

METHOD AND SYSTEM FOR SPECULATIVELY SOURCING CACHE MEMORY
DATA WITHIN A DATA-PROCESSING SYSTEM

BACKGROUND OF THE INVENTION

1. Technical Field

The present invention relates to a method and system for sharing cache memory data in general and, in particular, to a method and system for sharing cache memory data between a processing unit and an I/O device within a data-processing system. Still more particularly, the present invention relates to a method and system for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system.

2. Description of the Prior Art

A data-processing system, includes at least one processing unit, a system memory, and various I/O devices. A processing unit may include a processor core having multiple registers and execution units for carrying out program instructions. In addition, the processing unit may have one or more primary caches (i.e., level one or L1 caches), such as an instruction cache and/or a data cache, which are implemented utilizing high-speed memories. Further, the processing unit also may include additional caches, typically referred to as a secondary cache (i.e., level two or L2 cache) for supporting the primary caches such as those mentioned above.

Typically, the transfer of data from one processing unit to another processing unit or to an I/O device on a system bus without going through the system memory is

referred to as an intervention. An intervention protocol improves system performance by reducing the number of cases in which the system memory must be accessed in order to satisfy a read or read-with-intent-to-modify (RWITM) request by any one of the processing units or I/O devices within the system.

Broadly speaking, when there is an outstanding read/RWITM request by an I/O device, any one of the other processing units, attached to the system bus, that possesses the requested data within its cache(s) can source the data to the requesting I/O device. Under the traditional intervention protocol, the processing unit having the data residing in its cache will wait for a "combined" response from all processing units within the system before issuing a data bus request to source the data from its cache(s).

At the same time, the traditional intervention protocol also allows for a "retry" mechanism, and any read/RWITM request that could be satisfied by an intervention could also be interrupted by a "retry" from any one of the processing units on the system bus. If one processing unit responds with an intervention while another processing unit responds with a "retry," under a well-established rule, the retry response automatically overrules the intervention response. As a result, if there is an outstanding retry request by any one of the processing units on the system bus, the processing unit that contains the data will not issue a data bus request.

Consequently, it would be desirable to provide an improved sourcing scheme in which intervention data will be sourced to the requesting I/O device in such a way that is

less influenced by the "retries" from any of the processing units within the data-processing system.

SUMMARY OF THE INVENTION

In view of the foregoing, it is therefore an object of the present invention to provide an improved method and system for sharing cache memory data.

It is another object of the present invention to provide an improved method and system for sharing cache memory data between a processing unit and an I/O device within a data-processing system.

It is yet another object of the present invention to provide an improved method and system for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system.

In accordance with the method and system of the present invention, a data-processing system includes at least one processing unit, each having at least one cache memory and at least one intelligent I/O device. In response to a request for data by an intelligent I/O device within the data-processing system, an intervention response is issued from a processing unit within the data-processing system having the requested data. The requested data is then read from a cache memory within the processing unit before a combined response from all processing units within the data-processing system returns to the processing unit.

All objects, features, and advantages of the present invention will become apparent in the following detailed written description.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention itself, as well as a preferred mode of use, further objects, and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figure 1 is a block diagram of a data-processing system in which the present invention may be applicable;

Figure 2 is a block diagram of an exemplary data-processing system for illustrating a sourcing scheme under the prior art;

Figure 3 is a high-level logic flow diagram for illustrating a method for speculatively sourcing cache memory data from a processing unit to an I/O device within a data-processing system, in accordance with a preferred embodiment of the present invention.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

The present invention may be implemented in any data-processing system having at least one cache memory. Also, it is understood that the features of the present invention may be applicable in various multiprocessor data-processing systems, each processor having a primary cache and a secondary cache.

Referring now to the drawings and, in particular, to Figure 1, there is depicted a block diagram of a data-processing system 10 in which the present invention may be applicable. Data-processing system 10 includes multiple central processor units (CPUs) 11a-11n, and each of CPUs 11a-11n contains a primary cache. As shown, CPU 11a contains a primary cache 12a, while CPU 11n contains a primary cache 12n. Each of primary caches 12a-12n may be a sectorized cache.

Each of CPUs 11a-11n is coupled to each of secondary caches 13a-13n, respectively. Each of secondary caches 13a-13n also may be a sectorized cache. CPUs 11a-11n, primary caches 12a-12n, and secondary caches 13a-13n are connected to each other via an interconnect 15 to a system memory 14. Interconnect 15 can be either a bus or a switch. Also attached to interconnect 15 are intelligent I/O devices 16a-16n. These intelligent I/O devices 16a-16n have the ability to initiate data transfers to and from system memory 14. Intelligent I/O devices 16a-16n may include various adaptors utilized for communicating with another data-processing system via networks such as an Intranet or Internet.

As a preferred embodiment of the present invention, a CPU, a primary cache, and a secondary cache, such as CPU

11a, primary cache 12a, and secondary cache 13a as depicted in Figure 1, may be collectively known as a processing unit. Although a preferred embodiment of a data-processing system is described in Figure 1, it should be understood that the present invention can be practiced within a variety of system configurations. For example, each of CPUs 11a-11n may have more than two levels of cache memory.

With reference now to Table I, there is illustrated a number of established coherency responses from a processing unit under the prior-art intervention protocol. After an I/O device within the multiprocessor data-processing system makes a read or read-with-intent-to-modify (RWITM) request on a system bus, any processing unit within the system may issue one of the responses in accordance with Table I, after snooping.

Coherency Response	Priority	Definition
000	-	Reserved
001	3	Shared Intervention
010	-	Reserved
011	-	Reserved
100	1	Retry
101	2	Modified Intervention
110	4	Shared
111	5	Null or Clean

Table I

As depicted in Table I, the coherency responses take the form of a 3-bit snoop response signal, with the

definition of each coherency response set forth. These signals are encoded to indicate the snoop result after the address tenure. In addition, a priority value is associated with each response to allow a system logic to determine which of the coherency responses should take priority when formulating a single snoop response signal to be returned to all processing units all I/O devices on the system bus. For example, if a processing unit responds with a shared intervention response (priority 3), and another processing unit responds with a retry response (priority 1), then the processing unit with the retry response will take priority such that the system logic will return a retry coherency response to the requesting processing unit as well as to all other processing units that are attached on the system bus. This system logic may reside in various components within the system, such as a system control unit or a memory controller.

Several well-known mechanisms may be employed to ascertain which cache (of a processing unit) is the "owner" of the data that is being requested, and therefore entitled to source the data. Under the prior-art MESI protocol, if a cache holds the requested data in a "Modified" or an "Exclusive" state, that means this cache is the only one within the system which contains a valid copy of the data and is clearly the owner. If, however, a cache holds the requested data in a "Shared" state, that means the data must also be held in at least one other cache within the system. Thus, potentially, either one of the two or more caches can source the data. In such a case, several alternatives are available to determine which cache should perform sourcing.

Referring now to Figure 2, there is depicted a block diagram of an exemplary data-processing system for illustrating a sourcing scheme under the prior art. As

shown, for example, an intelligent I/O device 24 desires to make a read or RWITM request on a system bus 23, and L2 cache of processing unit 21 contains the data being requested by I/O device 24. Furthermore, the L2 cache within processing unit 20 is in an "Invalid" state, the L2 cache within processing unit 21 is in a "Modified" state, and the L2 cache within processing unit 22 does not contain the requested data. The subsequent sequence of actions will be taken by a respective L2 cache controller of each processing unit for performing the source intervention as dedicated by the prior art.

After I/O device 24 makes its read/RWITM request, the read/RWITM request is "snooped" from system bus 23 by processing unit 21, processing unit 22, and processing unit 23. An L2 cache directory lookup is performed in each of the processing units 21-23 to determine whether or not the requested data is resident in its L2 cache. Because processing unit 21 has the requested data, an intervention response will be issued by processing unit 21, and a finite state machine within processing unit 21 will be dispatched to control the following actions. If the data within the L2 cache of processing unit 21 is in a "Modified" state, a modified intervention coherency response will be issued by processing unit 21. Otherwise, if the data within the L2 cache of processing unit 21 is in a "Shared" or "Exclusive" state, a shared intervention coherency response will be issued by processing unit 21. Because the L2 cache within processing unit 20 is in an "Invalid" state and the L2 cache within processing unit 22 does not contain the requested data, each of processing units 20 and 22 will send a null coherency response.

After the issuance of the intervention response, processing unit 21 is pending for a combined response which basically includes, in this example, the coherency response from itself and from processing units 20, 22 and I/O device 24. If the returned combined response is a modified intervention coherency response, processing unit 21 may start sourcing the requested data from its L2 cache. If processing unit 20 and/or processing unit 22 request(s) a retry for whatever reason, the sourcing must yield to the retry request (i.e., the sourcing sequence will not proceed), under the established intervention protocol. For example, processing unit 22 may be in a snoop queue busy condition.

If the data in the L2 cache of processing unit 21 has not been modified or is not resident in the L1 cache (i.e., not L1 inclusive) since the snoop action has been initiated, processing unit 21 may begin to make a system bus request to the system bus arbiter (typically, the requested data must be read into a buffer by the L2 cache controller before the system bus request can begin). Otherwise, the L1 cache of processing unit 21 will be flushed and invalidated (i.e., forcing the L1 cache to "push" any modified data back to the L2 cache and invalidating the copy in the L1 cache) before any system bus request can be made. If the L1 cache of processing unit 21 is in a "Shared" state, however, only an invalidation to the L1 cache is required before making any data bus request.

Processing unit 21 then waits for a system bus grant to return. The actual data-sourcing to I/O device 24 will begin after the data bus grant is received. Once the sourcing has completed, the L2 cache of processing unit 21 will be changed from a "Modified" state to a "Shared" state for a read request and to an "Invalid" state for an RWITM

request. There is no change of state in the L2 cache of processing units 20 and 22.

Referring now to Figure 3, there is depicted a high-level logic flow diagram for speculatively sourcing cache memory data from a processing unit to an I/O device within a data-processing system, in accordance with a preferred embodiment of the present invention. Starting at block 30, a read/RWITM request is snooped from a system bus by all processing units within the system, as shown in block 31. An L2 cache directory lookup is performed to determine by each processing unit as to whether or not the requested data is resident in its L2 cache, as depicted in block 32. A null coherency response will be issued by all those processing units that do not possess the requested data (such as processing units 20 and 22 of Figure 2), as illustrated in block 33, and the process exits at block 99. On the other hand, an intervention coherency response will be issued by a processing unit that possesses the requested data (such as processing unit 21 of Figure 2), as shown in block 34.

After the issuance of the intervention coherency response, the intervening processing unit must perform certain cache housekeeping tasks, as depicted in block 35. These tasks include flushing and invalidating the data copy in the L1 cache of the intervening processing unit if the data copy in the L1 cache has been modified, or simply invalidating the data copy in the L1 cache of the intervening processing unit if the data copy in the L1 cache has not been modified.

Subsequently, the requested data is read from the L2 cache of the intervening processing unit, preferably, to a buffer, and a request for the system data bus is made to a

system bus arbiter, as illustrated in block 36. A determination is made as to whether or not the system data bus has been granted, as shown in block 37. If the system data bus has not been granted, another determination is made as to whether or not a combined coherency response has returned yet, as depicted in block 38. If the combined coherency response has not returned, the process returns to block 37.

However, if the system bus has been granted, a sourcing of the requested data from the intervening processing may begin by driving the requested data to the system bus, as illustrated in block 39. Another determination is made as to whether or not a combined coherency response has returned at this point already, as shown in block 40. If the combined coherency response has not returned yet, the process will keep waiting for the combined coherency response to return while continuing the sourcing of the requested data to the system bus.

After the combined coherency response has returned, a determination is made as to whether or not the combined coherency response is a "retry," as depicted in block 41. If the combined coherency response is a retry, then the system bus request (from block 36) will be cancelled if the system bus has not been granted yet, or the sourcing of the requested data will be aborted immediately, as illustrated in block 42. Even if the sourcing has already been completed at this point, the results will be discarded due to the retry coherency response. Otherwise, if the combined coherency response is not a retry, the sourcing of the requested data will continue, if it has not been completed, until its completion. Finally, the status of the L2 cache in the

intervening processing unit is updated accordingly, as shown in block 43, and the process exits at block 99.

As has been described, the present invention provides a method for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system. Specifically, the present disclosure describes a novel intervention implementation in which the requested data is read from the L2 cache of the intervening processing unit before the combined coherency response has returned.

The present invention has obvious performance advantages over the prior art because the delay between a read/RWITM request on the system bus and the sampling of the combined response can be several system bus clock cycles. Hence, by allowing the requested data to be read from the L2 cache of the intervening processing unit before the combined coherency response is received, the intervention latency is reduced tremendously and the overall system performance is significantly improved.

While the invention has been particularly shown and described with reference to a preferred embodiment, it will be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention.



CLAIMS

1. A method for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system, said processing unit includes at least one cache memory, said method comprising the steps of:

in response to a request for data by said intelligent I/O device within said data-processing system, issuing an intervention response by a processing unit having said requested data; and

reading said requested data from a cache memory within said processing unit prior to a combined response from all processing units within said data-processing unit returns to said processing unit.



2. The method for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system according to Claim 1, wherein said reading step further includes a step of reading said requested data from a cache memory within said processing unit by a cache controller.

3. The method for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system according to Claim 1, wherein said reading step further includes a step of reading said requested data from a cache memory within said processing unit to a buffer.

4. The method for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system according to Claim 1, wherein said request for data includes a read request or a read-with-intent-to-modify request.

5. The method for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system according to Claim 1, wherein said intervention response is a modified intervention response or a shared intervention response.

6. The method for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system according to Claim 1, wherein said



method further includes a step of stopping said reading step if said returned combined response is a retry.

7. The method for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system according to Claim 1, wherein said method further includes a step of requesting a system bus for sourcing of said requested data by said processing unit before the return of said combined response.

8. The method for speculatively sourcing cache memory data from a processing unit to an intelligent I/O device within a data-processing system according to Claim 7, wherein said method further includes a step of sourcing said requested data by said processing unit before the return of said combined response.



9. A processing unit having a cache memory capable of speculatively sourcing data to an intelligent I/O device within a data-processing system, said processing unit comprising:

means for issuing an intervention response from a processing unit within said data-processing system having a requested data, in response to a request for said requested data by said intelligent I/O device within said data-processing system; and

means for reading said requested data from a cache memory within said processing unit prior to a combined response from all processing units returns to said processing unit.



10. The processing unit having a cache memory capable of speculatively sourcing data to an intelligent I/O device according to Claim 9, wherein said means for reading is a cache controller.

11. The processing unit having a cache memory capable of speculatively sourcing data to an intelligent I/O device according to Claim 9, wherein said means for reading further includes a means for reading said requested data from a cache memory within said processing unit to a buffer.

12. The processing unit having a cache memory capable of speculatively sourcing data to an intelligent I/O device according to Claim 9, wherein said request for data includes a read request or a read-with-intent-to-modify request.

13. The processing unit having a cache memory capable of speculatively sourcing data to an intelligent I/O device according to Claim 9, wherein said intervention response from said processing unit is a modified intervention response or a shared intervention response.

14. The processing unit having a cache memory capable of speculatively sourcing data to an intelligent I/O device according to Claim 9, wherein said processing unit further includes a means for stopping said reading by said reading means if said returned combined response is a retry.



15. The processing unit having a cache memory capable of speculatively sourcing data to an intelligent I/O device according to Claim 9, wherein said processing unit further includes a means for requesting a system bus for sourcing of said requested data by said processing unit before the return of said combined response.

16. The processing unit having a cache memory capable of speculatively sourcing data to an intelligent I/O device according to Claim 15, wherein said processing unit further includes a means of sourcing said requested data by said processing unit before the return of said combined response.