

(51) International Patent Classification:

G06F 9/38 (2006.01)

(21) International Application Number:

PCT/US2012/035757

(22) International Filing Date:

30 April 2012 (30.04.2012)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/486,197 13 May 2011 (13.05.2011) US

13/220,286 29 August 2011 (29.08.2011) US

(71) Applicant (for all designated States except US):

ORACLE INTERNATIONAL CORPORATION [US/US];
500 Oracle Parkway, Redwood Shores, CA 94065 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): CHEUNG, Connie, Wai Mun [US/US]; 703 San Conrado Terrace, #6, Sunnyvale, CA 94085 (US). KONDAPANENI, Madhavi [IN/US]; 22460 Cupertino Rd., Cupertino, CA 95014 (US). LAM, Joann, Yin [US/US]; 51 Wawona Street, San Francisco, CA 94127 (US). SIVARAMAKRISHNAN, Ramaswamy [US/US]; 1323 Teresita Drive, San Jose, CA 95129 (US).

(74) Agents: SCHERER, Thomas, K. et al.; Osha . Liang LLP,

909 Fannin Street, Suite 3500, Houston, TX 77010 (US).

(81) Designated States (unless otherwise indicated, for every

kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every

kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

[Continued on next page]

(54) Title: REQUEST TO OWN CHAINING IN MULTI-SOCKETED SYSTEMS




FIG. 1

(57) Abstract: A method including: receiving multiple local requests to access the cache line; inserting, into an address chain, multiple entries corresponding to the multiple local requests; identifying a first entry at a head of the address chain; initiating, in response to identifying the first entry and in response to the first entry corresponding to a request to own the cache line, a traversal of the address chain; setting, during the traversal of the address chain, a state element identified in a second entry; receiving a foreign request to access the cache line; inserting, in response to setting the state element, a third entry corresponding to the foreign request into the address chain after the second entry; and relinquishing, in response to inserting the third entry after the second entry in the address chain, the cache line to a foreign thread after executing the multiple local requests.



-
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*
- Published:**
- *with international search report (Art. 21(3))*

REQUEST TO OWN CHAINING IN MULTI-SOCKETED SYSTEMS

BACKGROUND

[0001] A thread is the smallest unit of processing that can be scheduled for execution. Accordingly, a thread may be regarded as a sequence of instructions, including instructions to access (*i.e.*, read from and/or write to) cache lines. Many computer systems have two or more multithreaded processors concurrently executing multiple threads. Although, in general, by having multiple multithreaded processors the performance of the computer system improves, concurrent attempts to access the same cache line among threads can pose performance problems depending on the order in which the threads are granted access. This is due to the long latency associated with an inter-socket cache line transfer. Despite the existence of this inter-socket cache line transfer latency, computer system architectures with multiple multithreaded processors remain in wide-spread use.

SUMMARY

[0002] In general, in one aspect, the invention relates to a method for managing a cache line. The method comprises: receiving, from a plurality of local threads executing on a local multithreaded processor, a plurality of local requests to access the cache line; inserting, into an address chain associated with the cache line, a plurality of entries corresponding to the plurality of local requests; identifying a first entry of the plurality of entries at a head of the address chain; initiating, in response to identifying the first entry at the head and in response to the first entry corresponding to a request to own (RTO) the cache line, a traversal of the address chain; setting, during the traversal of the address chain, a state element identified in a second entry of the plurality of entries, wherein the state element corresponds to a serialization acknowledgement from a central ordering unit for a local request

corresponding to the second entry; receiving, from a foreign thread executing on a foreign multithreaded processor, a foreign request to access the cache line; inserting, in response to setting the state element, a third entry corresponding to the foreign request into the address chain after the second entry; and relinquishing, in response to inserting the third entry after the second entry in the address chain, the cache line to the foreign thread after executing the plurality of local requests.

[0003] In general, in one aspect, the invention relates to a non-transitory computer readable storage medium storing instructions for managing a cache line. The instructions comprise functionality to: receive, from a plurality of local threads executing on a local multithreaded processor, a plurality of local requests to access the cache line; insert, into an address chain associated with the cache line, a plurality of entries corresponding to the plurality of local requests; identify a first entry of the plurality of entries at a head of the address chain; initiate, in response to the first entry being identified at the head and in response to the first entry corresponding to a request to own (RTO) the cache line, a traversal of the address chain; set, during the traversal of the address chain, a state element identified in a second entry of the plurality of entries, wherein the state element corresponds to a serialization acknowledgement from a central ordering unit for a local thread corresponding to the second entry; receive, from a foreign thread executing on a foreign multithreaded processor, a foreign request to access the cache line; insert, in response to the state element being set, a third entry corresponding to the foreign request into the address chain after the second entry; and relinquish, in response to the third entry being inserted after the second entry in the address chain, the cache line to the foreign thread after executing the plurality of local requests.

[0004] In general, in one aspect, the invention relates to a system for managing a cache line. The system comprises: a plurality of local threads executing on a local multithreaded processor and issuing a plurality of local requests to access the cache line; a foreign thread executing on a foreign multithreaded processor

and issuing a foreign request to access the cache line; and a local cache storing the cache line and comprising an address chain storing a plurality of local entries corresponding to the plurality of local requests and a foreign entry corresponding to the foreign request, wherein the local cache is configured to set, during a traversal of the address chain, a state element identified in the plurality of local entries and corresponding to a serialization acknowledgement from a central ordering unit, wherein the foreign entry is inserted into the address chain after the plurality of local entries in response to the state element being set, and wherein the cache line is relinquished to the foreign thread after the plurality of local requests are executed in response to the foreign entry being inserted into the address chain after the plurality of local entries.

[0005] Other aspects of the invention will be apparent from the following description and the appended claims.

BRIEF DESCRIPTION OF DRAWINGS

[0006] FIG. 1 shows a system in accordance with one or more embodiments of the invention.

[0007] FIG. 2 shows cache lines in accordance with one or more embodiments of the invention.

[0008] FIG. 3 shows an address chain in accordance with one or more embodiments of the invention.

[0009] FIG. 4 shows a flowchart in accordance with one or more embodiments of the invention.

[0010] FIG. 5 shows an example in accordance with one or more embodiments of the invention.

[0011] FIG. 6 shows a computer system in accordance with one or more embodiments of the invention.

DETAILED DESCRIPTION

[0012] Specific embodiments of the invention will now be described in detail with reference to the accompanying figures. Like elements in the various figures are denoted by like reference numerals for consistency.

[0013] In the following detailed description of embodiments of the invention, numerous specific details are set forth in order to provide a more thorough understanding of the invention. However, it will be apparent to one of ordinary skill in the art that the invention may be practiced without these specific details. In other instances, well-known features have not been described in detail to avoid unnecessarily complicating the description.

[0014] In general, embodiments of the invention provide a system and method for managing cache lines. Specifically, a cache line is retained in a cache of a local multithreaded processor until all the local requests (*i.e.*, requests issued by threads executing on the local multithreaded processor) are processed. It is only after processing the local requests that the cache line is relinquished to a foreign request issued by a foreign thread executing on a foreign processor. In one or more embodiments of the invention, the local multithreaded processor includes an address chain corresponding to the cache line and storing local entries corresponding to the local requests. State elements identified by the local entries are deliberately set to force an entry corresponding to a foreign request to be inserted at the end (*i.e.*, tail) of the address chain, and thus be processed after the local requests are processed.

[0015] FIG. 1 shows a multi-socketed system (100) in accordance with one or more embodiments of the invention. As shown in FIG. 1, the multi-socketed system (100) has multiple components including a central ordering unit (102), a memory (142), and one or more multithreaded processors (*i.e.*, Multithreaded Processor A (105), Multithreaded Processor B (125)) executing one or more threads (*i.e.*, Thread A1 (107), Thread AN (109), Thread B1 (127), Thread BN (129)). Further, each multithreaded processor (105, 125) includes a cache (*i.e.*, Cache A (111), Cache B (131)) storing one or more cache lines (*i.e.*, Cache

Line A (113), Cache Line B (133)) and one or more address chains (*i.e.*, Address Chain A (115), Address Chain B (135)).

[0016] In one or more embodiments of the invention, the threads (107, 109, 127, 129) include instructions (*i.e.*, requests) to access (*i.e.*, read from and/or write to) one or more locations in the memory (142). These memory locations may store shared resources (*e.g.*, locks, global variables, etc.) that are accessible to multiple threads and/or private variables that are only accessible to a specific thread. In order to expedite execution of these instructions, the contents of the memory addresses may be locally stored in the caches (111, 131), which have faster access times than the memory (142). Specifically, each cache line (113, 133) corresponds to a block of locations in the memory (142) and is a local copy of the contents (*i.e.*, shared resources, private variables, etc.) in the block of locations.

[0017] In one or more embodiments of the invention, if a local cache (*i.e.*, Cache A (111)) is not storing the contents of a memory address requested by a local thread (*e.g.*, Thread A1 (107)), the corresponding cache line is transferred from a foreign cache (*i.e.*, Cache B (131)) storing the cache line (*i.e.*, the foreign thread relinquishes the cache line). However, if no foreign cache is storing the cache line, the cache line is retrieved from the memory (142).

[0018] In one or more embodiments of the invention, concurrent attempts to access the same cache line among threads (107, 109, 127, 129) can pose performance problems depending on the order in which the threads are granted access. This is due to the long latency associated with an inter-socket cache line transfer (*i.e.*, transferring between a foreign cache to/from a local cache).

[0019] For example, suppose Thread A1 (107) has write access to Cache Line A (113) containing a global variable, and Thread AN (109) has instructions to store into the same global variable. It may take 10ns-20ns for Thread AN (109) to get the write access from Thread A1 (107). However, if Thread B1 (127) has instructions to store into the same global variable, the inter-processor transfer

(*i.e.*, transfer of Cache Line A (113) to Cache B (131)) can take more than 100ns. To improve the overall latency seen by all the threads, it is more efficient to amortize the latency of bringing in a cache line into a multithreaded processor over more threads that are local to that multithreaded processor, before relinquishing it to threads that are foreign to that multithreaded processor.

[0020] FIG. 2 shows multiple cache lines (*i.e.*, Cache Line W (210), Cache Line Z (250)) in accordance with one or more embodiments of the invention. The cache lines (210, 250) may correspond to either of cache line A (113) or cache line B (133), discussed above in reference to FIG. 1.

[0021] As shown in FIG. 2, a cache line (210, 250) may include shared resources (*e.g.*, Global Variable W (212)) and/or private variables (*i.e.*, Private Variable Z1 (213), Private Variable Z2 (214)). Multiple threads may concurrently issue requests to access Cache Line W (210) in order to read from and/or write to Global Variable W (212). This is an example of true sharing (*i.e.*, multiple threads concurrently issuing requests to access the same byte(s) in a cache line). Alternatively, one thread may issue a request to access Cache Line Z (250) in order to read from and/or write to Private Variable Z1 (213), while a different thread may concurrently issue a request to access Cache Line Z (250) in order to read from and/or write to Private Variable Z2 (214). This is an example of false sharing (*i.e.*, multiple threads concurrently issuing requests to access non-overlapping bytes in a cache line). A single cache line may have both shared resources and private variables.

[0022] Referring back to FIG. 1, each cache (111, 131) includes one or more address chains (115, 135) corresponding to the stored cache lines (113, 133). The address chains (115, 135) control/track the order in which requests for cache lines (113, 133) are processed. Specifically, Address Chain A (115) dictates the order in which threads (107, 109, 127, 129) access Cache Line A (113). Similarly, Address Chain B (135) dictates the order in which threads (107, 109, 127, 129) access Cache Line B (133).

[0023] FIG. 3 shows an address chain (*i.e.*, Address Chain Q (310)) in accordance with one or more embodiments of the invention. Address Chain Q (310) is an example of Address Chain A (115), discussed above in reference to FIG. 1. As shown in FIG. 3, Address Chain Q (310) includes multiple local entries (*i.e.*, Local Entry A1 (302), Local Entry A2 (304), Local Entry A3 (306), etc.) and at least one foreign entry (*i.e.*, Foreign Entry B1 (308)). The local entries (302, 304, 306) correspond to requests by local threads (*e.g.*, Thread A1 (107), Thread AN (109)) (*i.e.*, local requests) to access Cache Line A (113). The foreign entry B1 (308) corresponds to a request by a foreign thread (*e.g.*, Thread B1 (127)) (*i.e.*, foreign request) to access Cache Line A (113).

[0024] In one or more embodiments of the invention, if multiple threads executing on Multithreaded Processor A (105) issue store requests to Cache Line A (113), local entries (302, 304, 306) corresponding to these “local” stores are created and inserted into Address Chain Q (310). The “local” stores are processed in the order that they are received. In other words, they are processed serially (*i.e.*, only one member of the chain is processed at any given time). In one or more embodiments of the invention, Address Chain Q (310) is structured such that if the oldest request in the chain (*i.e.*, the entry at the head (312)) is a Request To Own (RTO) (*e.g.*, a store, a SWAP operation, a compare-and-swap (CAS) operation, etc.), all local entries in the chain are guaranteed to complete before the cache line is relinquished to a foreign request. In other words, Foreign Entry B1 (308) corresponding to a foreign request by a foreign thread (*e.g.*, Thread B1 (127), Thread BN (129)) is pending behind the youngest local entry in the address chain (310) if the head (312) of the address chain (310) is an RTO, allowing all local chain members to complete before the foreign request gets processed.

[0025] As discussed above, the multi-socketed system (100) includes the central ordering unit (102). The central ordering unit (102) selects requests (*i.e.*, requests issued by the threads) to be active in the system (100), and

issues serialization acknowledgements to the threads selected to be active. In one or more embodiments of the invention, each local entry (302, 304, 306) includes a state element called `mb_acked`. The state element `mb_acked` is set when a serialization acknowledgement from the central ordering unit (102) is received. The serialization acknowledgement indicates that the request corresponding to the local entry is now active in the system. In one or more embodiments of the invention, only the request corresponding to the entry at the head of the address chain (310) can receive a serialization acknowledgement from the central ordering unit (102).

[0026] In one or more embodiments of the invention, `mb_acked` is set preemptively for younger members of the address chain (310), provided that the oldest member (*i.e.*, the entry at the head (312)) is an RTO. In other words, even if a serialization acknowledgement has not yet been received from the central ordering unit for a younger member of the address chain (310), the state element `mb_acked` will be set for the younger member provided an RTO is at the head of the address chain. Further, even if the original oldest member is not an RTO, when an RTO reaches the head of the chain, the address chain will be traversed to set all subsequent (*i.e.*, younger) members' `mb_acked` bits (*i.e.*, set `mb_acked` = 1).

[0027] In one or more embodiments of the invention, the `mb_acked` state element is used by a foreign request to determine a location in the address chain (310) where the corresponding foreign entry B1 (308) should be inserted. In one or more embodiments of the invention, the foreign request inserts the corresponding foreign entry B1 (308) after the youngest member with `mb_acked` set to 1 (*i.e.*, after Local Entry A3 (306)).

[0028] Those skilled in the art, having the benefit of this detailed description, will appreciate that if the `mb_acked` state were to be set only when a request receives a serialization acknowledgement from a central ordering unit (*i.e.*, Central Ordering Unit (102), discussed above in reference to FIG.1), since younger members of the address chain cannot send requests to the central

ordering unit (102) until they get to the head of the chain, the younger members might not have their respective `mb_acked` bits set when the foreign request arrives. Thus, in such scenarios, Foreign Entry B1 (308) may be inserted immediately following the oldest active request (*i.e.*, immediately following Local Entry A1 (302)) in the chain and transfer the cache line to the foreign cache when processing of Local Entry A1 (302) is complete.

[0029] FIG. 4 shows a flowchart in accordance with one or more embodiments of the invention. The process shown in FIG. 4 may be executed by one or more components of the multi-socketed system (100). One or more steps shown in FIG. 4 may be omitted, repeated, and/or performed in a different order among different embodiments of the invention. Accordingly, embodiments of the invention should not be considered limited to the specific number and arrangement of steps shown in FIG. 4.

[0030] Initially, multiple local requests to access a cache line are received from multiple local threads executing on a local multithreaded processor (STEP 405). The cache line may include a resource shared by the multiple local threads (*e.g.*, global variable) (*i.e.*, true sharing) and/or the cache line may include individual private variables that are not shared by the multiple local threads but are stored in non-overlapping bytes of the cache line (*i.e.*, false sharing).

[0031] In STEP 410, multiple local entries corresponding to the local requests are inserted into an address chain associated with the cache line. Each local entry identifies the corresponding local thread and a state element (*i.e.*, `mb_acked`). If the cache of the multithreaded processor does not already store the cache line, the cache line will need to be transferred from a foreign cache storing the cache line. However, if no foreign cache is storing the cache line, the cache line will need to be retrieved from main memory.

[0032] In STEP 415, it is determined that the local entry at the head of the address chain (*i.e.*, head entry) corresponds to a request to own (RTO) the

cache line. As discussed above, examples of requests to own include a store operation, a SWAP operation, a CAS operation, etc.

[0033] In STEP 420, the address chain is traversed in response to determining that the head entry corresponds to a RTO the cache line. In other words, the traversal of the address chain is triggered by a RTO being at the head of the address chain. During the traversal, the state element in the local entries in the address chain are set (*i.e.*, `mb_acked = 1`). The traversal may be executed before or after the serialization acknowledgement for the RTO head entry is received from a central ordering unit.

[0034] In STEP 425, a foreign request to access the cache line is received. The foreign request is issued by a foreign thread executing on a foreign multithreaded processor. A foreign entry for the foreign request is created and inserted into the address chain. Specifically, the foreign entry is inserted after the youngest local entry having its state element set (*i.e.*, `mb_acked = 1`). Accordingly, the foreign entry is effectively inserted at the tail of the address chain and all entries between the head of the address chain and the foreign entry are local entries.

[0035] In STEP 430, all the local requests are processed before the foreign request because the foreign entry is after the local entries in the address chain. Once the local requests are processed, the cache line is relinquished to the foreign thread corresponding to the foreign entry. Relinquishing the cache line to the foreign thread may include transferring the cache line from the cache of the local multithreaded processor to the cache of the foreign multithreaded processor executing the foreign thread. The relinquished cache line may be stored in a cache of the foreign processor.

[0036] FIG. 5 shows an example in accordance with one or more embodiments of the invention. Specifically, the example in FIG. 5 includes two scenarios: Scenario I (598) and Scenario II (599).

[0037] In Scenario I (598), multiple local store requests to access a cache line are issued by local threads. Assume the cache line is already stored in a cache of the local multithreaded processor executing the local threads. Accordingly, the address chain (501) for the cache line includes three entries (*i.e.*, Store X1 (505), Store X2 (507), Store X3 (509)) corresponding to the store requests. Further, each entry identifies the `mb_acked` state element. As shown in Scenario I (598), the `mb_acked` state element is set (*i.e.*, `mb_acked = 1`) for Store X1 (505) in the address chain (501). The `mb_acked` state element for Store X1 (505) is set in response to receiving a serialization acknowledgement for the corresponding request from the central ordering unit. However, the `mb_acked` state element in the remaining entries (507, 509) is not set (*i.e.*, `mb_acked = 0`).

[0038] Still referring to Scenario I (598), upon arrival of a foreign store request for the cache line, a foreign entry (*i.e.*, Store Y1 (511)) is created. As shown in the address chain (502), the foreign entry (*i.e.*, Store Y1 (511)) is inserted immediately following the youngest local entry having the `mb_acked` state element set (*i.e.*, immediately following Store X1 (505)). Accordingly, once processing of the local request corresponding to Store X1 (505) is finished, Store Y1 (511) is next in the address chain (502), and the cache line will be relinquished to the foreign multithreaded processor that issued the foreign request. The cache line will eventually need to be returned to the local multithreaded processor in order to process the remaining local entries (*i.e.*, Store X2 (507), Store X3 (509)). The relinquishing of the cache line and return of the cache line are time expensive operations.

[0039] In Scenario II (599), multiple local store requests to access a cache line are issued by local threads. Assume the cache line is already stored in a cache of the local multithreaded processor executing the local threads. Accordingly, the address chain (552) for the cache line includes three entries (*i.e.*, Store X1 (505), Store X2 (507), Store X3 (509)) corresponding to the store requests. Further, each entry identifies the `mb_acked` state element. Upon identifying

the head entry in the address chain (552) as a RTO, the address chain (552) is traversed and the `mb_acked` state elements of all subsequent (*i.e.*, younger) members in the address chain (552) are set, even though the younger members have not yet received a serialization acknowledgement. As shown in Scenario II (599), the `mb_acked` state element is set (*i.e.*, `mb_acked` = 1) for all entries in the address chain (552) (*i.e.*, Store X1 (505), Store X2 (507), Store X3 (509)).

[0040] Still referring to Scenario II (599), upon arrival of a foreign store request for the cache line, a foreign entry (*i.e.*, Store Y1 (511)) is created. As shown in the address chain (554), the foreign entry (*i.e.*, Store Y1 (511)) is inserted immediately following the youngest local entry having the `mb_acked` state element set (*i.e.*, immediately following Store X3 (509)). Accordingly, it is only after the processing of all the local requests (*i.e.*, Store X1 (505), Store X2 (507), and Store X3 (509)) that the cache line is relinquished to the foreign multithreaded processor that issued the foreign request. By processing all the local requests before the cache line is relinquished, the cache line is not repeatedly transferred between the local and foreign processors and the delays resulting from the repeated transfers of the cache line are reduced. Scenario II (599) may be implemented using the multi-socketed system (100), discussed above in reference to FIG. 1, and/or the process discussed above in reference to FIG. 4.

[0041] As yet another example, assume there exists a local multithreaded processor executing eight (8) local threads and a foreign processor executing eight (8) foreign threads. Further, assume that all eight local thread and all eight foreign threads want to make updates to the same cache line. Further still, assume that intra-socket latency (*i.e.*, latency to transfer access of cache line between local threads) is 20ns while inter-socket latency (*i.e.*, latency to transfer access between a local thread and a foreign thread) is 100ns, and initially it takes 100ns to fetch the cache line from memory to the cache in the local multithreaded processor, which issues the first request. In Scenario I (598), the cache line will get repeatedly transferred between the local

multithreaded processor and the foreign multithreaded processor, and the incremental delay between access grants will be 100ns, or in other words, $16 * 100\text{ns} = 1600\text{ns}$ of total time used for all 16 threads to make their respective updates. In Scenario II (599), all eight threads in the local multithreaded processor will get to make their updates to the cache line, taking $100\text{ns} + (7 * 20\text{ns}) = 240\text{ns}$ before relinquishing the cache line to the foreign multithreaded processor. All 8 threads in the foreign multithreaded processor will then be able to take their turns to update the cache line, again taking $100\text{ns} + (7 * 20\text{ns}) = 240\text{ns}$. The total time taken for all 16 threads to complete their updates is $240\text{ns} + 240\text{ns} = 480\text{ns}$ in Scenario II (599), vs. 1600ns in Scenario I (598).

[0042] Embodiments of the invention may be implemented on virtually any type of computer regardless of the platform being used. For example, as shown in FIG. 6, a computer system (600) includes one or more processor(s) (602) (such as a central processing unit (CPU), integrated circuit, etc.), associated memory (604) (*e.g.*, random access memory (RAM), cache memory, flash memory, *etc.*), a storage device (606) (*e.g.*, a hard disk, an optical drive such as a compact disk drive or digital video disk (DVD) drive, a flash memory stick, *etc.*), and numerous other elements and functionalities typical of today's computers (not shown). The computer system (600) may also include input means, such as a keyboard (608), a mouse (610), or a microphone (not shown). Further, the computer system (600) may include output means, such as a monitor (612) (*e.g.*, a liquid crystal display (LCD), a plasma display, or cathode ray tube (CRT) monitor). The computer system (600) may be connected to a network (614) (*e.g.*, a local area network (LAN), a wide area network (WAN) such as the Internet, or any other type of network) via a network interface connection (not shown). Those skilled in the art will appreciate that many different types of computer systems exist, and the aforementioned input and output means may take other forms. Generally speaking, the computer system (600) includes at least the minimal processing, input, and/or output means necessary to practice embodiments of the invention.

[0043] Further, in one or more embodiments of the invention, one or more elements of the aforementioned computer system (600) may be located at a remote location and connected to the other elements over a network. Further, embodiments of the invention may be implemented on a distributed system having a plurality of nodes, where each portion of the invention may be located on a different node within the distributed system. In one embodiment of the invention, the node corresponds to a computer system. Alternatively, the node may correspond to a processor with associated physical memory. The node may alternatively correspond to a processor or micro-core of a processor with shared memory and/or resources. Further, software instructions in the form of computer readable program code to perform embodiments of the invention may be stored, temporarily or permanently, on a tangible computer readable storage medium, such as a compact disc (CD), a diskette, a solid state memory device, a tape, memory, or any other non-transitory tangible computer readable storage device.

[0044] While the invention has been described with respect to a limited number of embodiments, those skilled in the art, having benefit of this disclosure, will appreciate that other embodiments can be devised which do not depart from the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the attached claims.

CLAIMS

What is claimed is:

1. A method for managing a cache line, comprising:

receiving, from a plurality of local threads executing on a local multithreaded processor, a plurality of local requests to access the cache line;
inserting, into an address chain associated with the cache line, a plurality of entries corresponding to the plurality of local requests;
identifying a first entry of the plurality of entries at a head of the address chain;
initiating, in response to identifying the first entry at the head and in response to the first entry corresponding to a request to own (RTO) the cache line, a traversal of the address chain;
setting, during the traversal of the address chain, a state element identified in a second entry of the plurality of entries, wherein the state element corresponds to a serialization acknowledgement from a central ordering unit for a local request corresponding to the second entry;
receiving, from a foreign thread executing on a foreign multithreaded processor, a foreign request to access the cache line;
inserting, in response to setting the state element, a third entry corresponding to the foreign request into the address chain after the second entry; and
relinquishing, in response to inserting the third entry after the second entry in the address chain, the cache line to the foreign thread after executing the plurality of local requests.

2. The method of claim 1, further comprising:
 - retrieving, in response to receiving at least one of the plurality of local requests, the cache line from a memory operatively connected to the local multithreaded processor; and
 - storing the cache line in a cache of the local multithreaded processor, wherein the cache line comprises a global variable for the plurality of local threads and the foreign thread.
3. The method of claim 1, wherein the cache line comprises a plurality of private variables for the plurality of local threads.
4. The method of claim 1, further comprising:
 - receiving, after initiating the traversal, a serialization acknowledgement from the central ordering unit for a local request corresponding to the first entry in the address chain.
5. The method of claim 1, wherein foreign request corresponds to a store operation.
6. The method of claim 1, wherein the RTO is a SWAP operation.
7. The method of claim 1, wherein relinquishing the cache line to the foreign thread comprises storing the cache line in a cache of the foreign multithreaded processor.
8. A non-transitory computer readable storage medium storing instructions for managing a cache line, the instructions comprising functionality to:
 - receive, from a plurality of local threads executing on a local multithreaded processor, a plurality of local requests to access the cache line;
 - insert, into an address chain associated with the cache line, a plurality of entries corresponding to the plurality of local requests;
 - identify a first entry of the plurality of entries at a head of the address chain;
 - initiate, in response to the first entry being identified at the head and in response to the first entry corresponding to a request to own (RTO) the cache line, a traversal of the address chain;

set, during the traversal of the address chain, a state element identified in a second entry of the plurality of entries, wherein the state element corresponds to a serialization acknowledgement from a central ordering unit for a local thread corresponding to the second entry;

receive, from a foreign thread executing on a foreign multithreaded processor, a foreign request to access the cache line;

insert, in response to the state element being set, a third entry corresponding to the foreign request into the address chain after the second entry; and

relinquish, in response to the third entry being inserted after the second entry in the address chain, the cache line to the foreign thread after executing the plurality of local requests.

9. The non-transitory computer readable storage medium of claim 8, the instructions further comprising functionality to:

retrieve, in response to receiving at least one of the plurality of local requests, the cache line from a memory operatively connected to the local multithreaded processor; and

store the cache line in a cache of the local multithreaded processor, wherein the cache line comprises a global variable for the plurality of local threads and the foreign thread.

10. The non-transitory computer readable storage medium of claim 8, wherein the cache line comprises a plurality of private variables for the plurality of local threads.

11. The non-transitory computer readable storage medium of claim 8, the instructions further comprising functionality to:

receive, after initiating the traversal, a serialization acknowledgement from the central ordering unit for a local request corresponding to the first entry in the address chain.

12. The non-transitory computer readable storage medium of claim 8, wherein foreign request corresponds to a store operation.
13. The non-transitory computer readable storage medium of claim 8, wherein the RTO is a SWAP operation.
14. The non-transitory computer readable storage medium of claim 8, wherein the instructions to relinquish the cache line to the foreign thread comprise functionality to:
- store the cache line in a cache of the foreign multithreaded processor.
15. A system for managing a cache line, comprising:
- a plurality of local threads executing on a local multithreaded processor and issuing a plurality of local requests to access the cache line;
 - a foreign thread executing on a foreign multithreaded processor and issuing a foreign request to access the cache line; and
 - a local cache storing the cache line and comprising an address chain storing a plurality of local entries corresponding to the plurality of local requests and a foreign entry corresponding to the foreign request,
- wherein the local cache is configured to set, during a traversal of the address chain, a state element identified in the plurality of local entries and corresponding to a serialization acknowledgement from a central ordering unit,
- wherein the foreign entry is inserted into the address chain after the plurality of local entries in response to the state element being set, and
- wherein the cache line is relinquished to the foreign thread after the plurality of local requests are executed in response to the foreign entry being inserted into the address chain after the plurality of local entries.
16. The system of claim 15, wherein the traversal is initiated in response to a head of the address chain being a request to own (RTO) the cache line.
17. The system of claim 16, wherein the RTO is a CAS operation.

18. The system of claim 15, further comprising:

a foreign cache configured to store the cache line after the cache line is relinquished to the foreign thread.

19. The system of claim 15, wherein the cache line comprises a plurality of private variables for the plurality of local threads.

20. The system of claim 15, further comprising:

a memory operatively connected to the local multithreaded processor and the foreign multithreaded processor,

wherein the cache line is retrieved from the memory in response to at least one of the plurality of local requests, and

wherein the cache line comprises a global variable for the plurality of local requests and the foreign request.




FIG. 1




FIG. 2




FIG. 3

4/6




FIG. 4




FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2012/035757

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06F9/38
 ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2004/162951 A1 (JACOBSON QUINN A [US] ET AL) 19 August 2004 (2004-08-19) the whole document -----	1-20
A	US 7 774 552 B1 (CYPHER ROBERT E [US] ET AL) 10 August 2010 (2010-08-10) the whole document -----	1-20
A	WO 2007/027365 A1 (INTEL CORP [US]; KUKANOV ALEXEY [RU]; ROBISON ARCH [US]) 8 March 2007 (2007-03-08) abstract -----	1,8,15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

11 July 2012

Date of mailing of the international search report

23/07/2012

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040,
 Fax: (+31-70) 340-3016

Authorized officer

Moraiti, Marina

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2012/035757

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004162951 A1	19-08-2004	EP 1593037 A1	09-11-2005
		US 2004162951 A1	19-08-2004
		US 2005262301 A1	24-11-2005
		WO 2004075052 A1	02-09-2004

US 7774552 B1	10-08-2010	NONE	

WO 2007027365 A1	08-03-2007	CN 101253482 A	27-08-2008
		EP 1920332 A1	14-05-2008
		US 2007067774 A1	22-03-2007
		US 2010293553 A1	18-11-2010
		WO 2007027365 A1	08-03-2007
