

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
15 April 2004 (15.04.2004)


PCT

(10) International Publication Number
WO 2004/031964 A2

- (51) International Patent Classification⁷: G06F 12/08
- (21) International Application Number: PCT/US2003/028934
- (22) International Filing Date: 18 September 2003 (18.09.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 10/261,642 30 September 2002 (30.09.2002) US
- (71) Applicant: ADVANCED MICRO DEVICES, INC. [US/US]; One AMD Place, Mail Stop 68, P.O. Box 3453, Sunnyvale, CA 94088-3453 (US).
- (72) Inventor: CONWAY, Patrick; 973 Dolores Avenue, Los Altos, CA 94024-6049 (US).
- (74) Agents: CAYWOOD, Michael et al.; Advanced Micro Devices, Inc., 5204 East Ben White Boulevard, Mail Stop 562, Austin, TX 78741 (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR REDUCING OVERHEAD IN A DATA PROCESSING SYSTEM WITH A CACHE



(57) Abstract: A data processor (120) recognizes a special data processing operation in which data will be stored in a cache (124) for one use only. The data processor (120) allocates a memory location to at least one cache line of the cache (124). A data producer such as a data communication driver program running on a central processing unit (122) then writes a data element to the allocated memory location. A data consumer (160) reads the data element by sending a READ ONCE request to a host bridge (130). The host bridge (130) provides the READ ONCE request to a memory controller (126), which reads the data from the cache (124) and de-allocates the at least one cache line without performing a writeback from the cache to a main memory (170). In one form the memory controller (126) de-allocates the at least one cache line by issuing a probe marking the next state of the associated cache line as invalid.

WO 2004/031964 A2



Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

METHOD AND APPARATUS FOR REDUCING OVERHEAD
IN A DATA PROCESSING SYSTEM WITH A CACHE

TECHNICAL FIELD

5

The present invention generally relates to data processing systems, and more specifically to data processing systems with caches.

BACKGROUND ART

10

A known way to increase the performance of a computer system is to include a local, high-speed memory known as a cache. A cache increases system performance because there is a high probability that once the central processing unit (CPU) accesses a data element at a particular address, its next access will be to an adjacent address. The cache fetches and stores data that is located adjacent to the requested piece of data from a slower, main memory or lower-level cache. In very high performance computer systems, several caches may be placed in a hierarchy. The cache that is closest to the CPU, known as the upper-level or "L1" cache, is the highest-level cache in the hierarchy and is generally the fastest. Other, generally slower caches are then placed in descending order in the hierarchy starting with the "L2" cache, etc., until the lowest level cache that is connected to main memory.

20

A cache follows certain policies when storing and discarding data. For example, many processors follow an "allocate-on-write" policy that dictates that the cache line corresponding to memory locations that are written by the CPU will be stored in the cache. Typically caches follow a policy known as least-recently-used (LRU) to determine which location to discard to make room for a new data element once all locations have been filled.

25

Caches typically contain multiple status bits to indicate the status of the cache line to maintain data coherency throughout the system. One common coherency protocol is known as the "MOESI" protocol. According to this protocol each cache line includes status bits to indicate which MOESI state the line is in, including bits that indicate that the cache line has been modified (M), that the cache line is exclusive (E) or shared (S), or that the cache line is invalid (I). The Owned (O) state indicates that the line is modified in one cache, that there may be shared copies in other caches and that the data in memory is stale.

30

In a typical configuration all caches are combined with the CPU in the same integrated circuit and main memory is located off-chip. Main memory is the slowest and least expensive memory in the system and may be constructed of inexpensive but relatively slow dynamic random access memory (DRAM) chips. This characteristic results in bottlenecks in accessing the off-chip DRAM main memory and it is desirable to avoid these bottlenecks whenever possible. Furthermore in recent years microprocessor speeds have increased faster than DRAM access speeds, compounding the bottleneck problem.

35

DISCLOSURE OF INVENTION

In accordance with the present invention, a method is provided for reducing overhead in a data processing system that includes a cache and a lower-level memory system. A memory location is first allocated to at least one cache line of the cache. A data producer writes a data element to the at least one cache line. A read once request is received from a data consumer. In response thereto, the data element is provided to the data consumer by reading the data element directly from the cache. The cache line is subsequently de-allocated without accessing the lower-level memory system.

In one form the invention comprises a data processor having a central processing unit including a cache, a host bridge, and a memory controller. The host bridge is coupled to the central processing unit and is adapted to receive a read once request for a data element from a data consumer. The memory controller has an input coupled to the host bridge and has an output coupled to the central processing unit. The memory controller issues a probe for the data element to the central processing unit in response to receiving the read once request. The probe indicates that a next state of at least one cache line containing the requested data element will be invalid. If the cache contains the data element, the central processing unit provides the data element to the host bridge and the cache marks the at least one cache line thereof invalid.

BRIEF DESCRIPTION OF DRAWINGS

The present invention will hereinafter be described in conjunction with the appended drawing figures, wherein like numerals denote like elements, and:

FIG. 1 illustrates a timing diagram associated with the flow of data and control signals in a data processing system known in the prior art;

FIG. 2 illustrates a timing diagram associated with the flow of data and control signals according to the present invention;

FIG. 3 illustrates a block diagram of a data processing system according to the present invention; and

FIG. 4 illustrates an encoding table of a HyperTransport packet that can be used to form a READ ONCE request packet.

MODE(S) FOR CARRYING OUT THE INVENTION

FIG. 1 illustrates a timing diagram 10 associated with the flow of data and control signals in a data processing system known in the prior art. As shown in FIG. 1 the vertical axis represents time in which time flows from top to bottom. FIG. 1 depicts three devices (nodes) as vertical lines distributed along the horizontal direction. A first node, labeled "NODE A", is a data consumer. For example, NODE A may correspond to an input/output device such as a transmitter of a data communications channel. For this example NODE A may have an associated direct memory access (DMA) controller to assist the flow of data. A second node, labeled "NODE B", is a lower-level memory system and includes main memory and possibly one or more lower-level caches. A third node labeled "NODE C" is a data producer having a cache, such as a central processing unit running an application program. In the data communications context, NODE C would produce data in the form of data frames using a circular buffer system in memory for eventual transmission by a data communication device coupled to NODE A.

The flow of data and control signals will now be described. First an application program running on NODE C, such as a communication protocol driver, creates a data frame. NODE C writes the data for the data frame by performing a store operation to a memory location. Since the address has not been allocated to its local cache, NODE C initiates a store request to main memory represented by an arrow 12. The memory controller at NODE B initiates a memory access during a time 14 and sends an invalidating probe 16 to all caches including NODE C. The invalidating probe informs all caches in the system that the memory controller intends to allocate the cache line in exclusive state to NODE C for coherency purposes. When the data is returned from the lower-level memory system the memory controller forwards it to the data processor's cache on NODE C as shown by an arrow 18. The processor then writes the data into the cache line during time period 20. This type of operation is a conventional allocate-on-write operation.

After the cache write is complete the application program on NODE C signals the creation of the data frame through a notification message 22 from NODE C to NODE A. NODE A responds with a DMA read request 24 back to NODE B. Node B initiates a memory access during a time 26 and sends a probe 28 to NODE C. Since this data is present in the cache at NODE C, the data processor returns the data to satisfy this request as shown by arrow 30. The data consumer at NODE A then writes the data to the communication channel's transmission first-in, first-out (FIFO) memory, transmit buffer, or the like during time period 32.

Later after the passage of an indeterminate amount of time based on other events in the system, the modified cache line stored in the cache at NODE C becomes the least-recently-used cache line and will be cast out. NODE C thus performs a writeback by sending the modified data to the lower-level memory system at NODE B shown as arrow 36. The memory controller at NODE B writes the data to the lower-level memory system during a time period 38.

Note that this conventional operation requires one read and one write to main memory and thus occupies the main memory bus during the time periods 14 and 38. For a system implementing a high-bandwidth communication protocol this operation would be repeated frequently, consuming a significant portion of the available bandwidth. In recent years main memory speeds have increased but not as fast as

processor speeds and a reduction in the main memory bus bandwidth required for this operation would be desirable.

To decrease main memory bandwidth the inventor realized that the communication operation is a special case in which the data can be operated upon solely within the processor's cache, avoiding the need for any access to main memory at all. Moreover since the data is to be used only once, the subsequent writeback is not needed. In order to avoid the writeback the inventor has devised a method and a corresponding data processing system to signal such a condition and prevent the writeback from occurring. It can also be implemented with a general-purpose cache that responds to invalidating probes without adding to the number of cache line states. Such a method can be understood with reference to FIG. 2, which illustrates a timing diagram 50 associated with the flow of data and control signals according to the present invention. As with respect to FIG. 1, the vertical axis represents time in which time flows from top to bottom, and the horizontal axis depicts the same three nodes described above.

The operation commences with a request 52 between the CPU and the memory controller known as a "PREFETCH ALLOCATE". This special cache control instruction causes the cache line to be allocated to the cache within NODE C without first reading the contents of that memory location from DRAM. The contents of the line are set to all zeros and the line is created in the M state. The memory controller updates its directory and responds with a target done signal 54. Then the data processor's driver program writes the data directly into the cache to construct the frame during a time period 56. Once the application program has constructed the frame, it sends a notify message 58 to the DMA controller in the input/output device at NODE A. The I/O device responds with a DMA read request 60 to NODE B. Node B initiates a memory access and sends a probe 62 to NODE C. Node C satisfies the read request by providing the data as indicated by an arrow 64. NODE A then writes the data to its internal FIFO or transmit buffer during time period 66 for subsequent transmission. When NODE A reads the data, the cache at NODE C de-allocates the cache line without the writeback that would normally be required. This operation allows the data processing system to avoid all accesses to the lower-level memory system by using the PREFETCH ALLOCATE instruction and removing (de-allocating) the data from the cache concurrently with the read.

There are at least three ways to de-allocate the cache line without writing back its contents to main memory. First the I/O controller at NODE A could add a READ ONCE attribute signal to indicate that the cache line is to be invalidated without being written back. Second the CPU at NODE C could decode the I/O address for the DMA read request and if the read falls within a predefined address window, invalidate the line. Third, the data processor's instruction set could be modified to add a "CFLUSH DEALLOCATE" command to de-allocate a single cache line. As described more fully below, the inventor prefers the first technique to take advantage of some available characteristics of the on-chip I/O controllers. This technique and an apparatus that implements it will now be more fully described with reference to FIGs. 3 and 4 below.

FIG. 3 illustrates a block diagram of a data processing system 100 according to the present invention. Data processing system 100 includes generally a data processor in the form of a single-chip microprocessor 120, an input/output (I/O) device 160, and a dynamic random access memory (DRAM) 170. Microprocessor 120 includes generally a central processing unit (CPU) 122, a memory controller 126, a crossbar switch labeled "XBAR" 128, and three host bridges 130, 140, and 150 each labeled "HT" for HyperTransport,

described more fully below. CPU 122 is a processor adapted to execute instructions of the so-called x86 instruction set. The x86 instruction set is based on the instruction set of the 8086 microprocessor first manufactured by Intel Corporation of Santa Clara, California. CPU 122 however includes many sophisticated functions for high-performance execution of x86 programs including pipelining and superscalar design. CPU
5 122 includes at least one cache 124 for storage of frequently used data. In a preferred form CPU actually includes two L1 caches, one for instructions and the other one for data, and an L2 cache that is shared by the instruction and data streams. In this form the shared L2 cache stores the READ ONCE data. However the invention is applicable to any CPU having at least one cache.

Memory controller 126 is the mechanism for data transfer between microprocessor 120 and DRAM
10 170. Memory controller 126 offloads the task of initiating and terminating memory accesses from CPU 122. It includes internal queues to allow efficient use of the external bus to DRAM 170. In other embodiments DRAM 170 could be replaced by a lower-level memory system including one or more additional caches and main memory, by static RAM, by nonvolatile memory, etc.

XBAR 128 is a switching/multiplexing circuit designed to couple together the buses internal to
15 microprocessor 120.

Host bridges 130, 140, and 150 are coupled to devices external to microprocessor 120 over output channels 132, 142, and 152, respectively, and over input channels 134, 144, and 154, respectively. Each of host bridges 130, 140, and 150 comply with the HyperTransport™ I/O Link Specification, Revision 1.03, ©
20 2001 HyperTransport Technology Consortium, and are able to achieve a throughput of 3.2 GB/sec when using a 1600 MHz data rate. HyperTransport technology is a packet-based link implemented on two independent unidirectional sets of wires. So for example host bridge 130 communicates with I/O device 160 over an output connection 132 and an input connection 134. The HyperTransport link is nominally point-to-point and connects two devices. Chains of HyperTransport links can also be used as an I/O channel, connecting I/O devices and bridges to a host system.

The HyperTransport link is designed to deliver a high-performance and scalable interconnect between
25 CPU, memory, and I/O devices. The HyperTransport link uses low swing differential signaling with on-die differential termination to achieve very high data rates. The HyperTransport link uses scalable frequency and data width to achieve scalable bandwidth.

I/O device 160 is an input/output device that functions as a consumer of data. In the illustrated
30 embodiment I/O device 160 is a controller that implements the local area network communication protocol standardized by the Institute of Electrical and Electronics Engineers (IEEE) under the auspices of the IEEE 802.3 committee, commonly referred to as "Ethernet". However the present invention is applicable to other data consumers as well.

In operation an application program running on CPU 122 prepares a frame for transmission and stores
35 the frame in cache 124. As described above with respect to FIG. 2, CPU 122 first performs a PREFETCH ALLOCATE command so that the cache lines are allocated to cache 124 without the necessity of accessing DRAM 170. I/O device 160 accesses this data prior to sending it out on the external Ethernet bus. According to the HyperTransport standard, I/O device 160 can request this data element by sending a block read request

packet over the HyperTransport link. HT block 130 receives the block read request and forwards it to memory controller 126 over XBAR 128. Memory controller 126 then probes the on-chip devices to determine whether they contain a valid copy of the requested data element. In this example cache 124 has a valid copy of this data. Cache 124 completes the access by forwarding the data to I/O device 160 through HT block 130.

5 In existing data processors cache 124 would follow a known policy such as the LRU replacement policy. At an appropriate time, when this data element becomes the least-recently-used data element and cache 124 needs to allocate a new cache line, cache 124 would write the data to DRAM 170 through memory controller 126. However in some circumstances such as the producer-consumer data communication example this data will be used only once by the system and the write to DRAM 170 would be superfluous. Thus
10 according to the present invention I/O device 160 sends a new packet type known as a READ ONCE packet over HyperTransport input channel 134. The READ ONCE packet preferably uses a reserved encoding of the command (Cmd) field of the read request packet, not currently defined in the HyperTransport standard, to signal that its use of the requested data is for a READ ONCE purpose. The READ ONCE request allows microprocessor 120 to suppress the DRAM access that would normally be required on writeback, and thus
15 conserves bandwidth on the main memory bus.

The mechanism for suppressing this future DRAM access will now be described in detail. In response to receiving this novel READ ONCE request, memory controller 126 issues probe with an indication that the next state of the cache line of any device having the requested data element will be invalid (I). Thus if it stores the data element, cache 124 will forward the requested data to memory controller 126 and mark
20 invalid the cache line in which the data element is located. By marking the cache line containing the READ ONCE data as invalid, microprocessor 122 avoids the need to write the data back to DRAM 170 when a valid cache line is cast out of cache 124.

FIG. 4 illustrates a diagram of a HyperTransport packet 200 that can be used to form a READ ONCE request packet. This packet structure is known as the request packet with address. The packet is four bytes
25 long with an additional four-byte address. According to the present invention I/O device 160 signals a READ ONCE request by inserting a new READ ONCE encoding into the command field as described further below.

The various fields associated with this packet will now be described. SeqID[3:0] (formed by concatenating bits 7 and 6 of Bit-time 0 with bits 6 and 5 of Bit-time 1) is used to tag groups of requests that were issued as part of an ordered sequence by a device and must be strongly ordered within a virtual channel.
30 All requests within the same I/O stream and virtual channel that have matching nonzero SeqID fields must have their ordering maintained. The SeqID value of 0x0 is reserved to mean that a transaction is not part of a sequence. Transactions with this value have no sequence-ordering restrictions, although they may be ordered for other reasons as described in Chapter 6 of the HyperTransport standard.

The Cmd[5:0] field is used to signal various commands in the HyperTransport protocol. The 2001
35 HyperTransport standard defines various encodings of the Cmd[5:0] field, and reserves other encodings. TABLE I shows the defined and reserved encodings:

TABLE I

Code	VChan	Command	Comments/Options	Packet Type
000000	-	NOP	Null packet. Contains flow control information.	Info
000001		Reserved-HOST		
000010	NPC	Flush	Flush posted writes	Request
000011 0001xx		Reserved-HOST		
x01xxx x01xxx	NPC or PC (bit 5)	Wr (sized)	Write request [5] Defines whether request is posted: 0: Nonposted 1: Posted [2] Defines the data length: 0: Byte 1: Doubleword [1] Defines bandwidth/latency requirements: 0: Normal 1: Isochronous [0] Indicates whether access requires host cache coherence (ignored if access is not to host memory): 0: Noncoherent 1: Coherent	Req/Addr/Data

TABLE I (Continued)

01xxxx 01xxxx	NPC	Rd (sized)	Read requests [3] Defines ordering requirements for response: 0: Response may not pass posted requests 1: Response may pass posted requests [2] Defines the data length: 0: Byte 1: Doubleword [1] Defines bandwidth/latency requirements: 0: Normal 1: Isochronous [0] Indicates whether access requires host cache coherence (ignored if access is not to host memory): 0: Noncoherent 1: Coherent	Req/Addr
100xxx		Reserved-I/O		
110000	R	RdResponse	Read Response	Resp/Data
110001 110010		Reserved-HOST		
Code	VChan	Command	Comments/Options	Packet Type
110011	R	TgtDone	Tell source of request that target is done.	Response
11010x		Reserved-HOST		
11010x		Reserved-HOST		
11011x		Reserved-I/O		
11100x		Reserved-HOST		
11100x		Reserved-HOST		
11100x		Reserved-HOST		
111010	PC	Broadcast	Broadcast message	Req/Addr
111011		Reserved-HOST		
111100	PC	Fence	Fence posted requests	Request
111101	NPC	Atomic RMW	Atomic Read-Modify-Write	Req/Addr/Data
111110		Reserved-I/O		
111111	-	Sync/Error	Link Synchronization and Error Packet	Info

PassPW indicates that that this packet is allowed to pass packets in the posted request channel of the same I/O stream. Otherwise, this packet must stay ordered behind them. This bit should be cleared to maintain the standard personal computer interface- (PCI-) like ordering model of HyperTransport technology. Systems that do not require PCI-like ordering may set PassPW for higher performance.

5 The UnitID[4:0] field serves to identify participants in a transaction. Since all packets are transferred either to or from the host bridge at the end of the fabric, either the source or destination node is implied. The value 0 is reserved for the UnitID of the host bridge. See Section 4.2 of the 2001 HyperTransport standard for more details on the use of UnitID. Nodes with multiple logical I/O streams can own multiple UnitID values.

10 The HyperTransport protocol reserves two bytes, transmitted during bit times 2 and 3, to be specific to the command.

 Addr[39:2] represents the doubleword address accessed by the request. Not all address bits are included in all request types. Where finer granularity is required, byte masks are used.

15 I/O device 160 preferably signals the READ ONCE request by using any one of the Cmd[5:0] field definitions that are marked "Reserved-I/O". In addition, however, since there are very few reserved command encodings for I/O devices, the new READ ONCE packet is preferably combined with other new packet formats using only one Reserved-I/O command encoding. This objective is accomplished by using the four-bit SeqID field to encode additional conditions to differentiate the between the other new packet types.

20 It should be apparent that the foregoing discussion of the READ ONCE signaling is specific to the HyperTransport protocol. However in other embodiments of the present invention other methods of signaling the READ ONCE condition may be used. For example, special encodings of bits fields, existing or additional input/output pins and the like may be used to signal the READ ONCE condition. The READ ONCE signaling may be accomplished on other I/O buses, such as PCI, PCI-X, PCIE, by using reserved or vendor specific fields in the I/O DMA read request packet.

25 While the invention has been described in the context of a preferred embodiment, various modifications will be apparent to those skilled in the art. For example a data processing system using the disclosed method may be applicable to systems interfacing directly to main memory as well as to a lower-level cache. Furthermore the invention is applicable to systems with one or more higher-level caches. Main memory may be constructed using DRAM, SRAM, nonvolatile memory, or any other type of memory. Furthermore several methods of de-allocating the cache line after it has been read once are possible.


30 Accordingly, it is intended by the appended claims to cover all modifications of the invention that fall within the true scope of the invention.

CLAIMS

What is claimed is:

1. In a data processing system (100) having a memory hierarchy including a cache (124) and a lower-level memory system (170), a method comprising the steps of:
5 allocating a memory location to at least one cache line of the cache (124);
writing a data element to said at least one cache line by a data producer (122); and
receiving a read-once request for said data element from a data consumer (160) and in
10 response thereto providing said data element to said data consumer (160) by reading
said data element directly from the cache (124), and subsequently de-allocating said
at least one cache line without accessing the lower-level memory system (170).
2. The method of claim 1 wherein said step of de-allocating comprises the step of marking said at least one cache line as invalid in response to said step of reading.
- 15 3. The method of claim 1 wherein said step of allocating comprises the step of allocating said memory location to at least one cache line of the cache (124) without accessing the lower level memory system (170).
4. The method of claim 1 wherein said step of writing said data element comprises the step of writing at
20 least a portion of a data communication frame.
5. A method for reducing overhead in a data processing system having a memory hierarchy including a cache and a lower-level memory system, comprising the steps of:
25 receiving a read once request for a data element;
probing for the presence of said data element in at least one cache line of the cache;
indicating in conjunction with said step of probing that a next state of said at least one cache
line is to be invalid; and
30 satisfying said read once request by either, providing data from said at least one cache line of
the cache and subsequently invalidating said at least one cache line if said data
element is stored in the cache, or fetching said data element from the lower-level
memory system if said data element is not stored in the cache.
6. The method of claim 5 wherein said step of providing data from said at least one cache line comprises
35 the steps of:
initiating an access to the lower-level memory system; and
subsequently aborting said access when said data element is found to be stored in the cache.

7. The method of either claim 1 or claim 5 wherein said step of receiving said read-once request further comprises the step of receiving said read-once request using a link substantially compatible with the HyperTransport™ I/O Link Specification, Revision 1.03.
- 5 8. A data processor (120) comprising:
- a central processing unit (122) including a cache (124);
 - a host bridge (130) coupled to said central processing unit (122) and adapted to receive a read once request for a data element from a data consumer (160); and
 - 10 a memory controller (126) having an input coupled to said host bridge (130) and having an output coupled to said central processing unit (122), wherein said memory controller (126) issues a probe for said data element to said central processing unit (122) in response to receiving said read once request, said probe indicating a next state of at least one cache line containing said data element will be invalid;
- 15 wherein if said cache (124) contains said data element, said central processing unit (122) provides said data element to said host bridge (130) and said cache (124) marks said at least one cache line invalid.
9. The data processor (120) of claim 8 wherein said host bridge (130) is coupled to an input/output device (160) using a link substantially compatible with the HyperTransport™ I/O Link Specification, Revision 1.03.
- 20
10. The data processor (120) of claim 8 wherein said central processing unit (122), said host bridge (130), and said memory controller (126) are combined into a single integrated circuit.






FIG. 2




FIG. 3

4 / 4

Bit-Time	7	6	5	4	3	2	1	0
0	SeqID[3:2]			Cmd[5:0]				
1	PassPW	SeqID[1:0]	UnitID[4:0]					
2	COMMAND-SPECIFIC							
3	COMMAND-SPECIFIC							
4	Addr[15:8]							
5	Addr[23:16]							
6	Addr[31:24]							
7	Addr[39:32]							

200

FIG. 4