

3.3 Functional Timing

3.3.1 Reset Timing

Figure 3-2 illustrates the required RESET timing for both a power-on reset and a reset that occurs during operation. The WM_RST, FLUSH# and QDUMP# inputs are sampled at

the falling edge of RESET to determine if the 6x86 CPU should enter built-in self-test, enable tri-state test mode or enable the scatter-gather interface pins, respectively. WM_RST, FLUSH# and QDUMP# must be valid at least two clocks prior to the RESET falling edge.




Figure 3-2. RESET Timing



3.3.2 Bus State Definition

The 6x86 CPU bus controller supports non-pipelined and pipelined operation as well as single transfer and burst bus cycles. During each CLK period, the bus controller exists in one of six states as listed in Table 3-12. Each of these bus states and its associated state transitions is illustrated in Figure 3-3, (Page 3-27) and listed in Table 3-13, (Page 3-28).

Table 3-12. 6x86 CPU Bus States

STATE	NAME	DESCRIPTION
Ti	Idle Clock	During Ti, no bus cycles are in progress. BOFF# and RESET force the bus to the idle state. The bus is always in the idle state while HLDA is active.
T1	First Bus Cycle Clock	During the first clock of a non-pipelined bus cycle, the bus enters the T1 state. ADS# is asserted during T1 along with valid address and bus cycle definition information.
T2	Second and Subsequent Bus Cycle Clock	During the second clock of a non-pipelined bus cycle, the bus enters the T2 state. The bus remains in the T2 state for subsequent clocks of the bus cycle as long as a pipelined cycle is not initiated. During T2, valid data is driven during write cycles and data is sampled during reads. BRDY# is also sampled during T2. The bus also enters the T2 state to complete bus cycles that were initiated as pipelined cycles but complete as the only outstanding bus cycle.
T12	First Pipelined Bus Cycle Clock	During the first clock of a pipelined cycle, the bus enters the T12 state. During T12, data is being transferred and BRDY# is sampled for the current cycle at the same time that ADS# is asserted and address/bus cycle definition information is driven for the next (pipelined) cycle.
T2P	Second and Subsequent Pipelined Bus Cycle Clock	During the second and subsequent clocks of a pipelined bus cycle where two cycles are outstanding, the bus enters the T2P state. During T2P, data is being transferred and BRDY# is sampled for the current cycle. However, valid address and bus cycle definition information continues to be driven for the next pipelined cycle.
Td	Dead Clock	The bus enters the Td state if a pipelined cycle was initiated that requires one idle clock to turn around the direction of the data bus. Td is required for a read followed immediately by a pipelined write, and for a write followed immediately by a pipelined read.




Figure 3-3. IBM 6x86 CPU Bus State Diagram



Table 3-13. Bus State Transitions

TRANSITION	CURRENT STATE	NEXT STATE	EQUATION
A	Ti	Ti	No Bus Cycle Pending.
B	Ti	T1	New or Aborted Bus Cycle Pending.
C	T1	T2	Always.
D	T2	T2	Not Last BRDY# and No New Bus Cycle Pending, or Not Last BRDY# and New Bus Cycle Pending and NA# Negated.
E	T2	T1	Last BRDY# and New Bus Cycle Pending and HITM# Negated.
F	T2	Ti	Last BRDY# and No New Bus Cycle Pending, or Last BRDY# and HITM# Asserted.
G	T2	T12	Not Last BRDY# and New Bus Cycle Pending and NA# Sampled Asserted.
H	T12	T2	Last BRDY# and No Dead Clock Required.
I	T12	Td	Last BRDY# and Dead Clock Required.
J	T12	T2P	Not Last BRDY#.
K	T2P	T2P	Not Last BRDY#.
L	T2P	T2	Last BRDY# and No Dead Clock Required.
M	T2P	Td	Last BRDY# and Dead Clock Required.
N	Td	T12	New Bus Cycle Pending and NA# Sampled Asserted.
O	Td	T2	No New Bus Cycle Pending, or New Bus Cycle Pending and NA# Negated.
P	Any State	Ti	RESET Asserted, or BOFF# Asserted.

3.3.3 Non-pipelined Bus Cycles

Non-pipelined bus operation may be used for all bus cycle types. The term “non-pipelined” refers to a mode of operation where the CPU allows only one outstanding bus cycle. In other words, the current bus cycle must complete before a second bus cycle is allowed to start.

3.3.3.1 Non-pipelined Single Transfer Cycles

Single transfer read cycles occur during non-cacheable memory reads, I/O read cycles, and special cycles. A non-pipelined single transfer read cycle begins with address and bus cycle definition information driven on the bus during the first clock (T1 state) of the bus cycle. The CPU then monitors the BRDY# input at the end of the second clock (T2 state). If BRDY# is asserted, the CPU reads the appropriate data and data parity lines and terminates the bus cycle. If BRDY# is not active, the CPU continues to sample the BRDY# input at the end of each subsequent cycle (T2 states). Each of the additional clocks is referred to as a wait state.

The CPU uses the data parity inputs to check for even parity on the active data lines. If the CPU detects an error, the parity check output (PCHK#) asserts during the second clock following the termination of the read cycle.

Figure 3-4 (Page 3-30) illustrates the functional timing for two non-pipelined single-transfer read cycles. Cycle 2 is a potentially cacheable cycle as indicated by the CACHE# output. Because this cycle is potentially cacheable, the CPU samples the KEN# input at the same clock edge that BRDY# is asserted. If KEN# is negated, the cycle terminates as shown in the diagram. If KEN# is asserted, the CPU converts this cycle into a burst cycle as described in the next section. NA# must be negated for non-pipelined operation. Pipelined bus cycles are described later in this chapter.



Functional Timing




Figure 3-4. Non-Pipelined Single Transfer Read Cycles

Single transfer write cycles occur for writes that are neither line replacement nor write-back cycles. The functional timing of two non-pipelined single transfer write cycles is shown in Figure 3-5. During a write cycle, the data and data parity lines are outputs and are driven

valid during the second clock (T2 state) of the bus cycle. Data and data parity remain valid during all wait states. If the write cycle is a write to a valid cache location in the “shared” state, the WB/WT# pin is sampled with BRDY#. If WB/WT# is sampled high, the cache line transitions from the “shared” to the “exclusive” state.




Figure 3-5. Non-Pipelined Single Transfer Write Cycles



3.3.3.2 Non-pipelined Burst Read Cycles

The 6x86 CPU uses burst read cycles to perform cache line fills. During a burst read cycle, four 64-bit data transfers occur to fill one of the CPU's 32-byte internal cache lines. A non-pipelined burst read cycle begins with address and bus cycle definition information driven on the bus during the first clock (T1 state) of the bus cycle. The CACHE# output is always active during a burst read cycle and is driven during the T1 clock.

The CPU then monitors the BRDY# input at the end of the second clock (T2 state). If BRDY# is asserted, the CPU reads the data and data parity and also checks the KEN# input. If KEN# is negated, the CPU terminates the bus cycle as a single transfer cycle. If KEN# is asserted, the CPU converts the cycle into a burst (cache line fill) by continuing to sample BRDY# at the end of each subsequent clock. BRDY# must be asserted a total of four times to complete the burst cycle.

WB/WT# is sampled at the same clock edge as KEN#. In conjunction with PWT and the on-chip configuration registers, WB/WT# determines the MESI state of the cache line for the current line fill.

Each time BRDY# is sampled asserted during the burst cycle, a data transfer occurs. The CPU reads the data and data parity busses and assigns the data to an internally generated burst address. Although the CPU internally generates the burst address sequence, only the first address of the burst is driven on the external address bus. System logic must predict the burst address sequence based on the first address. Wait states may be added to any transfer within a burst by delaying the assertion of BRDY# by the desired number of clocks.

The CPU checks even data parity for each of the four transfers within the burst. If the CPU detects an error, the parity check output (PCHK#) asserts during the second clock following the BRDY# assertion of the data transfer.

Figure 3-6 (Page 3-33) illustrates two non-pipelined burst read cycles. The cycles shown are the fastest possible burst sequences (2-1-1-1). NA# must be negated for non-pipelined operation as shown in the diagram. Pipelined bus cycles are described later in this chapter.

Figure 3-7 (Page 3-34) depicts a burst read cycle with wait states. A 3-2-2-2 burst read is shown.




Figure 3-6. Non-Pipelined Burst Read Cycles




Figure 3-7. Burst Cycle with Wait States

Burst Cycle Address Sequence.

The IBM 6x86 CPU provides two different address sequences for burst read cycles. The IBM 6x86 CPU burst cycle address sequence modes are referred to as “1+4” and “linear”. After reset, the CPU default mode is “1+4”.

In “1+4” mode, the CPU performs a single transfer read cycle prior to the burst cycle, if the desired first address is (...xx8). During this single transfer read cycle, the CPU reads the critical data. In addition, the 6x86 CPU sam-

ples the state of KEN#. If KEN# is active, the CPU then performs the burst cycle with the address sequence shown in Table 3-14 (Page 3-35). The IBM 6x86 CPU CACHE# output is not asserted during the single read cycle prior to the burst. Therefore, CACHE# must not be used to qualify the KEN# input to the processor. In addition, if KEN# is returned active for the “1” read cycle in the “1+4”, all data bytes supplied to the CPU must be valid. The CPU samples WB/WT# during the “1” read cycle, and does not resample WB/WT# during the following burst cycle. Figure 3-8 (Page 3-35) illustrates a “1+4” burst read cycle.

Table 3-14. "1+4" Burst Address Sequences

BURST CYCLE FIRST ADDRESS	SINGLE READ CYCLE PRIOR TO BURST	BURST CYCLE ADDRESS SEQUENCE
0	None	0-8-10-18
8	Address 8	0-8-10-18
10	None	10-18-0-8
18	Address 18	10-18-0-8




Figure 3-8. "1+4" Burst Read Cycle



The address sequences for the 6x86 CPU's linear burst mode are shown in Table 3-15. Operating the CPU in linear burst mode minimizes processor bus activity resulting in higher system performance. Linear burst mode can be enabled through the IBM 6x86 CPU CCR3 configuration register.

Table 3-15. Linear Burst Address Sequences

BURST CYCLE FIRST ADDRESS	BURST CYCLE ADDRESS SEQUENCE
0	0-8-10-18
8	8-10-18-0
10	10-18-0-8
18	18-0-8-10

3.3.3.3 Burst Write Cycles

Burst write cycles occur for line replacement and write-back cycles. Burst writes are similar to burst read cycles in that the CACHE# output is asserted and four 64-bit data transfers occur. Burst writes differ from burst reads in that the data and data parity lines are outputs rather than inputs. Also, KEN# and WB/WT# are not sampled during burst write cycles.

Data and data parity for the first data transfer are driven valid during the second clock (T2 state) of the bus cycle. Once BRDY# is sampled asserted for the first data transfer, valid data and data parity for the second transfer are driven during the next clock cycle. The same timing relationship between BRDY# and data applies for the third and fourth data transfers as well. Wait states may be added to any transfer within a burst by delaying the assertion of BRDY# by the required number of clocks.

As on burst read cycles, only the first address of a burst write cycle is driven on the external address bus. System logic must predict the remaining burst address sequence based on the first address. Burst write cycles always begin with a first address ending in 0 (signals A4-A0=0) and follow an ascending address sequence for the remaining transfers (0-8-10-18).

Figure 3-9 illustrates two non-pipelined burst write cycles. The cycles shown are the fastest possible burst sequences (2-1-1-1). As shown, an idle clock always exists between two back-to-back burst write cycles. Therefore, the second burst write cycle in a pair of back-to-back burst writes is always issued as a non-pipelined cycle regardless of the state of the NA# input.




Figure 3-9. Non-Pipelined Burst Write Cycles



3.3.4 Pipelined Bus Cycles

Pipelined addressing is a mode of operation where the CPU allows up to two outstanding bus cycles at any given time. Using pipelined addressing, the address of the first bus cycle is driven on the bus and while the CPU waits for the data for the first cycle, the address for a second bus cycle is issued. Pipelined bus cycles occur for all cycle types except locked cycles and burst write cycles.

Pipelined cycles are initiated by asserting NA#. The CPU samples NA# at the end of each T2, T2P and Td state. KEN# and WB/WT# are sampled at either the same clock as NA# is active, or at the same clock as the first BRDY# for that cycle, whichever occurs first. The CPU issues the next address a mini-

um of two clocks after NA# is sampled asserted.

The CPU latches the state of the NA# pin internally. Therefore, even if a new bus cycle is not pending internally at the time NA# was sampled asserted, the CPU still issues a pipelined bus cycle if an internal bus request occurs prior to completion of the current bus cycle. Once NA# is sampled asserted, the state of NA# is ignored until the current bus cycle completes. If two cycles are outstanding and the second cycle is a read, the CPU samples KEN# and WB/WT# for the second cycle when NA# is sampled asserted.

Figure 3-10 and Figure 3-11 (Page 3-39) illustrate pipelined single transfer read cycles and pipelined burst read cycles, respectively.




Figure 3-10. Pipelined Single Transfer Read Cycles




Figure 3-11. Pipelined Burst Read Cycles



3.3.4.1 Pipelined Back-to-Back Read/Write Cycles

Figure 3-12 depicts a read cycle followed by a pipelined write cycle. Under this condition, the data bus must change from an input for the read cycle to an output for the write cycle. In order to accomplish this transition without

causing data bus contention, the CPU automatically inserts a “dead” (Td) clock cycle. During the Td state, the data bus floats. The CPU then drives the write data onto the bus in the following clock. The CPU also inserts a Td clock between a write cycle and a pipelined read cycle to allow the data bus to smoothly transition from an output to an input.




Figure 3-12. Read Cycle Followed by Pipelined Write Cycle

3.3.5 Interrupt Acknowledge Cycles

The CPU issues interrupt acknowledge bus cycles in response to an active INTR input. Interrupt acknowledge cycles are single transfer cycles and always occur in locked pairs as shown in Figure 3-13. The CPU reads the interrupt vector from the lower eight bits of the data bus at the completion of the second inter-

rupt acknowledge cycle. Parity is not checked during the first interrupt acknowledge cycle.

M/IO#, D/C# and W/R# are always logic low during interrupt acknowledge cycles. Additionally, the address bus is driven with a value of 0000 0004h for the first interrupt acknowledge cycle and with a value of 0000 0000h for the second. A minimum of one idle clock always occurs between the two interrupt acknowledge cycles.

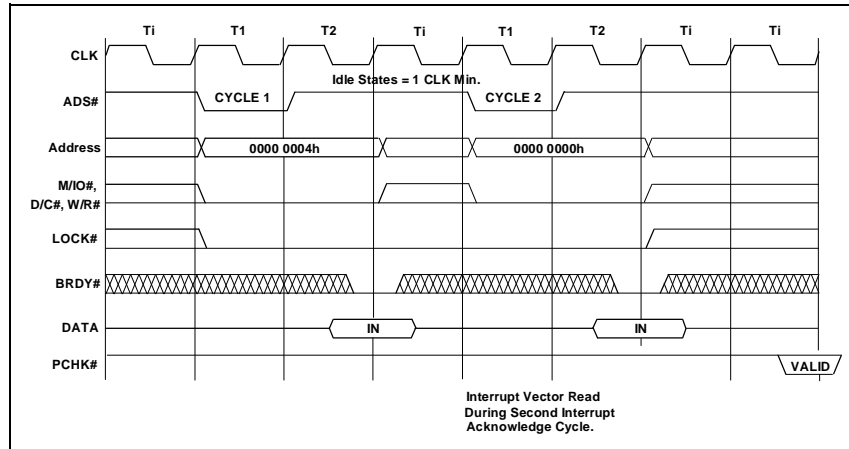


Figure 3-13. Interrupt Acknowledge Cycles



3.3.6 SMI# Interrupt Timing

The CPU samples the System Management Interrupt (SMI#) input at each clock edge. At the next appropriate instruction boundary, the CPU recognizes the SMI# and completes all pending write cycles. The CPU then asserts SMIACT# and begins saving the SMM header information to the SMM address space. SMIACT# remains asserted until after execution of a RSM instruction. Figure 3-14 illustrates the functional timing of the SMIACT# signal.

To facilitate using SMI# to power manage I/O peripherals, the 6x86 CPU implements a feature called I/O trapping. If the current bus cycle is an I/O cycle and SMI# is asserted a minimum of three clocks prior to BRDY#, the CPU immediately begins execution of the SMI service routine following completion of the I/O instruction. No additional instructions are executed prior to entering the SMI service routine. I/O trap timing requirements are shown in Figure 3-15 (Page 3-43).




Figure 3-14. SMIACT# Timing




Figure 3-15. SMM I/O Trap Timing

3.3.7 Cache Control Timing

3.3.7.1 Invalidating the Cache Using FLUSH#

The FLUSH# input forces the CPU to write-back and invalidate the entire contents of the on-chip cache. FLUSH# is sampled at each clock edge, latched internally and then recognized internally at the next instruction boundary. Once FLUSH# is recognized, the CPU issues a series of burst write cycles to write-back any "modified" cache lines. The cache lines are invalidated as they are written back. Following completion of the write-back cycles, the CPU issues a flush acknowledge special bus cycle.

The latency between when FLUSH# occurs and when the cache invalidation actually completes varies depending on:

- (1) the state of the processor when FLUSH# is asserted,
- (2) the number of modified cache lines,
- (3) the number of wait states inserted during the write-back cycles.

Figure 3-16 (Page 3-44) illustrates the sequence of events that occur on the bus in response to a FLUSH# request.




Figure 3-16. Cache Invalidation Using FLUSH#

3.3.7.2 EWBE# Timing

During memory and I/O write cycles, the 6x86™ CPU samples the external write buffer empty (EWBE#) input. If EWBE# is negated, the CPU does not write any data to “exclusive” or “modified” internal cache lines. After sampling EWBE# negated, the CPU continues to

sample EWBE# at each clock edge until it asserts. Once EWBE# is asserted, all internal cache writes are allowed. Through use of this signal, the external system may enforce strong write ordering when external write buffers are used. EWBE# functional timing is shown in Figure 3-17.




Figure 3-17. External Write Buffer Empty (EWBE#) Timing