 2-51
Figure 2-28. Selector Mechanism . . . 2-51
Figure 2-29. Paging Mechanisms . . . 2-53

LIST OF FIGURES (Continued)

Figure Name **Page Number**

Figure 2-30. Directory and Page Table Entry (DTE and PTE) Format 2-53
Figure 2-31. TLB Test Registers 2-55
Figure 2-32. Unified Cache 2-58
Figure 2-33. Cache Test Registers 2-59
Figure 2-34. Error Code Format 2-69
Figure 2-35. SMI Execution Flow Diagram 2-70
Figure 2-36. System Management Memory Address Space 2-71
Figure 2-37. SMM Memory Space Header. 2-72
Figure 2-38. SMHR Register 2-74
Figure 2-39. SMM and Suspend Mode State Diagram 2-81
Figure 2-40. FPU Tag Word Register 2-87
Figure 2-41. FPU Status Register 2-87
Figure 2-42. FPU Mode Control Register 2-88
Figure 3-1. 6x86MX Functional Signal Groupings 3-1
Figure 3-2. RESET Timing 3-23
Figure 3-3. 6x86MX CPU Bus State Diagram 3-25
Figure 3-4. Non-Pipelined Single Transfer Read Cycles 3-28
Figure 3-5. Non-Pipelined Single Transfer Write Cycles 3-29
Figure 3-6. Non-Pipelined Burst Read Cycles 3-31
Figure 3-7. Burst Cycle with Wait States 3-32
Figure 3-8. “1+4” Burst Read Cycle 3-33
Figure 3-9. Non-Pipelined Burst Write Cycles 3-35
Figure 3-10. Pipelined Single Transfer Read Cycles 3-36
Figure 3-11. Pipelined Burst Read Cycles 3-37
Figure 3-12. Read Cycle Followed by Pipelined Write Cycle 3-38
Figure 3-13. Interrupt Acknowledge Cycles 3-39
Figure 3-14. SMIACT# Timing 3-40
Figure 3-15. SMM I/O Trap Timing 3-41
Figure 3-16. Cache Invalidation Using FLUSH# 3-42
Figure 3-17. External Write Buffer Empty (EWBE#) Timing 3-43
Figure 3-18. Requesting Hold from an Idle Bus 3-44
Figure 3-19. Requesting Hold During a Non-Pipelined Bus Cycle. 3-45

LIST OF FIGURES (Continued)

Figure Name		Page Number
Figure 3-20.	Requesting Hold During a Pipelined Bus Cycle	3-46
Figure 3-21.	Back-Off Timing	3-47
Figure 3-22.	HOLD Inquiry Cycle that Hits on a Modified Line	3-49
Figure 3-23.	BOFF# Inquiry Cycle that Hits on a Modified Line	3-50
Figure 3-24.	AHOLD Inquiry Cycle that Hits on a Modified Line	3-51
Figure 3-25.	AHOLD Inquiry Cycle During a Line Fill	3-52
Figure 3-26.	APCHK# Timing	3-53
Figure 3-27.	Hold Inquiry that Hits on a Modified Data Line	3-54
Figure 3-28.	BOFF# Inquiry Cycle that Hits on a Modified Data Line	3-56
Figure 3-29.	Hold Inquiry that Misses the Cache While in SMM Mode	3-57
Figure 3-30.	AHOLD Inquiry Cycle During a Line Fill from SMM Memory	3-58
Figure 3-31.	SUSP# Initiated Suspend Mode	3-60
Figure 3-32.	HALT Initiated Suspend Mode	3-61
Figure 3-33.	Stopping CLK During Suspend Mode	3-62
Figure 4-1.	Drive Level and Measurement Points for Switching Characteristics	4-7
Figure 4-2.	CLK Timing and Measurement Points	4-8
Figure 4-3.	Output Valid Delay Timing	4-9
Figure 4-4.	Output Float Delay Timing	4-10
Figure 4-5.	Input Setup and Hold Timing	4-12
Figure 4-6.	TCK Timing Measurement Points	4-13
Figure 4-7.	JTAG Test Timings	4-14
Figure 4-8.	Test Reset Timing	4-14
Figure 5-1.	296-Pin SPGA Package Pin Assignments (Top View)	5-1
Figure 5-1.	296-Pin SPGA Package Pin Assignments (Bottom View)	5-2
Figure 5-2.	296-Pin SPGA Package	5-5
Figure 5-3	Typical HeatSink/Fan	5-8
Figure 6-1	Instruction Set Format	6-1

LIST OF TABLES

Table Name **Page Number**

Table 1-1. Register Renaming with WAR Dependency 1-2
Table 1-2. Register Renaming with WAR Dependency 1-7
Table 1-2. Register Renaming with WAW Dependency 1-8
Table 1-3. Example of Operand Forwarding 1-10
Table 1-4. Result Forwarding Example 1-11
Table 1-5. Example of Data Bypassing 1-12
Table 2-1. Initialized Register Controls 2- 2
Table 2-2. Segment Register Selection Rules 2-8
Table 2-3. EFLAGS Bit Definitions 2-10
Table 2-4. CR0 Bit Definitions 2-14
Table 2-5. Effects of Various Combinations of EM, TS and MP Bits 2-14
Table 2-6. CR4 Bit Definitions 2-15
Table 2-7. Segment Descriptor Bit Definitions 2-18
Table 2-8. TYPE Field Definitions with DT = 0 2-18
Table 2-9. TYPE Field Definitions with DT = 1 2-19
Table 2-10. Gate Descriptor Bit Definitions 2-20
Table 2-11. 6x86MX Configuration Registers 2-25
Table 2-12. CCR0 Bit Definitions 2-26
Table 2-13. CCR1 Bit Definitions 2-27
Table 2-14. CCR2 Bit Definitions 2-28
Table 2-15. CCR3 Bit Definitions 2-29
Table 2-16. CCR4 Bit Definitions 2-30
Table 2-17. CCR5 Bit Definitions 2-31
Table 2-18. CCR6 Bit Definitions 2-32
Table 2-19. ARR0 - ARR7 Registers Index Assignments 2-34
Table 2-20. Bit Definitions for SIZE Field 2-34
Table 2-21. RCR0 -RCR7 Bit Definitions 2-36
Table 2-22. Machine Specific Registers 2-38
Table 2-23. Counter Event Control Register Bit Definitions 2-40
Table 2-24. Event Type Register 2-41
Table 2-25. DR6 and DR7 Debug Register Field Definitions 2-45
Table 2-26. Memory Addressing Modes 2-49

LIST OF TABLES (Continued)

Table Name **Page Number**

Table 2-27. Directory and Page Table Entry (DTE and PTE) Bit Definitions 2-54
Table 2-28. CMD Field 2-54
Table 2-29. TLB Test Register Bit Definitions 2-56
Table 2-30. Cache Test Register Bit Definitions 2-59
Table 2-31. Cache Locking Operations 2-61
Table 2-32. Interrupt Vector Assignments 2-65
Table 2-33. Interrupt and Exception Priorities. 2-67
Table 2-34. Exception Changes in Real Mode 2-68
Table 2-35. Error Code Bit Definitions. 2-69
Table 2-36. SMM Memory Space Header 2-73
Table 2-37. SMHR Register 2-74
Table 2-38. SMM Instruction Set 2-75
Table 2-39. Requirements for Recognizing SMI# and SMINT 2-76
Table 2-40. Descriptor Types Used for Control Transfer. 2-84
Table 2-41. FPU Status Register Bit Definitions 2-87
Table 2-42. FPU Mode Control Register Bit Definitions 2-88
Table 2-43. Saturation Limits 2-90
Table 3-1. 6x86MX CPU Signals Sorted by Signal Name 3-2
Table 3-2. Clock Control 3-7
Table 3-3. Pins Sampled During RESET 3-7
Table 3-4. Signal States During RESET 3-8
Table 3-5. Byte Enable Signal to Data Bus Byte Correlation. 3-9
Table 3-6. Parity Bit to Data Byte Correlation. 3-10
Table 3-7. Bus Cycle Types. 3-12
Table 3-8. Effects of WB/WT# on Cache Line State. 3-16
Table 3-9. Signal States During Bus Hold. 3-17
Table 3-10. Signal States During Suspend Mode. 3-21
Table 3-11. 6x86MX CPU Bus States 3-24
Table 3-12. Bus State Transitions 3-26
Table 3-13. “1+4” Burst Address Sequences. 3-33
Table 3-14. Linear Burst Address Sequences. 3-34
Table 4-1. Pins Connected to Internal Pull-Up and Pull-Down Resistors 4-1

LIST OF TABLES (Continued)

Table Name **Page Number**

Table 4-2. Absolute Maximum Ratings 4-2
Table 4-3. Recommended Operating Conditions 4-3
Table 4-4. DC Characteristics (at Recommended Operating Conditions) 1 of 2 4-4
Table 4-5. DC Characteristics (at Recommended Operating Conditions) 2 of 2 4-5
Table 4-6. Power Dissipation 4-5
Table 4-7. Drive Level and Measurement Points for Switching Characteristics 4-7
Table 4-8. Clock Specifications 4-8
Table 4-9. Output Valid Delays, C_L = 50 pF, T_{case} = 0°C to 70 °C 4-9
Table 4-10. Output Float Delays, C_L = 50 pF, T_{case} = 0°C to 70 °C. 4-10
Table 4-11. Input Setup Times T_{case} = 0°C to 70 °C 4-11
Table 4-12. Input Hold Times T_{case} = 0°C to 70 °C. 4-11
Table 4-13. JTAG AC Specifications 4-13
Table 5-1. 296-Pin SPGA Package Signal Names Sorted by Pin Number 5-3
Table 5-2. 296-Pin SPGA Package Pin Numbers Sorted by Signal Name 5-4
Table 5-3. 296-Pin SPGA Package Dimensions 5-6
Table 5-4. Required θ_{CA} to Maintain 70°C Case Temperature. 5-7
Table 5-5. Heatsink/Fan Dimensions 5-8
Table 6-1. Instruction Set Format 6-1
Table 6-2. Instruction Fields 6-2
Table 6-3. Instruction Prefix Summary 6-3
Table 6-4. w Field Encoding 6-4
Table 6-5. d Field Encoding. 6-4
Table 6-6. s Field Encoding 6-5
Table 6-7. eee Field Encoding. 6-5
Table 6-8. mod r/m Field Encoding 6-6
Table 6-9. mod r/m Field Encoding Dependent on w Field 6-7
Table 6-10. reg Field 6-7
Table 6-11. sreg3 Field Encoding. 6-8
Table 6-12. sreg2 Field Encoding. 6-8
Table 6-13. ss Field Encoding 6-9
Table 6-14. index Field Encoding 6-9
Table 6-15. mod base Field Encoding 6-10

Table 6-16. CPUID Data Returned When EAX = 0 6-11
Table 6-17. CPUID Data Returned When EAX = 1 6-11
Table 6-18. CPU Clock Count Abbreviations 6-13
Table 6-19. Flag Abbreviations 6-13
Table 6-20. Action of Instruction on Flag 6-13
Table 6-21. 6x86MX CPU Instruction Set Clock Count Summary. 6-14
Table 6-22. FPU Clock Count Table Abbreviations 6-30
Table 6-23. 6x86MX FPU Instruction Set Summary. 6-31
Table 6-24. MMX Clock Count Table Abbreviations 6-37
Table 6-25. MMX Instruction Set Summary 6-38

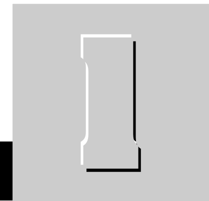

1. ARCHITECTURE OVERVIEW

The Cyrix 6x86MX™ processor is an enhanced 6x86 processor, that can process 57 new multimedia instructions compatible with MMX™ technology. The processor also operates at a higher frequency, contains a enlarged cache, a two-level TLB, and a improved branch target cache.

The 6x86MX processor is based on the proven 6x86 core that is superscalar in that it contains two separate pipelines that allow multiple instructions to be processed at the same time. The use of advanced processing technology and superpipelining (increased number of pipeline stages) allow the 6x86MX CPU to achieve high clocks rates.

Through the use of unique architectural features, the 6x86MX processor eliminates many data dependencies and resource conflicts, resulting in optimal performance for both 16-bit and 32-bit x86 software.

For maximum performance, the 6x86MX CPU contains two caches, a large unified 64 KByte 4-way set associative write-back cache and a small high-speed instruction line cache.

To provide support for multimedia operations, the cache can be turned into a scratchpad RAM memory on a line by line basis. The cache area set aside as scratchpad memory acts as a private memory for the CPU and does not participate in cache operations.

Within the 6x86MX processor there are two TLBs, the main L1 TLB and the larger L2 TLB. The direct-mapped L1 TLB has 16 entries and the 6-way associative L2 TLB has 384 entries.

The on-chip FPU has been enhanced to process MMX™ instructions as well as the floating point instructions. Both types of instructions execute in parallel with integer instruction processing. To facilitate FPU operations, the FPU features a 64-bit data interface, a four-deep instruction queue and a six-deep store queue.

The CPU operates using a split rail power design. The core runs on a 2.9 volt power supply, to minimize power consumption. External signal level compatibility is maintained by using a 3.3 volt power supply for the I/O interface.

For mobile systems and other power sensitive applications, the 6x86MX processor incorporates low power suspend mode, stop clock capability, and system management mode (SMM).

1.1 Major Differences Between the 6x86MX and the 6x86 Processors

The major differences between the 6x86MX and the 6x86 processors are summarized in Table 1-1.

Table 1-1. The 6x86MX Processor Versus the 6x86 Processor

<table>
<tr><th>FEATURE</th><th>6x86MX Processor</th><th colspan="2">6x86 Processor</th></tr>
<tr><td>Pinout</td><td>P55C</td><td colspan="2">P54C</td></tr>
<tr><td>Supply Voltage
Core
I/O</td><td>
2.9 V
3.3 V</td><td>6x86:
3.3 or 3.52 V
3.3 V</td><td>6x86L:
2.8 V
3.3 V</td></tr>
<tr><td>CPU Primary Cache</td><td>64 KBytes</td><td colspan="2">16 KBytes</td></tr>
<tr><td>Translation Lookaside Buffer (TLB)</td><td>L1: 16 entry
L2: 384 entry</td><td colspan="2">L1: 128 entry
Victim TLB: 8 entry</td></tr>
<tr><td>Branch Prediction</td><td>512 entry branch target cache
1024 entry branch history table</td><td colspan="2">256 entry branch target cache
512 entry branch history table</td></tr>
<tr><td>MMX</td><td>Yes</td><td colspan="2">No</td></tr>
<tr><td>Performance Monitor including Time Stamp Counter and Model Specific Registers</td><td>Yes</td><td colspan="2">No</td></tr>
<tr><td>Scratchpad RAM in Primary Cache</td><td>Yes</td><td colspan="2">No</td></tr>
<tr><td>Cacheable SMI Code/Data</td><td>Yes</td><td colspan="2">No</td></tr>
<tr><td>Clock Modes</td><td>2x, 2.5x, 3x, 3.5x</td><td colspan="2">2x, 3x</td></tr>
</table>

1.2 Major Functional Blocks

The 6x86MX processor consists of four major functional blocks, as shown in the overall block diagram on the first page of this manual:

- Memory Management Unit
- CPU Core
- Cache Unit
- Bus Interface Unit

The CPU contains the superpipelined integer unit, the BTB (Branch Target Buffer) unit and the FPU (Floating Point Unit).

The BIU (Bus Interface Unit) provides the interface between the external system board and the processor's internal execution units. During a memory cycle, a memory location is selected through the address lines (A31-A3 and BE7# -BE0#). Data is passed from or to memory through the data lines (D63-D0).

Each instruction is read into 256-Byte Instruction Line Cache. The Cache Unit stores the most recently used data and instructions to allow fast access to the information by the Integer Unit and FPU.

The CPU core requests instructions from the Cache Unit. The received integer instructions are decoded by either the X or Y processing pipelines within the superpipelined integer unit. If the instruction is a MMX or FPU instruction it is passed to the floating point unit for processing.

As required data is fetched from the 64-KByte unified cache. If the data is not in the cache it is accessed via the bus interface unit from main memory.

The Memory Management Unit calculates physical addresses including addresses based on paging.

Physical addresses are calculated by the Memory Management Unit and passed to the Cache Unit and the Bus Interface Unit (BIU).

1.3 Integer Unit

The Integer Unit (Figure 1-1) provides parallel instruction execution using two seven-stage integer pipelines. Each of the two pipelines, X and Y, can process several instructions simultaneously.

The Integer Unit consists of the following pipeline stages:

- Instruction Fetch (IF)
- Instruction Decode 1 (ID1)
- Instruction Decode 2 (ID2)
- Address Calculation 1 (AC1)
- Address Calculation 2 (AC2)
- Execute (EX)
- Write-Back (WB)

The instruction decode and address calculation functions are both divided into superpipelined stages.

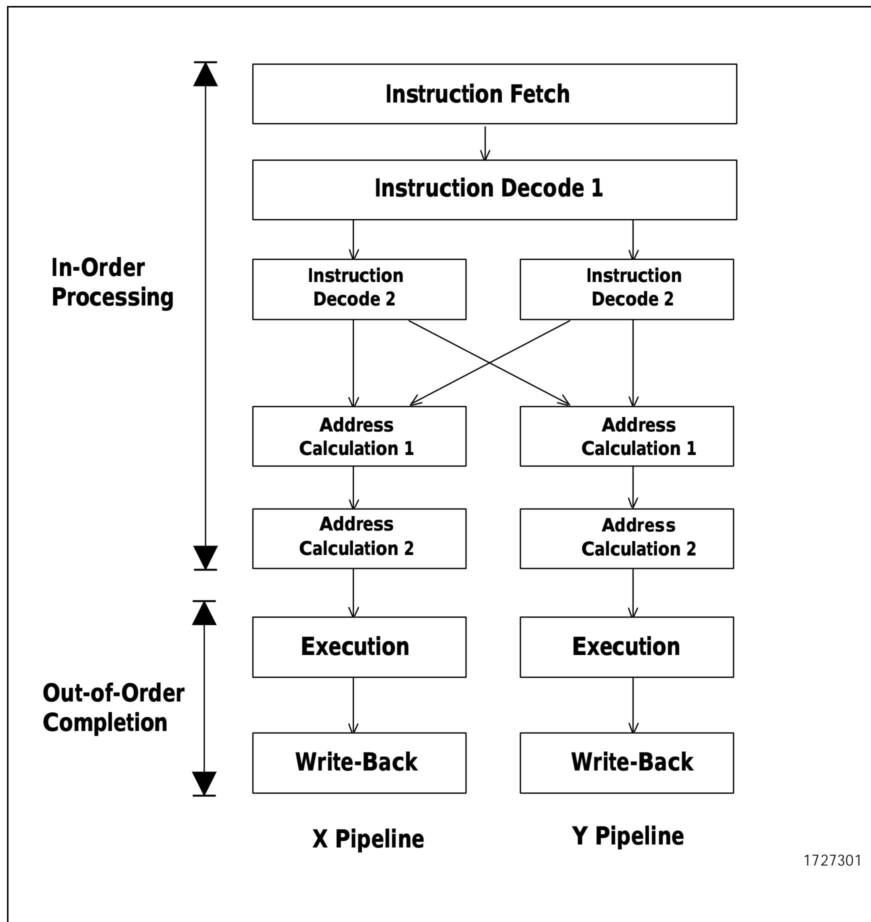

Figure 1-1. Integer Unit

1.3.1 Pipeline Stages

The **Instruction Fetch** (IF) stage, shared by both the X and Y pipelines, fetches 16 bytes of code from the cache unit in a single clock cycle. Within this section, the code stream is checked for any branch instructions that could affect normal program sequencing.

If an unconditional or conditional branch is detected, branch prediction logic within the IF stage generates a predicted target address for the instruction. The IF stage then begins fetching instructions at the predicted address.

The superpipelined **Instruction Decode** function contains the ID1 and ID2 stages. ID1, shared by both pipelines, evaluates the code stream provided by the IF stage and determines the number of bytes in each instruction. Up to two instructions per clock are delivered to the ID2 stages, one in each pipeline.

The ID2 stages decode instructions and send the decoded instructions to either the X or Y pipeline for execution. The particular pipeline is chosen, based on which instructions are already in each pipeline and how fast they are expected to flow through the remaining pipeline stages.

The **Address Calculation** function contains two stages, AC1 and AC2. If the instruction refers to a memory operand, the AC1 calculates a linear memory address for the instruction.

The AC2 stage performs any required memory management functions, cache accesses, and register file accesses. If a floating point instruction is detected by AC2, the instruction is sent to the FPU for processing.

The **Execute** (EX) stage executes instructions using the operands provided by the address calculation stage.

The **Write-Back** (WB) stage is the last IU stage. The WB stage stores execution results either to a register file within the IU or to a write buffer in the cache control unit.

1.3.2 Out-of-Order Processing

If an instruction executes faster than the previous instruction in the other pipeline, the instructions may complete out of order. All instructions are processed in order, up to the EX stage. While in the EX and WB stages, instructions may be completed out of order.

If there is a data dependency between two instructions, the necessary hardware interlocks are enforced to ensure correct program execution. Even though instructions may complete out of order, exceptions and writes resulting from the instructions are always issued in program order.

1.3.3 Pipeline Selection

In most cases, instructions are processed in either pipeline and without pairing constraints on the instructions. However, certain instructions are processed only in the X pipeline:

- Branch instructions
- Floating point instructions
- Exclusive instructions

Branch and floating point instructions may be paired with a second instruction in the Y pipeline.

Exclusive Instructions cannot be paired with instructions in the Y pipeline. These instructions typically require multiple memory accesses. Although exclusive instructions may not be paired, hardware from both pipelines is used to accelerate instruction completion. Listed below are the 6x86MX CPU exclusive instruction types:

- Protected mode segment loads
- Special register accesses (Control, Debug, and Test Registers)
- String instructions
- Multiply and divide
- I/O port accesses
- Push all (PUSHA) and pop all (POPA)
- Intersegment jumps, calls, and returns

1.3.4 Data Dependency Solutions

When two instructions that are executing in parallel require access to the same data or register, one of the following types of data dependencies may occur:

- Read-After-Write (RAW)
- Write-After-Read (WAR)
- Write-After-Write (WAW)

Data dependencies typically force serialized execution of instructions. However, the 6x86MX CPU implements three mechanisms that allow parallel execution of instructions containing data dependencies:

- Register Renaming
- Data Forwarding
- Data Bypassing

The following sections provide detailed examples of these mechanisms.

1.3.4.1 Register Renaming

The 6x86MX CPU contains 32 physical general purpose registers. Each of the 32 registers in the register file can be temporarily assigned as one of the general purpose registers defined by the x86 architecture (EAX, EBX, ECX, EDX, ESI, EDI, EBP, and ESP). For each register write operation a new physical register is selected to allow previous data to be retained temporarily. Register renaming effectively removes all WAW and WAR dependencies. The programmer does not have to consider register renaming as register renaming is completely transparent to both the operating system and application software.

Example #1 - Register Renaming Eliminates Write-After-Read (WAR) Dependency

A WAR dependency exists when the first in a pair of instructions reads a logical register, and the second instruction writes to the same logical register. This type of dependency is illustrated by the pair of instructions shown below:

X PIPE	Y PIPE
(1) MOV BX, AX	(2) ADD AX, CX
BX ← AX	AX ← AX + CX

Note: In this and the following examples the original instruction order is shown in parentheses.

In the absence of register renaming, the ADD instruction in the Y pipe would have to be stalled to allow the MOV instruction in the X pipe to read the AX register.

The 6x86MX CPU, however, avoids the Y pipe stall (Table 1-2). As each instruction executes, the results are placed in new physical registers to avoid the possibility of overwriting a logical register value and to allow the two instructions to complete in parallel (or out of order) rather than in sequence.

Table 1-2. Register Renaming with WAR Dependency

Instruction	Physical Register Contents					Action	
	Reg0	Reg1	Reg2	Reg3	Reg4	Pipe	
(Initial)	AX	BX	CX				
MOV BX, AX	AX		CX	BX		X	Reg3 ← Reg0
ADD AX, CX			CX	BX	AX	Y	Reg4 ← Reg0 + Reg2

Note: The representation of the MOV and ADD instructions in the final column of Table 1-2 are completely independent.

Example #2 - Register Renaming Eliminates Write-After-Write (WAW) Dependency

A WAW dependency occurs when two consecutive instructions perform writes to the same logical register. This type of dependency is illustrated by the pair of instructions shown below:

X PIPE	Y PIPE
(1) ADD AX, BX	(2) MOV AX, [mem]
AX ←AX + BX	AX ← [mem]

Without register renaming, the MOV instruction in the Y pipe would have to be stalled to guarantee that the ADD instruction in the X pipe would write its results to the AX register first.

The 6x86MX CPU uses register renaming and avoids the Y pipe stall. The contents of the AX and BX registers are placed in physical registers (Table 1-3). As each instruction executes, the results are placed in new physical registers to avoid the possibility of overwriting a logical register value and to allow the two instructions to complete in parallel (or out of order) rather than in sequence.

Table 1-3. Register Renaming with WAW Dependency

Instruction	Physical Register Contents				Action	
	Reg0	Reg1	Reg2	Reg3	Pipe	
(Initial)	AX	BX				
ADD AX, BX		BX	AX		X	Reg2 ← Reg0 + Reg1
MOV AX, [mem]		BX		AX	Y	Reg3 ← [mem]

Note: All subsequent reads of the logical register AX will refer to Reg 3, the result of the MOV instruction.

1.3.4.2 Data Forwarding

Register renaming alone cannot remove RAW dependencies. The 6x86MX CPU uses two types of data forwarding in conjunction with register renaming to eliminate RAW dependencies:

- Operand Forwarding
- Result Forwarding

Operand forwarding takes place when the first in a pair of instructions performs a move from register or memory, and the data that is read by the first instruction is required by the second instruction. The 6x86MX CPU performs the read operation and makes the data read available to both instructions simultaneously.

Result forwarding takes place when the first in a pair of instructions performs an operation (such as an ADD) and the result is required by the second instruction to perform a move to a register or memory. The 6x86MX CPU performs the required operation and stores the results of the operation to the destination of both instructions simultaneously.

Example #3 - Operand Forwarding Eliminates Read-After-Write (RAW) Dependency

A RAW dependency occurs when the first in a pair of instructions performs a write, and the second instruction reads the same register. This type of dependency is illustrated by the pair of instructions shown below in the X and Y pipelines:

X PIPE	Y PIPE
(1) MOV AX, [mem]	(2) ADD BX, AX
AX ← [mem]	BX ← AX + BX

The 6x86MX CPU uses operand forwarding and avoids a Y pipe stall (Table 1-4). Operand forwarding allows simultaneous execution of both instructions by first reading memory and then making the results available to both pipelines in parallel.

Table 1-4. Example of Operand Forwarding

Instruction	Physical Register Contents				Action	
	Reg0	Reg1	Reg2	Reg3	Pipe	
(Initial)	AX	BX				
MOV AX, [mem]		BX	AX		X	Reg2 ← [mem]
ADD BX, AX			AX	BX	Y	Reg3 ← [mem] + Reg1

Operand forwarding can only occur if the first instruction does not modify its source data. In other words, the instruction is a move type instruction (for example, MOV, POP, LEA). Operand forwarding occurs for both register and memory operands. The size of the first instruction destination and the second instruction source must match.

Example #4 - Result Forwarding Eliminates Read-After-Write (RAW) Dependency

In this example, a RAW dependency occurs when the first in a pair of instructions performs a write, and the second instruction reads the same register. This dependency is illustrated by the pair of instructions in the X and Y pipelines, as shown below:

X PIPE	Y PIPE
(1) ADD AX, BX	(2) MOV [mem], AX
AX ←AX + BX	[mem] ← AX

The 6x86MX CPU uses result forwarding and avoids a Y pipe stall (Table 1-5). Instead of transferring the contents of the AX register to memory, the result of the previous ADD instruction (Reg0 + Reg1) is written directly to memory, thereby saving a clock cycle.

Table 1-5. Result Forwarding Example

Instruction	Physical Register Contents			Action	
	Reg0	Reg1	Reg2	Pipe	
(Initial)	AX	BX			
ADD AX, BX		BX	AX	X	Reg2 ←Reg0 + Reg1
MOV [mem], AX		BX	AX	Y	[mem] ← Reg0 +Reg1

The second instruction must be a move instruction and the destination of the second instruction may be either a register or memory.

1.3.4.3 Data Bypassing

In addition to register renaming and data forwarding, the 6x86MX CPU implements a third data dependency-resolution technique called data bypassing. Data bypassing reduces the performance penalty of those memory data RAW dependencies that cannot be eliminated by data forwarding.

Data bypassing is implemented when the first in a pair of instructions writes to memory and the second instruction reads the same data from memory. The 6x86MX CPU retains the data from the first instruction and passes it to the second instruction, thereby eliminating a memory read cycle. Data bypassing only occurs for cacheable memory locations.

Example #1- Data Bypassing with Read-After-Write (RAW) Dependency

In this example, a RAW dependency occurs when the first in a pair of instructions performs a write to memory and the second instruction reads the same memory location. This dependency is illustrated by the pair of instructions in the X and Y pipelines as shown below:

X PIPE	Y PIPE
(1) ADD [mem], AX	(2) SUB BX, [mem]
[mem] ←[mem] + AX	BX ← BX - [mem]

The 6x86MX CPU uses data bypassing and stalls the Y pipe for only one clock by eliminating the Y pipe's memory read cycle (Table 1-6). Instead of reading memory in the Y pipe, the result of the previous instruction ([mem] + Reg0) is used to subtract from Reg1, thereby saving a memory access cycle.

Table 1-6. Example of Data Bypassing

Instruction	Physical Register Contents			Action	
	Reg0	Reg1	Reg2	Pipe	
(Initial)	AX	BX			
ADD [mem], AX	AX	BX		X	[mem] ← [mem] + Reg0
SUB BX, [mem]	AX		BX	Y	Reg2 ← Reg1 - {[mem] + Reg0}

1.3.5 Branch Control

Branch instructions occur on average every four to six instructions in x86-compatible programs. When the normal sequential flow of a program changes due to a branch instruction, the pipeline stages may stall while waiting for the CPU to calculate, retrieve, and decode the new instruction stream. The 6x86MX CPU minimizes the performance degradation and latency of branch instructions through the use of branch prediction and speculative execution.

1.3.5.1 Branch Prediction

The 6x86MX CPU uses a 512-entry, 4-way set associative Branch Target Buffer (BTB) to store branch target addresses. The 6x86MX CPU has 1024-entry branch history table. During the fetch stage, the instruction stream is checked for the presence of branch instructions. If an unconditional branch instruction is encountered, the 6x86MX CPU accesses the BTB to check for the branch instruction's target address. If the branch instruction's target address is found in the BTB, the 6x86MX CPU begins fetching at the target address specified by the BTB.

In case of conditional branches, the BTB also provides history information to indicate whether the branch is more likely to be taken or not taken. If the conditional branch instruction is found in the BTB, the 6x86MX CPU begins fetching instructions at the predicted target address. If the conditional branch misses in the BTB, the 6x86MX CPU predicts that the branch will not be taken, and instruction fetching continues with the next sequential instruction. The decision to fetch the taken or not taken target address is based on a four-state branch prediction algorithm.

Once fetched, a conditional branch instruction is first decoded and then dispatched to the X pipeline only. The conditional branch instruction proceeds through the X pipeline and is then resolved in either the EX stage or the WB stage. The conditional branch is resolved in the EX stage, if the instruction responsible for setting the condition codes is completed prior to the execution of the branch. If the instruction that sets the condition codes is executed in parallel with the branch, the conditional branch instruction is resolved in the WB stage.

Correctly predicted branch instructions execute in a single core clock. If resolution of a branch indicates that a misprediction has occurred, the 6x86MX CPU flushes the pipeline and starts fetching from the correct target address. The 6x86MX CPU prefetches both the predicted and the non-predicted path for each conditional branch, thereby eliminating the cache access cycle on a misprediction. If the branch is resolved in the EX stage, the resulting misprediction latency is four cycles. If the branch is resolved in the WB stage, the latency is five cycles.

Since the target address of return (RET) instructions is dynamic rather than static, the 6x86MX CPU caches target addresses for RET instructions in an eight-entry return stack rather than in the BTB. The return address is pushed on the return stack during a CALL instruction and popped during the corresponding RET instruction.

1.3.5.2 Speculative Execution

The 6x86MX CPU is capable of speculative execution following a floating point instruction or predicted branch. Speculative execution allows the pipelines to continuously execute instructions following a branch without stalling the pipelines waiting for branch resolution. The same mechanism is used to execute floating point instructions (see Section 1.6) in parallel with integer instructions.

The 6x86MX CPU is capable of up to four levels of speculation (i.e., combinations of four conditional branches and floating point operations). After generating the fetch address using branch prediction, the CPU checkpoints the machine state (registers, flags, and processor environment), increments the speculation level counter, and begins operating on the predicted instruction stream.

Once the branch instruction is resolved, the CPU decreases the speculation level. For a correctly predicted branch, the status of the checkpointed resources is cleared. For a branch misprediction, the 6x86MX processor generates the correct fetch address and uses the checkpointed values to restore the machine state in a single clock.

In order to maintain compatibility, writes that result from speculatively executed instructions are not permitted to update the cache or external memory until the appropriate branch is resolved. Speculative execution continues until one of the following conditions occurs:

1) A branch or floating point operation is decoded and the speculation level is already at four.
2) An exception or a fault occurs.
3) The write buffers are full.
4) An attempt is made to modify a non-checkpointed resource (i.e., segment registers, system flags).

1.4 Cache Units

The 6x86MX CPU employs two caches, the Unified Cache and the Instruction Line Cache (Figure 1-2, Page 1-15). The main cache is a 4-way set-associative 64-KByte unified cache. The unified cache provides a higher hit rate than using equal-sized separate data and instruction caches. While in Cyrix SMM mode both SMM code and data are cacheable.

The instruction line cache is a fully associative 256-byte cache. This cache avoids excessive conflicts between code and data accesses in the unified cache.

1.4.1 Unified Cache

The 64-KByte unified write-back cache functions as the primary data cache and as the secondary instruction cache. Configured as a four-way set-associative cache, the cache stores up to 64 KBytes of code and data in 2048 lines. The cache is dual-ported and allows any

two of the following operations to occur in parallel:

- Code fetch
- Data read (X pipe, Y pipeline or FPU)
- Data write (X pipe, Y pipeline or FPU)

The unified cache uses a pseudo-LRU replacement algorithm and can be configured to allocate new lines on read misses only or on read and write misses.

1.4.2 Instruction Line Cache

The fully associative 256-byte instruction line cache serves as the primary instruction cache. The instruction line cache is filled from the unified cache through the data bus. Fetches from the integer unit that hit in the instruction line cache do not access the unified cache. If an instruction line cache miss occurs, the instruction line data from the unified cache is transferred to the instruction line cache and the integer unit, simultaneously.

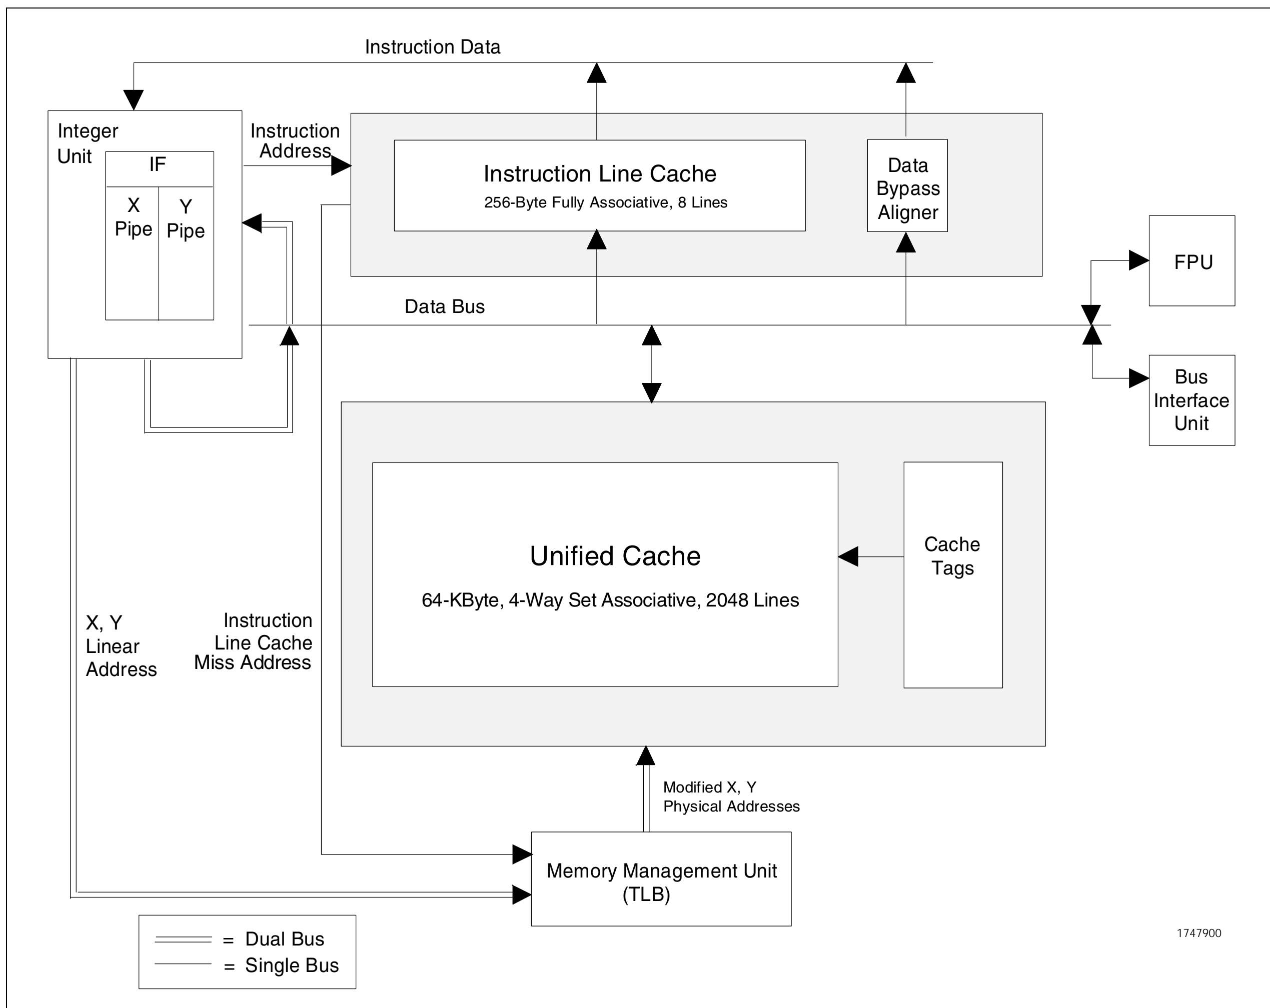

Figure 1-2. Cache Unit Operations

The instruction line cache uses a pseudo-LRU replacement algorithm. To ensure proper operation in the case of self-modifying code, any writes to the unified cache are checked against the contents of the instruction line cache. If a hit occurs in the instruction line cache, the appropriate line is invalidated.

1.5 Memory Management Unit

The Memory Management Unit (MMU), shown in Figure 1-3, translates the linear address supplied by the IU into a physical address to be used by the unified cache and the bus interface. Memory management procedures are x86 compatible, adhering to standard paging mechanisms.

Within the 6x86MX CPU there are two TLBs, the main L1 TLB and the larger L2 TLB. The 16-entry L1 TLB is direct mapped and holds 42 lines. The 384-entry L2 TLB is 6-way associative and hold 384 lines. The DTE is located in memory.

Scratch Pad Cache Memory

The 6x86MX CPU has the capability to "lock down" lines in the L1 cache on a line by line basis. Locked down lines are treated as private memory for use by the CPU. Locked down memory does not participate in hardware--cache coherency protocols.

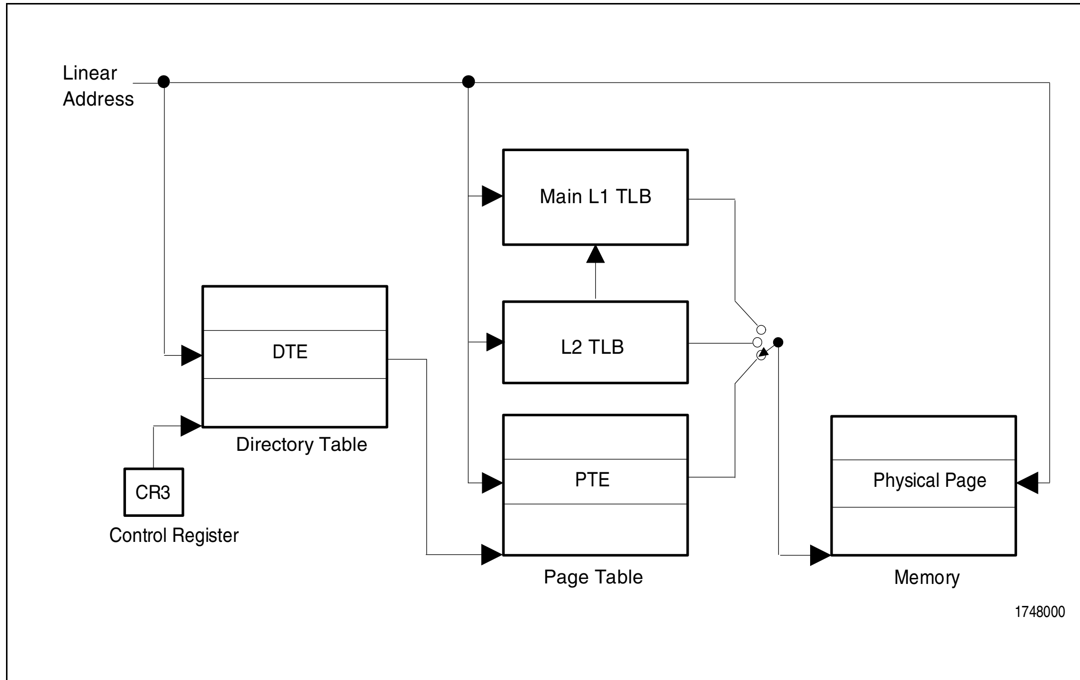

Figure 1-3. Paging Mechanism within the Memory Management Unit

Cache locking is controlled through use of the RDMSR and WRMSR instructions.

1.6 Floating Point Unit

The 6x86MX Floating Point Unit (FPU) processes floating point and MMX instructions. The FPU interfaces to the integer unit and the cache unit through a 64-bit bus. The 6x86MX FPU is x87 instruction set compatible and adheres to the IEEE-754 standard. Since most applications contain FPU instructions mixed with integer instructions, the 6x86MX FPU achieves high performance by completing integer and FPU operations in parallel.

FPU Parallel Execution

The 6x86MX CPU executes integer instructions in parallel with FPU instructions. Integer instructions may complete out of order with respect to the FPU instructions. The 6x86MX CPU maintains x86 compatibility by signaling exceptions and issuing write cycles in program order.

As previously discussed, FPU instructions are always dispatched to the integer unit's X pipeline. The address calculation stage of the X pipeline checks for memory management exceptions and accesses memory operands used by the FPU. If no exceptions are detected, the 6x86MX CPU checkpoints the state of the CPU and, during AC2, dispatches the floating point instruction to the FPU instruction queue. The 6x86MX CPU can then complete any subsequent integer instructions speculatively and out of order relative to the FPU instruction and relative to any potential FPU exceptions which may occur.

As additional FPU instructions enter the pipeline, the 6x86MX CPU dispatches up to four FPU instructions to the FPU instruction queue. The 6x86MX CPU continues executing speculatively and out of order, relative to the FPU queue, until the 6x86MX CPU encounters one of the conditions that causes speculative execution to halt. As the FPU completes instructions, the speculation level decreases and the checkpointed resources are available for reuse in subsequent operations. The 6x86MX FPU also uses a set of six write buffers to prevent stalls due to speculative writes.

1.7 Bus Interface Unit

The Bus Interface Unit (BIU) provides the signals and timing required by external circuitry. The signal descriptions and bus interface timing information is provided in Chapters 3 and 4 of this manual.