

DATA PROCESSING SYSTEMS WITH DELAYED CACHE WRITEDESCRIPTION

The present invention relates to data processing systems with delayed cache write providing an improvement for increasing the tolerance to slow memory components for a cache read miss without impacting wait state parameters.

Background information respecting the 80386, its characteristics and its use in microcomputer systems including cache memory subsystems are described in "Introduction to the 80386", April 1986 and the 80386 Hardware Reference Manual (1986). The characteristics and operating performance of the 82385 are described in the publication "82385 High Performance 32-Bit Cache Controller" (1987). All of these publications are produced by Intel Corporation.

In microcomputer systems, as in other computer systems, speed of operations is an important criterion which in most cases has to be balanced against system cost. Many features which were first introduced to speed up operations in mainframe computers and minicomputers are now finding their way into microcomputer systems. These include cache memory subsystems. The use of cache memory subsystems results in a multi-bus computer architecture. More particularly, in a microprocessor with a cache memory subsystem, a first bus, for convenience referred to as a CPU local bus, connects the microprocessor (for example an 80386), a cache control (which can include an 82385 cache controller) and a random access memory serving as the cache memory. The CPU local bus may be interconnected via a buffer with a second bus system, for convenience termed a system bus.

Other components can be connected to the system bus (either directly or indirectly) such as main memory, input- output equipment, read only memory, etc.

A prime reason for using a cache subsystem is to speed up memory operations by enabling the processor to read from the cache memory subsystem to the extent the required information is stored therein. The cache memory subsystem usually has a speed advantage over the main memory. To the extent that memory accesses can be restricted to the cache subsystem then, the processor does not require access to the system bus. This significantly frees up the system bus for use in other operations, for example input- output operations, DMA, etc. This is another advantage of microcomputer systems with cache memory subsystems.

In order to maintain orderly sequencing of the various operations that must be performed, operations are divided into units of time referred to as clock states. In a microcomputer system employing for example the 80386 processor and 82385 cache controller, the fastest memory operations require two clock cycles, each of two clock states. Other operations which require more than two clock cycles are referred to as having that number of wait states which is equal to the difference between the number of clock cycles required for the operation and the two clock cycles minimum (which is also referred to as a zero wait state).

Since speed of operations is an important criterion, it is an advantage to ensure that, to the extent feasible, operations can be fit into zero wait state operations as opposed to one or two wait state operations, etc.

Cache accesses are one type of operation which is a zero wait state operation.

While desirably as many memory accesses as possible are handled by the cache subsystem, it is of course necessary at times to access main memory. One principle used in operating a cache subsystem is that in the event of a read miss, i.e. a read operation in which the required information is not found in the cache subsystem, the information when read from main memory is immediately written to the cache subsystem. By using this principle, the information read is thereafter available in the cache subsystem (unless overwritten) so that subsequent accesses to the same information need not access main memory. Thus in the event of a read miss, two operations are required, (1) accessing main memory to read the required information to make it available to the processor and (2) writing the cache subsystem with the information just read from main memory.

It is a characteristic of the 82385 that, in the event of a read miss condition, the information which is to be accessed from main memory is necessarily rewritten to the cache subsystem prior to the time that it must be available to the processor. The chip manufacturer recognizes this condition and suggests one of two alternatives, i.e. either select sufficiently fast main memory so that the data can be accessed and made available to the 82385 within the time necessary for a fixed wait state operation, or increase the length of the cycle from what would normally be required by adding additional wait state(s) as necessary. While either alternative is practical, the first alternative impacts system cost since the faster memories required to meet the timing specifications of the 82385 necessarily are more expensive, and the second alternative introduces delay into any read miss operation by imposing one or more wait states.

It is an object of the present invention to eliminate the necessity for this choice so as to improve system tolerance to slower memory components without impacting wait state parameters for read miss operations.

According to the invention, there is provided a data processing system having a processor, a cache ~~memory~~, a cache-to-main memory interface and a cache controller adapted to generate separate control signals for required activities in respect of the cache, including cache write ~~enable~~ signals, and in respect of the cache interface in order to provide data access as required by the processor, wherein logic is interposed between the cache controller and the cache, receiving control signals ~~that relate~~ ^{enable} ~~relative~~ to both the cache and the cache interface, including cache write ~~enable~~ signals, to delay the same in the presence of received combinations of control signals indicating a cache read miss.

An embodiment of the invention will now be described, by way of example, with reference to the accompanying drawings, in which:-

Fig. 1 is an overall three-dimensional view of a typical microcomputer system employing the present invention;

Fig. 2 is a detailed block diagram of a majority of the components of a typical microcomputer system employing the present invention;

Fig. 3 is a detailed block diagram of a cache controller, a random access memory, a logic circuit and buffers embodying the present invention;

Fig. 4 is a timing diagram illustrating the different timing requirements of the cache controller and the microprocessor; and

Figs. 5A-5C show operation of the invention for read miss cache writes as well as cache writes not caused by a read miss.

Fig. 1 shows a typical microcomputer system in which the present invention can be employed. As shown, the microcomputer system 10 comprises a number of components which are interconnected together. More particularly, a system unit 30 is coupled to and drives a monitor

20 (such as a conventional video display). The system unit 30 is also coupled to input devices such as a keyboard 40 and a mouse 50. An output device such as a printer 60 can also be connected to the system unit 30. Finally, the system unit 30 may include one or more disk drives, such as the disk drive 70. As will be described below, the system unit 30 responds to input devices such as the keyboard 40 and the mouse 50, and input/output devices such as the disk drive 70 for providing signals to drive output devices such as the monitor 20 and the printer 60. Of course, those skilled in the art are aware that other conventional components can also be connected to the system unit 30 for interaction therewith. In this arrangement, the microcomputer system 10 includes (as will be more particularly described below) a cache memory subsystem such that there is a CPU local bus interconnecting a processor, a cache control and a cache memory which CPU local bus is coupled via a buffer to a system bus. The system bus is interconnected to and interacts with the I/O devices such as the keyboard 40, mouse 50, disk drive 70, monitor 20 and printer 60. Furthermore, the system unit 30 may also include a third bus comprising a Micro Channel (TM) bus for interconnection between the system bus and other (optional) input/output devices, memory, etc.

Fig. 2 is a high level block diagram illustrating the various components of a typical microcomputer system embodying with the present invention. A CPU local bus 230 (comprising data, address and control components) provides for the connection of a microprocessor 225 (such as an 80386), a cache control 260 (which may include an 82385 cache controller) and a random access cache memory 255. Also coupled on the CPU local bus 230 is a buffer 240. The buffer 240 is itself connected to the system bus 250, also comprising address, data and control components. The system bus 250 extends between the buffer 240 and a further buffer 253.

The system bus 250 is also connected to a bus control and timing element 265 and a DMA controller 325. An arbitration control bus 340

couples the bus control and timing element 265 and a central arbitration element 335. Memory 350 is also connected to the system bus 250. The memory 350 includes a memory control element 351, an address multiplexer 352 and a data buffer 353. These elements are interconnected with memory elements 360 through 364, as shown in Fig. 2.

A further buffer 254 is coupled between the system bus 250 and a planar bus 270. The planar bus 270 includes address data and control components, respectively. Coupled along the planar bus 270 are a variety of I/O adaptors and other components such as the display adaptor 275 (which is used to drive the monitor 20), a clock 280, additional random access memory 285, an RS 232 adaptor 290 (used for serial I/O operations), a printer adaptor 295 (which can be used to drive the printer 60), a timer 300, a diskette adaptor 305 (which cooperates with the disk drive 70), an interrupt controller 310 and read only memory 315. The buffer 253 provides an interface between system bus 250 and an optional feature bus such as the Micro-Channel (TM) bus 320 represented by the Micro-Channel (TM) sockets. Devices such as memory 331 may be coupled to the bus 320.

While data for cache writes may be derived from memory 350, such data may also be derived from other memory such as memory installed on the Micro-Channel (TM) bus.

In a conventional 80386/82385 microcomputer system, the user is faced with selecting one of two unattractive alternatives in connection with memory operations and particularly operations following a read miss.

In such a system, during a read cycle, the microprocessor 225 places an address on the address component of the CPU local bus 230. The cache control 260 responds to the address and determines whether the required information is contained in the cache memory 255. In the event that the information is found in the cache memory 255, the cache memory 255

is addressed, it places the data on the data component of the CPU local bus where it is available to the microprocessor 225. In the event that the information required is not in the cache memory 255, then the cache control 260 enables the buffer 240 so that the address can be passed from the CPU local bus 230 to the system bus 250. When the required address reaches the system bus 250 it is available to memory 350 and after a period of time required by the characteristics of the memory 350, the addressed data appears on the data component of the system bus 250. The data is coupled through the buffer 240 where it is then available both to the cache memory 255 and to the microprocessor 225. The information will be used in the cache memory 255 to be written therein so that in the event the same information is required, a further access to memory 350 is not required. Similar operations can occur with memory on the planar bus 270 or on the optional feature bus. In the case of the optional feature bus, address information is coupled through buffer 253 to memory on the optional feature bus. Data from such memory is coupled back through buffer 253, system bus 250 to the CPU local bus 230 via buffer 240.

As has been mentioned, however, the characteristics of the 82385 require that information extracted as a result of a read miss be available at the CPU local bus for writing into the cache memory 255 before it is required by the processor 225. In other words, the timing requirements placed on the memory 350 or other memory by the 82385 are more stringent than the timing requirements dictated by the 80386. The manufacturer of the 82385 thus suggests that the user can either:

- 1) select memory elements (such as the elements 361-364) in main memory 350 or other memory to be fast enough to meet the timing requirements of the 82385 with a specified wait state, or
- 2) ensure that the dual operation initiated by a read miss occupies an additional wait state.

As will be described below, the present invention eliminates this requirement to select one of two undesirable alternatives by in effect eliminating the stringent timing requirements dictated by the cache control 260 so that the timing requirements placed on main memory 350 or other memory are no more stringent than those dictated by the microprocessor 225.

In order to implement the invention, the cache control 260 has added to it several logic elements, beyond the 82385 chip as is illustrated in Fig. 3.

Fig. 3 shows, in block diagram form, the detailed components of the cache control 260 and the cache memory 255. More particularly, the cache memory 255 is coupled to the data element of the CPU local bus 230 and via the latch 255L to the address component of the CPU local bus 230.

The cache control 260 includes the 82385 cache controller as well as a logic element 261. Fig. 3 shows those output signals of the 82385 which are pertinent to cache write operations. Those include the cache latch enable (CALEN), the cache write enables including CWEA (for bank A) and CWEB (for bank B) and the chip select signals CS0, CS1, CS2 and CS3.

As shown in Fig. 3, logic 261 receives, as inputs, the CWEA and CWEB along with two clocking signals, CLK and CLK2 (the former is exactly half the rate of the latter) and BUSRD (indicating a system bus read). BUSRD becomes active for operations which take place on the system bus 250. Any memory (other than cache) read takes place (at least in part) on the system bus 250.

Logic 261 responds to its input signals and under the appropriate circumstance generates DCWEA (Delayed Cache Write Enable A) or DCWEB (Delayed Cache Write Enable B). More particularly, in the event CWEA

is active and BUSRD is active (thus unmistakably indicating a read miss), then at the appropriate time DCWEA will be generated. In a like fashion, in the presence of an active CWEB as well as an active BUSRD, then at the appropriate time DCWEB will be generated.

The logic 261 outputs DCWEA and DCWEB each provide an input to the associated logic gate 263A or 263B. Each of these gates also receives a corresponding input from the associated output of the 82385, i.e. gate 263A has its other input provided by CWEA and likewise 263B has its other input provided by CWEB.

In addition to the logic gates 263A and 263B, the cache control 260 further includes a buffer 262 comprising buffer elements 262a-262e, one for each of the signals CALEN, CS0, CS1, CS2 and CS3. As shown in Fig. 3, each element of the buffer 262 is a logic gate which is permanently, partially enabled (by one input being tied to the appropriate potential). The other input to the buffer element comes from the corresponding output of the 82385. The output of the buffer element 262a is input as the control input to the latch 255L. The outputs of the buffer elements 262b-262e are directly input to the cache memory 255 as CS0-CS3.

Before operation of the components of Fig. 3 is described, reference is made to Fig. 4 to illustrate appropriate timing.

Fig. 4 shows three timing diagrams. Each of the diagrams begins at the time a system bus operation occurs, i.e. when BUSRD becomes active. The line labeled 80386 shows that at a time MT1 following commencement of a read miss, data output from memory is valid. The time MT1 is the timing tolerance imposed on the memory by the requirements of the 80386.

On that line of Fig. 4 labeled 82385, the timing requirements imposed by the 82385 are illustrated. More particularly, the 82385 generates

CWEA and CWEB such that valid data must be available from memory by the time MT2, i.e. at a time before valid data is required by the 80386. Thus Fig. 4 shows the more stringent timing requirement (MT2) of the 82385 compared to the more lenient timing requirements (MT₁) of the 80386.

Fig. 4 also shows the effect of the logic 261. More particularly, the DCWEA and/or DCWEB outputs of the logic 261 track the more lenient timing requirements of the 80386. More particularly, the DCWEA and/or DCWEB are delayed by the "delay" shown in Fig. 4 as compared to the CWEA and CWEB. Thus, by this delay, the more lenient timing requirements of the 80386 are the only timing requirements imposed on main memory 350 in the event of a read miss. This enables more inexpensive memory components to be used as compared to the components which would have been required for the 82385 in order to complete a read miss within a specified number of wait states.

The internal components of the logic 261 implement the following logic equations:

$\text{/DCWEA} = \text{/BUSRD} \ \& \ \text{/CLK} \ \& \ \text{/CWEA}$

$\text{/DCWEB} = \text{/BUSRD} \ \& \ \text{/CLK} \ \& \ \text{/CWEB}$

where the operator "&" is a logical AND, the operator "/" represents negation and the other signal elements (aside from BUSRD) have already been defined.

Figs. 5A-5C are useful in explaining the operation of the elements shown in Fig. 3. More particularly, Fig. 5A shows a typical CWE signal. The signal has two transitions, a first low going transition and a second high going transition. The cache 255 is arranged to effect a write event on the high going transition of CWE such as is illustrated in Fig. 5A.

Fig. 5B shows typical CWE, DCWE and the output of the gate (either 263A or 263B) for a read miss condition. As shown on a first line, the CWE signal is generated by the 82385. Under a read miss condition, the logic element 261 generates the DCWE delayed from the CWE by an appropriate delay. The third line of Fig. 5B (labelled GATE) shows the output of the appropriate gate (263A or 263B). More particularly the output of the gate 263 has a low going transition produced as a result of the low going transition CWE. When the CWE signal undergoes a high going transition, the gate output remains low because of the low input from DCWE. It is only when DCWE goes high that the output of the gate goes high so that timing of the write event is occasioned by the high going transition of DCWE. As is seen in the third line of Fig. 5B, the write event is delayed by the delay interposed by logic element 261 as compared with the high going transition of CWE.

Fig. 5C shows the operation during a cache write other than one occasioned by a read miss. The first line of Fig. 5C (labelled CWE) shows the CWE signal as generated by the 82385. Because Fig. 5C illustrates a cache write not occasioned by the a read miss, the DCWE shows no transition at all (BUSRD remains inactive). Accordingly, the output of the gate (263A or 263B) is synchronous with the CWE such that the write event is not delayed at all.

In an embodiment of the invention actually constructed, the duration DELAY is on the order of 25 nanoseconds.

The logic equations which have been referenced above are reproduced immediately below. In this material the symbols have the following meanings associated with them:

<u>Symbol</u>	<u>Definition</u>
/	Negation
:=	A registered term, equal to
=	A combinatorial term, equal to
&	Logical AND
+	Logical OR

Logic Equations

$$\begin{aligned} \text{/BUSRD} &:= \text{BUSRD} \ \& \ \text{BUSCYC385} \ \& \ \text{/BADS} \ \& \ \text{/(BW/R)} \ \& \ \text{CLK} \\ &+ \text{BUSRD} \ \& \ \text{/PIPECYCLE385} \ \& \ \text{/(BW/R)} \ \& \ \text{CLK} \\ &+ \text{/BUSRD} \ \& \ \text{BREADY} \\ &+ \text{/BUSRD} \ \& \ \text{/MISS1} \ + \ \text{/BUSRD} \ \& \ \text{/CLK} \end{aligned} \quad (1)$$

$$\begin{aligned} \text{/BUSCYC385} &:= \text{BUSCYC385} \ \& \ \text{/BADS} \ \& \ \text{CLK} \\ &+ \text{BUSCYC385} \ \& \ \text{/PIPECYC385} \ \& \ \text{CLK} \\ &+ \text{BUSCYC385} \ \& \ \text{/BT2} \ \& \ \text{CLK} \\ &+ \text{/BUSCYC385} \ \& \ \text{BREADY} \ + \ \text{/BUSCYC385} \ \& \ \text{/CLK} \end{aligned} \quad (2)$$

$$\begin{aligned} \text{/PIPECYC385} &:= \text{PIPECYC385} \ \& \ \text{/BADS} \ \& \ \text{/BUSCYC385} \ \& \ \text{CLK} \ \& \ \text{/BREADY} \\ &+ \text{PIPECYC385} \ \& \ \text{/MISS1} \ \& \ \text{BT2} \ \& \ \text{/BUSCYC385} \ \& \ \text{CLK} \ \& \ \text{/BREADY} \\ &+ \text{/PIPECYC385} \ \& \ \text{/CLK} \end{aligned} \quad (3)$$

$$\begin{aligned} \text{/MISS1} &:= \text{MISS1} \ \& \ \text{BUSCYC385} \ \& \ \text{CPUNA} \ \& \ \text{/BADS} \ \& \ \text{/(BW/R)} \ \& \ \text{CLK} \ \& \ \text{NCA} \\ &+ \text{MISS1} \ \& \ \text{/BUSCYC385} \ \& \ \text{/BADS} \ \& \ \text{/(BW/R)} \ \& \ \text{CLK} \ \& \ \text{NCA} \ \& \ \text{/BREADY} \\ &+ \text{/MISS1} \ \& \ \text{/CLK} \\ &+ \text{/MISS1} \ \& \ \text{BREADY} \end{aligned} \quad (4)$$

/CPUNA:=/MISS1 & CLK & CPUNA & /NACACHE (5)
 + /MISS1 & CLK & CPUNA & /BREADY & /BUSCYC385
 + /CPUNA & /CLK + /CPUNA & /MISS1 & CLK
 + /CPUNA & CLK & BREADY
 + /CPUNA & BUSCYC385 & NACACHE & CLK

/BT2:=BUSCYC385 & PIPECYC385 & /BADS & CLK & BT2 (6)
 + BUSCYC385 & /PIPECYC385 & BADS & CLK & NACACHE & BT2
 + MISS1 & /BUSCYC385 & /BADS & /(BW/R) & CLK & NCA & /BREADY
 + /MISS1 & /BREADY & /BUSCYC385 & CLK
 + /BT2 & BREADY & NACACHE
 + /CLK & /BT2

In the foregoing logic equations the following signals are described or referred to in the cited Intel publications:

BADS

BREADY

(BW/R) actually referred to as BW/R, the
 parenthesis are used to indicate that the
 entire term is one signal

CLK

BADS, when active indicates a valid address on the system bus 250.
 BREADY is a ready signal from the system bus 250 to the CPU local bus 230. BW/R defines a system bus 250 Write or Read. CLK is a processor clocking signal which is in phase with the processor 225.

Equations (1)-(6) define:

BT2

BUSCYC385

BUSRD

CPUNA

MISS1

PIPECYC385 in terms of the defined signals, the signals described or referred to in the cited Intel publications and NCA and NACACHE. BT2 reflects the state of the system bus 250. The state BT2 is a state defined in the cited Intel publications.

BUSCYC385 also reflects the state of the system bus 250. It is high for bus states BTI, BT1, BT1P and low for bus states BT2, BT2P and BT2I (again these are bus states referenced in the cited Intel publications).

/BUSRD is active during reads occurring on the system bus 250.

CPUNA is a signal to the 80386 allowing pipelined operation.

MISS1 is active defining the first cycle in a double cycle for handling 64 bit read to cacheable devices.

PIPECYC385 is active during BT1P (which is a bus state referred to in the cited Intel publications).

NCA is a signal created by decoding the address component on the CPU local bus 230 to reflect, when active, a non-cacheable access. Cacheability is determined by a tag component (A31 to A17) and programmable information defining what tags (if any) refer to cacheable as opposed to non-cacheable addresses,

NACACHE is a signal similar to the BNA signal. BNA is a system generated signal requesting a next address from the CPU local bus 230, and is referenced in the cited Intel publications. NACACHE differs from BNA only in respect of the fact that BNA is created for 32K cache while NACACHE is created for a 64K cache. So long as the cache memory is 32K, as cited in the Intel publications the NACACHE signal referred to here could be replaced by the BNA signal. In an embodiment actually constructed the logic element 261 was in the form of a programmable logic array. It should be apparent that other and further logic devices can be used to perform the identical function.

CLAIMS

1. A data processing system having a processor, a cache ~~memory~~, a cache-to-main memory interface and a cache controller adapted to generate separate control signals for required activities in respect of the cache, including cache write^{enable} signals, and in respect of the cache interface in order to provide data access as required by the processor, wherein logic is interposed between the cache controller and the cache, receiving control signals ^{that relate} ~~relative~~ to both the cache and the cache interface, including cache write^{enable} signals, to delay the same in the presence of received combinations of control signals indicating a cache read miss.
2. A system as claimed in claim 1, wherein the cache is a multi-chip memory and the logic also connects chip select signals from the cache controller to the cache ~~memory~~ and further comprises buffer means for delaying transmission of the chip select signals, the buffer means having an input for each of the chip select signals and an output for each of the chip select signals, outputs of the buffer means being coupled to chip select terminals of the cache ~~memory~~.
3. A system as claimed in claim 2, wherein the buffer means further includes an input for a cache latch enable signal from the cache controller and an output coupled to an address latch of the cache ~~memory~~.
4. A system as claimed in any preceding claim, wherein the processor, cache, cache controller and interface are interconnected by a local bus, and the interface connects to a system bus to which the main memory is connected, the cache controller including first means responsive to a detected read miss condition for passing addressing information from the local bus to the system bus for addressing the



main memory and second means responsive to data on the system bus retrieved from the main memory for coupling the data to the local bus under control of a bus read control signal/^{so that}~~for~~ writing into the cache/^{can take} place when ^a~~the~~ delayed/^{cache}write enable signal.

5. A system as claimed in claim 4, wherein the system bus is coupled to an optional feature bus to which a memory device is coupled in such a way that the memory device can be addressed and data so accessed returned to the local bus via the system bus.

6. A system as claimed in claim 4 or claim 5, wherein the control signal combination to which the logic responds to generate ^{the}~~a~~ delayed/^{cache}write enable signal is the concurrence of ^{cache}~~a~~ write enable signal and a bus read signal.

7. A system as claimed in any one of the preceding claims, wherein the cache ~~memory~~ includes a first and a second memory bank and, correspondingly, first and second write enable inputs, the cache controller having a first/^{cache}write enable output corresponding to the first memory bank and a second/^{cache}write enable output corresponding to the second memory bank and the delay logic separately responding to the first or second ^{cache}write enable outputs of the cache controller to produce a first or second delayed/^{cache}write enable signal at the corresponding first or second write enable input.

8. A system as claimed in claim 7, wherein the logic includes a gate connecting to the first write enable input and responsive to a first cache write enable output from the cache controller and to a concurrent delayed/^{cache}write enable signal to enable the first bank and a corresponding second gate responsive to a second ^{cache}write enable output from the controller and a concurrent delayed/^{cache}write enable signal to enable the second bank via the second enable input.

