



(19) **United States**

(12) **Patent Application Publication**
COE et al.

(10) **Pub. No.: US 2019/0213135 A1**

(43) **Pub. Date: Jul. 11, 2019**

(54) **HARDWARE CONTROL OF CPU HOLD OF A CACHE LINE IN PRIVATE CACHE**

(52) **U.S. Cl.**

CPC *G06F 12/0891* (2013.01); *G06F 12/0831* (2013.01); *G06F 2212/283* (2013.01); *G06F 12/0811* (2013.01); *G06F 2212/621* (2013.01); *G06F 12/128* (2013.01)

(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION**,
Armonk, NY (US)

(72) Inventors: **DWIFUZI COE**, POUGHKEEPSIE, NY (US); **YAIR FRIED**, PETAH TIQWA (IL); **MARTIN RECKTENWALD**, SCHOENAICH (DE); **YOSSI SHAPIRA**, SHOHAM (IL)

(57) **ABSTRACT**

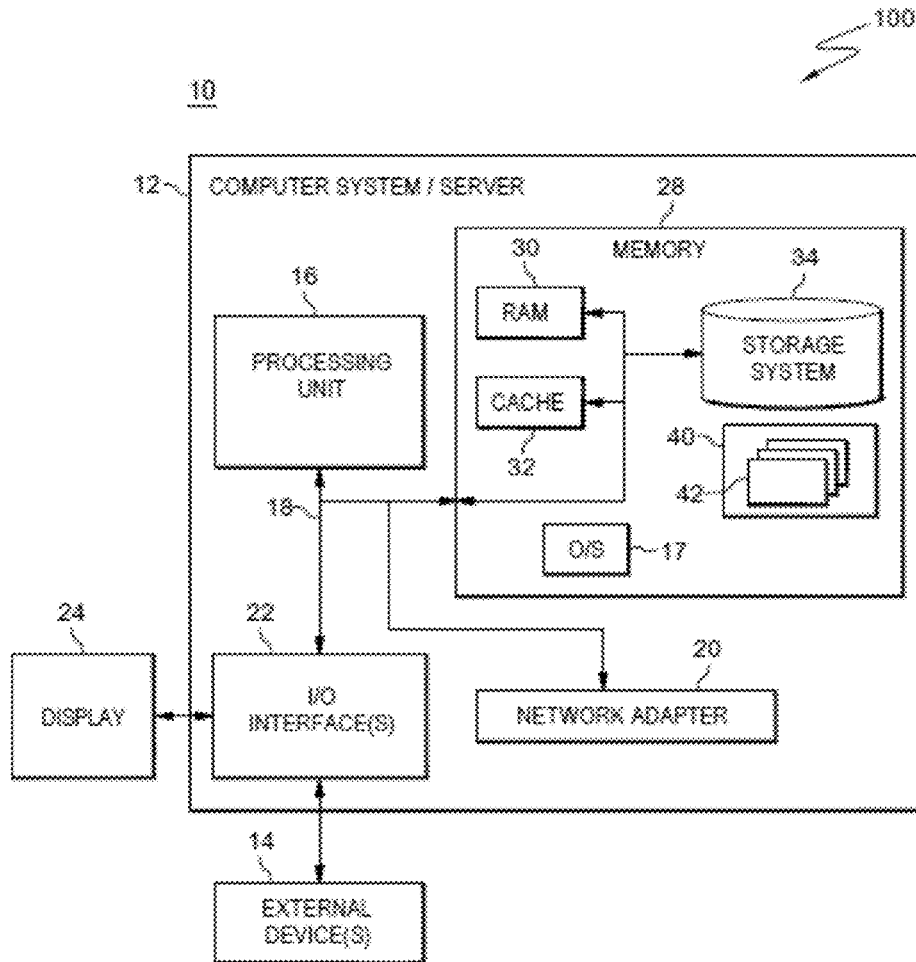
A memory access control includes a tracker configured to receive cache invalidate (XI) commands from the memory controller and to provide responses to the memory controller and an address storage element in the tracker that stores an address to be locked by one of the processing units. The system also includes a lock required, a cache invalidate (XI) tracker bit, a set input that upon receipt of a set command sets the lock required bit when a first condition is met, a first reset input that resets the lock required bit upon receipt of a reset command; and a second reset input that resets the XI tracker bit. The tracker rejects incoming XI commands from the memory controller when the lock required bit is set, allows incoming XI commands when the lock bit is not set and sets the XI tracker bit when a first incoming XI command is received.

(21) Appl. No.: **15/867,989**

(22) Filed: **Jan. 11, 2018**

Publication Classification

(51) **Int. Cl.**
G06F 12/0891 (2006.01)
G06F 12/0831 (2006.01)
G06F 12/128 (2006.01)
G06F 12/0811 (2006.01)



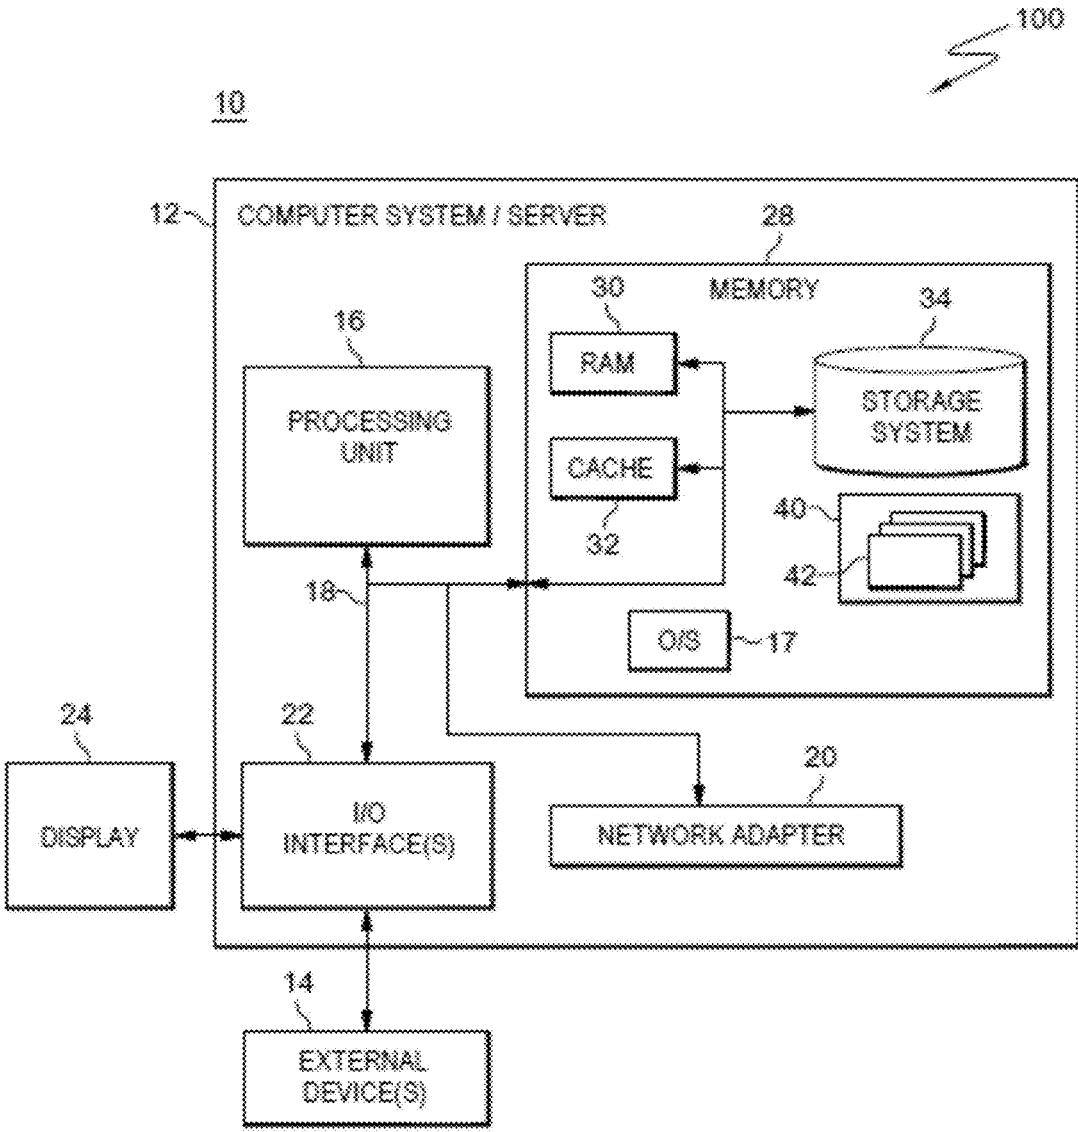


FIG. 1

200

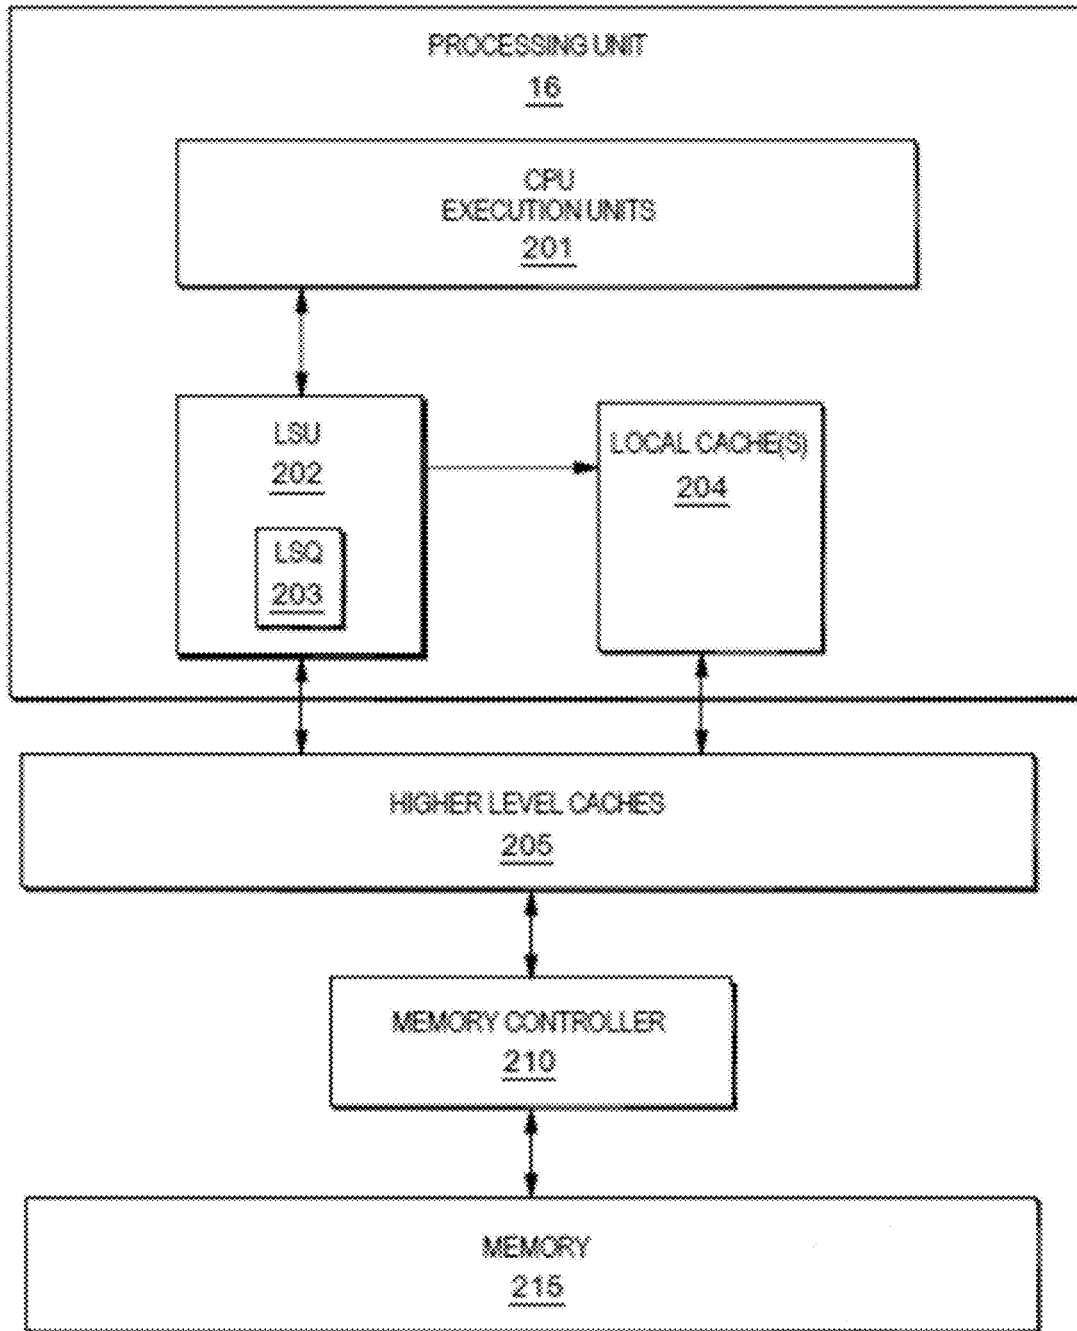


FIG. 2

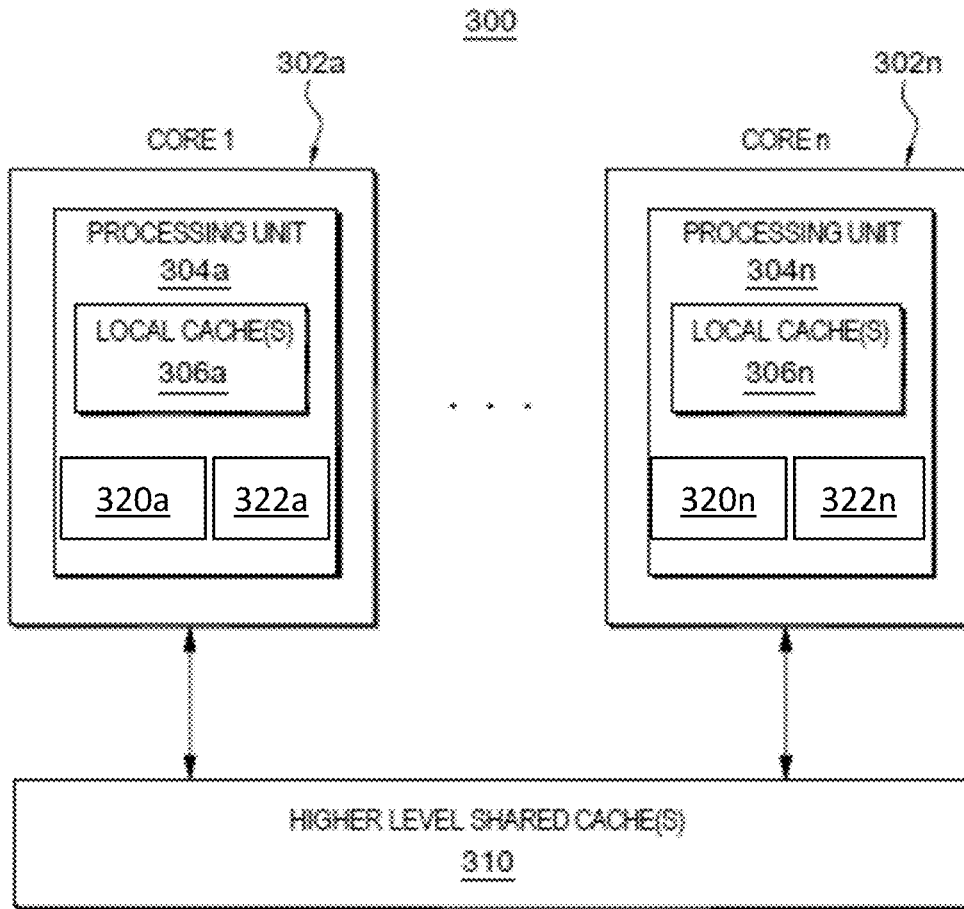


FIG. 3

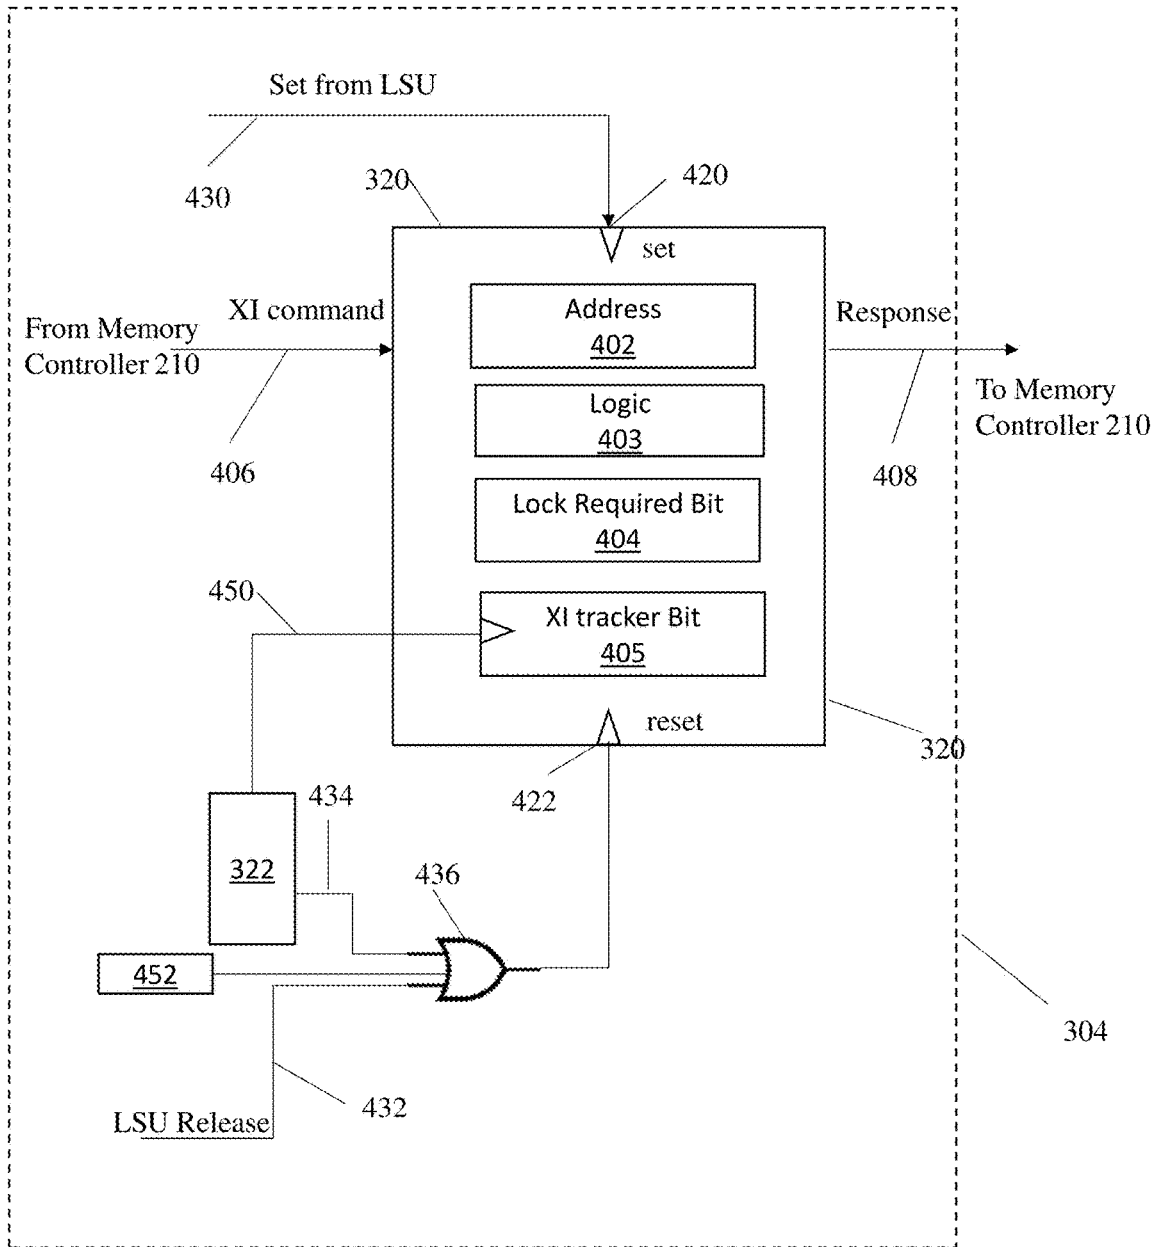


FIG. 4

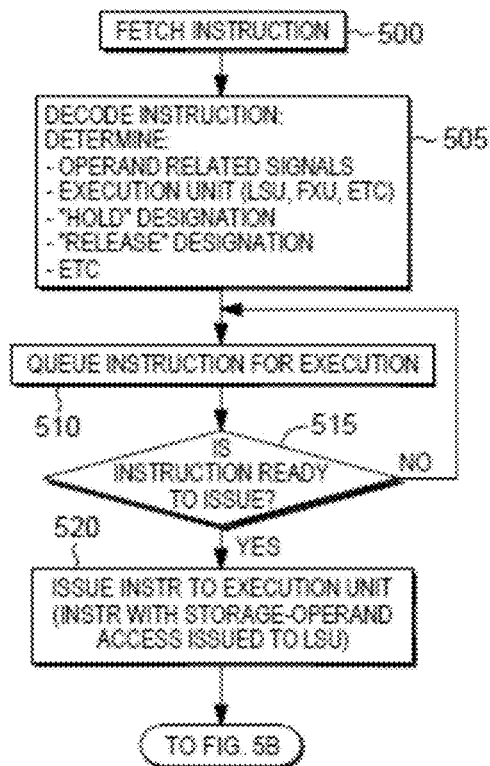


FIG. 5A

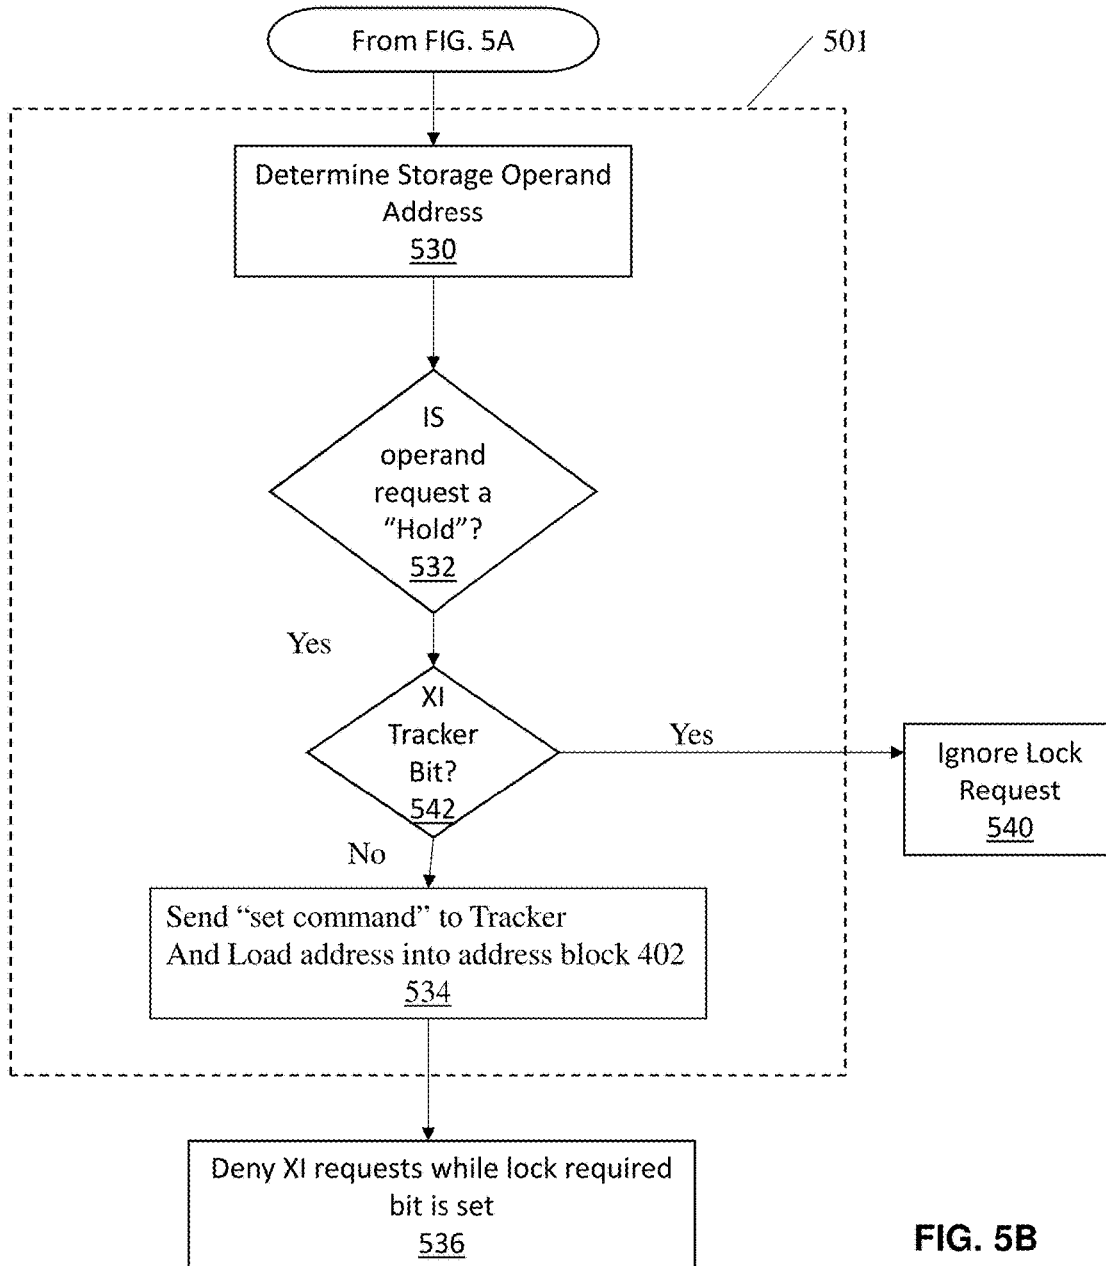


FIG. 5B



FIG. 6A

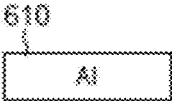


FIG. 6B

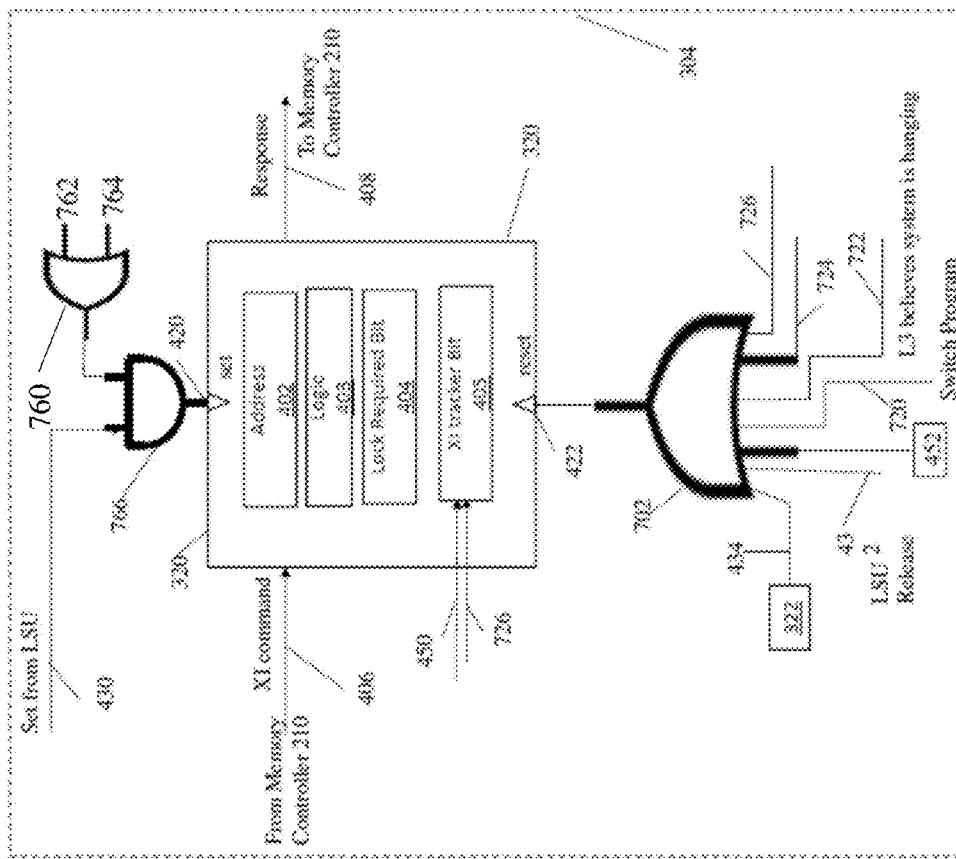


FIG. 7

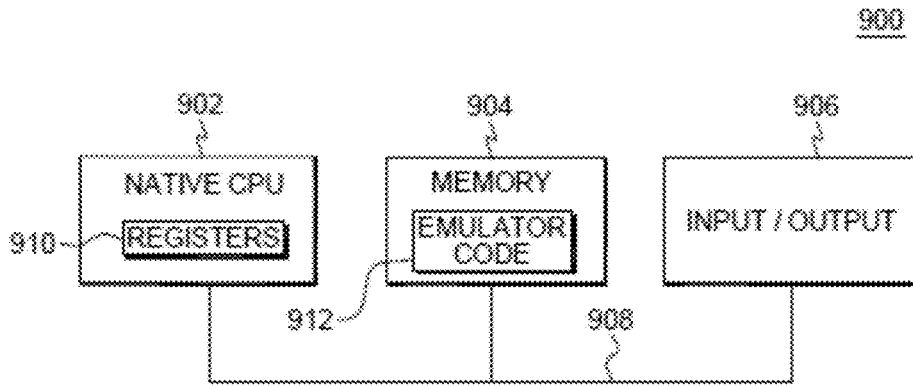


FIG. 8A

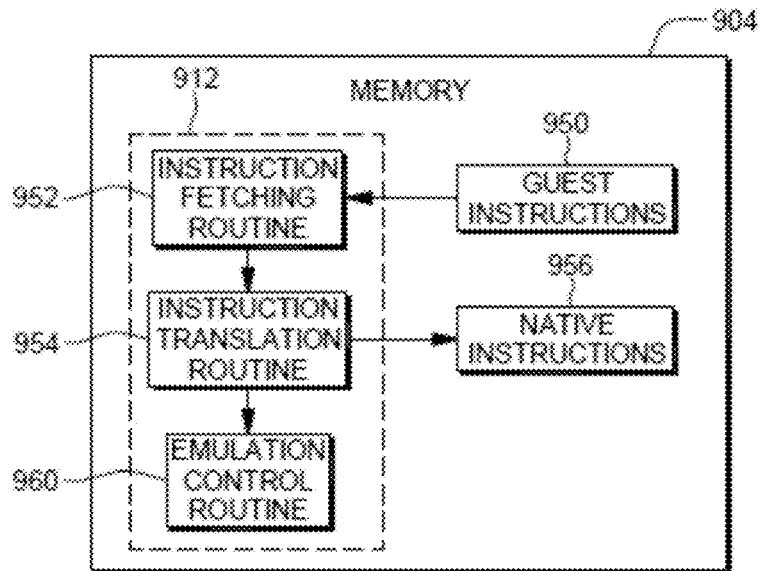


FIG. 8B

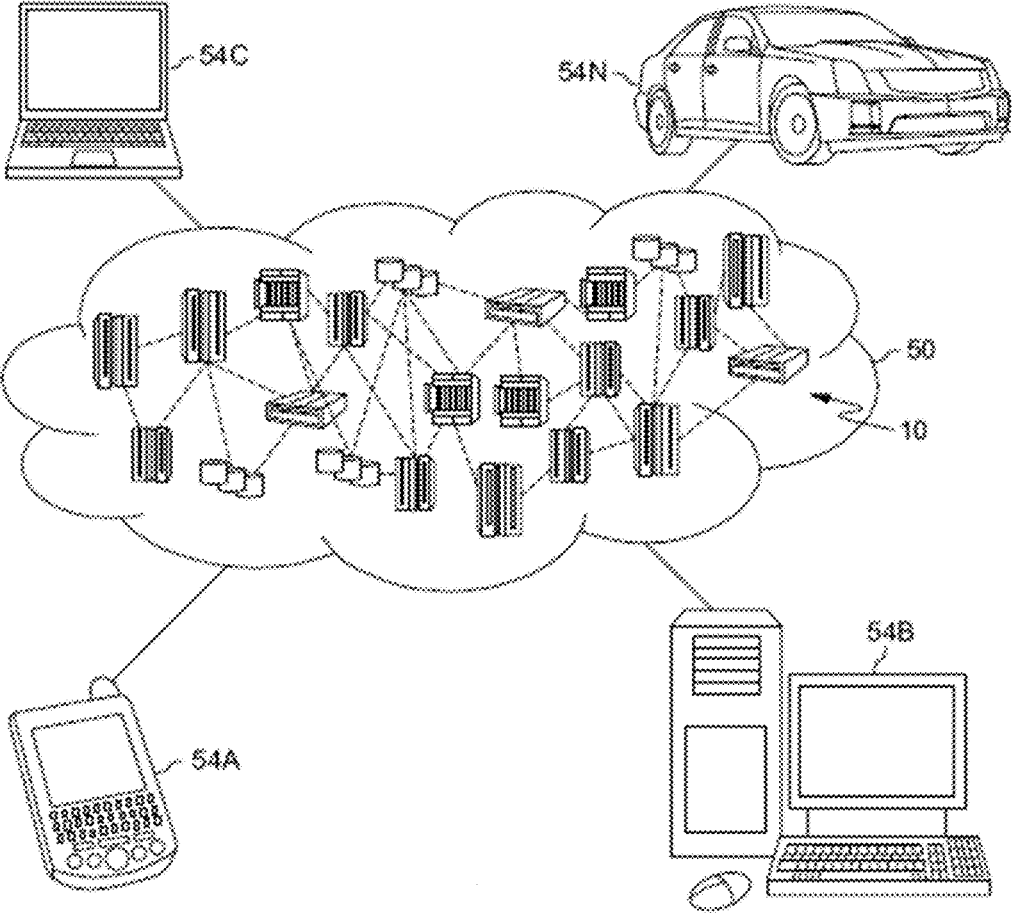


FIG. 9

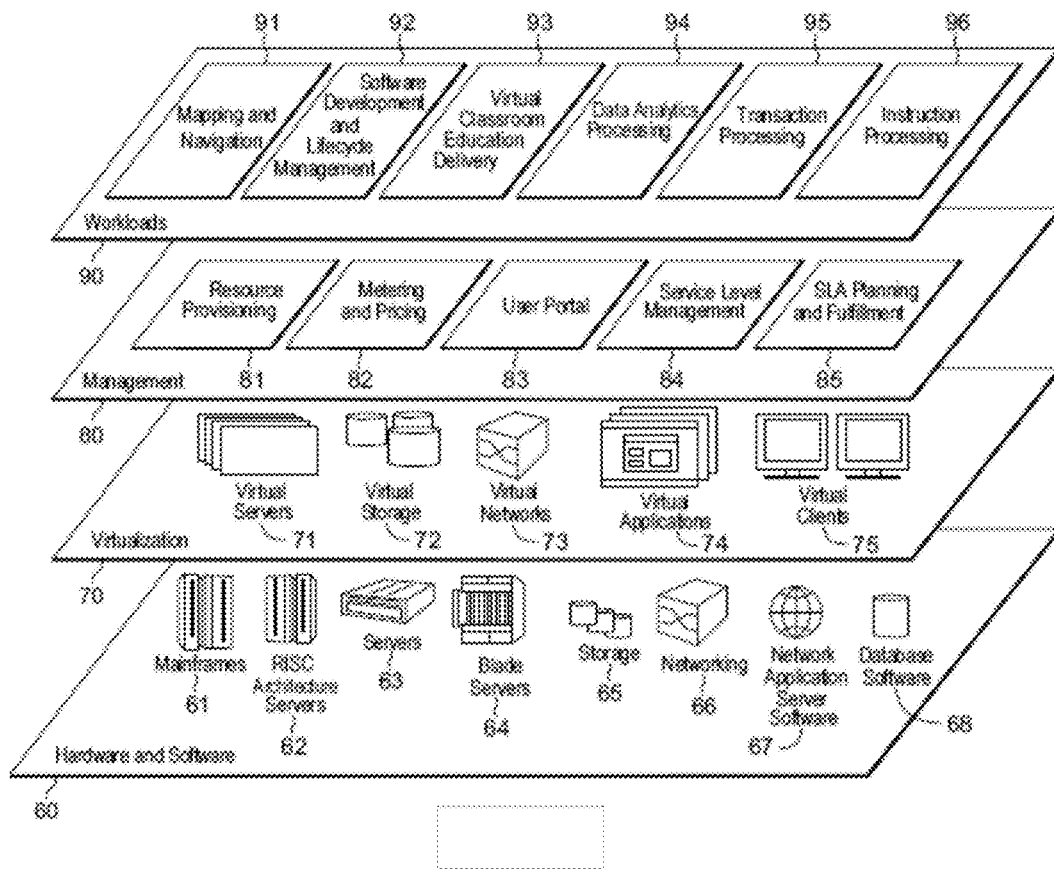


FIG. 10

HARDWARE CONTROL OF CPU HOLD OF A CACHE LINE IN PRIVATE CACHE

BACKGROUND

[0001] One or more aspects relate, in general, to locking a cache line and, in particular, to a hardware control of a next instruction access intent instruction.

[0002] In a multiprocessing computing environment, multiple processing units, such as central processing units (CPUs) or cores, are connected to share access to common storage, such as common main storage locations. The storage hierarchy across the configuration may include various cache levels private or local to each processing unit, various cache levels shared among several processing units, and main storage. When a processing unit is to update a storage location, a line of storage is transferred to the processing unit by the storage hierarchy, and the line is held exclusively within a cache level that is private to the processing unit performing the update. While a line of storage is held exclusively by a processing unit, other processing units are not permitted to update or read that line of storage, until after the line is released and no longer held exclusively by the processing unit.

[0003] Where several computer processes share access to an area of main storage, a semaphore (or lock) may be used to control access to that area of storage. The semaphore is typically a word or double word in main storage and the semaphore contents dictate which processing unit within the configuration currently is permitted access to the corresponding area of storage. When a process executing on a processing unit requires access to the storage area, that process updates the semaphore to indicate the shared area is in use, completes tasks that access the shared area, and then updates the semaphore again to indicate that the shared area is no longer in use.

[0004] While the above process works well in most instances, in some cases, a process or processor that has access can hang and not be able to release the line. In other cases, the processor that has the lock may release it but then reacquire it before a competing processor has a chance to request a lock.

SUMMARY

[0005] In one embodiment, a memory access control system including a tracker for use in computing environment including two or more processing units is disclosed. Each of the two or more processing units include a tracker that includes an address storage element that stores an address to be locked by one of the processing units, a lock required bit and a cache invalidate (XI) tracker bit. The tracker also includes a set input that upon receipt of a set command sets the lock required bit when a first condition is met, a first reset input that resets the lock required bit upon receipt of a reset command; a second reset input that resets the XI tracker bit. The tracker also includes logic that causes the tracker to reject incoming XI commands when the lock required bit is set, to allow incoming XI commands when the lock bit is not set and to set the XI tracker bit when a first incoming XI command is received.

[0006] In one embodiment, a memory access control system for use in computing environment including two or more processing units and at least one shared memory controlled by a memory controller is disclosed. The system

of this embodiment includes a tracker configured to receive cache invalidate (XI) commands from the memory controller and to provide responses to the memory controller and an address storage element in the tracker that stores an address to be locked by one of the processing units. The system also includes a lock required bit in the tracker, a cache invalidate (XI) tracker bit in the tracker a set input in the tracker that upon receipt of a set command sets the lock required bit when a first condition is met, a first reset input in the tracker that resets the lock required bit upon receipt of a reset command and a second reset input that resets the XI tracker bit. The tracker is configured to reject incoming XI commands from the memory controller when the lock required bit is set, to allow incoming XI commands when the lock bit is not set and to set the XI tracker bit when a first incoming XI command is received.

[0007] Additional features and advantages are realized through the techniques described herein. Other embodiments and aspects are described in detail herein and are considered a part of the claimed aspects.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] One or more aspects are particularly pointed out and distinctly claimed as examples in the claims at the conclusion of the specification. The foregoing and objects, features, and advantages of one or more aspects are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

[0009] FIG. 1 depicts one example of a computing environment to incorporate and use one or more aspects of the present invention;

[0010] FIG. 2 depicts a further example of a computing environment to incorporate and use one or more aspects of the present invention;

[0011] FIG. 3 depicts one example of a multiprocessing unit environment to incorporate and use one or more aspects of the present invention;

[0012] FIG. 4 depicts one embodiment of a tracker in accordance with one or more aspects of the present invention;

[0013] FIGS. 5A and 5B are flowcharts depicting a more detailed example of storage operand request-related processing, in accordance with one or more aspects of the present invention;

[0014] FIG. 6A illustrates one embodiment of an instruction format useful in accordance with one or more aspects of the present invention;

[0015] FIG. 6B illustrates one embodiment of an access intent control field of an instruction format such as depicted in FIG. 6A, in accordance with one or more aspects of the present invention;

[0016] FIG. 7 depicts another embodiment of a tracker in accordance with one or more aspects of the present invention

[0017] FIG. 8A depicts another example of a computing environment to incorporate and use one or more aspects of the present invention;

[0018] FIG. 8B depicts further details of the memory of FIG. 8A;

[0019] FIG. 9 depicts one embodiment of a cloud computing environment; and

[0020] FIG. 10 depicts one example of abstraction model layers.

DETAILED DESCRIPTION

[0021] As noted, processing unit features intended to enhance individual processing unit performance, such as pipelining, superscalar, out-of-order execution, and branch prediction, may result in speculative memory requests. Speculative requests to private storage locations can be advantageous, but excessive speculative requests to common storage locations can be detrimental to semaphore (or lock) processing.

[0022] Contention among processing units for a semaphore to control a common shared storage location increases as: the number of speculative memory requests increases by the processing units in the computing environment (or computing configuration), the number of processing units in the computing environment increases, or potentially based on memory subsystem topology, etc. As contention among processing units increases, it becomes increasingly more difficult for each individual processing unit to successfully obtain and hold a line of storage containing a particular shared memory location. Therefore, overall system performance may degrade.

[0023] For instance, as multiple processing units attempt to update a storage location corresponding to a semaphore, system controls may broadcast cache invalidate commands to all processing units in the computing environment, including the processing unit currently holding the line of storage, which may contain the semaphore in an exclusive state. Standard memory management protocols may cause the processing unit, currently executing a process owning the lock to release the exclusive ownership of the line, before the process can complete all tasks and before the process can update the lock to indicate that the shared area is no longer in use. In this case, the processing unit will need to later retrieve the line exclusively again in order to finish the process. In view of this, a cache line containing a semaphore may frequently be traversing the computing environment, without benefitting any individual process.

[0024] In some systems, a facility is provided for a program to signal a processing unit that a specific storage operand access corresponds to a common storage location containing a semaphore (or lock) and is to have an extended hold. When signaled, the processing unit may deviate from standard memory management protocols for the purpose of holding a cache line exclusively, throughout the computing environment, for a longer duration than conventional, thereby increasing the probability of completing the process being run on the processing unit before releasing the cache line. As a result, the number of times the cache line, containing the semaphore traverses the computing environment unproductively is reduced, thereby increasing system performance.

[0025] One embodiment of a computing environment to incorporate and use one or more aspects of the present invention is described with reference to FIG. 1. In one example, the computing environment may be based on the z/Architecture, offered by International Business Machines Corporation, Armonk, N.Y. One embodiment of the z/Architecture is described in “z/Architecture Principles of Operation,” IBM Publication No. SA22-7832-10, March 2015, which is hereby incorporated herein by reference in its entirety. Z/ARCHITECTURE is a registered trademark of International Business Machines Corporation, Armonk, N.Y., USA.

[0026] In another example, the computing environment may be based on the Power Architecture, offered by International Business Machines Corporation, Armonk, N.Y. One embodiment of the Power Architecture is described in “Power ISA™ Version 2.07B,” International Business Machines Corporation, Apr. 9, 2015, which is hereby incorporated herein by reference in its entirety. POWER ARCHITECTURE is a registered trademark of International Business Machines Corporation, Armonk, N.Y., USA.

[0027] The computing environment may also be based on other architectures, including, but not limited to, the Intel x86 architectures. Other examples also exist.

[0028] As shown in FIG. 1, a computing environment **100** includes, for instance, a node **10** having, e.g., a computer system/server **12**, which is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with computer system/server **12** include, but are not limited to, personal computer (PC) systems, server computer systems, thin clients, thick clients, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems, and distributed cloud computing environments that include any of the above systems or devices, and the like.

[0029] Computer system/server **12** may be described in the general context of computer system-executable instructions, such as program modules, being executed by a computer system. Generally, program modules may include routines, programs, objects, components, logic, data structures, and so on that perform particular tasks or implement particular abstract data types. Computer system/server **12** may be practiced in many computing environments, including but not limited to, distributed cloud computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed cloud computing environment, program modules may be located in both local and remote computer system storage media including memory storage devices.

[0030] As shown in FIG. 1, computer system/server **12** is shown in the form of a general-purpose computing device. The components of computer system/server **12** may include, but are not limited to, one or more processors or processing units **16**, a system memory **28**, and a bus **18** that couples various system components including system memory **28** to processor **16**.

[0031] Bus **18** represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus.

[0032] Computer system/server **12** typically includes a variety of computer system readable media. Such media may be any available media that is accessible by computer system/server **12**, and it includes both volatile and non-volatile media, removable and non-removable media.

[0033] System memory **28** can include computer system readable media in the form of volatile memory, such as random access memory (RAM) **30** and/or cache memory **32**. Computer system/server **12** may further include other removable/non-removable, volatile/non-volatile computer system storage media. By way of example only, storage system **34** can be provided for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a “hard drive”). Although not shown, a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a “floppy disk”), and an optical disk drive for reading from or writing to a removable, non-volatile optical disk such as a CD-ROM, DVD-ROM or other optical media can be provided. In such instances, each can be connected to bus **18** by one or more data media interfaces. As will be further depicted and described below, memory **28** may include at least one program product having a set (e.g., at least one) of program modules that are configured to carry out the functions of embodiments of the invention.

[0034] Program/utility **40**, having a set (at least one) of program modules **42**, may be stored in memory **28** by way of example, and not limitation, as well as an operating system **17**, one or more application programs, other program modules, and program data. Each of the operating system, one or more application programs, other program modules, and program data or some combination thereof, may include an implementation of a networking environment. Program modules **42** may generally carry out the functions and/or methodologies of one or more embodiments of the invention as described herein.

[0035] Computer system/server **12** may also communicate with one or more external devices **14** such as a keyboard, a pointing device, a display **24**, etc.; one or more devices that enable a user to interact with computer system/server **12**; and/or any devices (e.g., network card, modem, etc.) that enable computer system/server **12** to communicate with one or more other computing devices. Such communication can occur via Input/Output (I/O) interfaces **22**. Still yet, computer system/server **12** can communicate with one or more networks such as a local area network (LAN), a general wide area network (WAN), and/or a public network (e.g., the Internet) via network adapter **20**. As depicted, network adapter **20** communicates with the other components of computer system/server **12** via bus **18**. It should be understood that although not shown, other hardware and/or software components could be used in conjunction with computer system/server **12**. Examples, include, but are not limited to: microcode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

[0036] As an example, processing unit **16** may include a plurality of functional components used to execute instructions. These functional components may include, for instance, an instruction fetch component to fetch instructions to be executed; an instruction decode unit to decode the fetched instructions; instruction execution components to execute the decoded instructions; a memory access component to access memory for instruction execution, if necessary; and a write back component to provide the results of the executed instructions. One or more of these components may, in accordance with an aspect of the present invention, be used to implement deviations from typical memory management protocols with respect to a cache line corre-

sponding to a storage-operand request identified as hold, and resuming typical memory management protocols with respect to a cache line corresponding to a storage-operand request identified as release, as described further below.

[0037] Processing unit **16** also includes, in one embodiment, one or more registers to be used by one or more of the functional components.

[0038] Another example of a computing environment to incorporate and use one or more aspects of the present invention is described with reference to FIG. **2**. In one example, the computing environment may again be based on the above-referenced *z*/Architecture, offered by International Business Machines Corporation, Armonk, N.Y. In another example, the computing environment may be based on the above-referenced Power Architecture, offered by International Business Machines Corporation, Armonk, N.Y. The computing environment may also be based on other architectures, including, but not limited to, the Intel 64 and IA-32 architectures. Other examples also exist.

[0039] Referring to FIG. **2**, in one example, a computing environment **200** includes at least one processing unit (PU), such as a central processing unit (CPU), coupled to one or more higher level caches **205**. Central processing unit **16** includes, for instance, one or more execution units **201** to request and consume data. Execution units **201** are coupled to a load/store unit (LSU) **202**, which includes a load/store queue (LSQ) **203**. Load/store unit **202** issues memory access operations (loads and stores) that retrieve prefetched data or cause the data to be fetched from the memory subsystem or higher level caches. LSU **202** may be coupled to a local (or private) cache(s) **204** via, e.g., a translation mechanism, such as, e.g., a translation look-aside buffer (TLB) or an effective-to-real address translation table (ERAT).

[0040] Processing unit **16** is coupled to and communicates with a memory hierarchy that includes, for instance, local (or private) cache(s) **204**, which may include an L1 data cache; one or more higher level caches **205**, which may include a single level two (L2) cache or multiple other sequentially numbered higher levels, e.g., L3, L4; a memory **215**; and an associated memory controller **210**, which controls accesses to memory **215**. Local cache(s) **204** serves as a prefetch buffer for data (and/or data streams) that are prefetched. Local cache(s) **204** has a corresponding load miss queue (LMQ), which the cache utilizes to save information about ongoing prefetch requests.

[0041] Further, in one embodiment, local cache(s) **204** may have an associated cache directory, which may be implemented as part of the cache or be maintained separate therefrom. The cache directory may include a tag for each cache line in the local cache. The tag provides information about the cache line, including at least part of a real or absolute address of the data in the cache line, as well as status information, such as whether the data is being held exclusive or shared, etc.; the use of the data, such as whether it is most recently used, least recently used, or somewhere in between, etc.; and/or other status information. The tag may also include additional information regarding the data in the cache line. Each cache level may have an associated cache directory, either its own directory or a shared directory.

[0042] In addition to the example memory hierarchy, computing environment **200** may also include additional storage devices that form a part of the memory hierarchy from the perspective of processing unit **16**. The storage devices may be one or more electronic storage media, such

as a floppy disk, hard drive, CD-ROM, or DVD. CPU 102 communicates with each of the above devices within the memory hierarchy by various mechanisms, including via buses and/or direct channels, as examples.

[0043] The buses may represent one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus.

[0044] In a further embodiment, the processing disclosed herein may be used in a multiprocessor environment, such as the one described below with reference to FIG. 3. In one example, a multiprocessor environment 300 includes a plurality of cores 302a . . . 302n. Each core includes at least one hardware thread or processing unit 304a, 304n that includes a local (or private) cache, such as an L1 cache 306a, 306n. Each core 302a, 302n is coupled to zero or more higher level private caches (e.g., L2 cache(s)). Additionally, multiprocessor environment 300 includes zero or more shared caches 310, such as L3 and/or L4 caches, which are shared between the cores, and more particularly, shared among the processing units. Other topologies may also exist.

[0045] Each processing unit in the multiprocessor environment (or a subset of the processing units) may also include a query function to be used to query one or more of the caches, including one or more of the shared cache(s), if available, to obtain cache residency status and/or other information of selected data.

[0046] In accordance with one or more aspects of the present invention, a cache line tracker 320 is provided in one or more of the processing units 306a, 306n. The tracker 320 can include one or more logical operators or instructions that allow it to store a cache line address from, for example, higher level shared cache(s) 310 that the particular processor 306a has a lock. An example is instructive in this regard and is set forth with respect to FIGS. 3 and 4. It shall be understood that the tracker 320 could be included in any (e.g. processing unit 16) disclosed herein. Also, while shown in the processing units herein, the tracker 320 is not so limited and can be included in any location in any of the cores disclosed herein.

[0047] The tracker 320 includes a cache line address store (address) 402 that is currently being used for a particular instruction or that may be needed for next (or future) instruction that the processing unit 304 has currently locked. The address 402 can be implemented in hardware as an address store. This can be register in one embodiment.

[0048] When the lock is acquired, one or more timers discussed below may be started. The timers can be located in any location and may be implemented in software/hardware or a combination thereof and for clarity are shown, for example, generally by timers 322a, 322n in FIG. 3. In general operation, the lock is held for as long as possible to allow for as much of a particular operation to be completed. To that end, the tracker 320 includes a lock required bit 404 that is set (e.g., to a logical "1") while the lock is being held released (e.g., forced to a logical zero) when the lock is no longer needed. In some instances, the lock required bit 404 can be referred to or considered a "tracker valid bit" that indicates that the tracker 320 is to deny cache invalidate (XI)

requests as further described below. When this bit is not set, XI commands are allowed and the line is released.

[0049] In operation, when another processor requests access to the locked address, the memory controller 210 can send a request (such as cache invalidate (XI) request 406) to the processor that has locked the address. In embodiments disclosed herein, the tracker 320 will receive such a request and provide a response 408 that either allows another processor to lock the cache line or not. The response will depend on the status of the lock required bit. If the lock is still needed, any request will be denied (e.g., response 408 will be negative). If it is not, then the request will be granted. The response can be generated by tracker logic 403 in one embodiment. The logic 403 can include instructions or hardware causing it to provide the comparisons described herein and produce the responses 408 in the manner described herein.

[0050] It shall further be understood that the lock bit is required when the state of the cache line is deemed exclusive to the particular processor.

[0051] As discussed herein, upon determination that a cache line is to be held exclusively, a timer 322 can be set. The timer may be initiated to a set interval for continued holding of the corresponding cache line (e.g., setting the lock required bit 404) in the private cache exclusively after completing processing of the storage operand request, and discontinuing the holding of the corresponding cache line in private cache exclusively after reaching the set interval. In one or more embodiments, a length of the set interval may be dependent on an access type of the storage operand request. The access type may be, for instance, a fetch-type, a store-type or update-type access. In certain embodiments, the length of the set interval may be greater for a store-type access or an update-type access than for fetch-type access. By way of example, the set interval may be a time-based interval or a cycle-based interval, and the continuing hold of the corresponding cache line may be released upon the timer counting down (or up) to the set count interval.

[0052] As discussed, above, in one or more embodiments, any cache invalidate commands (e.g., from a system control of the computing environment) for the corresponding cache line in the private cache are rejected during the continuing hold of the corresponding cache line in private cache exclusively after completing processing of the storage operand request.

[0053] In one embodiment, the tracker 320 can also include an XI tracker bit 405. This bit is set when an XI command has been received and rejected. In one embodiment, this bit can serve to prevent the tracker 320 from dominating a particular address. Consider the following situation: The tracker 320 has locked a particular address 402 and received an XI command that it rejected. In normal operation, the requester may wait a certain amount time to make a subsequent request. During that time, the address 402 may be released. Now suppose that the same processing unit 320 again wants to lock the same address before the subsequent XI command is received. Without an indication that another processing unit has requested the address, the processing unit would just "reestablish" the same lock by never releasing the address. Certain logic in the tracker 320 can address such a situation as described further below.

[0054] The tracker 320 also includes a first reset input 422 that causes the lock required bit 404 to be reset and allow for access to the cache line address 402 upon the next request.

[0055] With further reference to FIGS. 5A and 5B operation of the system shown in FIGS. 3-5B will be described. In one aspect, some or all of the release steps described therein can be performed in hardware. Additional releases may also be provided as are suggested below. Regardless of the number of releases, embodiments herein may help alleviate or remove hang ups that can occur when a processor that has locked a particular address is not functioning properly or is abusing access to a particular cache line.

[0056] Referring first to FIG. 5A, a processing unit may fetch an instruction 500, and decode the instruction to determine, for instance, operand related signals, an execution unit (LSU, FXU, etc.), a “hold” designation, a “release” designation, etc. 505. As will be understood, the hold and release designations can be reflected in the value of the lock required bit 404. When bit 404 is “set” XI commands are denied. Otherwise, they are allowed. To that end, the tracker 320 includes a set input 420 that causes the lock required bit 404 to be set. The setting of the tracker 404 can include loading the locked address into address 402. The set input can receive a set command 430 from, for example, an LSU 202 or other controller.

[0057] The lock required bit 404 is shown as a separate element but in one embodiment it can be bit in the same hardware element as the address 402. The same is true of the XI tracker bit 405.

[0058] The instruction is queued for execution 510, and the processing unit determines whether the instruction is ready to issue 515. If “no”, then the instruction remains in the queue 510. Otherwise, the instruction is issued to the execution unit. In this example, the instruction is assumed to include a storage operand request or access and is issued to the load store unit (LSU) 501 (FIG. 5B) for processing.

[0059] Referring to FIG. 5B, LSU 501 determines the storage operand address 530, and requests storage operand access from the entire storage hierarchy (including private cache to the processing unit, such as L1 cache).

[0060] In the implementation depicted, LSU 501 determines whether the storage operand request is designated as hold 532. If “no”, then standard memory management protocols are employed in processing the storage operand request. If “yes”, then LSU 501 determines the type of storage access. In the example depicted, the type of storage access of interest may involve a fetch-type access, or alternatively, a store-type access or update-type access associated with the storage operand request.

[0061] If the corresponding cache line in private cache is held exclusive (EX), then timer 322 is initiated to facilitate holding of the cache line for a desired (or set) interval. Note that in one or more other embodiments, a cycle-based interval may be used rather than a time-based interval. Further, it shall be understood that the timer 322 can have be set based on the type of command. For instance, a first timer value (e.g., timer value) can be established for fetch type command and a second, longer timer value can be set for store or update types of commands.

[0062] By way of further example, one embodiment of an instruction format which may be used in accordance with one or more embodiments of the present invention is described below with reference to FIGS. 6A-7. In FIG. 6A, an instruction format is depicted known as a Next Instruction Access Intent (NIAI) instruction. This instruction format may include an operation code field 601 indicating that this is a next instruction access intent instruction. Addition-

ally, a reserved field 602 may be provided, which in one or more embodiments may contain zeros. The instruction format may also include fields or operands, such as an access intent control field 603 (I_1), and another access intent control field 604 (I_2).

[0063] For this instruction format, the term primary-access operand means the lowest numbered storage operand of an instruction. Similarly, the term secondary-access operand means the next lowest numbered storage operand of the instruction. These terms do not have the same meaning as the terms first operand and second operand, even though there may be some correlation.

[0064] Subject to the controls in the I_1 603 and I_2 604 fields, the CPU is signaled to future access intent for either or both the primary-access and the secondary-access operands of the next sequential instruction. The next sequential instruction may be the instruction retrieved and executed after the NIAI instruction. The I_1 603 field may contain a code to signal the CPU the access intent for the primary-access operand of the next sequential instruction. The I_2 604 field may contain a code to signal the CPU the access intent for the secondary-access operand of the next sequential instruction. When the next sequential instruction has only a single storage operand, the I_2 604 field may be ignored.

[0065] One embodiment of an instruction format may include a single access intent field. Another embodiment of an instruction format may include two access intent fields. Another embodiment of instruction format may include multiple access intent fields. Still another embodiment of instruction format may include a field or operand that contains a value specifying the number of subsequent next sequential instructions. The access intent control fields I_1 603 and I_2 604 may have format 610 such as depicted in FIG. 6B, where a set number of bits of the I_1 603 and I_2 604 fields contain an unsigned integer that is used as a code to signal the CPU the access intent for the corresponding operand of the next sequential instruction.

[0066] As described above, the hardware based tracker can be set or reset based on whether a command (such as a hold) is received and a timer, such as timer 322, expires or when the LSU issues a release command such as indication instruction 432 in FIG. 4. The expiration of the timer 322 is indicated by timeout signal 434. Methods/hardware for generation of the timeout signal 434 upon expiration of timer 322 are known in the art and generally understood to be included with timer 322. The actual signal could be generated, however, by the processing unit. As discussed further below, timer 322 can generate multiple time out signals. As such, in some instances, the timer 322 may be multiple timers.

[0067] According to one embodiment, the processing unit 304 can include an OR gate 436 that receives the timeout signal 434 and the release instruction 434. In the event either instruction is positive, the OR gate 436 will cause output a positive signal. The output of the OR gate is connected to the first reset input 422 of the tracker 320 and will cause the tracker 320 to be reset. In one embodiment, resetting the tracker only resets the lock required bit 404 and not the XI tracker bit 405. As will be discussed further below, additional optional signals may also be provided to the OR gate 436 to ensure that tracker 320 releases the address 402 by resetting (e.g., setting to “0”) the lock required bit 404.

[0068] In one or more embodiments, a processing unit behavior control is described herein to selectively continue

holding the corresponding cache line in private cache exclusively after completing processing of a storage operand request where the storage operand request is designated hold, and the corresponding cache line in private cache used for processing the storage operand request is owned exclusively by the processing unit as indicated by block 536 in FIG. 5B.

[0069] As noted, the continuing hold may include initiating a timer to facilitate the hold for a desired interval. Further, based on the access type being store-type or update-type, a subsequent storage operand request may be designated release for the corresponding cache line in order to discontinue the hold on the corresponding cache line in private cache exclusively. Identifying the storage operand request as “hold” or “release” may be accomplished, as noted above, in one or more implementations, utilizing a designated access intent code in one or more of the control fields of the access intent instruction.

[0070] Stated otherwise, when a storage operand designated hold is successfully processed, a primary result of the hardware based system including the OR gate 436 and the tracker 320 described herein is for the processing unit 304 to reject cache invalidate commands from system control that are associated with a specific cache line held exclusively by the processing unit. The processing unit 304 can cause the reject by giving control to the tracker or forwarding the output of the tracker to the memory controller 210. A byproduct of rejecting these commands is that the corresponding cache line is held exclusively throughout the storage hierarchy by the processing unit longer than typical protocols would permit. In such a case, upon receipt of a cache invalidate (XI) command from system control, the processing unit will no longer hold the cache line. That is, whenever the output of OR gate 436 is a logical “1” the cache line is released (e.g., response 408 allows the line to go to another processing unit) the next time an XI command 406 is received.

[0071] In one embodiment, when the counter 322 reaches a second preprogrammed level, the XI tracker bit 405 can be reset. This will avoid situations where the other processor no longer needs the cache line. This second reset input is provided by output 450 and resets only the tracker bit 405.

[0072] In one embodiment, when an XI command is received a second timer 452 is started. This timer is reset each time any instruction completes. When it reaches a tracker operational time limit, the timer 452 will reset the tracker. Consider for example a hypothetical case where a first processing unit obtained a first lock for a first location A. As above, any XI requests for line A will reject that request. Now assume that the first processing unit also needs to lock location B. If another processing unit has already locked location B and now wants to lock line A, a stalemate can be created wherein the first processing unit is waiting for line B while holding line A and the second processing unit is waiting for line A and holding line B. The second timer 452 can detect such a case as it starts counting when some second processing units wants the line held by the first processing unit in the tracker, and it continues ticking as long as the first processing unit does not complete any instruction. This will suggest that the first processing unit is stuck and the tracker should be reset.

[0073] FIG. 7 shows another embodiment of a circuit that includes a tracker 320. This embodiment works in the same manner as described above but includes one or more

optional additional signals that cause the tracker 320 to reset. Each of these optional additional signals will be provided, along with signals 450, 434 and 432, to an OR gate 702 (or a functional equivalent thereof) to provide a signal to the first reset input 422.

[0074] One of the additional signals is “switch program” signal 720 that indicates that the processor is no longer working on the program that requested the hold. This signal can be generated by the processing unit 304. In such a case, the tracker 320 is reset. Another additional signal is an L3 hang signal 722 that indicates that the L3 cache believes a line has been held for too long. The tracker 320 can also be reset when the processing unit itself or another processor determines that the processor that includes the tracker 320 hung (line 724). Finally, the tracker 320 and the XI tracker bit 405 can be reset when L2 cache determines that it became the LRU line as indicated by line 726. In operation the private L1 and L2 caches in each processing unit have limited space, so sometimes when the processing unit fetches a new line from memory it must evict one of the existing lines in its cache, just for the space it occupies. Usually the line chosen to be evicted is the Least Recently Used (LRU), so the process of evicting a line to make space to a new line is nicknamed LRUing, and the line about to be evicted is the “LRU line”. So if the line held in the tracker is evicted, the tracker needs to be reset as the line is no longer being used in the processing unit.

[0075] Also shown in FIG. 7 is logic to control setting the tracker 320. This logic can include a OR gate 760 to ensure that only trusted software can set the tracker or that the processor has issued a configuration “enable all” command (line 762). The skilled artisan will understand how to determine the value of the trusted software input 764.

[0076] The output of OR gate 760 can be provided along with LSU set signal 430 at AND gate 766 to set tracker 320.

[0077] Another embodiment of a computing environment to incorporate and use one or more aspects is described with reference to FIG. 8A. In this example, a computing environment 900 includes, for instance, a native central processing unit (CPU) 902, a memory 904, and one or more input/output devices and/or interfaces 906 coupled to one another via, for example, one or more buses 908 and/or other connections. As examples, computing environment 900 may include a PowerPC processor or a pSeries server offered by International Business Machines Corporation, Armonk, N.Y.; an HP Superdome with Intel Itanium II processors offered by Hewlett Packard Co., Palo Alto, Calif.; and/or other machines based on architectures offered by International Business Machines Corporation, Hewlett Packard, Intel, Oracle, or others.

[0078] Native central processing unit 902 includes one or more native registers 910, such as one or more general purpose registers and/or one or more special purpose registers used during processing within the environment. These registers include information that represent the state of the environment at any particular point in time.

[0079] Moreover, native central processing unit 902 executes instructions and code that are stored in memory 904. In one particular example, the central processing unit executes emulator code 912 stored in memory 904. This code enables the computing environment configured in one architecture to emulate another architecture. For instance, emulator code 912 allows machines based on architectures other than the z/Architecture, such as PowerPC processors,

pSeries servers, HP Superdome servers or others, to emulate the z/Architecture and to execute software and instructions developed based on the z/Architecture.

[0080] Further details relating to emulator code **912** are described with reference to FIG. **9B**. Referring collectively to FIGS. **8A** and **8B**, guest instructions **950** stored in memory **904** include software instructions (e.g., correlating to machine instructions) that were developed to be executed in an architecture other than that of native CPU **902**. For example, guest instructions **950** may have been designed to execute on a z/Architecture processor, but instead, are being emulated on native CPU **902**, which may be, for example, an Intel Itanium II processor. In one example, emulator code **912** includes an instruction fetching routine **952** to obtain one or more guest instructions **950** from memory **904**, and to optionally provide local buffering for the instructions obtained. It also includes an instruction translation routine **954** to determine the type of guest instruction that has been obtained and to translate the guest instruction into one or more corresponding native instructions **956**. This translation includes, for instance, identifying the function to be performed by the guest instruction and choosing the native instruction(s) to perform that function.

[0081] Further, emulator code **912** includes an emulation control routine **960** to cause the native instructions to be executed. Emulation control routine **960** may cause native CPU **902** to execute a routine of native instructions that emulate one or more previously obtained guest instructions and, at the conclusion of such execution, return control to the instruction fetch routine to emulate the obtaining of the next guest instruction or a group of guest instructions. Execution of the native instructions **956** may include loading data into a register from memory **904**; storing data back to memory from a register; or performing some type of arithmetic or logic operation, as determined by the translation routine.

[0082] Each routine is, for instance, implemented in software, which is stored in memory and executed by native central processing unit **902**. In other examples, one or more of the routines or operations are implemented in firmware, hardware, software or some combination thereof. The registers of the emulated processor may be emulated using registers **910** of the native CPU or by using locations in memory **904**. In embodiments, guest instructions **950**, native instructions **956** and emulator code **912** may reside in the same memory or may be disbursed among different memory devices.

[0083] As used herein, firmware includes, e.g., the microcode, millicode and/or macrocode of the processor. The term includes, for instance, the hardware-level instructions and/or data structures used in implementation of higher level machine code. In one embodiment, firmware may include, for instance, proprietary code that is typically delivered as microcode that includes trusted software or microcode specific to the underlying hardware and controls operating system access to the system hardware.

[0084] A guest instruction **950** that is obtained, translated and executed is, for instance, one or more of the instructions described herein. The instruction, which is of one architecture (e.g., the z/Architecture), is fetched from memory, translated and represented as a sequence of native instructions **956** of another architecture (e.g., PowerPC, pSeries, Intel, etc.). These native instructions may then be executed.

[0085] Although various examples are provided, variations are possible without departing from a spirit of the

claimed aspects. For example, values that are included in registers and/or fields used by the instruction may, in other embodiments, be in other locations, such as memory locations, etc. Many other variations are possible.

[0086] One or more aspects may relate to cloud computing.

[0087] It is understood in advance that although this disclosure includes a detailed description on cloud computing, implementation of the teachings recited herein are not limited to a cloud computing environment. Rather, embodiments of the present invention are capable of being implemented in conjunction with any other type of computing environment now known or later developed.

[0088] Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g. networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines, and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models, and at least four deployment models.

[0089] Characteristics are as follows:

[0090] On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with the service's provider.

[0091] Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

[0092] Resource pooling: the provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter).

[0093] Rapid elasticity: capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

[0094] Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized service.

[0095] Service Models are as follows:

[0096] Software as a Service (SaaS): the capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

[0097] Platform as a Service (PaaS): the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

[0098] Infrastructure as a Service (IaaS): the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

[0099] Deployment Models are as follows:

[0100] Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

[0101] Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

[0102] Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

[0103] Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

[0104] A cloud computing environment is service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability. At the heart of cloud computing is an infrastructure comprising a network of interconnected nodes. One such node is node 10 depicted in FIG. 1.

[0105] Computing node 10 is only one example of a suitable cloud computing node and is not intended to suggest any limitation as to the scope of use or functionality of embodiments of the invention described herein. Regardless, cloud computing node 10 is capable of being implemented and/or performing any of the functionality set forth hereinabove.

[0106] Referring now to FIG. 9, illustrative cloud computing environment 50 is depicted. As shown, cloud computing environment 50 comprises one or more computing nodes 10 with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or cellular telephone 54A, desktop computer 54B, laptop computer 54C, and/or automobile computer system 54N may communicate. Nodes 10 may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as Private, Community, Public, or Hybrid clouds as described hereinabove, or a combination thereof. This allows cloud computing environment 50 to offer infrastructure, platforms and/or

software as services for which a cloud consumer does not need to maintain resources on a local computing device. It is understood that the types of computing devices 54A-N shown in FIG. 7 are intended to be illustrative only and that computing nodes 10 and cloud computing environment 50 can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

[0107] Referring now to FIG. 10, a set of functional abstraction layers provided by cloud computing environment 50 (FIG. 9) is shown. It should be understood in advance that the components, layers, and functions shown in FIG. 11 are intended to be illustrative only and embodiments of the invention are not limited thereto. As depicted, the following layers and corresponding functions are provided:

[0108] Hardware and software layer 60 includes hardware and software components. Examples of hardware components include mainframes 61; RISC (Reduced Instruction Set Computer) architecture based servers 62; servers 63; blade servers 64; storage devices 65; and networks and networking components 66. In some embodiments, software components include network application server software 67 and database software 68.

[0109] Virtualization layer 70 provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers 71; virtual storage 72; virtual networks 73, including virtual private networks; virtual applications and operating systems 74; and virtual clients 75.

[0110] In one example, management layer 80 may provide the functions described below. Resource provisioning 81 provides dynamic procurement of computing resources and other resources that are utilized to perform tasks within the cloud computing environment. Metering and Pricing 82 provide cost tracking as resources are utilized within the cloud computing environment, and billing or invoicing for consumption of these resources. In one example, these resources may comprise application software licenses. Security provides identity verification for cloud consumers and tasks, as well as protection for data and other resources. User portal 83 provides access to the cloud computing environment for consumers and system administrators. Service level management 84 provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment 85 provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA.

[0111] Workloads layer 90 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include: mapping and navigation 91; software development and lifecycle management 92; virtual classroom education delivery 93; data analytics processing 94; transaction processing 95; and instruction processing 96.

[0112] The present invention may be a system, a method, and/or a computer program product at any possible technical detail level of integration. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

[0113] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0114] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0115] Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may

execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[0116] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0117] These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0118] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0119] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[0120] In addition to the above, one or more aspects may be provided, offered, deployed, managed, serviced, etc. by a service provider who offers management of customer environments. For instance, the service provider can create, maintain, support, etc. computer code and/or a computer

infrastructure that performs one or more aspects for one or more customers. In return, the service provider may receive payment from the customer under a subscription and/or fee agreement, as examples. Additionally or alternatively, the service provider may receive payment from the sale of advertising content to one or more third parties.

[0121] In one aspect, an application may be deployed for performing one or more embodiments. As one example, the deploying of an application comprises providing computer infrastructure operable to perform one or more embodiments.

[0122] As a further aspect, a computing infrastructure may be deployed comprising integrating computer readable code into a computing system, in which the code in combination with the computing system is capable of performing one or more embodiments.

[0123] As yet a further aspect, a process for integrating computing infrastructure comprising integrating computer readable code into a computer system may be provided. The computer system comprises a computer readable medium, in which the computer medium comprises one or more embodiments. The code in combination with the computer system is capable of performing one or more embodiments.

[0124] Although various embodiments are described above, these are only examples. For example, computing environments of other architectures can be used to incorporate and use one or more embodiments. Further, different instructions, instruction formats, instruction fields and/or instruction values may be used. Many variations are possible.

[0125] Further, other types of computing environments can benefit and be used. As an example, a data processing system suitable for storing and/or executing program code is usable that includes at least two processors coupled directly or indirectly to memory elements through a system bus. The memory elements include, for instance, local memory employed during actual execution of the program code, bulk storage, and cache memory which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution.

[0126] Input/Output or I/O devices (including, but not limited to, keyboards, displays, pointing devices, DASD, tape, CDs, DVDs, thumb drives and other memory media, etc.) can be coupled to the system either directly or through intervening I/O controllers. Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modems, and Ethernet cards are just a few of the available types of network adapters.

[0127] The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting. As used herein, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms “comprises” and/or “comprising”, when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components and/or groups thereof.

[0128] The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below, if any, are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of one or more embodiments has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art. The embodiment was chosen and described in order to best explain various aspects and the practical application, and to enable others of ordinary skill in the art to understand various embodiments with various modifications as are suited to the particular use contemplated.

What is claimed is:

1. A memory access control system for use in a computing environment including two or more processing units, each of the two or more processing units including a tracker, the tracker on each processing unit including:

an address storage element that stores an address to be locked by one of the processing units;

a lock required bit;

a cache invalidate (XI) tracker bit;

a set input that upon receipt of a set command sets the lock required bit when a first condition is met;

a first reset input that resets the lock required bit upon receipt of a reset command;

a second reset input that resets the XI tracker bit; and

logic that causes the tracker to reject incoming XI commands when the lock required bit is set, to allow incoming XI commands when the lock required bit is not set, and to set the XI tracker bit when a first incoming XI command is received.

2. The system of claim 1, further comprising:

a first timer, the first timer connected to the tracker such that it resets the lock required bit after expiration of a first threshold.

3. The system of claim 2, wherein the first timer is connected to the tracker such that it resets the XI tracker bit after expiration of a second threshold, the second threshold being longer than the first threshold.

4. The system of claim 2, further comprising:

a second timer, the second timer connected to the tracker such that it resets the lock required bit and the XI tracker bit after expiration of a tracker threshold limit.

5. The system of claim 4, wherein the system includes an OR gate connected on its inputs to an output of the first timer, an output of the second timer, and a signal from a processing unit indicating that the cache line can be released and on its output to the first reset input of the tracker.

6. The system of claim 2, wherein the system includes an OR gate connected on its inputs to an output of the first timer and a signal from a processing unit indicating that the cache line can be released and on its output to the first reset input of the tracker.

7. The system of claim 6, wherein the OR gate includes one or more additional inputs.

8. The system of claim 7, wherein the one or more additional inputs include one or more of:

a switch program signal that indicates that a first processing unit is no longer working on a program that requested the hold;

an L3 cache hang signal that indicates that an L3 cache believes a line has been held for too long; and

an external reset signal indicating that the first processing unit itself or another processing unit determines that the first processing unit that includes the tracker is hung.

9. The system of claim 7, wherein the one or more additional inputs include an L2 cache signal received when an L2 cache determines that the address in the address storage element was evicted, wherein the XI tracker bit is also reset upon receipt of the L2 cache signal.

10. The system of claim 1, wherein the logic determines when the first condition is met, wherein the first condition is met when the XI tracker bit is not set.

11. The system of claim 1, wherein the set command is produced by a signal from a load store unit.

12. The system of claim 1, wherein the set command is produced by a logical combination of a signal from a load store unit and a signal produced by one or more of an enable signal and a trusted software input.

13. A memory access control system for use in a computing environment including two or more processing units and at least one shared memory controlled by a memory controller, the system including:

a tracker configured to receive cache invalidate (XI) commands from the memory controller and to provide responses to the memory controller, the tracker including:

an address storage element that stores an address to be locked by one of the processing units;

a lock required bit;

a cache invalidate (XI) tracker bit;

a set input that upon receipt of a set command sets the lock required bit when a first condition is met;

a first reset input that resets the lock required bit upon receipt of a reset command; and

a second reset input that resets the XI tracker bit;

wherein the tracker is configured to reject incoming XI commands from the memory controller when the lock required bit is set, to allow incoming XI commands when the lock required bit is not set, and to set the XI tracker bit when a first incoming XI command is received.

14. The system of claim 13, further comprising: a first timer,

wherein the first timer is connected to the tracker such that it resets the lock required bit after expiration of a first threshold; and

wherein the first timer is connected to the tracker such that it resets the XI tracker bit after expiration of a second threshold, the second threshold being longer than the first threshold.

15. The system of claim 14, further comprising:

a second timer, the second timer connected to the tracker such that it resets the lock required bit and the XI tracker bit after expiration of a tracker threshold limit.

16. The system of claim 15, wherein the system includes an OR gate connected on its inputs to an output of the first timer, an output of the second timer, and a signal from a processing unit indicating that the cache line can be released and on its output to the first reset input of the tracker.

17. The system of claim 14, wherein the system includes an OR gate connected on its inputs to an output of the first timer and a signal from a processing unit indicating that the cache line can be released and on its output to the first reset input of the tracker.

18. The system of claim 17, wherein the OR gate includes one or more additional inputs, wherein the one or more additional inputs include one or more of:

a switch program signal that indicates that a first processing unit is no longer working on a program that requested the hold;

an L3 cache hang signal that indicates that an L3 cache believes a line has been held for too long; and

an external reset signal indicating that the first processing unit itself or another processing unit determines that the first processing unit that includes the tracker is hung.

19. The system of claim 17, wherein the OR gate also receives an L2 cache signal received when an L2 cache determines that the address in the address storage element was evicted, wherein the XI tracker bit is also reset upon receipt of the L2 cache signal.

20. The system of claim 13, wherein the logic determines when the first condition is met, wherein the first condition is met when the XI tracker bit is not set.

* * * * *