The Intel Microprocessors

8086/8088, 80186/80188, 80286, 80386, 80486 Pentium, Pentium Pro Processor, Pentium II, Pentium 4, and Core2 with 64-bit Extensions

Architecture, Programming, and Interfacing

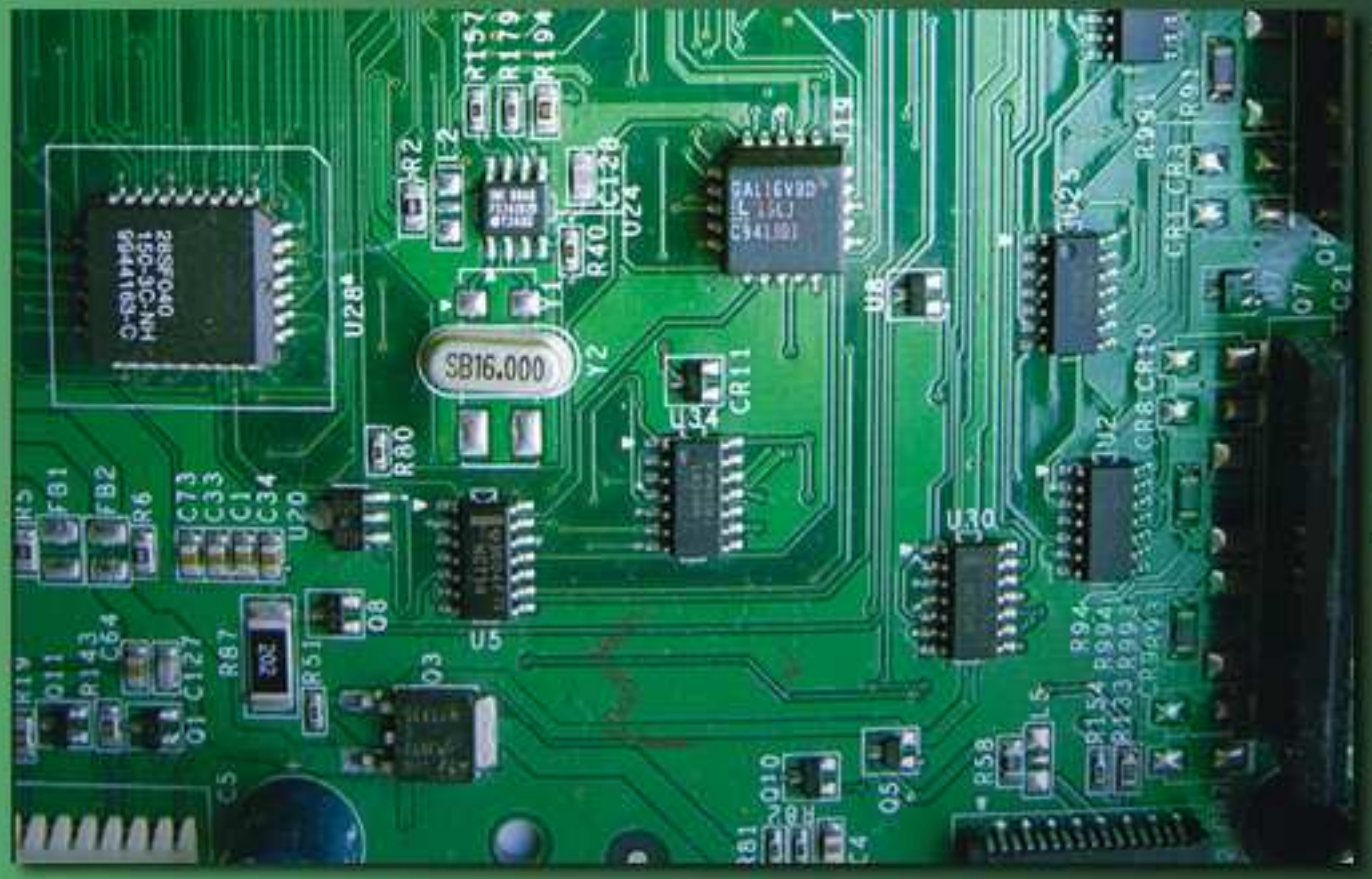

EIGHTH EDITION

Barry B. Brey

PEARSON

Chapter 18: The Pentium and Pentium Pro Microprocessors

Introduction

- The Pentium microprocessor signals an improvement to the architecture found in the 80486 microprocessor.
- 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.

18–1 INTRO TO THE PENTIUM PROCESSOR

- the pin-out of the Pentium processor, packaged in a huge 237-pin PGA (pin grid array)
- Pentium is available in two versions:
  - the full-blown Pentium
  - P24T version called the Pentium OverDrive
- P24T version contains a 32-bit data bus, for insertion into 80486 machines which contain the P24T socket

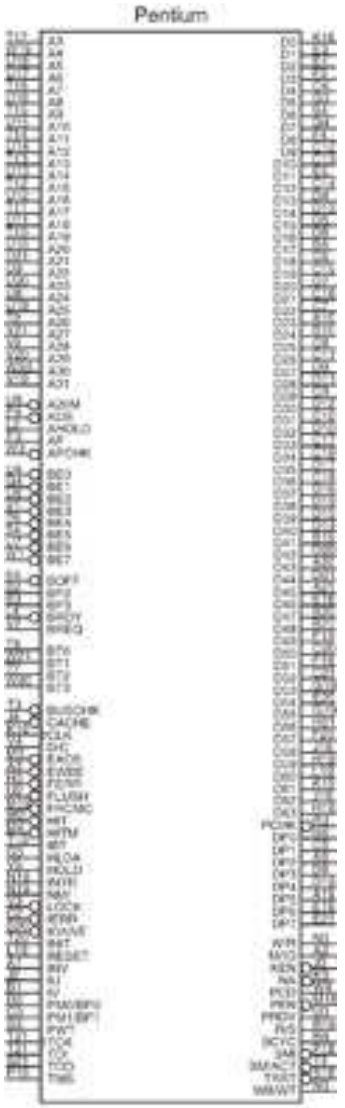

Figure 18–1 The pin-out of the Pentium microprocessor.

- Each Pentium output pin is capable of providing 4.0 mA of current at logic 0 level and 2.0 mA at logic 1 level.
  - compared to the 2.0 mA available 8086-80286
- Each input pin represents a small load requiring only 15 μA of current.
- In some systems, except the smallest, these current levels require bus buffers.

The Memory System

- Memory system for the Pentium is 4G bytes.
  - just as in 80386DX and 80486 processors
- Pentium uses a 64-bit data bus to address memory in eight 512M byte banks.
- Memory is divided into eight banks with each bank storing byte-wide data with a parity bit.
- The Pentium, like 80486, employs internal parity generation and checking logic for the memory system's data bus information.

- Pentium systems do not use parity checks, because ECC is available
- 64-bit-wide memory is important to double-precision floating-point data
- Pentium is able to retrieve floating-point data with one read cycle, instead of two as in 80486

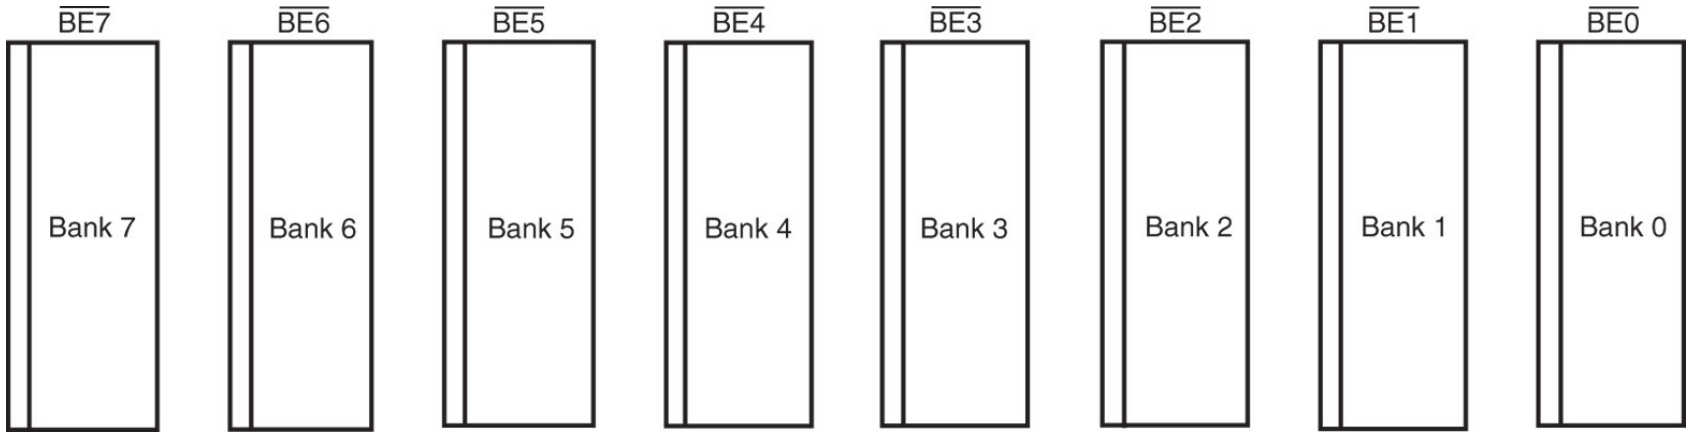

Figure 18–2 The 8-byte-wide memory banks of the Pentium microprocessor.

- The memory system is numbered in bytes, from byte 00000000H to byte FFFFFFFFH.
- Memory selection is accomplished with bank enable signals ($\overline{BE7}$-$\overline{BE0}$).
- The banks allow Pentium to access any single byte, word, doubleword, or quadword with one memory transfer cycle.
- A new feature is ability to check and generate parity for the address bus (A_{31}–A_5).

- The AP pin provides the system with parity information and the $\overline{\text{APCHK}}$ indicates a bad parity check for the address bus.
- 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

Input/Output System

- Input/output system of Pentium is completely compatible with earlier Intel processors.
- I/O port number appears on lines A_{15}–A_3
- Bank enable signals select actual memory banks used for the I/O transfer.
- I/O privilege information is added to the TSS segment when Pentium is in protected mode.
- If blocked I/O is accessed, Pentium generates a type 13 interrupt, I/O privilege violation.

System Timing

- Pentium nonpipelined memory cycle consists of two clocking periods: T_1 and T_2.
- Figure 18–3 shows basic nonpipelined read
- 66 MHz Pentium is capable of 33 million memory transfers per second.
  - assuming memory can operate at that speed
- W/$\overline{R}$ becomes valid if $\overline{ADS}$ is logic 0 at the positive edge of the clock (end of T_1).
  - clock must qualify cycle as read or write

Figure 18–3 The nonpipelined read cycle for the Pentium microprocessor.

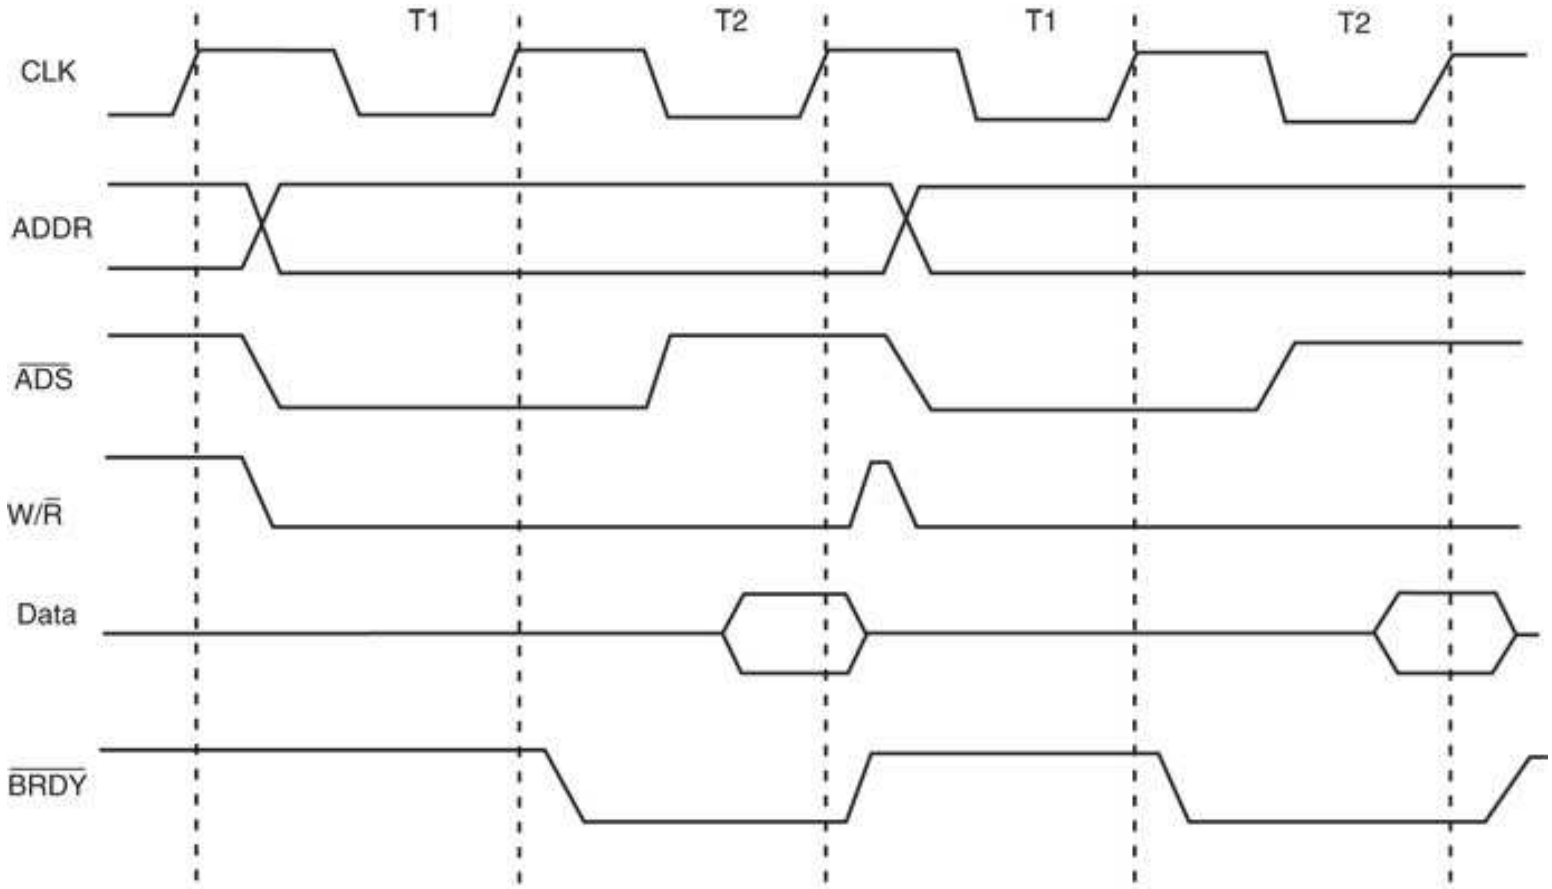

- During T_1, the processor issues the $\overline{ADS}$, W/$\overline{R}$ address, and M/$\overline{IO}$ signals.
- In order to qualify W/$\overline{R}$ & generate appropriate $\overline{MRDC}$ and $\overline{MWTC}$ signals, we use a flip-flop to generate the W/$\overline{R}$ signal.
- A two-line-to-one-line multiplexer then generates memory and I/O control signals.
- See Fig 18–4 for a circuit that generates I/O control and memory signals for Pentium.

Figure 18–4 A circuit that generates the memory and I/O control signals.

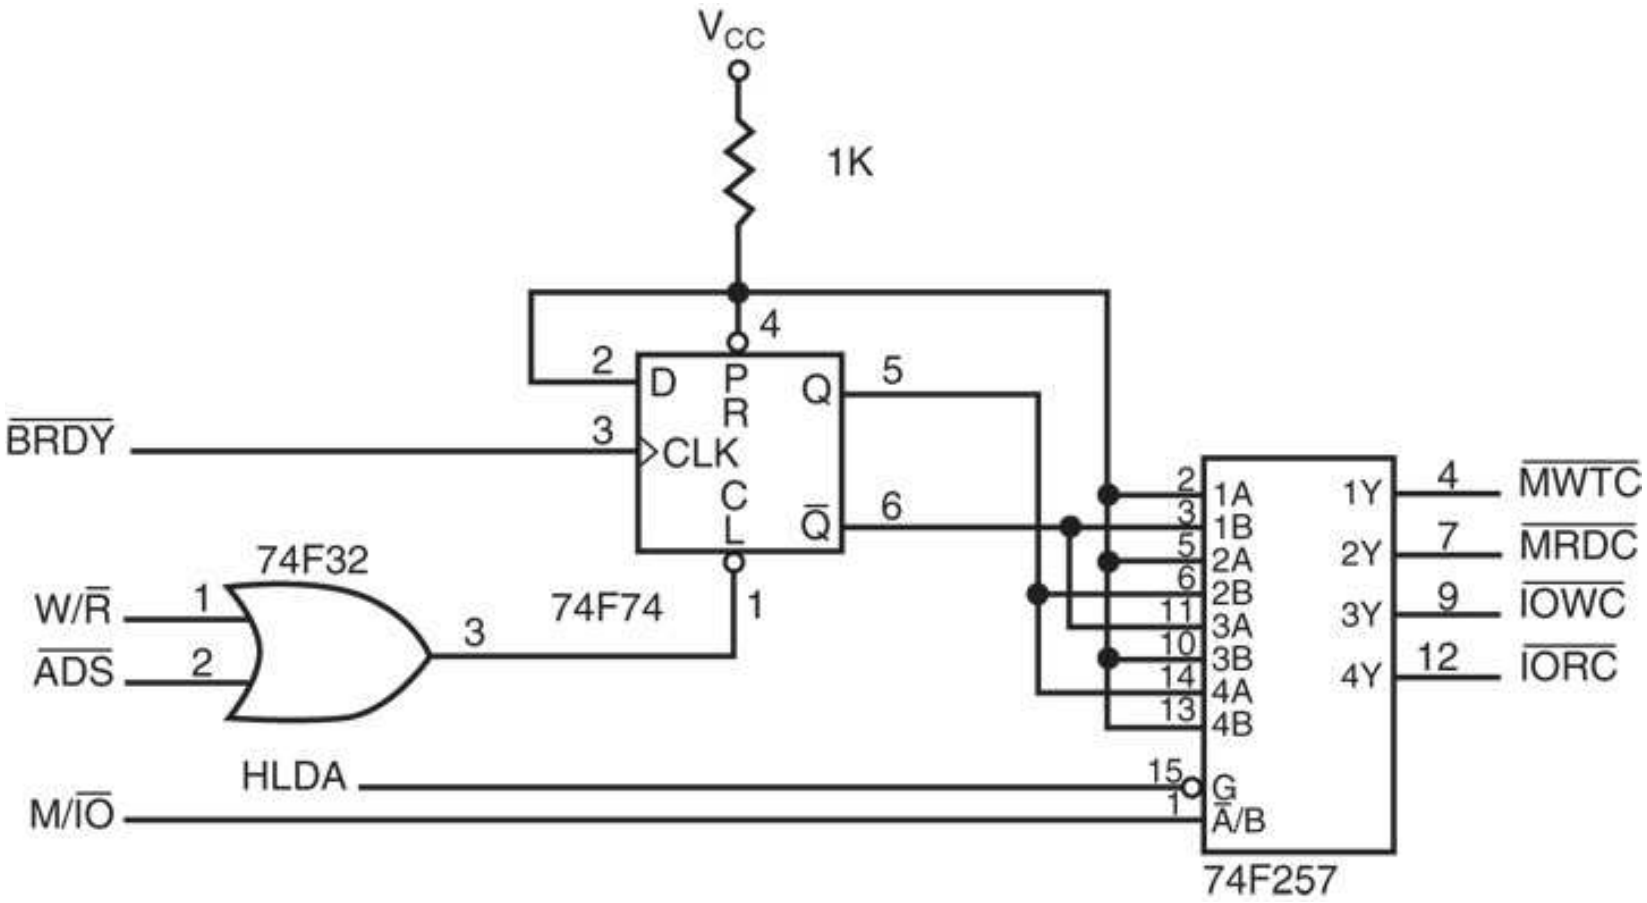

- During T_2, the data bus is sampled in synchronization with the end of T_2 at the positive transition of the clock pulse.
- Setup time before the clock is given as 3.8 ns.
- Hold time after the clock is given as 2.0 ns.
  - the data window around the clock is 5.8 ns
- The address appears on the 8.0 ns maximum after the start of T_1.

- Pentium at 66 MHz allows 30.3 ns (two clocking periods), minus the address delay of 8.0 ns and minus data setup 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 to allow access to SRAM.
  - not DRAM without inserting wait states
- SRAM is normally found in the form of an external level 2 cache.

- Wait states are inserted by controlling the $\overline{\text{BRDY}}$ input to the Pentium.
- $\overline{\text{BRDY}}$ must become logic 0 by the end of T_2
  - or additional T_2 states are inserted in the timing
- Inserting wait states lengthens timing to allow additional time for the memory to access data.
- In the timing shown, access time has been lengthened so 60 ns DRAM can be used.
  - requires insertion of four wait states of 15.2 ns (one clocking period) each to lengthen access time to 79.5 ns

– a read cycle timing diagram that contains wait states for slower memory.

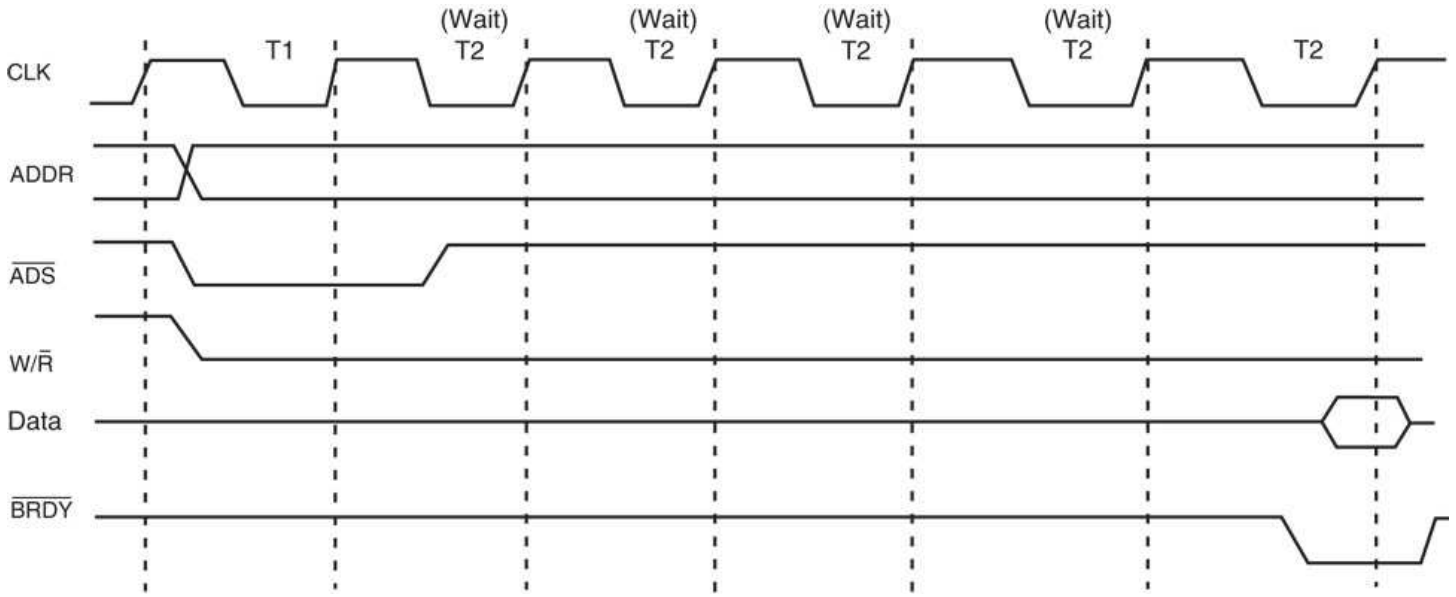

Figure 18–5 The Pentium timing diagram with four wait states inserted for an access time of 79.5 ns.

- $\overline{\text{BRDY}}$ is a synchronous signal generated by using the system clock.
- Figure 18–6 illustrates a circuit to generate $\overline{\text{BRDY}}$ for inserting any number of wait states into the Pentium timing diagram.
- The $\overline{\text{ADS}}$ signal is delayed between 0 and 7 clocking periods by the 74Fl61 shift register to generate the $\overline{\text{BRDY}}$ signal.
- The exact number of wait states is selected by the 74F151 eight-line-to-one-line multiplexer.

Figure 18–6 A circuit that generates wait states by delaying ADS. This circuit is wired to generate four wait states.

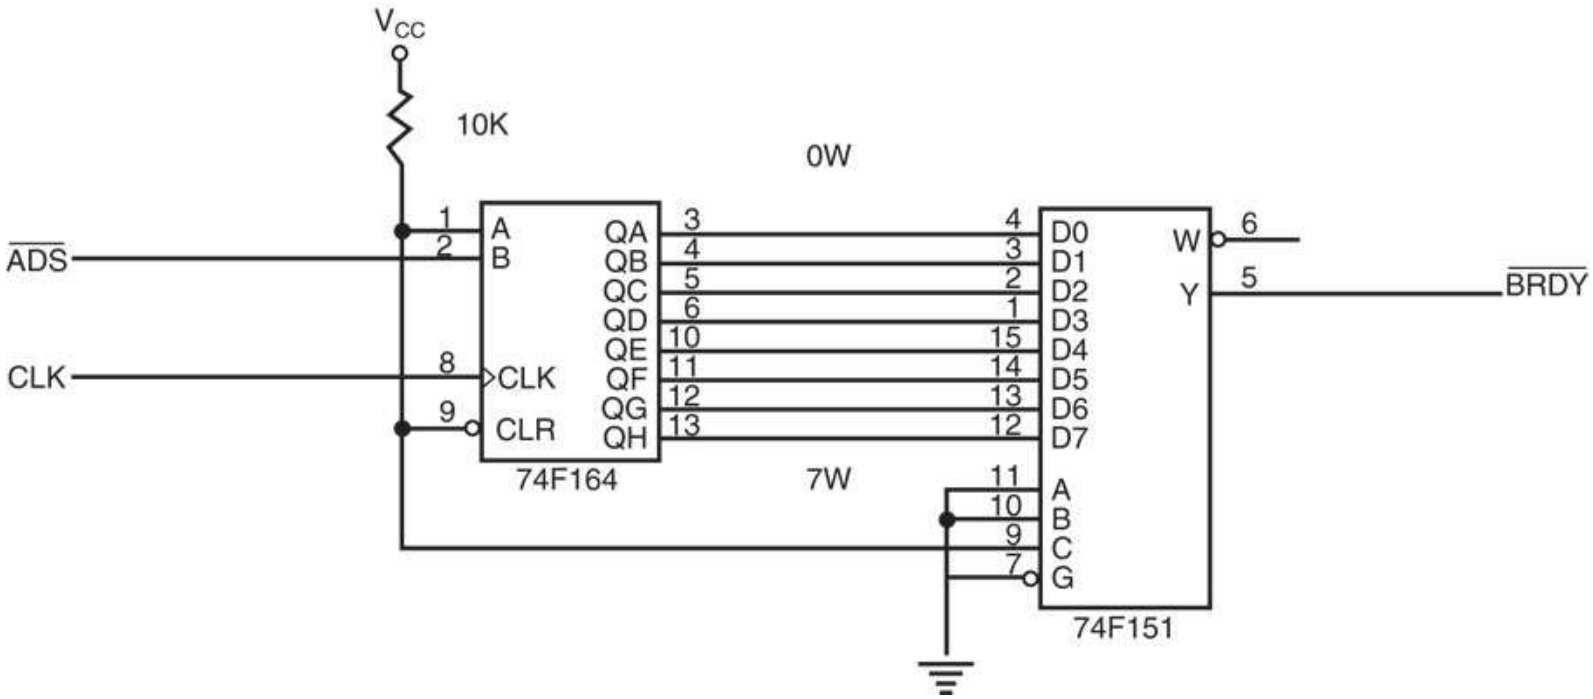

- The burst cycle is a more efficient method of reading memory data.
  - burst cycle in Pentium transfers four 64-bit numbers per cycle in five clocking periods
- A burst without wait states requires the memory system transfer data every 15.2 ns.
- With a level 2 cache, this speed is no problem as long as the data are read from the cache.
- If the cache does not contain the data, wait states must be inserted, reducing throughput.

Figure 18–7 The Pentium burst cycle operation that transfers four 64-bit data between the microprocessor and memory.

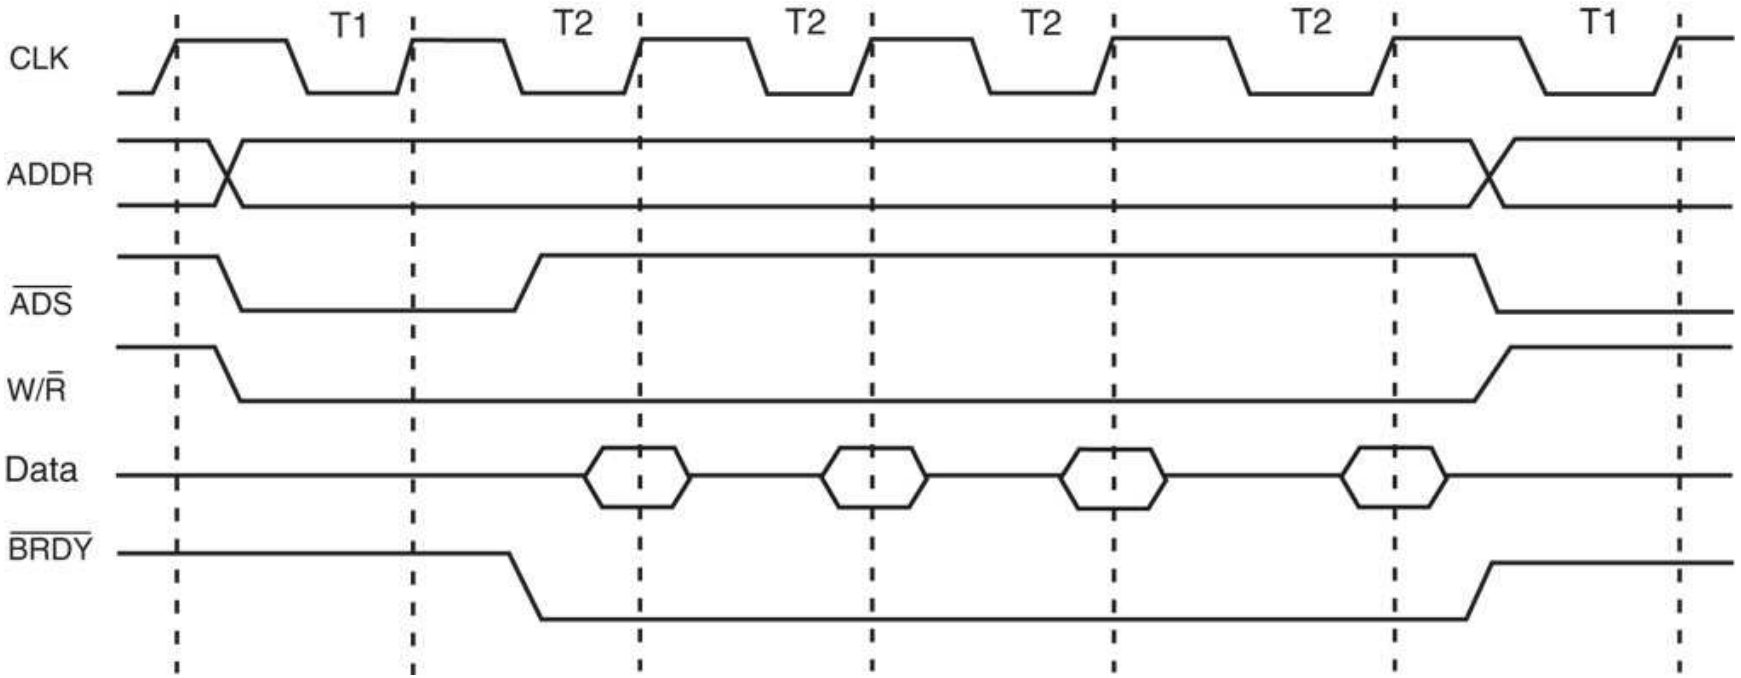

– wait states can be inserted to allow more time to the memory system for accesses

Branch Prediction Logic

- Used by Pentium to reduce time required for a branch caused by internal delays.
- Delays are minimized when a branch is encountered, because the processor begins prefetch instruction at the branch address.
- If the branch prediction logic errs, the branch requires an extra three clocking periods to execute.
  - in most cases, branch prediction is correct and no delay ensues

Cache Structure

- Pentium contains two 8K-byte cache memories instead of one as in 80486.
  - one 8K-byte instruction cache stores only instructions
  - another 8K-byte cache stores data used by instructions
- In the 80486 unified cache, a data-intensive program quickly filled the cache, and slowed execution speed.
  - this cannot occur in Pentium separate caches

Superscalar Architecture

- Pentium has three execution units.
  - One executes floating-point instructions
  - the other two (U-pipe and V-pipe) execute integer instructions
- This means it is possible to execute three instructions simultaneously.
  - because the floating-point unit is also used for MMX instructions, Pentium can simultaneously execute two integers and one MMX instruction.

18–5 INTRO TO THE PENTIUM PRO

- Before a processor can be used in a system, the function of each pin must be understood.
- This section details the operation of each pin, along with the external memory system and I/O structures of the Pentium Pro processor.
- Figure 18–12 illustrates the pin-out of the Pentium Pro processor, packaged in an immense 387-pin PGA (pin grid array).

Figure 18–12 The pin-out of the Pentium Pro microprocessor.

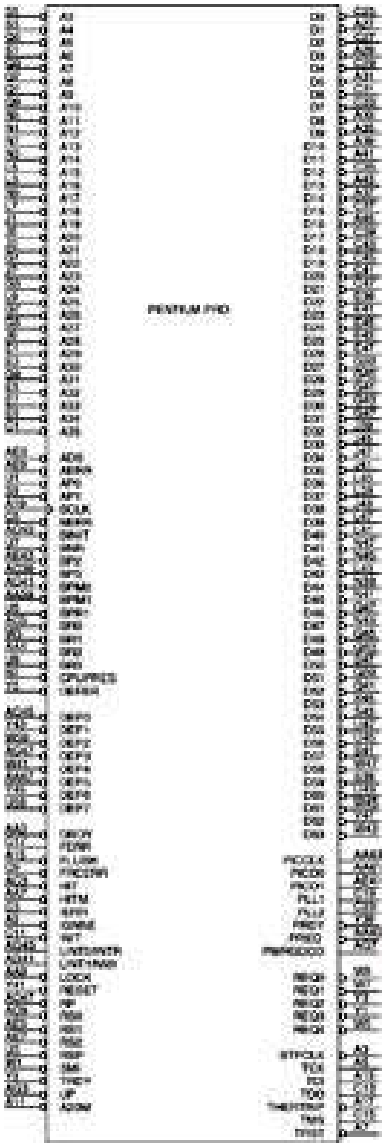

- Pentium Pro is available in two versions
- one version with 256K level 2 cache
- the other contains a 512K level 2 cache
- the most notable difference in the pin-out of Pentium Pro is the provision for a 36-bit address bus
- allows access to 64G bytes of memory

- Pro requires a +3.3 V or +2.7 V power supply.
  - current maximum is 9.9 A at 150 MHz
  - maximum power dissipation of 26.7 W
- A heat sink with good airflow is required.
- Multiple V_{cc} and V_{ss} connections must all be connected for proper operation.
- There are some pins labeled N/C (no connection) that must not be connected.
- Each output pin is capable of providing an ample 48.0 mA of current at a logic 0 level.
- Each input pin is a 15 μA current load.

Internal Structure of Pentium Pro

- Pentium Pro is structured differently than earlier processors.
  - which contained an execution unit, and cached bus interface unit buffering the execution unit
- Figure 18–13 shows a block diagram of the internal structure of the Pentium Pro processor.

Figure 18–13 The internal structure of the Pentium Pro microprocessor.

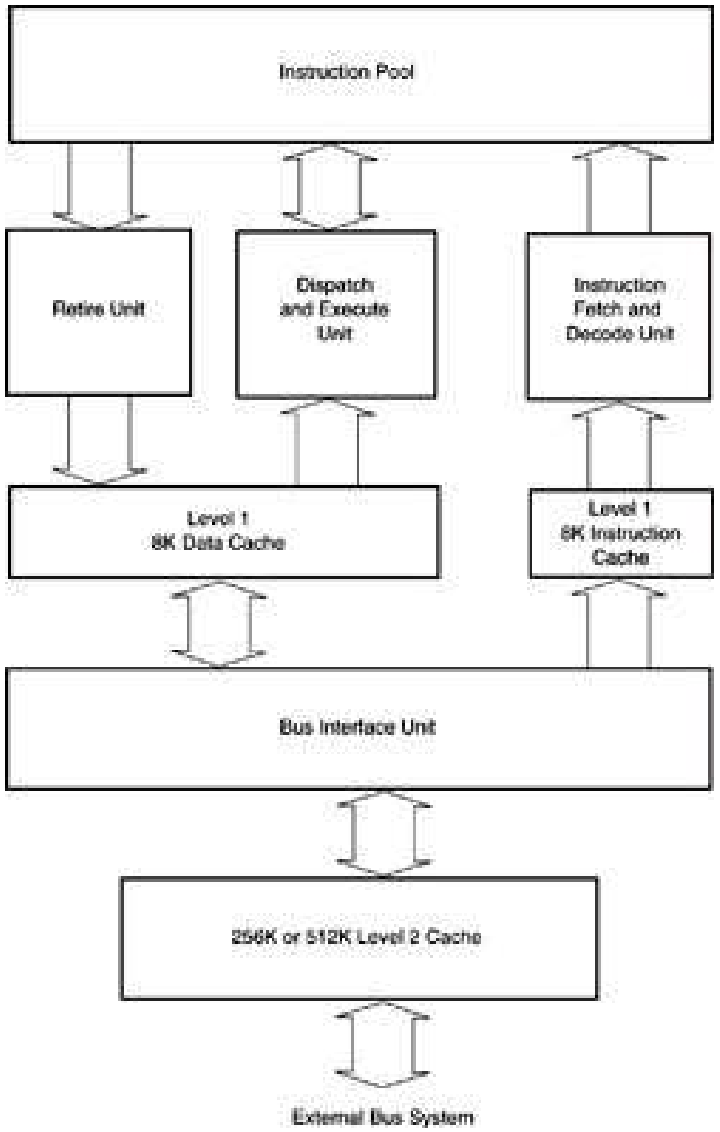

- system buses, which communicate to memory and I/O, connect to an internal level 2 cache
- often on the main board in most other systems
- level 2 cache in the Pro is either 256K or 512K bytes
- integration of the level 2 cache speeds processing and reduces number of components in a system

- The bus interface unit (BIU) controls access to system buses through the level 2 cache.
  - generates memory address and control signals, and passes/fetches data or instructions to an 8K level 1 data cache or a level 1 instruction cache
- The instruction cache is connected to the instruction fetch and decode unit (IFDU).
  - separate instruction decoders decode three instructions simultaneously
  - outputs are passed to the instruction pool, where they remain until the dispatch and execution unit or retire unit obtains them

- The dispatch and execute unit (DEU) retrieves instructions from the instruction pool when complete, and executes them.
- DEU has three instruction execution units:
  - two for processing integer instructions
  - one for floating-point instructions
- Reservation station (RS) schedules up to five events and can process four simultaneously.
- Retire unit (RU) removes instructions that have been executed from the instruction pool.
  - three decoded instructions per clock pulse

Figure 18–14 The Pentium Pro dispatch and execution unit (DEU).

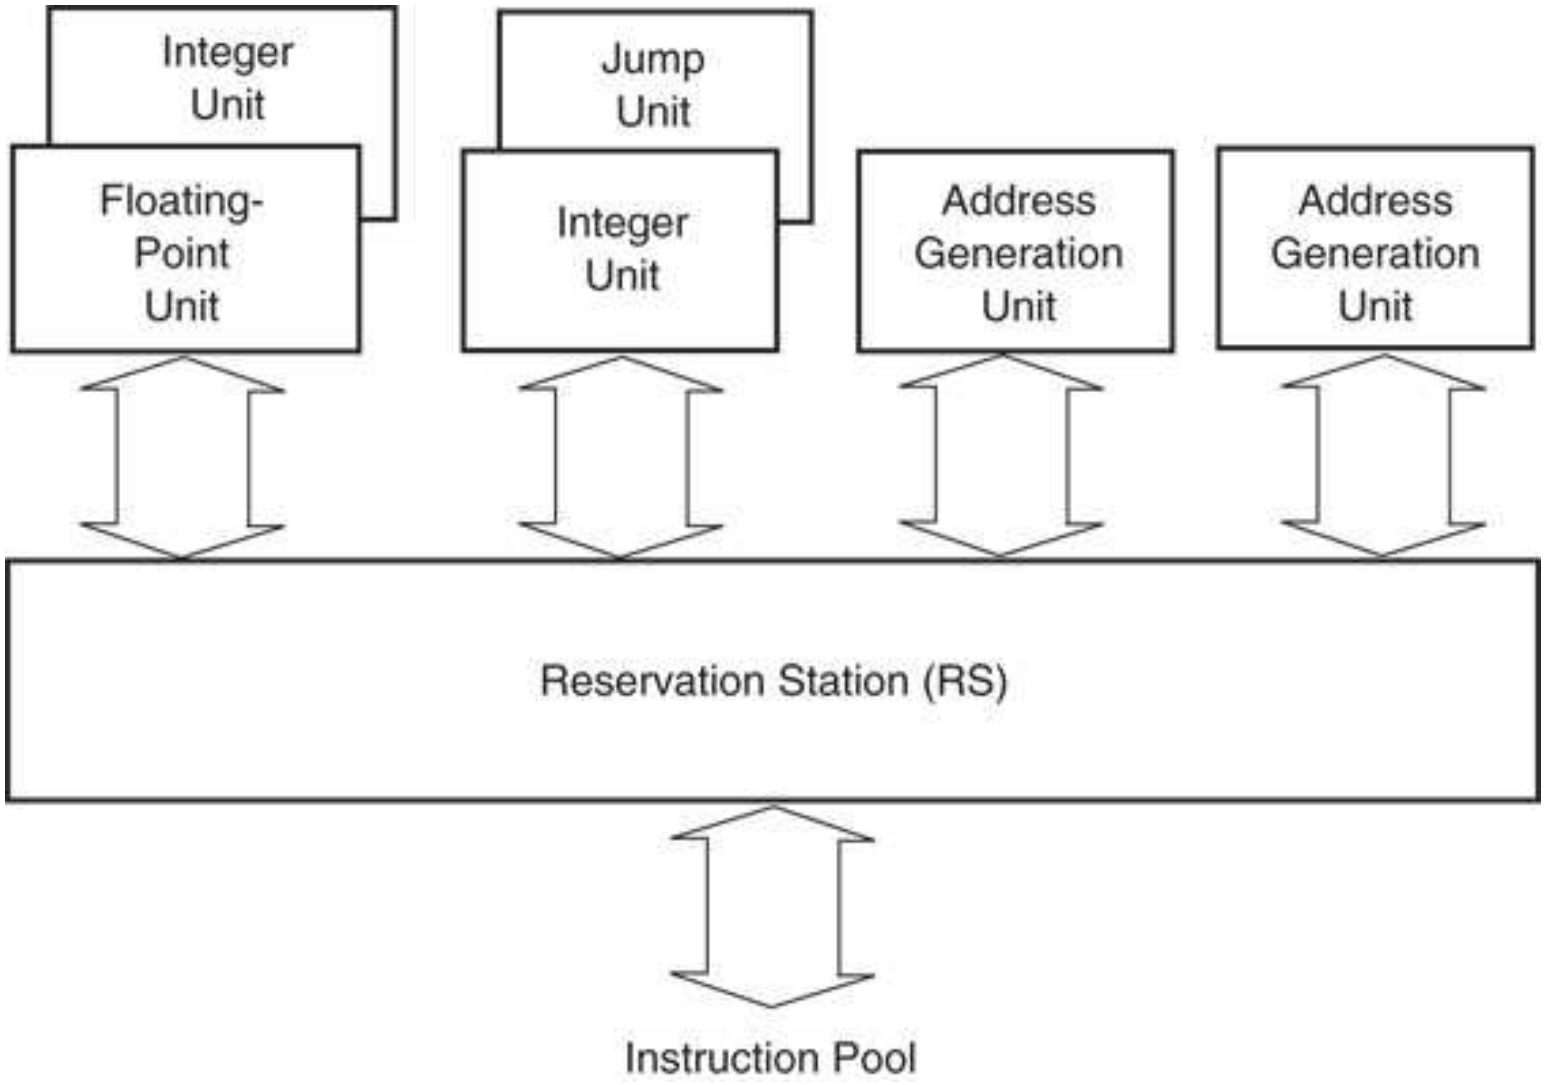

The Memory System

- The memory system for Pentium Pro is 4G bytes in size, just as in 80386DX–Pentium processors
- 2M paging is new to the Pentium Pro to allow memory above 4G to be accessed.
  - access between 4G and 64G is made possible by additional address signals A_{32}–A_{35}

– Pro uses a 64-bit data bus to address memory in eight banks that each contain 8G bytes of data.
– additional memory is enabled with bit position 5 of CR_4 , accessible when 2M paging is enabled

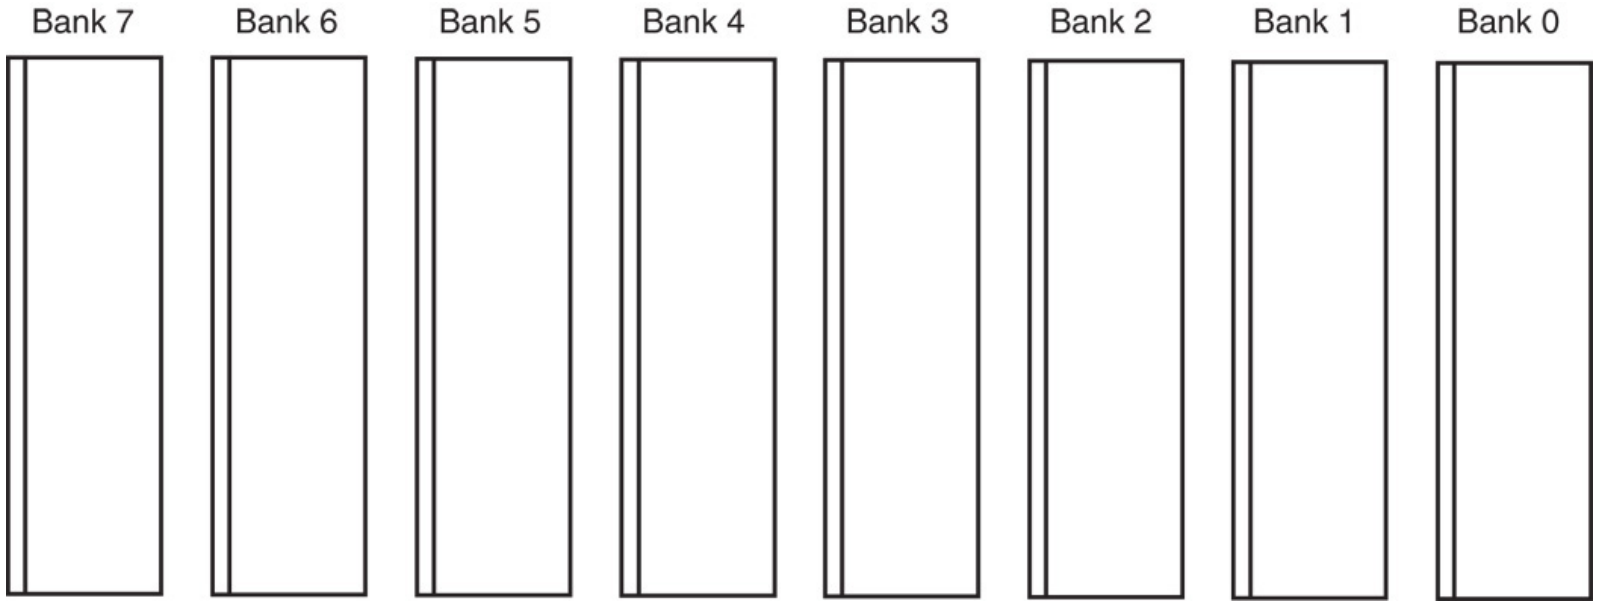

Figure 18–15 The eight memory banks in the Pentium Pro system. Note that each bank is 8 bits wide and 8G long if 36-bit addressing is enabled.

- Pentium Pro, like 80486 and Pentium, uses internal parity generation and checking logic for the memory system data bus information.
  - 64-bit-wide memory is important to double-precision floating-point data
- In the Pentium Pro processor, bank enable signals are presented on the address bus (A_{15}–A_8) during the second clock cycle of a memory or I/O access.
  - must be extracted from the address bus to access memory banks

- Separate banks allow Pro to access any single byte, word, doubleword, or quadword with one memory transfer cycle.
- Often eight separate write strobes for writing to the memory system.
- Memory write information is provided on the request lines from the processor during the second clock phase of a memory or I/O access.
- Pro is also able to check and generate parity for the address bus during certain operations.

- New to Pro is a built-in error-correction circuit (ECC) allowing correction of a one-bit error and the detection of a two-bit error.
  - to detect/correct, memory must have room for an extra 8-bit number stored with each 64-bit number
- The extra 8 bits store an error-correction code that allows Pro to correct any single-bit error.
- A 1M $\times$ 64 is a 64M SDRAM without ECC, and a 1M $\times$ 72 is an SDRAM with EEC.
- ECC code is far more reliable than the old parity scheme, now rarely used.

Input/Output System

- The input/output system of Pentium Pr is completely compatible with earlier Intel processors.
- I/O port number appears on address lines A_{15}–A_3 with the bank enable signals used to select the actual memory banks used for the I/O transfer.

System Timing

- The basic Pentium Pro memory cycle consists of two sections.
  - the address phase and the data phase
- During address phase, Pro sends the address (T_1) & control signals (T_2) to memory and I/O.
- The control signals include the ATTR lines (A_{31}–A_{24}), DID lines (A_{23}–A_{16}), bank enable signals (A_{15}–A_8), and the EXF lines (A_7–A_3).
- See Figure 18–16 for the basic timing cycle.

Figure 18–16 The basic Pentium Pro timing.

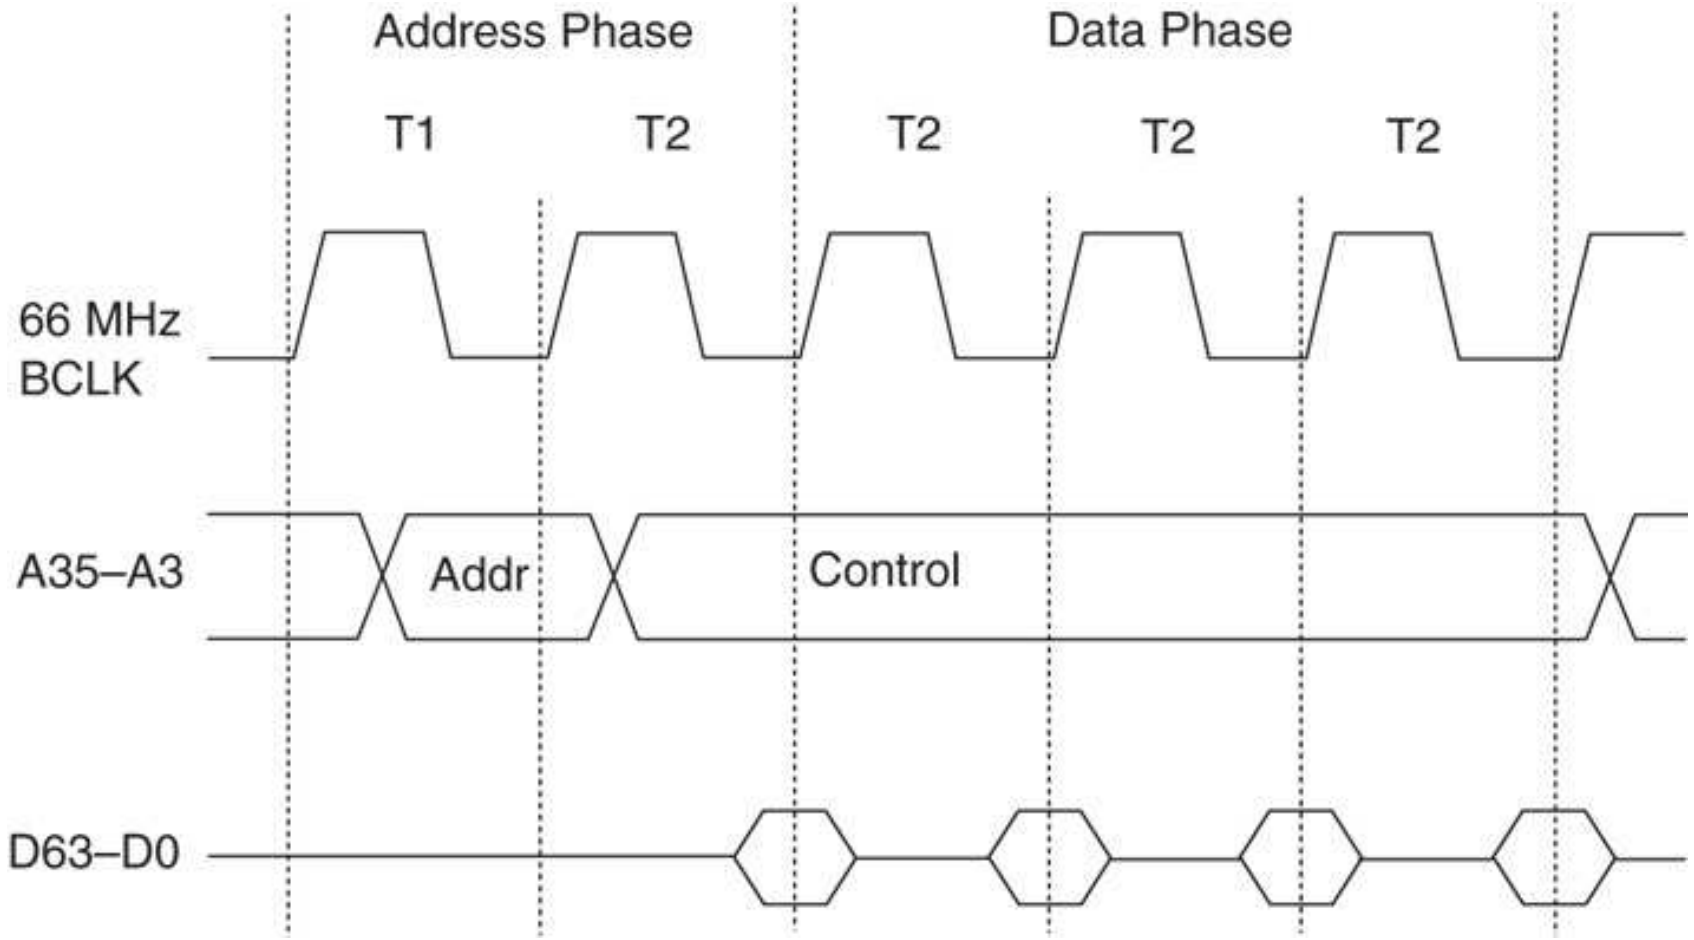

- Cycle type appears on the request pins.
- During the data phase, four 64-bit-wide numbers are fetched/written to the memory.
- The 66 MHz Pentium Pro is capable of 33 million memory transfers per second.
  - assuming memory can operate at that speed
- Setup time before the clock is given as 5.0 ns; hold time after the clock is given as 1.5 ns.
  - the data window around the clock is 6.5 ns
  - the address appears on the 8.0 ns maximum after the start of T_1

- The Pentium Pro 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 an SRAM, but not to any DRAM.
  - without inserting wait states into the timing