

FIGURE 15-28 Memory map for an AT-style clone.

EXAMPLE 15-4

```

.MODEL SMALL
.386P
.DATA
;page directory
;
0000 00000004  PDIR  DD    4
;
;page table 0
;
0004 0400 [  TAB0  DD    1024 DUP (?)
           00000000
           ]
0000          .CODE
           .STARTUP
0010 66| B8 00000000  MOV  EAX,0
0016 8C C8          MOV  AX,CS
0018 66| C1 E0 04    SHL  EAX,4
001C 66| 05 00000004 R  ADD  EAX,OFFSET TAB0
0022 66| 25 FFFF000  AND  EAX,0FFFF000H
0028 66| 83 C0 07    ADD  EAX,7
002C 66| A3 0000 R    MOV  PDIR,EAX          ;address page table 0

```

502 CHAPTER 15 THE 80386 AND 80486 MICROPROCESSORS

```
0030 B9 0100      MOV     CX,256
0033 BF 0004 R    MOV     DI,OFFSET TAB0
0036 8C D8       MOV     AX,DS
0038 8E C0       MOV     ES,AX
003A 66| B8 00000007 MOV     EAX,7
                .REPEAT                ;remap 00000H-09FFFH
0040 66| AB       STOSD                ;to 00000H-09FFFH
0042 66| 05 00001000 ADD     EAX,4096
                .UNTILCXZ
004A 66| B8 00102007 MOV     EAX,0102007H
0050 B9 0010     MOV     CX,16
                .REPEAT                ;remap 0A000H-0AFFFFH
0053 66| AB       STOSD                ;to 102000H-11FFFFH
0055 66| 05 00001000 ADD     EAX,4096
                .UNTILCXZ
005D 66| B8 00000000 MOV     EAX,0
0063 8C D8       MOV     AX,DS
0065 66| C1 E0 04   SHL     EAX,4
0069 66| 05 00000000 R ADD     EAX,OFFSET PDIR ;load CR3 with page directory
006F 0F 22 D8     MOV     CR3,EAX

                ;additional software to remap other areas of memory
                end
```

15-7 INTRODUCTION TO THE 80486 MICROPROCESSOR

The 80486 microprocessor is a highly integrated device, containing well over 1.2 million transistors. Located within this device circuit are a memory-management unit (MMU), a complete numeric coprocessor that is compatible with the 80387, a high-speed level-one cache memory that contains 8K bytes of space, and a full 32-bit microprocessor that is upward-compatible with the 80386 microprocessor. The 80486 is currently available as a 25 MHz, 33 MHz, 50 MHz, 66 MHz, or 100 MHz device. Note that the 66 MHz version is double-clocked and the 100 MHz version is triple-clocked. In 1990, Intel demonstrated a 100 MHz version (not double-clocked) of the 80486 for *Computer Design* magazine, but it has yet to be released. Advanced Micro Devices (AMD) has produced a 40 MHz version that is also available in an 80 MHz (double-clocked) and a 120 MHz (triple-clocked) form. The 80486 is available as an 80486DX or an 80486SX. The only difference between these devices is that the 80486SX does not contain the numeric coprocessor, which reduces its price. The 80487SX numeric coprocessor is available as a separate component for the 80486SX microprocessor.

This section details the differences between the 80486 and 80386 microprocessors. These differences are few, as shall be seen. The most notable differences apply to the cache memory system and parity generator.

Pin-out of the 80486DX and 80486SX Microprocessors

Figure 15-29 illustrates the pin-out of the 80486DX microprocessor, a 168-pin PGA. The 80486SX, also packaged in a 168-pin PGA, is not illustrated because only a few differences exist. Note that pin B15 is NMI on the 80486DX and pin A15 is NMI on the 80486SX. The only other differences are that pin A15 is IGNNE on the 80486DX (not present on the 80486SX), pin C14 is FERR on the 80486DX, and pins B15 and C14 on the 80486SX are not connected.

When connecting the 80486 microprocessor, all Vcc and Vss pins must be connected to the power supply for proper operation. The power supply must be capable of supplying 5.0 V \pm 10 percent, with up to 1.2 A of surge current for the 33 MHz version. The average supply current is 650 mA for the 33 MHz version. Intel has also produced a 3.3 V version that requires an average of 500 mA at a triple-clock speed of 100 MHz. Logic 0 outputs allow up to 4.0 mA of current, and logic 1 outputs allow up to 1.0 mA. If larger currents are required, as they often




FIGURE 15-29 The pin-out of the 80486. (Courtesy of Intel Corporation.)

are, then the 80486 must be buffered. Figure 15-30 shows a buffered 80486DX system. In the circuit shown, only the address, data, and parity signals are buffered.

Pin Definitions.

- A31-A2** **Address outputs** A31-A2 provide the memory and I/O with the address during normal operation; during a cache line invalidation, A31-A4 are used to drive the microprocessor.
- A20M** **Address bit 20 mask** causes the 80486 to wrap its address around from location 000FFFFFFH to 00000000H, as does the 8086 microprocessor. This provides a memory system that functions like the 1M-byte real memory system in the 8086 microprocessor.
- ADS** **Address data strobe** becomes a logic 0 to indicate that the address bus contains a valid memory address.
- AHOLD** **Address hold input** causes the microprocessor to place its address bus connections at their high-impedance state, with the remainder of the buses staying active. It is often used by another bus master to gain access for a cache-invalidation cycle.




FIGURE 15-30 An 80486 microprocessor showing the buffered address, data, and parity buses.

BE3-BE0	Byte enable outputs select a bank of the memory system when information is transferred between the microprocessors and its memory and I/O space. The BE3 signal enables D31-D24, BE2 enables D23-D16, BE1 enables D15-D8, and BE0 enables D7-D0.
BLAST	The burst last output shows that the burst bus cycle is complete on the next activation of the BRDY signal.
BOFF	The back-off input causes the microprocessor to place its buses at their high-impedance state during the next clock cycle. The microprocessor remains in the bus hold state until the BOFF pin is placed at a logic 1 level.
BRDY	The burst ready input is used to signal the microprocessor that a burst cycle is complete.
BREQ	The bus request output indicates that the 80486 has generated an internal bus request.
BS8	The bus size 8 input causes the 80486 to structure itself with an 8-bit data bus to access byte-wide memory and I/O components.
BS16	The bus size 16 input causes the 80486 to structure itself with a 16-bit data bus to access word-wide memory and I/O components.
CLK	The clock input provides the 80486 with its basic timing signal. The clock input is a TTL-compatible input that is 25 MHz to operate the 80486 at 25 MHz.
D31-D0	The data bus transfers data between the microprocessor and its memory and I/O system. Data bus connections D7-D0 are also used to accept the interrupt vector type number during an interrupt acknowledge cycle.
D/C	The data/control output indicates whether the current operation is a data transfer or control cycle. See Table 15-3 for the function of D/C, M/I/O, and W/R.
DP3-DP0	Data parity I/O provides even parity for a write operation and check parity for a read operation. If a parity error is detected during a read, the PCHK output becomes a logic 0 to indicate a parity error. If parity is not used in a system, these lines must be pulled high to +5.0 V or to 3.3 V in a system that uses a 3.3 V supply.
EADS	The external address strobe input is used with AHOLD to signal that an external address is used to perform a cache-invalidation cycle.
FERR	The floating-point error output indicates that the floating-point coprocessor has detected an error condition. It is used to maintain compatibility with DOS software.
FLUSH	The cache flush input forces the microprocessor to erase the contents of its 8K-byte internal cache.
HLDA	The hold acknowledge output indicates that the HOLD input is active and that the microprocessor has placed its buses at their high-impedance state.

TABLE 15-3 Bus cycle identification.

M/I/O	D/C	W/R	Bus Cycle Type
0	0	0	Interrupt acknowledge
0	0	1	Halt/special
0	1	0	I/O read
0	1	1	I/O write
1	0	0	Opcode fetch
1	0	1	Reserved
1	1	0	Memory read
1	1	1	Memory write

HOLD	The hold input requests a DMA action. It causes the address, data, and control buses to be placed at their high-impedance state and also, once recognized, causes HLDA to become a logic 0.
IGNNE	The ignore numeric error input causes the coprocessor to ignore floating-point errors and to continue processing data. This signal does not affect the state of the FERR pin.
INTR	The interrupt request input requests a maskable interrupt, as it does in all other family members.
KEN	The cache enable input causes the current bus to be stored in the internal cache.
LOCK	The lock output becomes a logic 0 for any instruction that is prefixed with the lock prefix.
M/IO	Memory/IO defines whether the address bus contains a memory address or an I/O port number. It is also combined with the W/R signal to generate memory, and I/O read and write control signals.
NMI	The non-maskable interrupt input requests a type 2 interrupt.
PCD	The page cache disable output reflects the state of the PCD attribute bit in the page table entry or the page directory entry.
PCHK	The parity check output indicates that a parity error was detected during a read operation on the DP3–DP0 pins.
PLOCK	The pseudo-lock output indicates that the current operation requires more than one bus cycle to perform. This signal becomes a logic 0 for arithmetic coprocessor operations that access 64- or 80-bit memory data.
PWT	The page write through output indicates the state of the PWT attribute bit in the page table entry or the page directory entry.
RDY	The ready input indicates that a non-burst bus cycle is complete. The RDY signal must be returned, or the microprocessor places wait states into its timing until RDY is asserted.
RESET	The reset input initializes the 80486, as it does in other family members. Table 15–4 shows the effect of the RESET input on the 80486 microprocessor.
W/R	Write/read signals that the current bus cycle is either a read or a write.

Basic 80486 Architecture

The architecture of the 80486DX is almost identical to the 80386. Added to the 80386 architecture inside the 80486DX is a math coprocessor and an 8K-byte level 1 cache memory. The 80486SX is almost identical to an 80386 with an 8K-byte cache, but no numeric coprocessor. Figure 15–31 illustrates the basic internal structure of the 80486 microprocessor. If this is compared to the architecture of the 80386, no differences are observed. The most prominent difference between the 80386 and the 80486 is that almost half of the 80486 instructions execute in one clocking period instead of the two clocking periods for the 80386 to execute similar instructions.

As with the 80386, the 80486 contains eight general-purpose 32-bit registers: EAX, EBX, ECX, EDX, EBP, EDI, ESI, and ESP. These registers may be used as 8-, 16-, or 32-bit data registers or to address a location in the memory system. The 16-bit registers are the same set as found in the 80286 and are assigned: AX, BX, CX, DX, BP, DI, SI, and SP. The 8-bit registers are AH, AL, BH, BL, CH, CL, DH, and DL.

In addition to the general-purpose registers, the 80486 also contains the same segment registers as the 80386, which are CS, DS, ES, SS, FS, and GS. Each are 16-bits wide, as in all earlier versions of the family.

The IP (instruction pointer) addresses the program located within the 1M byte of memory in combination with CS, or as EIP (extended instruction pointer) to address a program at any location within the 4G-byte memory

TABLE 15-4 The effect of the RESET signal.

Register	Initial Value with Self Test	Initial Value without Self Test
EAX	00000000H	?
EDX	00000400H + ID*	00000400H + ID*
EFLAGS	00000002H	00000002H
EIP	0000FFF0H	0000FFF0H
ES	0000H	0000H
CS	F000H	F000H
DS	0000H	0000H
SS	0000H	0000H
FS	0000H	0000H
GS	0000H	0000H
IDTR	base = 0, limit = 3FFH	base = 0, limit = 3FFH
CR0	60000010H	60000010H
DR7	00000000H	00000000H

*Note: Revision ID number is supplied by Intel for revisions to the microprocessor.

system. In protected mode operation, the segment registers function to hold selectors, as they did in the 80286 and 80386 microprocessors.

The 80486 also contains the global, local, and interrupt descriptor table register and memory management unit, as in the 80386. Although these registers are not illustrated in Figure 15-31, they are present as they are in the 80386. (The function of the MMU and its paging unit was described earlier in the chapter.)

The extended flag register (EFLAG) is illustrated in Figure 15-32. As with other family members, the rightmost flag bits perform the same functions for compatibility. The only new flag bit is the AC (**alignment check**), used to indicate that the microprocessor has accessed a word at an odd address or a doubleword stored at a non-doubleword boundary. Efficient software and execution require that data be stored at word or doubleword boundaries.

80486 Memory System

The memory system for the 80486 is identical to the 80386 microprocessor. The 80486 contains 4G bytes of memory, beginning at location 00000000H and ending at location FFFFFFFFH. The major change to the memory system is internal to the 80486 in the form of an 8K-byte cache memory, which speeds the execution of instructions and the acquisition of data. Another addition is the parity checker/generator built into the 80486 microprocessor.

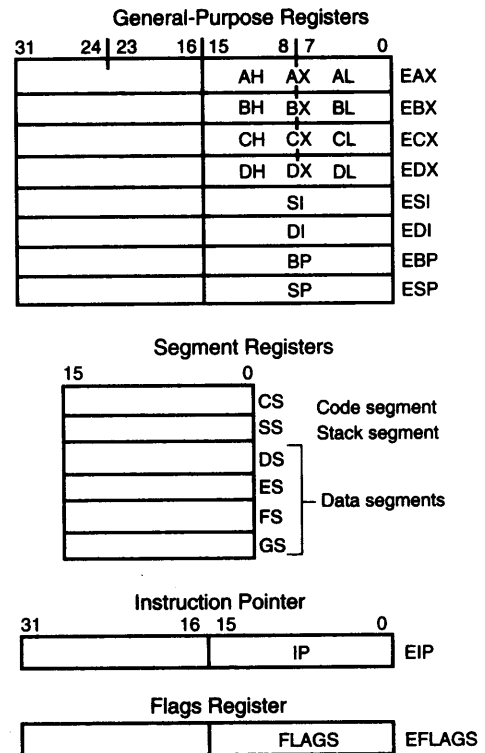
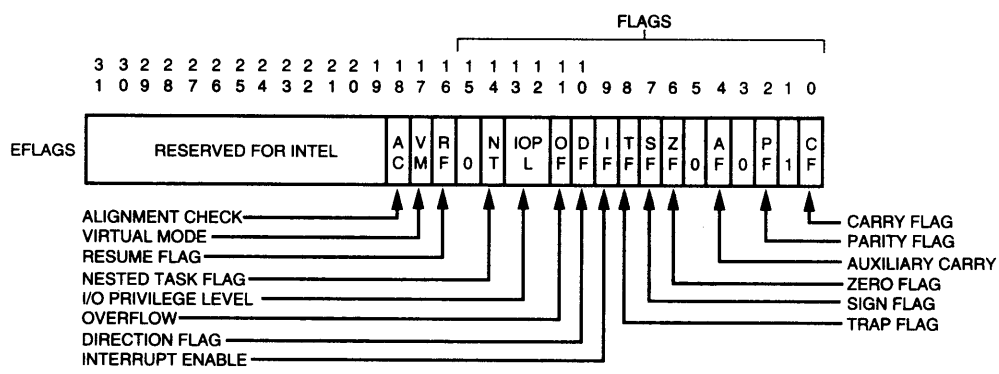


FIGURE 15-31 The internal programming model of the 80486. (Courtesy of Intel Corporation.)



Note: 0 indicates Intel Reserved: do not define.

240440-6

FIGURE 15-32 The EFLAG register of the 80486. (Courtesy of Intel Corporation.)




FIGURE 15-33 The organization of the 80486 memory, showing parity.

Parity Checker/Generator. Parity is often used to determine if data are correctly read from a memory location. To facilitate this, Intel has incorporated an internal parity generator/detector. Parity is generated by the 80486 during each write cycle. Parity is generated as even parity, and a parity bit is provided for each byte of memory. The parity check bits appear on pins DP0–DP3, which are also parity inputs as well as outputs. These are typically stored in memory during each write cycle and read from memory during each read cycle.

On a read, the microprocessor checks parity and generates a parity check error, if it occurs, on the PCHK pin. A parity error causes no change in processing unless the user applies the PCHK signal to an interrupt input. Interrupts are often used to signal a parity error in DOS-based computer systems. Figure 15-33 shows the organization of the 80486 memory system that includes parity storage. Note that this is the same as for the 80386, except for the parity bit storage. If parity is not used, Intel recommends that the DP0–DP3 pins be pulled-up to +5.0 V.

Cache Memory. The cache memory system caches (stores) data used by a program and also the instructions of the program. The cache is organized as a 4-way set associative cache, with each location (line) containing 16 bytes or four doublewords of data. The cache operates as a write-through cache. Note that the cache changes only if a




FIGURE 15-34 Control register zero (CR0) for the 80486 microprocessor.

miss occurs. This means that data written to a memory location not already cached are not written to the cache. In many cases, much of the active portion of a program is found completely inside the cache memory. This causes execution to occur at the rate of one clock cycle for many of the instructions that are commonly used in a program. About the only way that these efficient instructions are slowed is when the microprocessor must fill a line in the cache. Data are also stored in the cache, but it has less of an impact on the execution speed of a program because data are not referenced repeatedly as many portions of a program are.

Control register 0 (CR0) is used to control the cache with two new control bits not present in the 80386 microprocessor. (See Figure 15-34 for CR0 in the 80486 microprocessor.) The CD (cache disable) and NW (non-cache write-through) bits are new to the 80486 and are used to control the 8K-byte cache. If the CD bit is a logic 1, all cache operations are inhibited. This setting is used only for debugging software and normally remains cleared. The NW bit is used to inhibit cache write-through operations. As with CD, cache write-through is inhibited only for testing. For normal program operation, $CD = 0$ and $NW = 0$.

Because the cache is new to the 80486 microprocessor and the cache is filled by using burst cycles not present on the 80386, some detail is required to understand bus-filling cycles. When a bus line is filled, the 80486 must acquire four 32-bit numbers from the memory system to fill a line in the cache. Filling is accomplished with a burst cycle. The burst cycle is a special memory where four 32-bit numbers are fetched from the memory system in five clocking periods. This assumes that the speed of the memory is sufficient and that no wait states are required. If the clock frequency of the 80486 is 33 MHz, we can fill a cache line in 167 ns, which is very efficient considering that a normal, non-burst 32-bit memory read operation requires two clocking periods.

Memory Read Timing. Figure 15-35 illustrates the read timing, for the 80486, for a non-burst memory operation. Note that two clocking periods are used to transfer data. Clocking period T1 provides the memory




FIGURE 15-35 The non-burst read timing for the 80486 microprocessor.

address and control signals, and clocking period T2 is where the data are transferred between the memory and the microprocessor. Note that the RDY must become a logic 0 to cause data to be transferred and to terminate the bus cycle. Access time for a non-burst access is determined by taking two clocking periods, minus the time required for the address to appear on the address bus connection, minus a setup time for the data bus connections. For the 20 MHz version of the 80486, two clocking periods require 100 ns minus 28 ns for address setup time, and 6 ns for data setup time. This yields a non-burst access time of 100 ns – 34 ns, or 66 ns. Of course, if decoder time and delay times are included, the access time allowed the memory is even less for no wait-state operation. If a higher frequency version of the 80486 is used in a system, memory access time is still less.

The 80486 33 MHz, 66 MHz, and 100 MHz processors all access bus data at a 33 MHz rate. In other words, the microprocessor may operate at 100 MHz, but the system bus operates at 33 MHz. Notice that the non-burst access timing for the 33 MHz system bus allows 60 ns – 24 ns = 36 ns. It is obvious that wait states are required for operation with standard DRAM memory devices.

Figure 15–36 illustrates the timing diagram for filling a cache line with four 32-bit numbers using a burst. Note that the addresses (A31–A4) appear during T1 and remain constant throughout the burst cycle. Also, note that A2 and A3 change during each T2 after the first to address four consecutive 32-bit numbers in the memory system. As mentioned, cache fills using bursts require only five clocking periods (one T1 and four T2s) to fill a cache line with four doublewords of data. Access time using a 20 MHz version of the 80486 for the second and subsequent doublewords is 50 ns – 28 ns – 5 ns, or 17 ns, assuming no delays in the system. To use burst mode transfers, we need high-speed memory. Because DRAM memory access times are 40 ns at best, we are forced to use SRAM for burst cycle transfers. The 33 MHz system allows an access time of 30 ns – 19 ns – 5 ns, or 6 ns for the second and subsequent bytes. If an external counter is used in place of address bits A2 and A3, the 19 ns can be eliminated and the access time becomes 30 ns – 5 ns, or 25 ns, which is enough time for even the slowest SRAM connected to the system as a cache. This circuit is often called a *synchronous burst mode cache* if SRAM cache is used with the system. Note that the BRDY pin acknowledges a burst transfer rather than the RDY pin, which acknowledges a normal memory transfer.

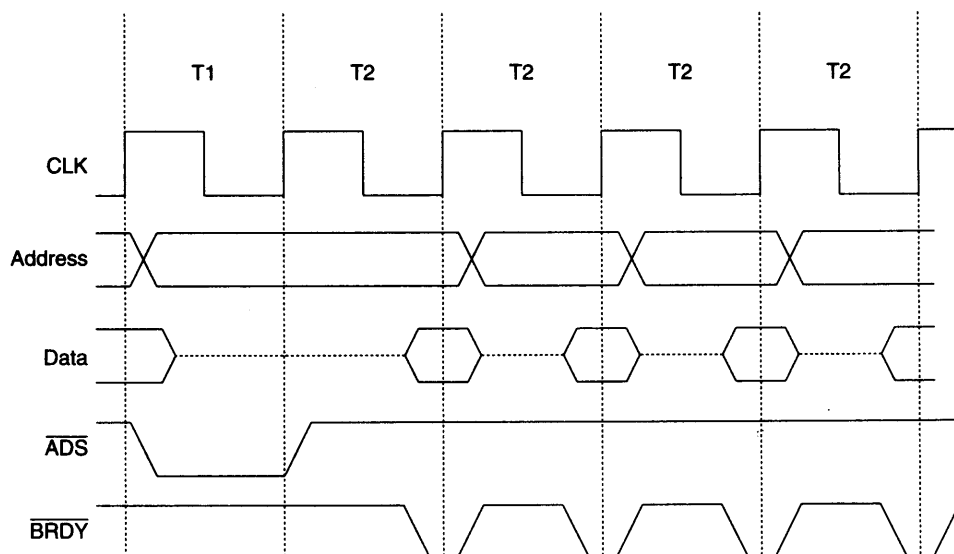


FIGURE 15–36 A burst cycle reads four doublewords in five clocking periods.

The cache data register (TR3) is used to access either the cache fill buffer for a write test operation or the cache read buffer for a cache read test operation. This register is a window into the 8K-byte cache memory located within the 80486 and is used for testing the cache. In order to fill or read a cache line (128 bits wide), TR3 must be written or read four times.

The contents of the set select field in TR5 determine which internal cache line is written or read through TR3. The 7-bit test field selects one of the 128 different 16 byte-wide cache lines. The entry select bits of TR5 select an entry in the set or the 32-bit location in the fill/read buffer. The control bits in TR5 enable the fill buffer or read buffer operation (00), perform a cache write (01), perform a cache read (10), or flush the cache (11).

The cache status register (TR4) holds the cache tag, LRU bits, and a valid bit. This register is loaded with the tag and valid bit before a cache write operation; and contains the tag, valid bit, LRU bits, and four valid bits on a cache test read.

The cache is tested each time that the microprocessor is reset if the AHOLD pin is high for two clocks prior to the RESET pin going low. This causes the 80486 to completely test itself with a built-in self-test or BIST. The BIST uses TR3, TR4, and TR5 to completely test the internal cache. Its outcome is reported in register EAX. If EAX is a zero, the microprocessor, coprocessor, and cache have passed the self-test. The value of EAX can be tested after a reset to determine if an error is detected. In most cases, we do not directly access the test registers unless we wish to perform our own tests on the cache or TLB.

15-8 SUMMARY

1. The 80386 microprocessor is an enhanced version of the 80286 microprocessor and includes a memory-management unit that is enhanced to provide memory paging. The 80386 also includes 32-bit extended registers, and a 32-bit address and data bus. A scaled-down version of the 80386DX with a 16-bit data and 24-bit address bus is available as the 80386SX microprocessor. The 80386EX is a complete AT-style personal computer on a chip.
2. The 80386 has a physical memory size of 4G bytes that can be addressed as a virtual memory with up to 64T bytes. The 80386 memory is 32 bits wide, and it is addressed as bytes, words, or doublewords.
3. When the 80386 is operated in the pipelined mode, it sends the address of the next instruction or memory data to the memory system prior to completing the execution of the current instruction. This allows the memory system to begin fetching the next instruction or data before the current is completed. This increases access time, thus reducing the speed of the memory.
4. A cache memory system allows data that are frequently read to be accessed in less time because they are stored in high-speed semiconductor memory. If data are written to memory, they are also written to the cache, so the most current data are always present in the cache.
5. The I/O structure of the 80386 is almost identical to the 80286, except that I/O can be inhibited when the 80386 is operated in the protected mode through the I/O bit protection map stored with the TSS.
6. The register set of the 80386 contains extended versions of the registers introduced on the 80286 microprocessor. These extended registers include: EAX, EBX, ECX, EDX, EBP, ESP, EDI, ESI, EIP, and EFLAGS. In addition to the extended registers, two supplemental segment registers (FS and GS) are added. Debug registers and control registers handle system debugging tasks and memory management in the protected mode.
7. The instruction set of the 80386 is enhanced to include instructions that address the 32-bit extended register set. The enhancements also include additional addressing modes that allow any extended register to address memory data. Scaling has been added so that an index register can be multiplied by 1, 2, 4, or 8. New instruction types include bit scan, string moves with sign- or zero-extension, set byte upon condition, and double-precision shifts.
8. In the 80386 microprocessor, interrupts have been expanded to include additional predefined interrupts in the interrupt vector table. These additional interrupts are used with the memory-management system.

9. The 80386 memory manager is similar to the 80286, except that the physical addresses generated by the MMU are 32 bits wide instead of 24 bits wide. The 80386 MMU is also capable of paging.
10. The 80386 is operated in the real mode (8086 mode) when it is reset. The real mode allows the microprocessor to address data in the first 1M byte of memory. In the protected mode, the 80386 addresses any location in its 4G bytes of physical address space.
11. A descriptor is a series of eight bytes that specify how a code or data segment is used by the 80386. The descriptor is selected by a selector that is stored in one of the segment registers. Descriptors are used only in the protected mode.
12. Memory management is accomplished through a series of descriptors, stored in descriptor tables. To facilitate memory management, the 80386 uses three descriptor tables: the global descriptor table (GDT), the local descriptor table (LDT), and the interrupt descriptor table (IDT). The GDT and LDT each hold up to 8192 descriptors; the IDT holds up to 256 descriptors. The GDT and LDT describe code and data segments as well as tasks. The IDT describes the 256 different interrupt levels through interrupt gate descriptors.
13. The TSS (task state segment) contains information about the current task and the previous task. Appended to the end of the TSS is an I/O bit protection map that inhibits selected I/O port addresses.
14. The memory paging mechanism allows any 4K-byte physical memory page to be mapped to any 4K-byte linear memory page. For example, memory location 00A00000H can be assigned memory location A0000000H through the paging mechanism. A page directory and page tables are used to assign any physical address to any linear address. The paging mechanism can be used in the protected mode or the virtual mode.
15. The 80486 microprocessor is an improved version of the 80386 microprocessor that contains an 8K-byte cache and an 80387 arithmetic coprocessor; it executes many instructions in one clocking period.
16. The 80486 microprocessor executes a few new instructions that control the internal cache memory and allow addition (XADD) and comparison (CMPXCHG) with an exchange and a byte swap (BSWAP) operation. Other than these few additional instructions, the 80486 is 100 percent upward-compatible with the 80386 and 80387.
17. A new feature found in the 80486 is the BIST (built-in self-test) that tests the microprocessor, coprocessor, and cache at reset time. If the 80486 passes the test, EAX contains a zero.
18. Additional test registers are added to the 80486 to allow the cache memory to be tested. These new test registers are TR3 (cache data), TR4 (cache status), and TR5 (cache control). Although we seldom use these registers, they are used by BIST each time that a BIST is performed after a reset operation.

15-9 QUESTIONS AND PROBLEMS

1. The 80386 microprocessor addresses _____ bytes of physical memory when operated in the protected mode.
2. The 80386 microprocessor addresses _____ bytes of virtual memory through its memory-management unit.
3. Describe the differences between the 80386DX and the 80386SX.
4. Draw the memory map of the 80386 when operated in the
 - (a) protected mode
 - (b) real mode
5. How much current is available on various 80386 output pin connections? Compare these currents with the currents available at the output pin connection of an 8086 microprocessor.
6. Describe the 80386 memory system, and explain the purpose and operation of the bank selection signals.
7. Explain the action of a hardware reset on the address bus connections of the 80386.
8. Explain how pipelining lengthens the access time for many memory references in the 80386 microprocessor-based system.

9. Briefly describe how the cache memory system functions.
10. I/O ports in the 80386 start at I/O address _____ and extend to I/O address _____.
11. What I/O ports communicate data between the 80386 and its companion 80387 coprocessor?
12. Compare and contrast the memory and I/O connections found on the 80386 with those found in earlier microprocessors.
13. If the 80386 operates at 20 MHz, what clocking frequency is applied to the CLK2 pin?
14. What is the purpose of the BS16 pin on the 80386 microprocessor?
15. What two additional segment registers are found in the 80386 programming model that are not present in the 8086?
16. List the extended registers found in the 80386 microprocessor.
17. List each 80386 flag register bit and describe its purpose.
18. Define the purpose of each of the control registers (CR0, CR1, CR2, and CR3) found within the 80386.
19. Define the purpose of each 80386 debug register.
20. The debug registers cause which level of interrupt?
21. Describe the operation of the bit scan forward instruction.
22. Describe the operation of the bit scan reverse instruction.
23. Describe the operation of the SHRD instruction.
24. Form an instruction that accesses data in the FS segment at the location indirectly addressed by the DI register. The instruction should store the contents of EAX into this memory location.
25. What is scaled index addressing?
26. Is the following instruction legal? `MOV AX,[EBX+ECX]`
27. Explain how the following instructions calculate the memory address:
 - (a) `ADD [EBX+8*ECX],AL`
 - (b) `MOV DATA[EAX+EBX],CX`
 - (c) `SUB EAX,DATA`
 - (d) `MOV ECX,[EBX]`
28. What is the purpose of interrupt type number 7?
29. Which interrupt vector type number is activated for a protection privilege violation?
30. What is a double interrupt fault?
31. If an interrupt occurs in the protected mode, what defines the interrupt vectors?
32. What is a descriptor?
33. What is a selector?
34. How does the selector choose the local descriptor table?
35. What register is used to address the global descriptor table?
36. How many global descriptors can be stored in the GDT?
37. Explain how the 80386 can address a virtual memory space of 64T bytes when the physical memory contains only 4G bytes of memory.
38. What is the difference between a segment descriptor and a system descriptor?
39. What is the task state segment (TSS)?
40. How is the TSS addressed?
41. Describe how the 80386 switches from the real mode to the protected mode.
42. Describe how the 80386 switches from the protected mode to the real mode.
43. What is virtual 8086 mode operation of the 80386 microprocessor?
44. How is the paging directory located by the 80386?
45. How many bytes are found in a page of memory?
46. Explain how linear memory address D0000000H can be assigned to physical memory address C0000000H with the paging unit of the 80386.
47. What are the differences between an 80386 and 80486 microprocessor?
48. What is the purpose of the FLUSH input pin on the 80486 microprocessor?
49. Compare the register set of the 80386 with the 80486 microprocessor.

50. What differences exist in the flags of the 80486 when compared to the 80386 microprocessor?
51. Which pins are used for parity checking on the 80486 microprocessor?
52. The 80486 microprocessor uses _____ parity.
53. The cache inside the 80486 microprocessor is ____ K bytes.
54. A cache line is filled by reading _____ bytes from the memory system.
55. What is an 80486 burst?
56. Define the term *cache-write through*.
57. What is a BIST?
58. Can 80486 caching be disabled by software? Explain your answer.
59. Explain how the XADD EBX,EDX instruction operates.
60. The CMPXCHG CL,AL instruction compares CL with AL. What else occurs when this instruction executes?
61. Compare the INVD instruction with the WBINVD instruction.
62. What is the purpose of the PCD bit in the page table directory or page table entry?
63. Does the PWT bit in the page table directory or page table entry affect the on-chip cache?

CHAPTER 16

The Pentium and Pentium Pro Microprocessors

INTRODUCTION

The Pentium microprocessor signals an improvement to the architecture found in the 80486 microprocessor. The changes include an improved cache structure, a wider data bus width, a faster numeric coprocessor, a dual integer processor, and branch prediction logic. The cache has been reorganized to form two caches that are each 8K bytes in size, one for caching data, and the other for instructions. The data bus width has been increased from 32 bits to 64 bits. The numeric coprocessor operates at about five times faster than the 80486 numeric coprocessor. A dual-integer processor often allows two instructions per clock. Finally, the branch prediction logic allows programs that branch to execute more efficiently. Notice that these changes are internal to the Pentium, which makes software upward-compatible from earlier Intel 80X86 microprocessors. A later improvement to the Pentium was the addition of the MMX instructions.

The Pentium Pro is a still faster version of the Pentium, and it contains a modified internal architecture that can schedule up to five instructions for execution and an even faster floating-point unit. The Pentium Pro also contains a 256K-byte or 512K-byte level two cache in addition to the 16K-byte (8K for data and 8K for instruction) level one cache. The Pentium Pro includes error correction circuitry (ECC) to correct a one bit error and indicate a two bit error. Also added are four additional address lines, giving the Pentium Pro access to an astounding 64G bytes of directly addressable memory space.

CHAPTER OBJECTIVES

Upon completion of this chapter, you will be able to:

1. Contrast the Pentium and Pentium Pro with the 80386 and 80486 microprocessors.
2. Describe the organization and interface of the 64-bit wide Pentium memory system and its variations.
3. Contrast the changes in the memory-management unit and paging unit when compared to the 80386 and 80486 microprocessors.
4. Detail the new instructions found with the Pentium microprocessor.
5. Explain how the superscaler dual integers units improve performance of the Pentium microprocessor.
6. Describe the operation of the branch prediction logic.
7. Detail the improvements in the Pentium Pro when compared with the Pentium.
8. Explain how the dynamic execution architecture of the Pentium Pro functions.

16-1 INTRODUCTION TO THE PENTIUM MICROPROCESSOR

Before the Pentium or any other microprocessor can be used in a system, the function of each pin must be understood. This section of the chapter details the operation of each pin, along with the external memory system and I/O structures of the Pentium microprocessor.

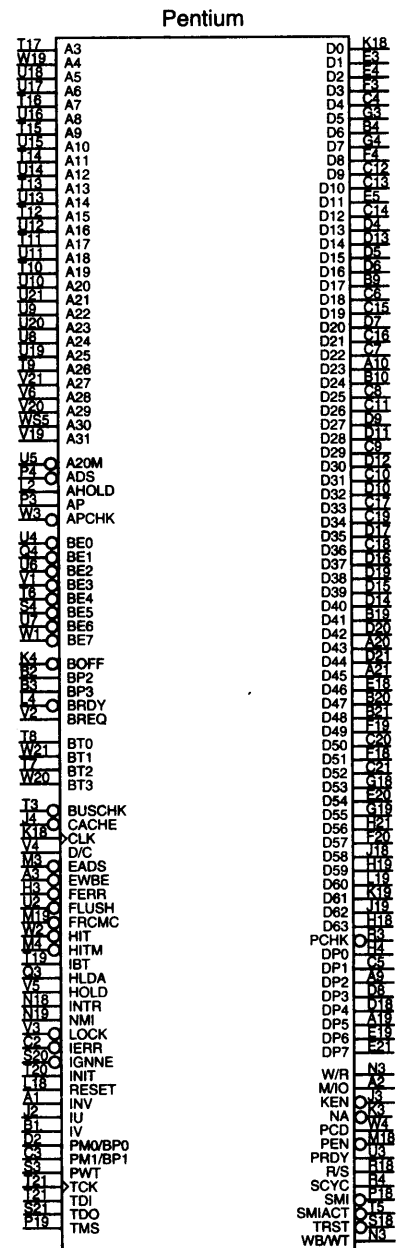
Figure 16-1 illustrates the pin-out of the Pentium microprocessor, which is packaged in a huge 237-pin PGA (pin grid array). Currently, the Pentium is available in two versions: the full-blown Pentium and the P24T version called the Pentium OverDrive. The P24T version contains a 32-bit data bus, compatible for insertion into 80486 machines, which contains the P24T socket. The P24T version also comes with a fan built into the unit. The most notable difference in the pin-out of the Pentium, when compared to earlier 80486 microprocessors, is that there are 64 data bus connections instead of 32, which require a larger physical footprint.

As with earlier versions of the Intel family of microprocessors, the early versions of the Pentium require a single +5.0 V power supply for operation. The power supply current averages 3.3 A for the 66 MHz version of the Pentium, and 2.91 A for the 60 MHz version. Because these currents are significant, so are the power dissipations of these microprocessors: 13 W for the 66 MHz version and 11.9 W for the 60 MHz version. The current versions of the Pentium, 90 MHz and above, use a 3.3 V power supply with reduced current consumption. At present, a good heat sink with considerable airflow is required to keep the Pentium cool. The Pentium contains multiple VCC and VSS connections that must all be connected to +5.0 V or +3.3 V and ground for proper operation. Some of the pins are labeled N/C (no connection) and must not be connected. The latest versions of the Pentium have been improved to reduce the power dissipation. For example, the 233 MHz Pentium requires 3.4 A or current, which is only slightly more than the 3.3 A required by the early 66 MHz version.

Each Pentium output pin is capable of providing 4.0 mA of current at a logic 0 level and 2.0 mA at a logic 1 level. This represents an increase in drive current, compared to the 2.0 mA available on earlier 8086, 8088, and 80286 output pins. Each input pin represents a small load requiring only 15 μ A of current. In some systems, except the smallest, these current levels require bus buffers.

The function of each Pentium group of pins follows:

- A20** The address **A20 mask** is an input that is asserted in the real mode to signal the Pentium to perform address wraparound, as in the 8086 microprocessor, for use of the HIMEM.SYS driver.



A31–A3	Address bus connections address any of the $512K \times 64$ memory locations found in the Pentium memory system. Note that A0, A1, and A2 are encoded in the bus enable (BE7–BE0), described elsewhere, to select any or all of the eight bytes in a 64-bit wide memory location.
ADS	The address data strobe becomes active whenever the Pentium has issued a valid memory or I/O address. This signal is combined with the W/R and M/IO signals to generate the separate read and write signals present in the earlier 8086–80286 microprocessor-based systems.
AHOLD	Address hold is an input that causes the Pentium to hold the address and AP signals for the next clock.
AP	Address parity provides even parity for the memory address on all Pentium-initiated memory and I/O transfers. The AP pin must also be driven with even parity information on all inquire cycles in the same clocking period as the EADS signal.
APCHK	Address parity check becomes a logic 0 whenever the Pentium detects an address parity error.
BE7–BE0	Bank enable signals select the access of a byte, word, doubleword, or quadword of data. These signals are generated internally by the microprocessor from address bits A0, A1, and A2.
BOFF	The back-off input aborts all outstanding bus cycles and floats the Pentium buses until BOFF is negated. After BOFF is negated, the Pentium restarts all aborted bus cycles in their entirety.
BP[3:2] and PM/BP[1:0]	The breakpoint pins BP3–BP0 indicate a breakpoint match when the debug registers are programmed to monitor for matches. The performance monitoring pins PM1 and PM0 indicate the settings of the performance monitoring bits in the debug mode control register.
BRDY	The burst ready input signals the Pentium that the external system has applied or extracted data from the data bus connections. This signal is used to insert wait states into the Pentium timing.
BREQ	The bus request output indicates that the Pentium has generated a bus request.
BT3–BT0	The branch trace outputs provide bits 2–0 of the branch target linear address and the default operand size on BT3. These outputs become valid during a branch trace special message cycle.
BUSCHK	The bus check input allows the system to signal the Pentium that the bus transfer has been unsuccessful.
CACHE	The cache output indicates that the current Pentium cycle can cache data.
CLK	The clock is driven by a clock signal that is at the operating frequency of the Pentium. For example, to operate the Pentium at 66 MHz, we apply a 66 MHz clock to this pin.
D63–D0	Data bus connections transfer byte, word, doubleword, and quadword data between the microprocessor and its memory and I/O system.
D/C	Data/control indicates that the data bus contains data for or from memory or I/O when a logic 1. If D/C is a logic 0, the microprocessor is either halted or executing an interrupt acknowledge.
DP7–DP0	Data parity is generated by the Pentium and detects its eight memory banks through these connections.
EADS	The external address strobe input signals that the address bus contains an address for an inquire cycle.

EWBE	The external write buffer empty input indicates that a write cycle is pending in the external system.
FERR	The floating-point error is comparable to the ERROR line in the 80386 and shows that the internal coprocessor has erred.
FLUSH	The flush cache input causes the cache to flush all write-back lines and invalidate its internal caches. If the FLUSH input is a logic 0 during a reset operation, the Pentium enters its test mode.
FRCMC	The functional redundancy check is sampled during a reset to configure the Pentium in the master (1) or checker mode (0).
HIT	Hit shows that the internal cache contains valid data in the inquire mode.
HITM	Hit modified shows that the inquire cycle found a modified cache line. This output is used to inhibit other master units from accessing data until the cache line is written to memory.
HOLD	Hold requests a DMA action.
HLDA	Hold acknowledge indicates that the 80386 is currently in a hold condition.
IBT	Instruction branch taken indicates that the Pentium has taken an instruction branch.
IERR	The internal error output shows that the Pentium has detected an internal parity error or functional redundancy error.
IGNNE	The ignore numeric error input causes the Pentium to ignore a numeric coprocessor error.
INIT	The initialization input performs a reset without initializing the caches, write-back buffers, and floating-point registers. This may not be used to reset the microprocessor in lieu of RESET after power-up.
INTR	The interrupt request is used by external circuitry to request an interrupt.
INV	The invalidation input determines the cache line state after an inquiry.
IU	The U-Pipe instruction complete output shows that the instruction in the U-pipe is complete.
IV	The V-Pipe instruction complete output shows that the instruction in the V-pipe is complete.
KEN	The cache enable input enables internal caching.
LOCK	LOCK becomes a logic 0 whenever an instruction is prefixed with the LOCK: prefix. This is most often used during DMA accesses.
M/IO	Memory/IO selects a memory device when a logic 1 or an I/O device when a logic 0. During the I/O operation, the address bus contains a 16-bit I/O address on address connections A15–A3.
NA	Next address indicates that the external memory system is ready to accept a new bus cycle.
NMI	The non-maskable interrupt requests a non-maskable interrupt, just as on the earlier versions of the microprocessor.
PCD	The page cache disable output shows that the internal page caching is disabled by reflecting the state of the CR3 PCD bit.
PCHK	The parity check output signals a parity check error for data read from memory or I/O.
PEN	The parity enable input enables the machine check interrupt or exception.
PRDY	The probe ready output indicates that the probe mode has been entered for debugging.

PWT	The page write-through output shows the state of the PWT bit in CR3.
R/S	This pin is provided for use with the Intel Debugging Port and causes an interrupt.
RESET	Reset initializes the Pentium, causing it to begin executing software at memory location FFFFFFF0H. The Pentium is reset to the real mode and the leftmost 12 address connections remain logic 1s (FFFH) until a far jump or far call is executed. This allows compatibility with earlier microprocessors. See Table 16–1 for the state of the Pentium after a hardware reset.
SCYC	The split cycle output signals a misaligned LOCKed bus cycle.
SMI	The system management interrupt input causes the Pentium to enter the system management mode of operation.
SMIACK	The system management interrupt active output shows that the Pentium is operating in the system management mode.
TCK	The testability clock input selects the clocking function in accordance to the IEEE 1149.1 Boundary Scan interface.
TDI	The test data input is used to test data clocked into the Pentium with the TCK signal.
TDO	The test data output is used to gather test data and instructions shifted out of the Pentium with TCK.
TMS	The test mode select input controls the operation of the Pentium in test mode.
TRST	The test reset input allows the test mode to be reset.
W/R	Write/read indicates that the current bus cycle is a write when a logic 1 or a read when a logic 0.
WB/WT	Write-back/write-through selects the operation for the Pentium data cache.

The Memory System

The memory system for the Pentium microprocessor is 4G bytes in size, just as in the 80386DX and 80486 microprocessors. The difference lies in the width of the memory data bus. The Pentium uses a 64-bit data bus to address

TABLE 16–1 State of the Pentium after a RESET.

<i>Register</i>	<i>RESET Value</i>	<i>RESET + BIST Value</i>
EAX	0	0 (if test passes)
EDX	0500XXXXH	0500XXXXH
EBX, ECX, ESP, EBP, ESI, and EDI	0	0
EFLAGS	2	2
EIP	0000FFF0H	0000FFF0H
CS	F000H	F000H
DS, ES, FS, GS, and SS	0	0
GDTR and TSS	0	0
CR0	60000010H	60000010H
CR2, CR3, and CR4	0	0
DR0–DR3	0	0
DR6	FFFF0FF0H	FFFF0FF0H
DR7	00000400H	00000400H

Notes: BIST = built-in self-test; XXXX = Pentium version number.




FIGURE 16-2 The 8-byte wide memory banks of the Pentium microprocessor.

memory organized in eight banks that each contain 512M bytes of data. See Figure 16-2 for the organization of the Pentium physical memory system.

The Pentium memory system is divided into eight banks that each store a byte of data with a parity bit. The Pentium, like the 80486, employs internal parity generation and checking logic for the memory system's data bus information. (Note that most Pentium systems do not use parity checks, but it is available.) The 64-bit wide memory is important to double-precision floating-point data. Recall that a double-precision floating-point number is 64 bits wide. Because of the change to a 64-bit wide data bus, the Pentium is able to retrieve floating-point data with one read cycle, instead of two as in the 80486. This causes the Pentium to function at a higher throughput than an 80486. As with earlier 32-bit Intel microprocessors, the memory system is numbered in bytes from byte 00000000H to byte FFFFFFFFH.

Memory selection is accomplished with the bank enable signals (BE7-BE0). These separate memory banks allow the Pentium to access any single byte, word, doubleword, or quadword with one memory transfer cycle. As with earlier memory selection logic, we often generate eight separate write strobes for writing to the memory system.

A new feature added to the Pentium is its capability to check and generate parity for the address bus (A31-A5) during certain operations. The AP pin provides the system with parity information and the APCHK indicates a bad parity check for the address bus. The Pentium takes no action when an address parity error is detected. The error must be assessed by the system and the system must take appropriate action (an interrupt), if so desired.

How is a 32-bit memory system connected to the Pentium? The Pentium can function with a 32-bit wide memory system by using a multiplexer to convert the 64-bit data bus to a 32-bit data bus. Figure 16-3 shows a set of bi-directional multiplexers (bi-directional buffers are used as multiplexers) that are used to convert the Pentium's 64-bit data bus into a 32-bit data bus. Care must be taken when using this arrangement because software could access a doubleword that crosses the boundary between the lower and upper halves of the data bus. All doublewords must be stored at doubleword boundaries. Note that a doubleword boundary is an address that is divisible by 4.

Input/Output System

The input/output system of the Pentium is completely compatible with earlier Intel microprocessors. The I/O port number appears on address lines A15-A3 with the bank enable signals used to select the actual memory banks used for the I/O transfer.

Beginning with the 80386 microprocessor, I/O privilege information is added to the TSS segment when the Pentium is operated in the protected mode. Recall that this allows I/O ports to be selectively inhibited. If the blocked I/O location is accessed, the Pentium generates a type 13 interrupt to signal an I/O privilege violation.

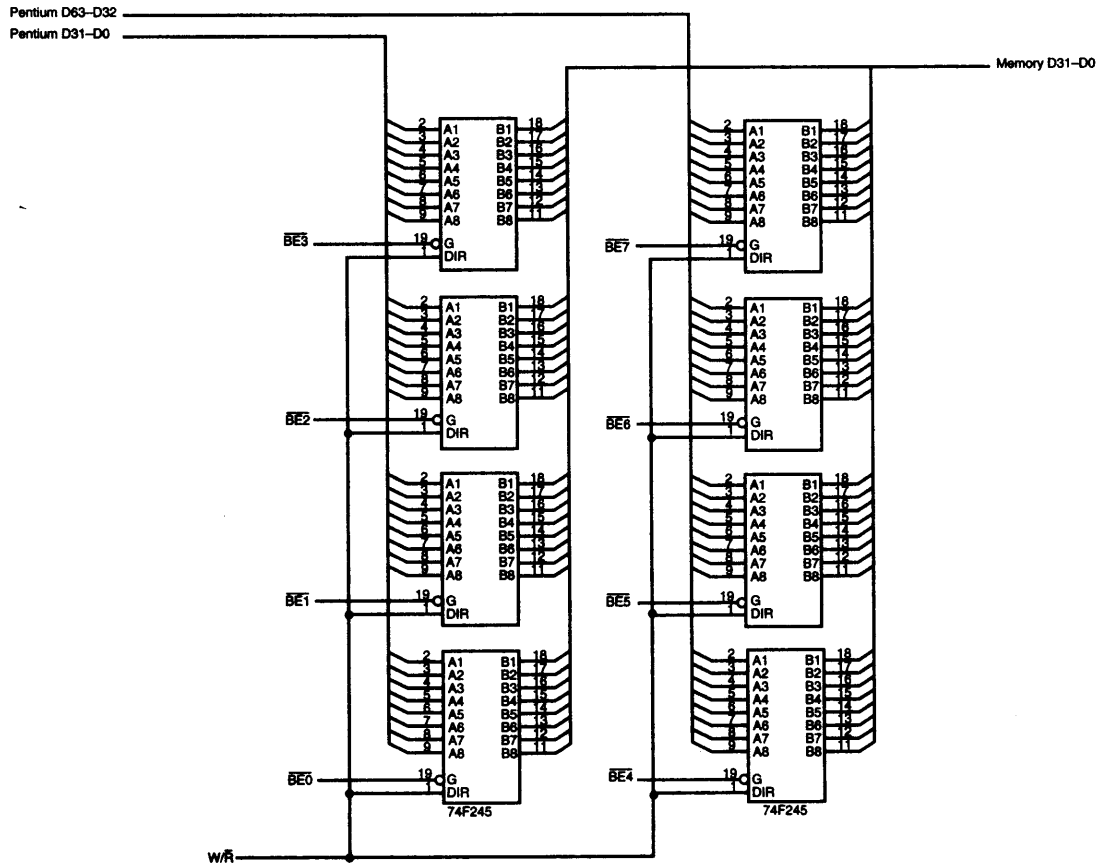


FIGURE 16-3 A circuit that generates a 32-bit memory data bus from the 64-bit Pentium data bus.

System Timing

As with any microprocessor, the system timing signals must be understood in order to interface the microprocessor. This portion of the text details the operation of the Pentium through its timing diagrams and shows how to determine memory access times.

The basic Pentium, non-pipelined memory cycle consists of two clocking periods: T1 and T2. See Figure 16-4 for the basic non-pipelined read cycle. Notice from the timing diagram that the 66 MHz Pentium is capable of 33 million memory transfers per second. This assumes that the memory can operate at that speed.

Also notice from the timing diagram that the W/R signal becomes valid if ADS is a logic 0 at the positive edge of the clock (end of T1). This clock must be used to qualify the cycle as a read or a write.

During T1, the microprocessor issues the ADS, W/R, address, and M/IO signals. In order to qualify the W/R signal and generate appropriate MRDC and MWTC signals, we use a flip-flop to generate the W/R signal. Then we use a 2 line-to-1 line multiplexer to generate the memory and I/O control signals. See Figure 16-5 for a circuit that generates the memory and I/O control signals for the Pentium microprocessor.

During T2, the data bus is sampled in synchronization with the end of T2 at the positive transition of the clock pulse. The setup time before the clock is given as 3.8 ns, and the hold time after the clock is given as 2.0 ns. This means that the data window around the clock is 5.8 ns. The address appears on the 8.0 ns maximum after the




FIGURE 16-4 The non-pipelined read cycle for the Pentium microprocessor.




FIGURE 16-5 A circuit that generates the memory and I/O control signals.

start of T1. This means that the Pentium microprocessor operating at 66 MHz allows 30.3 ns (two clocking periods), minus the address delay time of 8.0 ns and minus the data setup time of 3.8 ns. Memory access time without any wait states is $30.3 - 8.0 - 3.8$, or 18.5 ns. This is enough time to allow access to a SRAM, but not to any DRAM without inserting wait states into the timing. The SRAM is normally found in the form of an external level 2 cache.




FIGURE 16-6 The Pentium timing diagram with four wait states inserted for an access time of 79.5 ns.

Wait states are inserted into the timing by controlling the BRDY input to the Pentium. The BRDY signal must become a logic 0 by the end of T2 or additional T2 states are inserted into the timing. See Figure 16-6 for a read cycle timing diagram that contains wait states for slower memory. The effect of inserting wait states into the timing is to lengthen the timing, allowing additional time to the memory to access data. In the timing shown, the access time has been lengthened so that standard 60 ns DRAM can be used in a system. Note that this requires the insertion of four wait states of 15.2 ns (one clocking period) each to lengthen the access time to 79.5 ns. This is enough time for the DRAM and any decoder in the system to function.

The BRDY signal is a synchronous signal generated by using the system clock. Figure 16-7 illustrates a circuit that can be used to generate BRDY for inserting any number of wait states into the Pentium timing diagram. You may recall a similar circuit inserting wait states into the timing diagram of the 80386 microprocessor. The ADS signal is delayed between 0 and 7 clocking periods by the 74F161 shift register to generate the BRDY signal. The exact number of wait states is selected by the 74F151 8 line-to-1 line multiplexer. In this example, the multiplexer selects the 4-wait output from the shift register.

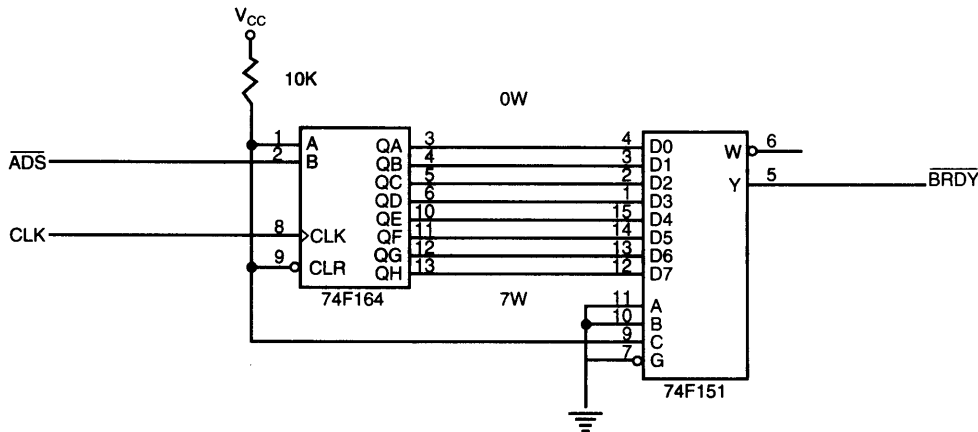


FIGURE 16-7 A circuit that generates wait states by delaying ADS. This circuit is wired to generate four wait states.




FIGURE 16-8 The Pentium burst cycle operation that transfers four 64-bit data between the microprocessor and memory.

A more efficient method of reading memory data is via the burst cycle. The burst cycle in the Pentium transfers four 64-bit numbers per burst cycle in five clocking periods. A burst without wait states requires that the memory system transfers data every 15.2 ns. If a level 2 cache is in place, this speed is no problem as long as the data are read from the cache. If the cache does not contain the data, then wait states must be inserted, which will reduce the data throughput. See Figure 16-8 for the Pentium burst cycle transfer without wait states. As before, wait states can be inserted to allow more time to the memory system for accesses.

Branch Prediction Logic

The Pentium microprocessor uses a branch prediction logic to reduce the time required for a branch caused by internal delays. These delays are minimized because when a branch instruction (short or near only) is encountered, the microprocessor begins pre-fetch instruction at the branch address. The instructions are loaded into the instruction cache, so when the branch occurs, the instructions are present and allow the branch to execute in one clocking period. If for any reason the branch prediction logic errs, the branch requires an extra three clocking periods to execute. In most cases, the branch prediction is correct and no delay ensues.

Cache Structure

The cache in the Pentium has been changed from the one found in the 80486 microprocessor. The Pentium contains two 8K-byte cache memories instead of one as in the 80486. There is an 8K-byte data cache and an 8K-byte instruction cache. The instruction cache stores only instructions, while the data cache stores data used by instructions.

In the 80486 with its unified cache, a program that was data-intensive quickly filled the cache, allowing little room for instructions. This slowed the execution speed of the 80486 microprocessor. In the Pentium, this cannot occur because of the separate instruction cache.

Superscaler Architecture

The Pentium microprocessor is organized with three execution units. One executes floating-point instructions, and the other two (U-pipe and V-pipe) execute integer instructions. This means that it is possible to execute three instructions simultaneously. For example, the FADD ST,ST(2) instruction, MOV EAX,10H instruction, and MOV EBX,12H instruction can all execute simultaneously because none of these instructions depend on each other. The FADD ST,ST(2) instruction is executed by the coprocessor; the MOV EAX,10H is executed by the U-pipe; and

the MOV EBX,12H instruction is executed by the V-pipe. Because the floating-point unit is also used for MMX instructions, if available, the Pentium can execute two integers and one MMX instruction simultaneously.

Software should be written to take advantage of this feature by looking at the instructions in a program, and then modifying them when cases are discovered in which dependent instructions can be separated by non-dependent instructions. These changes can result in up to a 40 percent execution speed improvement in some software. Make sure that any new compiler or other application package takes advantage of this new superscaler feature of the Pentium.

16-2 SPECIAL PENTIUM REGISTERS

The Pentium is essentially the same microprocessor as the 80386 and 80486, except that some additional features and changes to the control register set have occurred. This section highlights the differences between the 80386 control register structure and the flag register.

Control Registers

Figure 16-9 shows the control register structure for the Pentium microprocessor. Note that a new control register CR4 has been added to the control register array.

This section of the text only explains the new Pentium components in the control registers. See Figure 15-14 for a description and illustration of the 80386 control registers. Following is a description of the new control bits and new control register CR4:

- CD** **Cache disable** controls the internal cache. If *CD* = 1, the cache will not fill with new data for cache misses, but it will continue to function for cache hits. If *CD* = 0, misses will cause the cache to fill with new data.
- NW** **Not write-through** selects the mode of operation for the data cache. If *NW* = 1, the data cache is inhibited from cache write-through.
- AM** **Alignment mask** enables alignment checking when set. Note that alignment checking only occurs for protected mode operation when the user is at privilege level 3.




FIGURE 16-9 The structure of the Pentium control registers.

WP	Write protect protects user level pages against supervisor level write operations. When $WP = 1$, the supervisor can write to user level segments.
NE	Numeric error enables standard numeric coprocessor error detection. If $NE = 1$, the FERR pin becomes active for a numeric coprocessor error. If $NE = 0$, any coprocessor error is ignored.
VME	Virtual mode extension enables support for the virtual interrupt flag in protected mode. If $VME = 0$, virtual interrupt support is disabled.
PVI	Protected mode virtual interrupt enables support for the virtual interrupt flag in protected mode.
TSD	Time stamp disable controls the RDTSC instruction.
DE	Debugging extension enables I/O breakpoint debugging extensions when set.
PSE	Page size extension enables 4M-byte memory pages when set.
MCE	Machine check enable enables the machine checking interrupt.

The Pentium contains new features that are controlled by CR4 and a few bits in CR0. These new features are explained in later sections of the text.

EFLAG Register

The extended flag (EFLAG) register has been changed in the Pentium microprocessor. Figure 16-10 pictures the contents of the EFLAG register. Note that four new flag bits have been added to this register to control or indicate conditions about some of the new features in the Pentium.

Following is a list of the four new flags and the function of each:

ID The **identification flag** is used to test for the CPUID instruction. If a program can set and clear the ID flag, the processor supports the CPUID instruction.


VIP **Virtual interrupt pending** indicates that a virtual interrupt is pending.

VIF **Virtual interrupt** is the image of the virtual interrupt flag IF used with VIP.

AC **Alignment check** indicates the state of the AM bit in control register 0.

Built-In Self-Test (BIST)

The built-in self-test (BIST) is accessed on power-up by placing a logic 1 on INIT while the RESET pin changes from 1 to 0. The BIST tests 70 percent of the internal structure of the Pentium in approximately 150 μ s. Upon completion of the BIST, the Pentium reports the outcome in register EAX. If $EAX = 0$, the BIST passed and the Pentium is ready for operation. If EAX contains any other value, the Pentium has malfunctioned and is faulty.



Note: The blank bits in the flag register are reserved for future use and must not be defined.

FIGURE 16-10 The structure of the Pentium EFLAG register.

16-3 PENTIUM MEMORY MANAGEMENT

The memory-management unit within the Pentium is upward-compatible with the 80386 and 80486 microprocessors. Many of the features of these earlier microprocessors are basically unchanged in the Pentium. The main change is in the paging unit and a new system memory-management mode.

Paging Unit

The paging mechanism functions with 4K-byte memory pages or with a new extension available to the Pentium with 4M byte-memory pages. As detailed in Chapters 1 and 15, the size of the paging table structure can become large in a system that contains a large memory. Recall that to fully repage 4G bytes of memory, the microprocessor requires slightly over 4M bytes of memory just for the page tables. In the Pentium, with the new 4M-byte paging feature, this is dramatically reduced to just a single page table. The new 4M-byte page sizes are selected by the PSE bit in control register 0.

The main difference between 4K paging and 4M paging is that in the 4M paging scheme there is no page table entry in the linear address. See Figure 16-11 for the 4M paging system in the Pentium microprocessor. Pay close attention to the way the linear address is used with this scheme. Notice that the leftmost 10 bits of the linear address select an entry in the page directory (just as with 4K pages). Unlike 4K pages, there are no page tables; instead, the page directory addresses a 4M-byte memory page.

Memory-Management Mode

The system memory-management mode (SMM) is on the same level as protected mode, real mode, and virtual mode, but it is provided to function as a manager. The SMM is not intended to be used as an application or a system-level feature. It is intended for high-level system functions such as power management and security, which most Pentiums use during operation.




FIGURE 16-11 The linear address 00200001H repaged to memory location 01000002H in 4M-byte pages. Note that there are no page tables.

Access to the SMM is accomplished via a new external hardware interrupt applied to the SMI pin on the Pentium. When the SMM interrupt is activated, the processor begins executing system-level software in an area of memory called the *system management RAM*, or *SMMRAM*, called the *SMM state dump record*. The SMI interrupt disables all other interrupts that are normally handled by user applications and the operating system. A return from the SMM interrupt is accomplished with a new instruction. RSM returns from the memory-management mode interrupt and returns to the interrupted program at the point of the interruption.

The SMM interrupt calls the software, initially stored at memory location 38000H, using *CS* = 3000H and *EIP* = 8000H. This initial state can be changed using a jump to any location within the first 1M byte of memory. An environment similar to real-mode memory addressing is entered by the management mode interrupt, but it is different because, instead of being able to address the first 1M of memory, SMM mode allows the Pentium to treat the memory system as a flat, 4G-byte system.

In addition to executing software that begins at location 38000H, the SMM interrupt also stores the state of the Pentium in what is called a *dump record*. The dump record is stored at memory locations 3FFA8H through 3FFFFH, with an area at locations 3FE00H through 3FEF7H that is reserved by Intel. The dump record allows a Pentium-based system to enter a sleep mode and reactivate at the point of program interruption. This requires that the SMMRAM be powered during the sleep period. Many laptop computers have a separate battery to power the SMMRAM for many hours during sleep mode. Table 16-2 lists the contents of the dump record.

The Halt auto restart and I/O trap restarts are used when the SMM mode is exited by the RSM instruction. These data allow the RSM instruction to return to the halt-state or return to the interrupt I/O instruction. If neither a halt nor an I/O operation is in effect upon entering the SMM mode, the RSM instruction reloads the state of the machine from the state dump and returns to the point of interruption.

The SMM mode can be used by the system before the normal operating system is placed in the memory and executed. It can also periodically be used to manage the system, provided that normal software doesn't exist at location 38000H-3FFFFH. If the system relocates the SMRAM before booting the normal operating system, it becomes available for use in addition to the normal system.

The base address of the SMM mode SMRAM is changed by modifying the value in the state dump base address register (locations 3FEF8H through 3F3FBH) after the first memory-management mode interrupt. When the first RSM instruction is executed, returning control back to the interrupted system, the new value from these locations changes the base address of the SMM interrupt for all future uses. For example, if the state dump base address is changed to 000E8000H, all subsequent SMM interrupts use locations E8000H-EFFFFH for the Pentium state dump. These locations are compatible with DOS and Windows.

TABLE 16-2 Pentium SMM state dump record.

Offset Address	Register
FFFCH	CR0
FFF8H	CR3
FFF4H	EFLAGS
FFF0H	EIP
FFECH	EDI
FFE8H	ESI
FFE4H	EBP
FFE0H	ESP
FFDCH	EBX
FFD8H	EDX
FFD4H	ECX
FFD0H	EAX
FFCCH	DR6
FFC8H	DR7
FFC4H	TR
FFC0H	LDTR
FFBCH	GS
FFB8H	FS
FFB4H	DS
FFB0H	SS
FFACH	CS
FFA8H	ES
FF04H-FFA7H	Reserved
FF02H	Halt auto restart
FF00H	I/O trap restart
FEFCH	SMM revision identifier
FEF8H	State dump base
FE00H-FEF7H	Reserved

Note: The offset addresses are initially located at base address 00003000H.

16-4 NEW PENTIUM INSTRUCTIONS

The Pentium contains only one new instruction that functions with normal system software; the remainder of the new instructions are added to control the memory-management mode feature and serializing instructions. Table 16-3 lists the new instructions added to the Pentium instruction set.

The **CMPXCHG8B** instruction is an extension of the **CMPXCHG** instruction added to the 80486 instruction set. The **CMPXCHG8B** instruction compares the number 64-bit stored in **EDX** and **EAX** with the contents of a 64-bit memory location or register pair. For example, the **CMPXCHG8B DATA1** instruction compared the eight bytes stored in memory location **DATA1** with the 64-bit number in **EDX** and **EAX**. If **DATA1** equals **EDX:EAX**, the 64-bit number stored in **ECX:EBX** is stored in memory location **DATA1**. If they are not equal, the contents of **DATA1** are stored into **EDX:EAX**. Note that the zero flag bit indicates that the contents of **EDX:EAX** were equal or not equal to **DATA1**.

The **CPUID** instruction reads the CPU identification code and other information from the Pentium. Table 16-4 shows different information returned from the **CPUID** instruction for various input values for **EAX**. To use the **CPUID**

TABLE 16-3 New Pentium instructions.

<i>Instruction</i>	<i>Function</i>
CMPXCHG8B	Compare and exchange eight bytes
CPUID	Return the CPU identification code
RDTSC	Read time stamp counter
RDMSR	Read model specific register
WRMSR	Write model specific register
RSM	Return from system management interrupt

TABLE 16-4 **CPUID** instruction execution.

<i>Input Value (EAX)</i>	<i>Result after CPUID Executes</i>
0	EAX = 1 for all microprocessors EBX-EDX-ECX = vendor identification
1	EAX (bits 3-0) = Stepping ID EAX (bits 7-4) = Model EAX (bits 11-8) = Family EAX (bits 13-12) = Type EAX (bits 31-14) = Reserved EDX (bit 0) = CPU contains FPU EDX (bit 1) = Enhanced 8086 virtual mode supported EDX (bit 2) = I/O breakpoints supported EDX (bit 3) = Page size extensions supported EDX (bit 4) = Time stamp counter TSC supported EDX (bit 5) = Pentium-style MSR supported EDX (bit 6) = Reserved EDX (bit 7) = Machine check exception supported EDX (bit 8) = CMPXCHG8B supported EDX (bit 9) = 3.3 V microprocessor EDX (bits 10-31) = Reserved

instruction, first load EAX with the input value and then execute CPUID. The information is returned in the registers indicated in the table.

If a 0 is placed in EAX before executing the CPUID instruction, the microprocessor returns the vendor identification in EBX, EDX, and ECX. For example, the Intel Pentium returns “GenuineIntel” in ASCII code with the “Genu” in the EBX, “ineI” in EDX, and “ntel” in ECX. The EDX register returns information if EAX is loaded with a 1 before executing the CPUID instruction.

Example 16-1 illustrates a short program that reads the vendor information with the CPUID instruction. It then displays it on the video screen using the DISP macro. Note that this program works with a Pentium or any of the other Pentium clones that are on the market. It also works with the Pentium Pro microprocessor and later versions of the 80486 microprocessor.

EXAMPLE 16-1

```

                                .MODEL TINY
                                .586                                ;select the Pentium
0000                                .CODE
                                DISP MACRO                            ;;display character macro
                                    MOV AH, 2
                                    MOV DL, BL
                                    INT 21H
                                    SHR EBX, 8
                                    ENDM
                                .STARTUP
0100 66| B8 00000000 MOV EAX, 0
0106 0F A2 CPUID ;get ID from Pentium
0108 66| 52 PUSH EDX
                                DISP ;display first 4 letters
                                DISP
                                DISP
                                DISP
0132 66| 5B POP EBX ;display next 4 letters
                                DISP
                                DISP
                                DISP
                                DISP
015C 66| 8B D9 MOV EBX, ECX ;display last 4 letters
                                DISP
                                DISP
                                DISP
                                DISP
                                .EXIT
                                END

```

The RDTSC instruction reads the time-stamp counter into EDX:EAX. The time-stamp counter counts CPU clocks from the time the microprocessor is reset, where the time stamp counter is initialized to an unknown count. Because this is a 64-bit count, a 100 MHz microprocessor can accumulate a count of over 5800 years before the time-stamp counter rolls over. This instruction functions only in real mode or privilege level 0 in protected mode. If you are using a DOS shell from Windows or operating with a memory manager, this instruction will not function and will cause a general protection error. Windows uses privilege level 0 and operates with level 0 as a protected level.

Example 16-2 shows a macro sequence that times events to the microsecond. It makes use of the Pentium time stamp to time events and returns the elapsed time in microseconds in the EAX register. Note that there are three parameters associated with the macro. The first parameter passes the location of a quadword memory location used to store an image of the time stamp clock in the memory system. Make sure that this is defined with the DQ directive. The second parameter passes the clock frequency of the Pentium in MHz to the macro and determines whether it is to start the clock or return the elapsed time. If the final parameter is START, the clock is

started; if it is READ, the elapsed time is returned in EAX in microseconds. Note that this macro can be used only in real mode or privilege level 0 of protected mode.

EXAMPLE 16-2

```

EVENT MACRO WHERE, SPEED, OPER
    PUSH    EDX                                ;save registers
    PUSH    ECX
    IFIDN  <OPER>, <START>
        RDTSC
        MOV    DWORD PTR WHERE, EAX          ;;save current count
        MOV    DWORD PTR WHERE+4, EDX
    ENDIF
    IFIDN  <OPER>, <READ>
        RDTSC
        SUB    EAX, DWORD PTR WHERE          ;;form difference
        SBB    EDX, DWORD PTR WHERE+4
        MOV    ECX, SPEED                    ;;convert to microseconds
        DIV    ECX
    ENDIF
    POP     ECX
    POP     EDX
ENDM

```

The RDMSR and WRMSR instructions allow the model-specific registers to be read or written. The model-specific registers are unique to the Pentium and are used to trace, check performance, test, and check for machine errors. Both instructions use ECX to convey the register number to the microprocessor and use EDX:EAX for the 64-bit wide read or write. Note that the register addresses are 0H–13H. See Table 16-5 for a list of the Pentium model-specific registers and their contents. As with the RDTSC instruction, these model-specific registers operate only in the real or privilege level 0 of protected mode.

TABLE 16-5 The Pentium model-specific registers.

<i>Address (ECX)</i>	<i>Size</i>	<i>Function</i>
00H	64-bits	Machine check exception address
01H	5-bits	Machine check exception type
02H	14-bits	TR1 parity reversal test register
03H	—	—
04H	4-bits	TR2 instruction cache end bits
05H	32-bits	TR3 cache data
06H	32-bits	TR4 cache tag
07H	15-bits	TR4 cache control
08H	32-bits	TR6 TLB command
09H	32-bits	TR7 TLB data
0AH	—	—
0BH	32-bits	TR9 BTB tag
0CH	32-bits	TR10 BTB target
0DH	12-bits	TR11 BTB control
0EH	10-bits	TR12 new feature control
0FH	—	—
10H	64-bits	Time stamp counter (can be written)
11H	26-bits	Events counter selection and control
12H	40-bits	Events counter 0
13H	40-bits	Events counter 1

Never use an undefined value in ECX before using the RDMSR or WRMSR instructions. If $ECX = 0$ before the read or write machine-specific register instruction, the value returned EDX:EAX is the machine check exception address. If $ECX = 1$, the value is the machine check exception type; if $ECX = 0EH$, the test register 12 (TR12) is accessed. Note that these are internal registers designed for in-house testing. The contents of these registers are proprietary to Intel and should not be used during normal programming.

The RMS instruction returns from a memory-management mode interrupt. The memory-management mode interrupt is explained in Section 16-3.

16-5 INTRODUCTION TO THE PENTIUM PRO MICROPROCESSOR

Before this or any other microprocessor can be used in a system, the function of each pin must be understood. This section of the chapter details the operation of each pin, along with the external memory system and I/O structures of the Pentium Pro microprocessor.

Figure 16-12 illustrates the pin-out of the Pentium Pro microprocessor, which is packaged in an immense 387-pin PGA (pin grid array). Currently, the Pentium Pro is available in two versions: one version contains a 256K level 2 cache; the other contains a 512K level 2 cache. The notable difference in the pin-out of the Pentium Pro when compared to earlier Pentiums is that there are provisions for a 36-bit address bus, which allows access to 64G bytes of memory. This is meant for future use because no system today contains anywhere near that amount of memory.

As with most recent versions of the Pentium microprocessor, the Pentium Pro requires a single +3.3 V or +2.7 V power supply for operation. The power supply current is a maximum of 9.9 A for the 150 MHz version of the Pentium Pro, which also has a maximum power dissipation of 26.7 W. At present, a good heat sink with considerable airflow is required to keep the Pentium Pro cool. As with the Pentium, the Pentium Pro contains multiple VCC and VSS connections that must all be connected for proper operation. The Pentium Pro contains VCCP pins (primary VCC) that connect to +3.1 V, VCCS (secondary VCC) pins that connect to +3.3 V, and VCC5 (standard VCC) pins that connect to +5.0 V. There are some pins labeled N/C (no connection) and must not be connected.

Each Pentium Pro output pin is capable of providing an ample 48.0 mA of current at a logic 0 level. This represents a considerable increase in drive current, compared to the 2.0 mA available on earlier microprocessor output pins. Each input pin represents a small load, requiring only 15 μ A of current. Because of the 48.0 mA of drive current available on each output, only an extremely large system requires bus buffers.

Internal Structure of the Pentium Pro

The Pentium Pro is structured differently than earlier microprocessors. Early microprocessors contained an execution unit and a bus interface unit with a small cache buffering the execution unit for the bus interface unit. This structure was modified in later microprocessors, but the modifications were just additional stages within the microprocessors. The Pentium architecture is also a modification, but more significant than earlier microprocessors. Figure 16-13 shows a block diagram of the internal structure of the Pentium Pro microprocessor.

The system buses, which communicate to the memory and I/O, connect to an internal level 2 cache that is often on the main board in most other microprocessor systems. The level 2 cache in the Pentium Pro is either 256K bytes or 512K bytes. The integration of the level 2 cache speeds processing and reduces the number of components in a system.

The bus interface unit (BIU) controls the access to the system buses through the level 2 cache, as it does in most other microprocessors. Again, the difference is that the level 2 cache is integrated. The BIU generates the memory address and control signals, and passes and fetches data or instructions to either a level 1 data cache or a level 1 instruction cache. Each of these are 8K bytes in size at present and may be made larger in future versions of the microprocessor. Earlier versions of the Intel microprocessor contained a unified cache that held both instructions and data. The implementation of separate caches improves performance because data-intensive programs no longer fill the cache with data.

The instruction cache is connected to the instruction fetch and decode unit (IFDU). Although not shown, the IFDU contains three separate instruction decoders that decode three instructions simultaneously. Once decoded, the outputs of the three decoders are passed to the instruction pool, where they remain until the dispatch and execution unit or retire unit obtains them. Also included within the IFDU is a branch prediction logic section that looks ahead in code sequences that contain conditional jump instructions. If a conditional jump is located, the branch prediction logic tries to determine the next instruction in the flow of a program.

Once decoded instructions are passed to the instruction pool, they are held for processing. The instruction pool is a content-addressable memory, but Intel never states its size in the literature.

The dispatch and execute unit (DEU) retrieves decoded instructions from the instruction pool when they are complete, and then executes them. The internal structure of the DEU is illustrated in Figure 16-14. Notice that the DEU contains three instruction execution units: two for processing integer instructions and one for floating-point instructions. This means that the Pentium Pro can process two integer instructions and one floating-point instruction simultaneously. The Pentium also contains three execution units, but the architecture is different because the Pentium does not contain a jump execution unit or address generation units, as does the Pentium Pro. The reservation station (RS) can schedule up to five events for execution and process four simultaneously.

The last internal structure of the Pentium Pro is the retire unit (RU). The RU checks the instruction pool and removes decoded instructions that have been executed. The RU can remove three decoded instructions per clock pulse.

Pin Connections

The number of pins on the Pentium Pro has increased from the 237 pins on the Pentium to 387 pins on the Pentium Pro. Following is a description of each pin or grouping of pins:

A20M

The **address A20 mask** is an input that is asserted in the real mode to signal the Pentium Pro to perform address wraparound, as in the 8086 microprocessor, for use of the HIMEM.SYS driver.

A35-A3

Address bus connections address any of the 1G × 64 memory locations found in the Pentium Pro memory system.

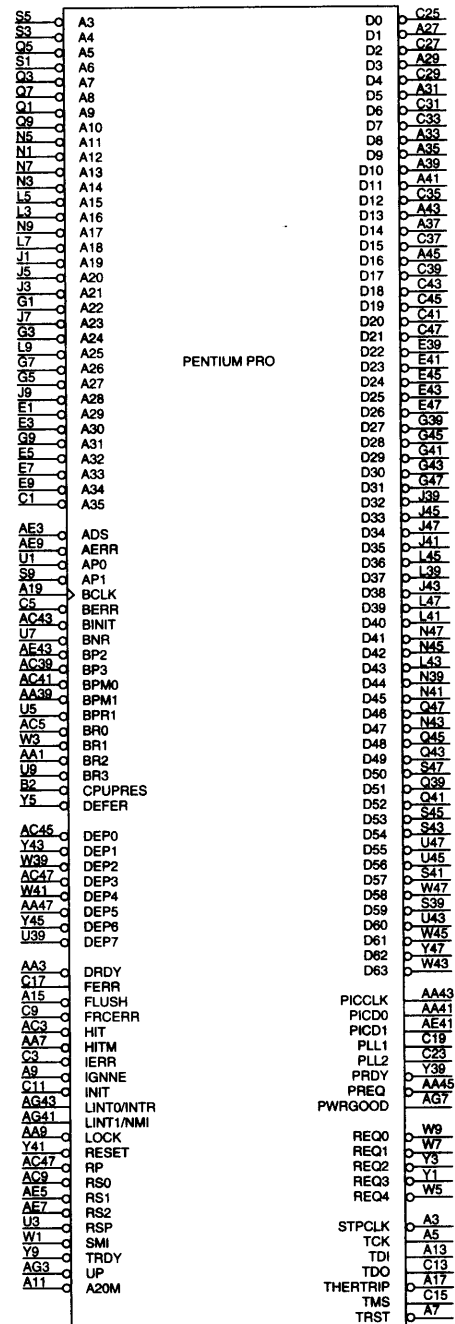


FIGURE 16-12 The pin-out of the Pentium Pro microprocessor.


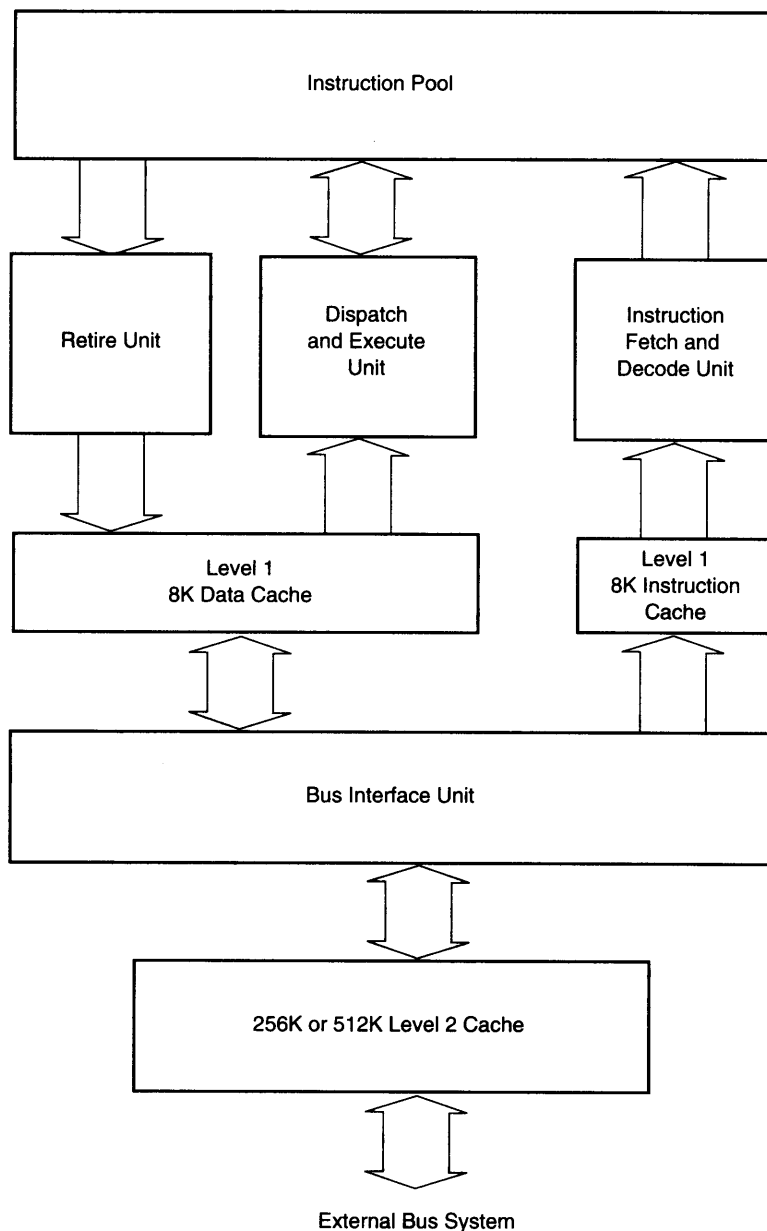


FIGURE 16-13 The internal structure of the Pentium Pro microprocessor.

ADS

The **address data strobe** becomes active whenever the Pentium Pro has issued a valid memory or I/O address.

AP0, AP1

Address parity provides even parity for the memory address on all Pentium Pro-initiated memory and I/O transfers. The AP0 output provides parity for address connections A23-A3, and the AP1 output provides parity for address connections A35-A24.




FIGURE 16-14 The Pentium Pro dispatch and execution unit (DEU).

ASZ0,ASZ1

Address size inputs are driven to select the size of the memory access. Table 16-6 illustrates the size of the memory access for the binary bit patterns on these two inputs to the Pentium Pro.

TABLE 16-6 The memory size as dictated by ASZ pins.

ASZ1	ASZ0	Memory Size
0	0	0-4G
0	1	4G-64G
1	X	Reserved

BCLK

The bus clock input determines the operating frequency of the Pentium Pro microprocessor. For example, if BCLK is 66 MHz, various internal clocking speeds are selected by the logic levels applied to the pins in Table 16-7. A BCLK frequency of 66 MHz runs the system bus at 66 MHz.

TABLE 16-7 The BCLK signal and its effect on the Pentium Pro clock speed.

LINT1/NMI	LINT0/INT R	IGNNE	A20M	Ratio	Speed with BCLK = 50 MHz	Speed with BCLK = 66 MHz
0	0	0	0	2	100 MHz	133 MHz
0	0	0	1	4	200 MHz	266 MHz
0	0	1	0	3	150 MHz	200 MHz
0	0	1	1	5	250 MHz	333 MHz
0	1	0	0	5/2	125 MHz	166 MHz
0	1	0	1	9/2	225 MHz	300 MHz
0	1	1	0	7/2	175 MHz	233 MHz
0	1	1	1	11/2	275 MHz	366 MHz
1	1	1	1	2	100 MHz	133 MHz

BERR	The bus error input/output either signals a bus error along or is asserted by an external device to cause a machine check interrupt or a non-maskable interrupt.
BINIT	Bus initialization is active on power-up to initialize the bus system.
BNR	Block next request is used to halt the system in a multiple microprocessor system.
BP3, BP2	The break point status outputs indicate the status of the Pentium Pro break points.
BPM1, BPM0	The break point monitor outputs indicate the status of the breakpoints and programmable counters.
BPRI	The priority agent bus request is an input that causes the microprocessor to cease bus requests.
BR3–BR0	The bus request inputs allow up to four Pentium Pro microprocessors to coexist on the same bus system.
BREQ3–BREQ0	Bus request signals are used for multiple microprocessors on the same system bus.
D63–D0	Data bus connections transfer byte, word, doubleword, and quadword data between the microprocessor and its memory and I/O system.
DBSY	Data bus busy is asserted to indicate that the data bus is busy transferring data.
DEFER	The defer input is asserted during the snoop phase to indicate that the transaction cannot be guaranteed in-order completion.
DEN	The defer enable signal is driven to the bus on the second phase of a request phase.
DEP7–DEP0	Data bus ECC protection signals provide error-correction codes for correcting a single-bit error and detecting a double-bit error.
FERR	The floating-point error , comparable to the ERROR line in the 80386, shows that the internal co-processor has erred.
FLUSH	The flush cache input causes the cache to flush all write-back lines and invalidate its internal caches. If the FLUSH input is a logic 0 during a reset operation, the Pentium enters its test mode.
FRCERR	Functional redundancy check error is used if two Pentium Pro microprocessors are configured in a pair.
HIT	Hit shows that the internal cache contains valid data in the inquire mode.
HITM	Hit modified shows that the inquire cycle found a modified cache line. This output is used to inhibit other master units from accessing data until the cache line is written to memory.
IERR	Internal error output shows that the Pentium Pro has detected an internal parity error or functional redundancy error.
IGNNE	The ignore numeric error input causes the Pentium Pro to ignore a numeric coprocessor error.
INIT	The initialization input performs a reset without initializing the caches, write-back buffers, and floating-point registers. This input may not be used to reset the microprocessor in lieu of RESET after power-up.
INTR	The interrupt request is used by external circuitry to request an interrupt.
LEN	Length signals (bit 0 and 1) indicate the size of the data transfer, as illustrated in Table 16–8.

TABLE 16–8 The LEN bits show the size of a data transfer.

<i>LEN1</i>	<i>LEN0</i>	<i>Data Transfer Size</i>
0	0	0–8 bytes
0	1	16 bytes
1	0	32 bytes
1	1	Reserved

TABLE 16–9 The function of the request signals in the first clocking period.

REQ4	REQ3	REQ2	REQ1	REQ0	Function
0	0	0	0	0	Deferred reply
0	0	0	0	1	Reserved
0	1	0	0	0	Case 1*
0	1	0	0	1	Case 2*
1	0	0	0	0	I/O read
1	0	0	0	1	I/O write
X	X	0	1	0	Memory read
X	X	0	1	1	Memory write
X	X	1	0	0	Memory code read
X	X	1	1	0	Memory data read
X	X	1	X	1	Memory write

*Note: See Table 18–10 for the second clock pulse for these codes.

TABLE 16–10 The second clock pulse and the request signals as they apply to case 1 and 2 from Table 18–9.

Case	REQ4	REQ3	REQ2	REQ1	REQ0	Function
1	X	X	X	0	0	Interrupt acknowledge
1	X	X	X	0	1	Special transactions
1	X	X	X	1	X	Reserved
2	X	X	X	0	0	Branch trace message
2	X	X	X	0	1	Reserved
2	X	X	X	1	X	Reserved

- LINT** The **local interrupt** inputs function as NMI and INTR, and also set the clock divider frequency on reset.
- LOCK** **LOCK** becomes a logic 0 whenever an instruction is prefixed with the LOCK: prefix. This is most often used during DMA accesses.
- NMI** The **non-maskable interrupt** requests a non-maskable interrupt, as it did on the earlier versions of the microprocessor.
- PICCLK** The **clock signal** input is used for synchronous data transfers.
- PICD** The **processor interface serial data** is used to transfer bi-directional serial messages between Pentium Pro microprocessors.
- PWRGOOD** **Power good** is an input that is placed at a logic 1 level when the power supply and clock have stabilized.
- REQ** **Request** signals (bits 0–4) define the type of data-transfer operation, as illustrated in Tables 16–9 and 16–10.
- RESET** **Reset** initializes the Pentium Pro, causing it to begin executing software at memory location FFFFFFF0H. The Pentium Pro is reset to the real mode and the leftmost 12 address connections remain logic 1s (FFFH) until a far jump or far call is executed. This allows compatibility with earlier microprocessors.

TABLE 16-11 The operation of the Pentium Pro in response to the RS inputs.

<i>RS2</i>	<i>RS1</i>	<i>RS0</i>	<i>Function</i>	<i>HITM</i>	<i>DEFER</i>
0	0	0	Idle state	X	X
0	0	1	Retry	0	1
0	1	0	Defer	0	1
0	1	1	Reserved	0	1
1	0	0	Hard failure	X	X
1	0	1	Normal, no data	0	0
1	1	0	Implicit write-back	1	X
1	1	1	Normal with data	0	0

RP	Request parity provides a means of requesting that the Pentium Pro checks parity.
RS	The response status inputs cause the Pentium Pro to perform the functions listed in Table 16-11.
RSP	The response parity input applies a parity error signal from an external parity checker.
SMI	The system management interrupt input causes the Pentium Pro to enter the system management mode of operation.
SMMEM	The system memory-management mode signal becomes a logic 0 whenever the Pentium Pro is executing in the system memory-management mode interrupt and address space.
SPCLK	The split lock signal is placed at a logic 0 level to indicate that the transfer will contain four locked transactions.
STPCLK	Stop clock causes the Pentium Pro to enter the power-down state when placed at a logic 0 level.
TCK	The testability clock input selects the clocking function in accordance with the IEEE 1149.1 Boundary Scan interface.
TDI	The test data input is used to test data clocked into the Pentium with the TCK signal.
TDO	The test data output is used to gather test data and instructions shifted out of the Pentium with TCK.
TMS	The test mode select input controls the operation of the Pentium in test mode.
TRDY	The target ready input is asserted when the target is ready for a data transfer operation.

The Memory System

The memory system for the Pentium Pro microprocessor is 4G bytes in size, just as in the 80386DX-Pentium microprocessors, but access to an area between 4G and 64G is made possible by additional address signals A32-A35. The Pentium Pro uses a 64-bit data bus to address memory organized in eight banks that each contain 8G bytes of data. Note that the additional memory is enabled with bit position 5 of CR4 and is accessible only when 2M paging is enabled. Note also that 2M paging is new to the Pentium Pro to allow memory above 4G to be accessed. More information is presented on Pentium Pro paging later in this chapter. Refer to Figure 16-15 for the organization of the Pentium Pro physical memory system.

The Pentium Pro memory system is divided into eight banks that each store a byte of data with a parity bit. Note that most Pentium and Pentium Pro microprocessor-based systems forgo the use of the parity bit. The




FIGURE 16–15 The eight memory banks in the Pentium Pro system. Note each bank is 8-bits wide and 8G long if 36-bit addressing is enabled.

Pentium Pro, like the 80486 and Pentium, employs internal parity generation and checking logic for the memory system data bus information. The 64-bit wide memory is important to double-precision floating-point data. Recall that a double-precision floating-point number is 64 bits wide. As with earlier Intel microprocessors, the memory system is numbered in bytes from byte 00000000H to byte FFFFFFFFH. This nine-digit hexadecimal address is employed in a system that addresses 64G of memory.

Memory selection is accomplished with the bank enable signals (BE7–BE0). In the Pentium Pro microprocessor, the bank enable signals are presented on the address bus (A15–A8) during the second clock cycle of a memory or I/O access. These must be extracted from the address bus to access memory banks. The separate memory banks allow the Pentium Pro to access any single byte, word, doubleword, or quadword with one memory transfer cycle. As with earlier memory selection logic, we often generate eight separate write strobes for writing to the memory system. Note that the memory write information is provided on the request lines from the microprocessor during the second clock phase of a memory or I/O access.

A new feature added to the Pentium and Pentium Pro is the capability to check and generate parity for the address bus during certain operations. The AP pin (Pentium) or pins (Pentium Pro) provides the system with parity information, and the APCHK (Pentium) or AP pins (Pentium Pro) indicate a bad parity check for the address bus. The Pentium Pro takes no action when an address-parity error is detected. The error must be assessed by the system, and the system must take appropriate action (an interrupt) if so desired.

New to the Pentium Pro is a built-in error correction circuit (ECC) that allows the correction of a one-bit error and the detection of a two-bit error. To accomplish the detection and correction of errors, the memory system must have room for an extra 8-bit number that is stored with each 64-bit number. The extra eight bits are used to store an error-correction code that allows the Pentium Pro to automatically correct any single-bit error. A 1M × 64 is a 64M SDRAM without ECC, and a 1M × 72 is an SDRAM with ECC support. The ECC code is much more reliable than the old parity scheme, which is rarely used in modern systems. The only drawback of the ECC scheme is the additional cost of SDRAM that is 72-bits wide.

Input/Output System

The input/output system of the Pentium Pro is completely compatible with earlier Intel microprocessors. The I/O port number appears on address lines A15–A3 with the bank-enable signals used to select the actual memory banks used for the I/O transfer.

Beginning with the 80386 microprocessor, I/O privilege information is added to the TSS segment when the Pentium is operated in the protected mode. Recall that this allows I/O ports to be selectively inhibited. If the blocked I/O location is accessed, the Pentium Pro generates a type-13 interrupt to signal an I/O privilege violation.

System Timing

As with any microprocessor, the system timing signals must be understood in order to interface the microprocessor. This portion of the text details the operation of the Pentium Pro through its timing diagrams and shows how to determine memory access times.

The basic Pentium Pro memory cycle consists of two sections: the address phase and the data phase. During the address phase, the Pentium Pro sends the address (T1) to the memory and I/O system, and also the control signals (T2). The control signals include the ATTR lines (A31–A24), the DID lines (A23–A16), the bank enable signals (A15–A8), and the EXF lines (A7–A3). See Figure 16–16 for the basic timing cycle. The type of memory cycle appears on the request pins. During the data phase, four 64-bit wide numbers are fetched or written to the memory. This operation is most common because data from the main memory are transferred between the internal 256K or 512K write-back cache and the memory system. Operations that write a byte, word, or doubleword, such as I/O transfers, use the bank selection signals and have only one clock in the data transfer phase. Notice from the timing diagram that the 66 MHz Pentium Pro is capable of 33 million memory transfers per second. (This assumes that the memory can operate at that speed.)

The setup time before the clock is given as 5.0 ns and the hold time after the clock is given as 1.5 ns. This means that the data window around the clock is 6.5 ns. The address appears on the 8.0 ns maximum after the start of T1. This means that the Pentium Pro microprocessor operating at 66 MHz allows 30 ns (two clocking periods), minus the address delay time of 8.0 ns and also minus the data setup time of 5.0 ns. Memory access time without any wait states is $30 - 8.0 - 5.0$, or 17.0 ns. This is enough time to allow access to a SRAM, but not to any DRAM without inserting wait states into the timing.

Wait states are inserted into the timing by controlling the TRDY input to the Pentium Pro. The TRDY signal must become a logic 0 by the end of T2; otherwise, additional T2 states are inserted into the timing. Note that 60 ns DRAM requires the insertion of four wait states of 15 ns (one clocking period) each to lengthen the access time to 77 ns. This is enough time for the DRAM and any decoder in the system to function. Because many EPROM memory devices require an access time of 100 ns, EPROM requires the addition of seven wait states to lengthen the access time to 122 ns.




FIGURE 16–16 The basic Pentium Pro timing.

3. The data bus on the Pentium is 64-bits wide and contains eight byte-wide memory banks selected with bank enable signals (BE0–BE7).
4. Memory access time, without wait states, is only about 18 ns in the 66 MHz Pentium. In many cases, this short access time requires wait states that are introduced by controlling the BRDY input to the Pentium.
5. The superscaler structure of the Pentium contains three independent processing units: a floating-point processor and two integer processing units labeled U and V by Intel.
6. The cache structure of the Pentium is modified to include two caches. One 8K × 8 cache is designed as an instruction cache; the other 8K × 8 cache is a data cache. The data cache can be operated as either a write-through or a write-back cache.
7. A new mode of operation called the *system memory-management* (SMM) mode has been added to the Pentium. The SMM mode is accessed via the system memory-management interrupt applied to the SMI input pin. In response to SMI, the Pentium begins executing software at memory location 38000H.
8. New instructions include the CMPXCHG8B, RSM, RDMSR, WRMSR, and CPUID. The CMPXCHG8B instruction is similar to the 80486 CMPXCHG instruction. The RSM instruction returns from the system memory-management interrupt. The RDMSR and WRMSR instructions read or write to the machine-specific registers. The CPUID instruction reads the CPU identification code from the Pentium.
9. The built-in self-test (BIST) allows the Pentium to be tested when power is first applied to the system. A normal power-up reset activates the RESET input to the Pentium. A BIST power-up reset activates INIT and then deactivates the RESET pin. EAX is equal to a 00000000H in the BIST passes.
10. A new proprietary Intel modification to the paging unit allows 4M-byte memory pages instead of the 4K-byte pages. This is accomplished by using the page directory to address 1024 pages that each contain 4M of memory.
11. The Pentium Pro is an enhanced version of the Pentium microprocessor that contains not only the level 1 caches found inside the Pentium, but also the level 2 cache of 256K or 512K found on most main boards.
12. The Pentium Pro operates by using the same 66 MHz bus speed as the Pentium and the 80486. It uses an internal clock generator to multiply the bus speed by various factors to obtain higher internal execution speeds.
13. The only significant software difference between the Pentium Pro and earlier microprocessors is the addition of the FCMOV and CMOV instructions.
14. The only hardware difference between the Pentium Pro and earlier microprocessors is the addition of 2M paging and four extra address lines that allow access to a memory address space of 64G bytes.
15. Error correction code has been added to the Pentium Pro, which corrects any single-bit error and detects any two-bit error.

16-8 QUESTIONS AND PROBLEMS

1. How much memory is accessible to the Pentium microprocessor?
2. How much memory is accessible to the Pentium Pro microprocessor?
3. The memory data bus width is _____ in the Pentium.
4. What is the purpose of the DP0–DP7 pins on the Pentium?
5. If the Pentium operates at 66 MHz, what frequency clock signal is applied to the CLK pin?
6. What is the purpose of the BRDY pin on the Pentium?
7. What is the purpose of the AP pin on the Pentium?
8. How much memory access time is allowed by the Pentium, without wait states, when it is operated at 66 MHz?
9. What Pentium pin is used to insert wait states into the timing?
10. A wait state is an extra ___ clocking period.
11. Explain how two integer units allow the Pentium to execute two non-dependent instructions simultaneously.

12. How many caches are found in the Pentium and what are their sizes?
13. How wide is the Pentium memory data sample window for a memory read operation?
14. Can the Pentium execute three instructions simultaneously?
15. What is the purpose of the SMI pin?
16. What is the system memory-management mode of operation for the Pentium?
17. How is the system memory-management mode exited?
18. Where does the Pentium begin to execute software for an SMI interrupt input?
19. How can the system memory-management unit dump address be modified?
20. Explain the operation of the CMPXCHG8B instruction.
21. What information is returned in register EAX after the CPUID instruction executes with an initial value of 0 in EAX?
22. What new flag bits are added to the Pentium microprocessor?
23. What new control register is added to the Pentium microprocessor?
24. Describe how the Pentium accesses 4M pages.
25. Explain how the time-stamp clock functions and how it can be used to time events.
26. Contrast the Pentium with the Pentium Pro microprocessor.
27. Where are the bank enable signals found in the Pentium Pro microprocessor?
28. How many address lines are found in the Pentium Pro system?
29. What changes have been made to CR4 in the Pentium Pro and for what purpose?
30. Compare access times in the Pentium system with the Pentium Pro system.
31. What is ECC?
32. What type of SDRAM must be purchased to use ECC?

CHAPTER 17

The Pentium II, Pentium III, and Pentium 4 Microprocessors

INTRODUCTION

The Pentium II, Pentium III, and Pentium 4 microprocessors may well signal the end to the evolution of the 32-bit architecture with the advent of the Itanium¹ microprocessor from Intel. The Itanium is a 64-bit architecture microprocessor. The Pentium II, Pentium III, and Pentium 4 architectures are extensions of the Pentium Pro architecture, with some differences. The most notable difference is that the internal cache from the Pentium Pro architecture has been moved out of the microprocessor in the Pentium II. Another major change is that the Pentium II is not available in integrated circuit form. Instead, the Pentium II is found on a small plug-in circuit board along with the level 2 cache chip. Various versions of the Pentium II are available. The Celeron² is a version of the Pentium II that does not contain the level 2 cache on the Pentium II circuit board. The Xeon³ is an enhanced version of the Pentium II that contains up to a 2M-byte cache on the circuit board.

Similar to the Pentium II, early Pentium III microprocessors were packaged in a cartridge instead of an integrated circuit. More recent versions, such as the Coppermine, are again packaged in an integrated circuit (370 pins). The Pentium III Coppermine, like the Pentium Pro, contains an internal cache. The Pentium 4 is packaged in a larger integrated circuit, with 421 pins. The Pentium 4 also uses physically smaller transistors, which makes it much smaller and faster than the Pentium III. Intel to date has released versions of the Pentium 4 that operate at frequencies over 2 GHz with a limit of possibly 10 GHz at some future date.

CHAPTER OBJECTIVES

Upon completion of this chapter, you will be able to:

1. Detail the differences between the Pentium II, Pentium III, and Pentium 4 and prior Intel microprocessors.
2. Explain how the architectures of the Pentium II, Pentium III, and Pentium 4 improve system speed.
3. Explain how the basic architecture of the computer system has changed by using the Pentium II, Pentium III, and Pentium 4 microprocessors.
4. Detail the changes to the CMOV instruction.
5. Describe the operation of the SYSENTER and SYSEXIT instructions.
6. Describe the operation of the FXSAVE and FXRSTOR instructions.

¹ Itanium is a registered trademark of Intel Corporation.

² The Celeron is a registered trademark of Intel Corporation.

³ Xeon is a registered trademark of Intel Corporation.

17-1 INTRODUCTION TO THE PENTIUM II MICROPROCESSOR

Before the Pentium II or any other microprocessor can be used in a system, the function of each pin must be understood. This section of the chapter details the operation of each pin, along with the external memory system and I/O structures of the Pentium II microprocessor.

Figure 17-1 illustrates the basic outline of the Pentium II microprocessor's slot 1 connector and the signals used to interface to the chipset. Figure 17-2 shows a simplified diagram of the components on the cartridge, and the placement of the Pentium II cartridge and bus components in the typical Pentium II system. There are 242 pins on the slot 1 connector for the microprocessor. (These connections are a reduction in the number of pins found on the Pentium and the Pentium II microprocessors.) The Pentium II is packaged on a printed circuit board instead of the integrated circuits of the past Intel microprocessors. The level 1 cache is 32K-bytes as it was in the Pentium Pro, but the level 2 cache is no longer inside the integrated circuit. Intel changed the architecture so that a level 2 cache could be placed very closely to the microprocessor. This change makes the microprocessor less expensive and still allows the level 2 cache to operate efficiently. The Pentium level 2 cache operates at one-half the microprocessor clock frequency, instead of the 66 MHz of the Pentium microprocessor. A 400 MHz Pentium II has a cache speed of 200 MHz. Currently, the Pentium II is available in three versions. The first is the full-blown Pentium II, which is the Pentium II for the slot 1 connector. The second is the Celeron, which is like the Pentium II, except that the slot 1 circuit board does not contain a level 2 cache; the level 2 cache in the Celeron system is located on the main board and operates at 66 MHz. The most recent version is the Xeon, which, because it uses a level 2 cache of 512K, 1M, or 2M, represents a significant speed improvement over the Pentium II. The Xeon's level 2 cache operates at the clock frequency of the microprocessor. A 400 MHz Xeon has a level 2 cache speed of 400 MHz, which is twice the speed of the regular Pentium II.

The early versions of the Pentium II require a 5.0 V, 3.3 V and variable voltage power supply for operation. The main variable power supply voltages vary from 3.5 V to as low as 1.8 V at the microprocessor. The power-supply current averages 14.2 A to 8.4 A, depending on the operating frequency and voltage of the Pentium II. Because these currents are significant, so is the power dissipation of these microprocessors. At present, a good heat sink with considerable airflow is required to keep the Pentium II cool. Luckily, the heat sink and fan are built into the Pentium II cartridge. The latest versions of the Pentium II have been improved to reduce the power dissipation.

Each Pentium II cartridge output pin is capable of providing at least 36 mA of current at a logic 0 level on the signal connections. Some of the output control signals provide only 14 mA of current. Another change to the Pentium II is that the outputs are open-drain and require an external pull-up resistor for proper operation.

The function of each Pentium II group of pins follows:

A20	Address A20 mask is an input that is asserted in the real mode to signal the Pentium II to perform address wraparound, as in the 8086 microprocessor, for use of the HIMEM.SYS driver.
A35-A3	Address buses , which are active low connections, address any of the memory locations found in the Pentium II memory system. Note that A0, A1, and A2 are encoded in the bus enable (BE7-BE0), which are generated by the chipset, to select any or all of the eight bytes in a 64-bit wide memory location.
ADS	Address data strobe is an input that is activated to indicate to the Pentium II that the system is ready to perform a memory or I/O operation. This signal causes the microprocessor to provide the address to the system.
AERR	Address error is an input used to cause the Pentium II to check for an address parity error if it is activated.
AP1-0	Address parity inputs indicate an address parity error.
BCLK	Bus clock is an input that sets the bus clock frequency. This is either 66 MHz or 100 MHz in the Pentium II.

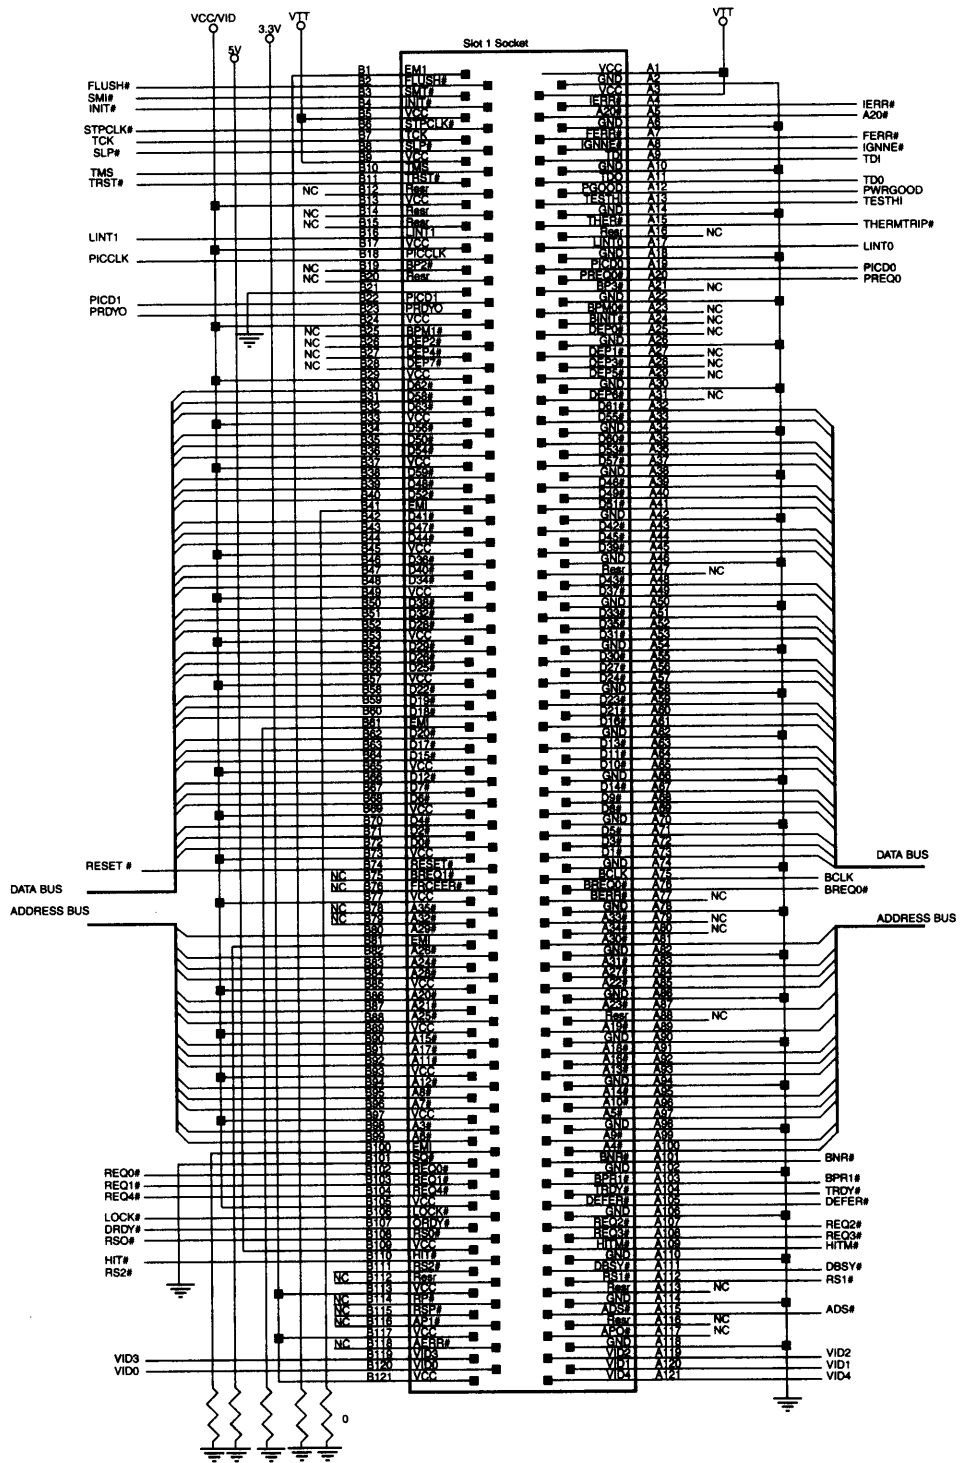



FIGURE 17-1 The pin-out of the slot 1 connector showing the connections to the system.



* The bus speed is 1/2 Pentium speed or the same as the Pentium speed in the Xeon.

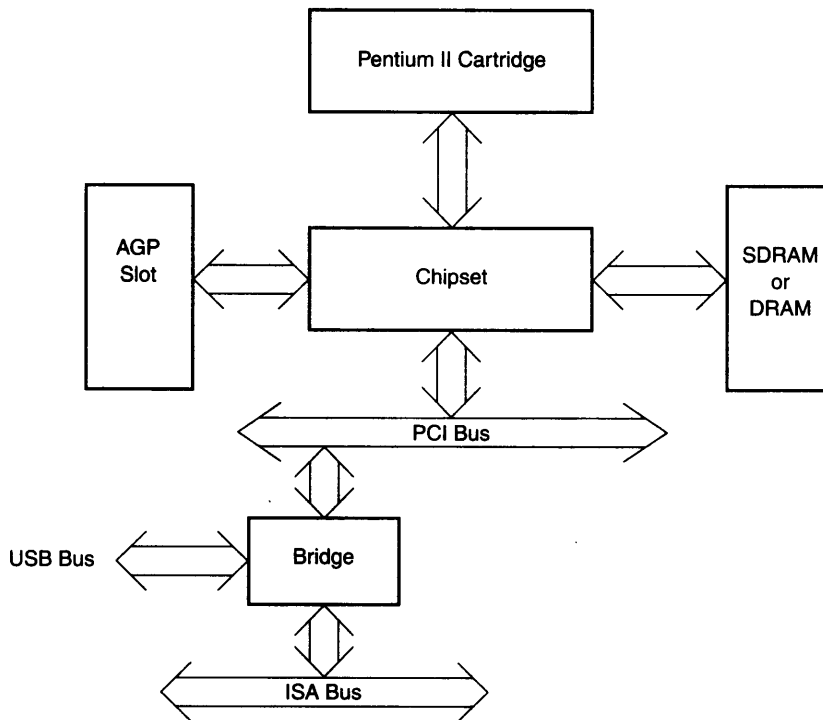


FIGURE 17-2 The structure of the Pentium II cartridge and the structure of the Pentium II system.

- BERR** **Bus error** is asserted to indicate that an error has occurred on the bus system.
- BINT** **Bus initialization** is a logic 0 during system reset or initialization. It is an input to indicate that a bus error has occurred and the system needs to be reinitialized.
- BNR** **Bus not ready** is an input used to insert wait states into the timing for the Pentium II. Placing a logic 0 in this pin causes the Pentium II to enter stall states or wait states.
- BP[3:2] and PM/BP[1:0]** The **breakpoint pins** BP3–BP0 indicate a breakpoint match when the debug registers are programmed to monitor for matches. The **performance monitoring pins** PM1 and PM0 indicate the settings of the performance monitoring bits in the debug mode control register.

BPRI	The bus priority request input is used to request the system bus from the Pentium II.
BR0-1	Bus requests indicate that the Pentium II has generated a bus request. During initialization, the BR0 pin must be activated.
BSEL	Bus select is currently not used by the Pentium II and must be connected to ground for proper operation.
D63-D0	Data bus connections transfer byte, word, doubleword, and quadword data between the microprocessor and its memory and I/O system.
DEFER	Defer is used to indicate that the external system cannot complete the bus cycle.
DEP7-DEP0	Data EEC pins are used in the error-correction scheme of the Pentium II and normally connect to an extra 8-bit memory section.
DRDY	Data ready is activated to indicate that the system is presenting valid data to the Pentium II.
EMI	Electro-magnetic interference must be grounded to prevent the Pentium II from generating or receiving noise.
FERR	Floating-point error , comparable to the ERROR line in the 80386, shows that the internal coprocessor has erred.
FLUSH	The flush cache input causes the cache to flush all write-back lines and invalidate its internal caches. If the FLUSH input is a logic 0 during a reset operation, the Pentium enters its test mode.
FRCERR	Functional redundancy check is sampled during a reset to configure the Pentium II in the master (1) or checker (0) mode.
HIT	Hit shows that the internal cache contains valid data in the inquire mode.
HITM	Hit modified shows that the inquire cycle found a modified cache line. This output is used to inhibit other master units from accessing data until the cache line is written to memory.
IERR	The internal error output shows that the Pentium II has detected an internal error or functional redundancy error.
IGNNE	The ignore numeric error input causes the Pentium II to ignore a numeric coprocessor error.
INIT	The initialization input performs a reset without initializing the caches, write-back buffers, and floating-point registers. This input may not be used to reset the microprocessor in lieu of RESET after power-up.
INTR	Interrupt request is used by external circuitry to request an interrupt.
LINT1, LINT0	Local APIC interrupt signals must connect the appropriate pins of all APIC bus agents. When the APIC is disabled, the LINT0 signal becomes INTR, a maskable interrupt request signal; LINT1 becomes NMI, a non-maskable interrupt.
LOCK	LOCK becomes a logic 0 whenever an instruction is prefixed with the LOCK: prefix. This is most often used during DMA accesses.
NMI	Non-maskable interrupt requests a non-maskable interrupt as it did on the earlier versions of the microprocessor.
PICCLK	Must be $\frac{1}{4}$ the frequency of BLCK.
PICD1-PICD0	Used for serial messages between the Pentium II and APIC.
PM1-PM0	Performance monitor signals are used to test the performance of the Pentium II.
PRDY	The probe ready output indicates that the probe mode has been entered for debugging.

TABLE 17-1 State of the Pentium II after a RESET.

<i>Register</i>	<i>RESET Value</i>	<i>RESET + BIST Value</i>
EAX	0	0 (if test passes)
EDX	0500XXXXH	0500XXXXH
EBX, ECX, ESP, EBP, ESI, and EDI	0	0
EFLAGS	2	2
EIP	0000FFF0H	0000FFF0H
CS	F000H	F000H
DS, ES, FS, GS, and SS	0	0
GDTR and TSS	0	0
CR0	60000010H	60000010H
CR2, CR3, and CR4	0	0
DR0–DR3	0	0
DR6	FFFF0FF0H	FFFF0FF0H
DR7	00000400H	00000400H

Notes: BIST = built-in self-test, XXXX = Pentium II version number.

PREQ	The probe request is used to request debugging.
PWRGOOD	An input that indicates that the system power supply is operational.
REQ4–REQ0	Request signals communicate commands between bus controllers and the Pentium II.
RESET	Reset initializes the Pentium II, causing it to begin executing software at memory location FFFFFFF0H or 000FFFF0. The A35–A32 address bits are set as logic 0s during the reset operation. The Pentium II is reset to the real mode and the leftmost 12 address connections remain logic 1s (FFFH) until a far jump or far call is executed. This allows compatibility with earlier microprocessors. See Table 17-1 for the state of the Pentium II after a hardware reset.
RP	Request parity is used to request parity.
RS2–RS0	Request status inputs are used to request the current status of the Pentium II.
RSP	The response parity input is activated to request parity.
SLOT0CC	The slot occupied output is a logic 0 if slot zero contains either a Pentium II or a dummy terminator.
SLP	Sleep is an input that, when inserted in the stop grant state, causes the Pentium II to enter the sleep state.
SMI	The system management interrupt input causes the Pentium II to enter the system management mode of operation.
STPCLK	The stop clock input causes the Pentium II to enter the low-power stop grant state.
TCK	The testability clock input selects the clocking function in accordance with the IEEE 1149.1 Boundary Scan interface.
TDI	The test data input is used to test data clocked into the Pentium II with the TCK signal.
TDO	The test data output is used to gather test data and instruction shifted out of the Pentium II with TCK.
TESTHI	Test high is an input that must be connected to +2.5 V through a 1K–10K Ω resistor for proper Pentium II operation.