

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
6 July 2006 (06.07.2006)

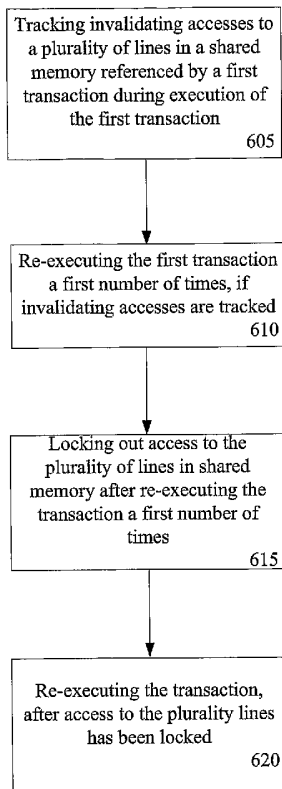
PCT

(10) International Publication Number  
WO 2006/071969 A1

- (51) International Patent Classification:  
G06F 9/52 (2006.01) G06F 9/46 (2006.01)
- (21) International Application Number:  
PCT/US2005/047376
- (22) International Filing Date:  
23 December 2005 (23.12.2005)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
11/027,623 29 December 2004 (29.12.2004) US
- (71) Applicant (for all designated States except US): INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): KOTTAPALLI, Sailesh [IN/US]; 330 Elan Village Lane #224, San Jose, CA 95134 (US). CRAWFORD, John, H. [US/US];
- 20128 Chateau Drive, Saratoga, CA 95070 (US). VAID, Kushagra [IN/US]; 507 White Chapel Drive, San Jose, CA 95136 (US).
- (74) Agents: VINCENT, Lester J. et al.; BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, 12400 Wilshire Boulevard, 7th Floor, Los Angeles, CA 90025 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

[Continued on next page]

(54) Title: TRANSACTION BASED SHARED DATA OPERATIONS IN A MULTIPROCESSOR ENVIRONMENT



(57) Abstract: The apparatus and method described herein are for handling shared memory Accesses between multiple processors utilizing lock-free synchronization through transactional-execution. A transaction demarcated in software is speculatively executed. During execution invalidating remote accesses/requests to addresses loaded from and to be written to share memory are tracked by a transactional buffer. If an invalidating access is encountered, the transaction is re-executed. After a pre-determined number of times re-executing the transaction, the transaction may be re-executed non-speculatively with locks/semaphores.

WO 2006/071969 A1



FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,  
RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,  
GN, GQ, GW, ML, MR, NE, SN, TD, TG).

— *before the expiration of the time limit for amending the  
claims and to be republished in the event of receipt of  
amendments*

**Published:**

— *with international search report*

*For two-letter codes and other abbreviations, refer to the "Guid-  
ance Notes on Codes and Abbreviations" appearing at the begin-  
ning of each regular issue of the PCT Gazette.*

**TRANSACTION BASED SHARED DATA OPERATIONS IN A MULTIPROCESSOR  
ENVIRONMENT**

**FIELD**

[0001] This invention relates to the field of integrated circuits and, in particular, to shared data operations between multiple integrated circuits, cores, and threads.

**BACKGROUND**

[0002] Advances in semi-conductor processing and logic design have permitted an increase in the amount of logic that may be present on integrated circuit devices. As a result, computer system configurations have evolved from a single or multiple integrated circuits in a system to multiple cores and multiple logical processors present on individual integrated circuits. An integrated circuit typically comprises a single processor die, where the processor die may include any number of cores or logical processors.

[0003] As an example, a single integrated circuit may have one or multiple cores. The term core usually refers to the ability of logic on an integrated circuit to maintain an independent architecture state, where each independent architecture state is associated with dedicated execution resources. Therefore, an integrated circuit with two cores typically

comprises logic for maintaining two separate and independent architecture states, each architecture state being associated with its own execution resources, such as low-level caches, execution units, and control logic. Each core may share some resources, such as higher level caches, bus interfaces, and fetch/decode units.

**[0004]** As another example, a single integrated circuit or a single core may have multiple logical processors for executing multiple software threads, which is also referred to as a multi-threading integrated circuit or a multi-threading core. Multiple logical processors usually share common data caches, instruction caches, execution units, branch predictors, control logic, bus interfaces, and other processor resources, while maintaining a unique architecture state for each logical processor. An example of multi-threading technology is Hyper-Threading Technology (HT) from Intel® Corporation of Santa Clara, California, that enables execution of threads in parallel using a single physical processor.

**[0005]** Current software has the ability to run individual software threads that may schedule execution on a plurality of cores or logical processors in parallel. The ever increasing number of cores and logical processors on integrated circuits enables more software threads to be executed. However, the increase in the number of software threads that

may be executed simultaneously have created problems with synchronizing data shared among the software threads.

[0006] One common solution to accessing shared data in multiple core or multiple logical processor systems comprises the use of locks to guarantee mutual exclusion across multiple accesses to shared data. As an example, if a first software thread is accessing a shared memory location, the semaphore guarding the shared memory location is locked to exclude any other software threads in the system from accessing the shared memory location until the semaphore guarding the memory location is unlocked.

[0007] However, as stated above, the ever increasing ability to execute multiple software threads potentially results in false contention and a serialization of execution. False contention occurs due to the fact that semaphores are commonly arranged to guard a collection of data, which, depending on the granularity of sharing supported by the software, may cover a very large amount of data. For this reason, semaphores act as contention "amplifiers" in that there may be contention by multiple software threads for the semaphores, even though the software threads are accessing totally independent data items. This leads to situations where a first software thread locks a semaphore guarding a data location that a second software thread may safely access without disrupting the execution of the first software thread. Yet, since the first

software thread locked the semaphore, the second thread must wait until the semaphore is unlocked, resulting in serialization of an otherwise parallel execution.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0008] The present invention is illustrated by way of example and not intended to be limited by the figures of the accompanying drawings.

[0009] Figure 1 illustrates an integrated circuit having N cores and M logical processors in each of the N cores.

[0010] Figure 2 illustrates an embodiment of an integrated circuit for implementing transactional execution.

[0011] Figure 3 illustrates an embodiment of the transaction buffer shown in Figure 3.

[0012] Figure 4 illustrates a transaction demarcated in software code, the software code shown compiled into a first and a second embodiment of compiled code.

[0013] Figure 5 illustrates an embodiment of transaction execution in a system.

[0014] Figure 6 illustrates an embodiment of a flow diagram for a method of executing a transaction.

[0015] Figure 7 illustrates an embodiment of the code flow for transactional execution.

#### **DETAILED DESCRIPTION**

[0016] In the following description, numerous specific details are set forth such as a specific number of physical/logical processors, specific transaction buffer fields, and specific processor logic and implementations in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that these specific details need not be employed to practice the present invention. In other instances, well known components or methods, such well-known functional blocks of a microprocessor, etc., have not been described in detail in order to avoid unnecessarily obscuring the present invention.

[0017] The apparatus and method described herein are for handling shared memory accesses between multiple software threads utilizing lock-free synchronization through transactional-execution. It is readily apparent to one skilled in the art, that the method and apparatus disclosed herein may be implemented in any level computer system, such as personal digital assistants, mobile platforms, desktop platforms, and server platforms, as well as with any number of integrated circuits, cores, or logical processors. For example, a multiprocessor system with four

integrated circuits may use the method and apparatus herein described to manage shared accesses to a memory shared by any four of the integrated circuits.

**[0018]** In **Figure 1** integrated circuit 105, which may implement transactional execution, is shown. In one embodiment, integrated circuit 105 is a microprocessor capable of operating independently from other microprocessors. Alternatively, integrated circuit 105 is a processing element that operates in conjunction with a plurality of processing elements.

**[0019]** Integrated circuit 105 illustrates first core 110, second core 115, and Nth core 120. A core, as used herein, refers to any logic located on an integrated circuit capable to maintain an independent architecture state, wherein each independently maintained architecture state is associated with at least some dedicated execution resources. Execution resources may include arithmetic logic units (ALUs), floating-point units (FPUs), register files, operand registers for operating on single or multiple integer and/or floating-point data operands in serial or parallel, and other logic for executing code. Moreover, a plurality of cores may share access to other resources, such as high-level caches, bus interface and control logic, and fetch/decode logic.



[0020] As an illustrative example, integrated circuit 105 has eight cores, each core associated with a set of architecture state registers, such as general-purpose registers, control registers, advanced programmable interrupt control (APIC) registers, machine state registers (MSRs), or registers for storing the state of an instruction pointer, to maintain an independent architecture state. Furthermore, each set of architecture state registers are exclusively associated with individual execution units.

[0021] Integrated circuit 105 also illustrates core 110 comprising first logical processor 125, second logical processor 130, and Mth logical processor 135. A logical processor, as used herein, refers any logic located on an integrated circuit capable to maintain an independent architecture state, wherein the independently maintained architecture states share access to execution resources. As above, each logical processor has a set of architecture state registers to maintain an independent architecture state; however, each of the architecture states share access to the execution resources. Consequently, on any single integrated circuit there may be any number of cores and/or any number of logical processors. For the purpose of illustration, the term processor will be referred to in discussing the operation of a core and/or a logical processor when discussing the apparatus and method used for transactional execution.

[0022] Referring to **Figure 2**, an embodiment of an integrated circuit is depicted to illustrate a specific implementation of transactional execution. However, it is readily apparent that the method and apparatus described in reference to **Figure 2** may be implemented in any level system, such as the system depicted in **Figure 5**. In one embodiment, integrated circuit 205 is capable of out-of-order speculative, where instructions are able to be executed in an order that is different that given in a program. Alternatively, processor 205 is capable of in-order execution, where the instructions are issued and executed in original program order.

[0023] Integrated circuit 205 may comprise any number of processors, which may be cores or logical processors. For instance, integrated circuit 205 has eight cores, each core having two logical processors, which would allow for execution of 16 software threads on integrated circuit 205 at one time. Consequently, integrated circuit 205 is typically referred to as a multi-threading multi-core processor. In **Figure 2**, integrated circuit 205 is depicted individually, as to not obscure the invention; yet, integrated circuit 205 may operate individually or in cooperation with other processors.

[0024] Integrated circuit 205 may also include, but is not required to include, any one or any combination of the following, which are not specifically depicted: a data path, an instruction path, a virtual memory address translation unit (a translation buffer), an arithmetic logic unit (ALU), a floating point calculation unit capable of executing a single instruction or multiple instructions, as well as capable to operate on single or multiple data operands in serial or in parallel, a register, an interrupt controller, an advanced programmable interrupt controller (APIC), a pre-fetch unit, an instruction re-order unit, and any other logic that is be used for fetching or executing instructions and operating on data.

[0025] Integrated circuit 205 illustrates front-end 210. Front-end 210 is shown as including instruction fetch 215, instruction decode 220, and branch predication 225. Front-end 210 is not limited to only including the logic shown, but may also include other logic, such as external data interface 265 and a low-level instruction cache. Front-end 210 fetches and decodes instructions to be executed by integrated circuit 205. As shown, front-end 210 also includes branch prediction logic 225 to predict instructions to be fetched and decoded. Front-end 210 may fetch and decode fixed length instructions, variable length instructions, macro-instructions, or instructions having individual operations.

[0026] An instruction usually includes multiple operations to be performed on data operands and is commonly referred to as a macro-instruction, while the individual operations to be executed are commonly referred to as micro-operations. However, an instruction may also refer to a single operation. Therefore, a micro-operation, as used herein, refers to any single operation to be performed by integrated circuit 205, while an instruction refers to a macro-instruction, a single operation instruction, or both. As an example, an add macro-instruction includes a first micro-operation to read a first data operand from a first associated address, a second micro-operation to read a second data operand from a second associated address, a third micro-operation to add the first and the second data operand to obtain a result, and a fourth micro-operation to store the result in a register location.

[0027] Transactional execution typically includes grouping a plurality of instructions or operations into a transaction or a critical section of code. In one embodiment, hardware in integrated circuit 205 groups macro-operations into transactions. Identifying transactions in hardware includes several factors, such as usage of lock acquire and lock releases, nesting of transactions, mutual exclusion of non-speculative memory operations, and overlay of memory ordering requirements over constructs

used to build transactions. In another embodiment, transactions are demarcated in software. Software demarcation of transactions is discussed in more detail in reference to **Figure 5**.

**[0028]** Integrated circuit 205 further comprises execution units 275 and register file 270 to execute the groups of macro-operations, also referred to as transactions and critical sections. Unlike traditional locking techniques, transactional execution usually entails speculatively executing a transaction/critical section and postponing state updates until the end of speculative execution, when the final status of the transaction is determined. As an example, a critical section is identified by front-end 210, speculatively executed, and then retired by retirement logic 235 only if remote agents, such as another core or logical processor have not made an invalidating request to the memory locations accessed during execution of the critical section.

**[0029]** As illustrative examples, remote agents include memory updating devices, such as another integrated circuit, processing element, core, logical processor, or any processor/device that is not scheduled to execute or is not executing the pending transaction. Typically, invalidating requests comprise requests/accesses by a remote agent to memory locations manipulated by micro-operations within the

transaction, requests to lock a semaphore guarding the memory locations manipulated by micro-operations within the transaction, or requests by a remote agent for ownership of memory locations manipulated by micro-operations within the transaction. Invalidating requests will be discussed in more detail in reference to **Figure 3**.

**[0030]** If at the end of executing the transaction/critical section the results are deemed inconsistent or invalid, then the transaction/critical section is not retired and the state updates are not committed to registers or memory. Additionally, if the transaction is not retired, then two options for re-executing the transaction include: (1) speculatively re-executing the transaction as previously executed or (2) non-speculatively re-executing the transaction utilizing locks/semaphores.

**[0031]** Speculative execution of transactions may include memory updates and register state updates. In one embodiment, integrated circuit 205 is capable of holding and merging speculative memory and register file state updates to ensure transaction execution results are valid and consistent before updating memory and the register file. As an illustrative example, integrated circuit 205 holds all instructions/micro-operations results identified as part of the same transaction in a speculative/temporary state for an arbitrary period of time. To accomplish

the holding and merging of speculative memory and register file state updates, special register checkpoint hardware and operand bypass logic is used to store the speculative results in temporary registers.

[0032] In another embodiment, integrated circuit 205 is capable of decoupling register state updates and instruction retirement from memory updates. In this embodiment, speculative updates are committed to register file 370 before speculation is resolved; however, the memory updates are buffered until after the transaction is retired. Therefore, one potential advantage is each individual instruction or micro-operation within a transaction may be retired immediately after execution. Furthermore, the decoupling of the register state update and the memory update potentially reduces the extra registers for storage of speculative results before committing to architectural register file 270.

[0033] However in this embodiment, speculatively updating register file 270 entails treating each update to register file 270 as a speculative update. Register re-use and allocation policies may account for updates to register file 270 as being speculative updates. As an illustrative example, input registers that are used for buffering data for transactions are biased against receiving new data during the pendency of commitment of the transaction. In this example, input registers used

during the transaction are biased against receiving new data; therefore, if the speculative execution fails or needs to be re-started, the input register set is usually able to be re-used without re-initialization, as other registers that are not part of the input register set would be used first.

[0034] In another example, if input registers receive new data during speculative execution or pendency of commitment of the transaction, the state of the input registers re-used are stored in a separate storage area, such as another register. The storage of the input register's original contents allows the input registers to be reloaded with their original contents in case of an execution failure or initiation of re-execution. The processor temporarily storing a registers contents and then re-loading upon re-execution is typically referred to as spilling and refilling.

[0035] The consistency of memory accesses to a shared memory, such as cache 240, within a transaction/critical section may be tracked to ensure memory locations read from still have the same information and memory locations to be updated/written-to have not been read or updated by another agent. As a first example, a memory access is a load operation that reads/loads data, a data operand, a data line, or any contents of a



memory location. As a second example, a memory access includes a memory update, store, or write operation.

[0036] In one embodiment, transaction buffer 265 tracks accesses to lines of data, such as cache lines 245, 250, and 255, in shared memory, such as cache 240. As an illustrative example, cache lines 245-255 comprise a line of data, an associated physical address, and a tag. The associated physical address references a memory location external to integrated circuit 205 or a memory location located on integrated circuit 205.

[0037] Turning to **Figure 3**, an embodiment of transaction buffer 265 is illustrated. Transaction buffer 265 may include transaction tracking logic to track invalidating requests/accesses by remote agents to each address loaded from and each address to be written to a shared memory within a transaction. As illustrative examples, remote agents include other processing elements, such as another logical processor, core, integrated circuit, processing element, or any processor/device that is not scheduled to execute or is not executing the pending transaction.

[0038] In one embodiment, transaction buffer 265 includes a load table 305 and a store/write buffer 325 to track the loads/reads and the stores/writes, respectively, during execution of a pending transaction. Here, the load table 305 stores a load entry, such as load entry 307, to

correspond to each line of data loaded/read from a shared memory during execution of a pending transaction/critical section. In one embodiment, load entry comprises a representation of a physical address 310 and an invalidating access field (IAF) 315. As first example, representation of physical address 310 includes the actual physical address used to reference the memory location. As a second example, the representation includes a coded version or a portion of the physical address, such as a tag value, to reference the loaded data line, along with length/size information. The length of loaded data may be implicit in the design; therefore, no specific reference to length/size of the data loaded is required. In one embodiment, the implicit length/size of loaded data is a single cache line.

[0039] As an illustrative example, IAF 315 has a first value when load entry 307 is first stored in load table 305 and is changed to a second value when a remote agent makes an invalidating access or invalidating access request to the memory location referenced by physical address 310. For instance, an invalidating request/access constitutes a remote agent writing to the memory location referenced by physical address 310 during execution of the pending critical section, where physical address 310 represents a memory location that was read from during execution of the pending critical section. As a simplified example, IAF 315 is initialized to a

first logical value of 1 upon storing load entry 307, load entry 307 comprising physical address 310, which references a memory location loaded from during execution of a critical section. If a remote agent, writes to the memory location referenced by physical address 310 during execution of the pending critical section, then IAF 315 field is changed to a second value of 0 to represent that a remote agent made an invalidating access to the memory location referenced by load entry 307.

**[0040]** In one embodiment, load table 305 may also be used to track invalidating lock/semaphore requests made by remote agents. When a transaction is executed, a semaphore or separate load entry, such as load entry 307 is used to track a semaphore for the transaction. A semaphore variable may be tracked using a common load operation for the semaphore variable, the load operation being tracked in a similar manner as discussed above. In fact, a semaphore load entry, such as load entry 307, to track invalidating requests to the semaphore comprises physical address field 310 and IAF 315. Physical address field 310 may comprise a representation of a physical address that the semaphore value is stored at.

**[0041]** Analogous to the operation of creating a load entry explained above, IAF 315 is loaded with a first value upon storing semaphore load entry 307 in load table 305 to track a locking

variable/semaphore for the current transaction. If a remote agent requests or acquires a lock with the semaphore, referenced by the physical address 310, during execution of the pending transaction, then IAF 315 is set to a second value to represent that a remote agent requested/obtained a lock on the transaction during execution. It is apparent that multiple agents may track a lock; however, the invalidation is performed when one of the agents acquires an actual lock.

[0042] Load table 305 is not limited to the embodiment shown in **Figure 4**. As an example, transaction buffer 265 determines which load entries, such as load entry 307, are empty (entries not used by the current transaction and may have default or garbage data) and which load entries are full (entries created by the current transaction). Here, a counter may be used to keep track of an allocation pointer that references the current load entry. Alternatively, another field, such as an allocation tracking field (ATF), is present in each load entry to track whether that load entry is empty or full. As an example, load entry 307 has an ATF with a first value, such as a logical 1, to represent an empty load entry that has not been created by the current transaction. The ATF in load entry 307 is changed to a second value, such as a logical 0, when load entry 307 is created by the current transaction.

[0043] In another embodiment, the size/length of the data line loaded/read is not implicit, but rather, another field, such as a length field, is present in load table 305 to establish the length/size of the data loaded. Load table 305 may be an advanced load address table (ALAT) known in the art for tracking speculative loads.

[0044] Referring again to **Figure 3**, store write buffer 325 stores a write entry, such as write entry 327, to correspond to each line of data or partial line of data to be written to/updated within a shared memory during execution of a pending transaction/critical section. For example, write entry 327 comprises a representation of a physical address 330, an invalidating access field (IAF) 335, and a data hold field 340. As a first example, representation of physical address 330 includes the actual physical address used to reference a memory location to be written to at the end or during execution of a pending critical section. As a second example, the representation includes a coded version or a portion of the physical address, such as a tag value, to reference a data line to be written to at the end of execution a pending critical section.

[0045] For the above example, IAF 335 has a first value when write entry 327 is first stored in write table 325 and is changed to a second value when an invalidating access to a memory location reference by physical

address 330 is made by a remote agent. In one embodiment, an invalidating access constitutes a remote agent writing to the memory location referenced by physical address 330 during execution of the pending critical section. Additionally, an invalidating access constitutes a remote agent reading from physical address 330 during execution of the pending critical section. Another invalidating access may constitute a remote agent gaining ownership of the memory location referenced by physical address 330. As a simplified example, IAF 335 is initialized to a first logical value of 1 upon storing write entry 327. If a remote agent reads or writes to the memory location referenced by physical address 330 during execution of the pending critical section, then IAF 325 is changed to a second logical value of 0 to represent that a remote agent has made an invalidating access to the memory location referenced by write entry 327.

**[0046]** Write entry 327 further illustrates data hold field 340 to buffer/hold the speculative data to be written. Data hold field 340 may also be used to track which portion of a tracked line of data contains new data versus which portion has not been targeted by the speculative store. Tracking the changed portions may aid in merging speculative data to actual memory locations later during the commitment process.

[0047] In one embodiment, ownership of a line to be written to, from a store operation, is gained upon execution and retirement of the individual operation within a transaction. As an alternative to pre-fetching ownership, at the retirement of each individual write/store micro-operation, the ownership of the physical address to be written to is not gained until the end of the transaction before transaction retirement. In either embodiment, at the end of the transaction, if ownership was relinquished during execution of the transaction, then the transaction is not retired (fails), because an invalidating access was made. Once the transaction is to be retired, ownership of each line to be written to is not relinquished until after all of the memory updates have been committed. If a remote agent requests ownership of a line during retirement, the request may be queued and held pending until after all of the memory updates/writes have been committed.

[0048] Write table 325 is not limited to what is shown in **Figure 4**. It may, for example, include a pinning field, not depicted, to block snoops from remote agents to a shared memory, such as a cache, when set. The pinning field of a write entry is set to a first value to allow snoops to a corresponding physical address and set to a second value when a cache line is pinned to block snoops to the cache line by remote agents. A

pinning field may be especially useful during the commit process to block snoops and to disallow any ownership changes. As stated above, any requests for ownership from a remote agent may be queued until after the transaction has been committed. One exemplary method to implement the pinning field is to block snoops for a predetermined length of time, when the pinning field is set, wherein the predetermined length of time is based on the number of store buffers present.

**[0049]** Write table 325 may also include a length field, such as the length field discussed in reference to load table 305 above, for storing the length of speculative data to be written. Any amount of other fields or combinations of fields may be included in store table/buffer 325. For instance, a remote agent field is used to track a processor ID or other ID to identify the remote agent that made an invalidating access.

**[0050]** Transaction buffer 265 may be implemented in hardware or firmware. In another instance, transaction buffer 365 is implemented in software and executed by integrated circuit 205. In yet another example, transaction buffer is implemented in microcode.

**[0051]** After executing all the micro-operations within a critical section/transaction, a transaction is typically committed, if no invalidating accesses occurred during execution of a pending critical section. After



retirement, the transaction is typically committed in an atomic manner. As an example, atomically writing/committing a pending critical section includes writing each and every data line buffered during execution of a critical section to a shared memory.

[0052] In one embodiment, a pending transaction is retired by retirement logic 235, shown in **Figure 2**, after checking transaction buffer 265 for invalidating accesses that were tracked during execution of the pending critical section. As an example, for a pending transaction to be retired, each load entry IAF stored in load table 305 and each write entry IAF stored in store table/buffer 325, which is associated with the pending transaction is checked. Additionally, any load entries that were created to track a lock variable or a semaphore for the pending transaction are also checked to ensure no invalidating access was made by a remote agent requesting the lock or the semaphore. If no invalidating accesses are discovered then the transaction retirement is granted and the store buffers are pinned. Once pinned and retirement is granted, which is done simultaneously, the memory updates may be performed in a serial fashion. Once completed, the "pin" status is removed, the line is relinquished, and the transaction is considered committed.

[0053] As a simplified example, a transaction includes a micro-operation to read from location 0001 and write the value 1010 to location 0002. When executing the first micro-operation, load table 305 would store load entry 307 comprising physical address field 310, which represents location 0001, and IAF 315 with a first value 1. When executing the second micro-operation store table 325 would store write entry 327 comprising physical address 330, which represents location 0002, IAF 335 with a first value of 1, and 1010 in data field 340. Additionally, the load and write entries may further comprise size/length information or other fields described above. If a remote agent writes to location 0001 during execution or while the transaction is still pending, then IAF 315 is set to the second value of 0 to represent an invalidating access was made. Upon trying to retire the transaction, IAF 315 represents an invalidating access, so the transaction would not be retired and the value 1010 would not be written to location 0002. However, if no remote agent writes to location 0001 and no remote agents reads/writes to location 0002 as represented by 1's in IAF 315 and 335, then the transaction is retired and the value 1010 is written to location 0002.

[0054] After determining an invalidating access occurred during the pending transaction, therefore, not retiring the transaction, there are a

number of options. The first option includes re-executing the transaction. As discussed above, the input registers are either (1) re-initialized to their original state, if they received new data during pendency of the transaction or (2) are already present in their original state, if they received no new data during pendency of the transaction. Consequently, the transaction is speculatively re-executed in the same manner as before. A second option includes speculatively re-executing the transaction using a back-off algorithm in conjunction with the remote agent that made the invalidating access. As an example, an exponential back-off algorithm is used to attempt to complete the transaction without the remote agent contending for the same data. Another option includes using a software non-blocking mechanism, known in the art, to re-execute the transaction. A fourth option includes re-executing the transaction non-speculatively with locks/semaphores after re-executing the transaction speculatively a predetermined number of times. The semaphores effectively locking the addresses to be read from and written to during the transaction.

[0055] The fourth option, utilizing locks/semaphores as a failure mechanism, may be implemented in hardware, software, or a combination of hardware for executing software. For instance, in software implemented lockout mechanism, a semaphore is used for locking access

to any granularity of memory locations. Each processor that wants to access a certain memory location contends for the semaphore guarding that location. If the semaphore is set to a first value representing no lock, then the first processor flips the semaphore to a second value representing that address/memory location is locked. Flipping the semaphore to the second value ensures through software that the processor, who flipped the semaphore, gets exclusive access to that memory location, and likely a range of memory locations guarded by that semaphore. Integrated circuit 205 may have separate lockout logic 260 to invoke/execute the semaphores in software or may simply use existing execution logic to execute/invoke the software lockouts. The semaphore may be software implemented; therefore, it the semaphore may be present in system memory (not depicted).

**[0056]** As another example of implementing lockout logic 260, shown in **Figure 2**, lockout logic 260 or software executed on lockout logic 260 uses a lockout mechanism for preventing at least one remote agent access to designated lines of a shared memory. In one embodiment, the lockout logic includes a lock bit. As a first example, in hardware, the lock bit is in a register or in the cache line. As a second example, the lock bit is

represented in software that is executed on lockout logic 260 and present in system memory.

[0057] When the lock bit has a first value access to predetermined or designated lines of shared memory is allowed. However, when the lock bit has a second value access to the designated lines of shared memory is prevented. The lock bit may be present in cache 240, in the lockout logic 260, any other memory in processor 205, or system memory. Any granularity of data lines may be locked by a single semaphore or by setting a single bit. As an example, 2<sup>8</sup> lines are locked by the setting of a single locking bit.

[0058] As an example of the use of semaphores as a fail safe mechanism, a transaction is executed a first number of time, such as five times, but during each execution a remote agent makes an invalidating access to an address that was read from during execution of the transaction, such as illustrative address 0001. Looping through the transaction code a sixth time, an execution threshold of six is met. Once the threshold or predetermined number of executions is met, a semaphore is used for executing the transaction.

[0059] In a software implementation, a semaphore guarding address 0001 is contended for. If address 0001 is not currently locked by

the semaphore, then the semaphore is flipped in value to represent that it is currently locked. The transaction is then re-executed non-speculatively.

**[0060]** As an alternative, in a hardware implementation, a locking circuit, such as locking circuit 263, which may consists of a single transistor or any number of transistors, sets a locking bit associated with address 0001 to a second value preventing remote agents access at least to address 0001 during the sixth execution of the transaction.

**[0061]** Locking of data lines is not limited to the use of semaphores or a locking bit, but includes any method or apparatus for preventing access to lines of data, whether implemented in hardware or software. As another example, a tri-state device is used to prevent interconnect access to lines of data.

**[0062]** Turning to **Figure 4**, an example of a transaction demarcated in software is shown. As stated above, a transaction typically includes a group of instructions/micro-operations to be executed. Therefore, a transaction declaration may be any method of demarcating a transaction. In **Figure 4**, transaction 410 has examples of some operations, such as read memory, perform operations, and update/write to memory. Transaction 410 is demarcated by transaction declaration/identifier 405, which is depicted as Atomic {...};. However, a transaction declaration is not so

limited. As a simple example, a pair of brackets grouping a plurality of operations or instructions is a transaction declaration/identifier to identify the bounds of a transaction/critical section.

**[0063]** An instance of transaction declaration 405 compiled is shown in compiled example 415. Transaction 430's bounds are identified by transaction identifier 425; therefore, a processor executing the transaction is able to identify the micro-operations that make up a transaction/critical section from the identifier. Another instance of transaction declaration 405 compiled is shown in compiled example 425. In this instance, transaction declaration 435 identifies the bounds of transaction 440.

**[0064]** To step through this example, lines 1 through 3 identify transactional execution, sets predicates Px to 1 and Py to 0, initializes a count variable to 0 in Rm, and the threshold of the count in Rn. Predicates typically include one type or path of execution when the predicate has one value and another type or path of execution when the predicate has another value. In lines 4-9, the count variable is initialized to a number representing the amount of times the transaction is to be executed speculatively, the count variable is then compared to a threshold or otherwise evaluated to see if the locking predicate should be set to execute the transaction with locks/semaphores (non-speculatively), the count

variable is decremented, or incremented depending on the design, to represent the amount of times the transaction has been executed, and the transaction is started. Lines 10 through 12 include any amount of operations within a critical section in transaction 440. Finally, line 14 includes a check instruction for probing the transaction tracking logic/buffer, discussed above, for invalidating accesses made by a remote agent during the execution of the transaction.

**[0065]** Turning to **Figure 5**, an embodiment of a system using transactional execution is shown. Microprocessors 505 and 510 are illustrated, however, the system may have any number of physical microprocessors, each physical microprocessor having any number of cores or any number of logical processors utilizing transactional execution. As an example, microprocessors 505 and 510 each have a plurality of cores present on their die, each core having a plurality of threads resulting in multi-threading cores. In one embodiment, micro-processor 505 and 510 are capable of out-of-order speculative and non-speculative execution. In another embodiment, microprocessor 505 and 510 are capable of only in-order execution.

**[0066]** Microprocessors 505 and 510 have caches 507 and 512. In one embodiment, caches 507 and 512 store recently fetched data and/or



instructions from system memory 530. In this embodiment, cache 507 and cache 512 would cache data private to their respective microprocessors. Memory 530 may be a shared memory that transactional execution is used to access. In another embodiment, any memory present in the system accessed during a transaction is a shared memory. For example, if microprocessors 505 and 510 accessed a higher level shared cache, not depicted in **Figure 5**.

**[0067]** Microprocessors 505 and 510 are shown coupled to memory controller 520 by interconnect 515. Memory controller is coupled to graphics device 540 by interconnects 535, respectively. In one embodiment, graphics device 540 is integrated in memory controller 520. Memory controller is also coupled to system memory 530 by interconnect 525. System memory 530 may be any type of access memory used in a system. In one embodiment, system memory 530 is a random access memory (RAM) device such as a static random access memory (SRAM), a dynamic random access memory (DRAM), a single data rate (SDR) RAM, a double data rate (DDR) RAM, any other multiple data rate RAM, or any other type of access memory.

**[0068]** Input/Output (I/O) controller 550 is coupled to memory controller 545 through interconnect 545. I/O controller 550 is coupled to

storage 560, network interface 565, and I/O devices 570 by interconnect 555. In one embodiment, storage 560 is a hard-drive. In another embodiment storage 560 is a disk drive. In yet another embodiment, storage 560 is any static storage device in the system. In one embodiment, network interface 565 interfaces with a local area network (LAN). In another embodiment, network interface 565 interfaces with a larger network, such as the internet. Input/output devices 570 may include any user input or system related output devices, such as a keyboard, mouse, monitor, or printer.

[0069] Referring next to **Figure 6**, an embodiment of a flow diagram for a method of executing a transaction is illustrated. In block 605, during execution of a first transaction, invalidating accesses to a plurality of lines in a shared memory referenced by the first transaction are tracked.

[0070] In one example, a transaction buffer is used to track the invalidating accesses. The transaction buffer includes a load table and a store table/buffer. The load table tracking invalidating accesses to addresses loaded from during execution of the first transaction. Invalidating accesses to addresses/memory locations loaded from include a remote agent, such as a processor, core, thread, or logical processor, not

scheduled to execute the first transaction, writing to an address or memory location loaded from during execution of the first transaction. Additionally, the load table may include a lockout mechanism entry to track invalidating accesses to a semaphore or other lockout mechanism during execution of the transaction. In this example, an invalidating access to the lockout mechanism includes a remote agent requesting or obtaining a lock on an address guarded/locked by the lockout mechanism.

[0071] The store table/buffer working similarly to the load table tracks invalidating accesses to addresses or memory locations that are to be written to upon commitment of the transaction. An invalidating access here may include a remote agent either reading from or writing to the aforementioned addresses or memory locations.

[0072] In block 610, the first transaction is re-executed a first number of times, if invalidating accesses are tracked. Therefore, if an invalidating access is tracked during execution of the first transaction, the first transaction is merely re-executed. However, if the first transaction has been re-executed a predetermined number of times, which may be represented by a count variable in software or logic within a processor, the plurality of lines in shared memory referenced by the first transaction are locked. Locking may occur through a software implemented lockout

mechanism, such as a semaphore, which locks out or gives exclusive access to one processor the plurality of lines. Locking may also occur through hardware utilizing lockout logic to physically lockout access to the plurality of lines referenced by the first transaction.

[0073] In block 620, the transaction is re-executed again, after access to the plurality of lines has been locked. Therefore, the processor, which may be a core or a logical processor that was re-executing the transaction speculatively, but failing to commit the results because invalidating accesses were tracked, would have exclusive access to the plurality of lines referenced by the first transaction. Consequently, the first transaction may be executed non-speculatively, since exclusive access is available to the executing processor.

[0074] Turning now to Figure 7, an embodiment of the code flow for transactional execution is shown. In block 705, a group of micro-operations, which when grouped together may span multiple instructions or macro-operations, are executed. As above, in block 710, invalidating accesses to shared memory locations associated with each load and store micro-operation are tracked.

[0075] In block 715, the execution of the first group of micro-operations is looped through until (1) no invalidating accesses are tracked

or (2) the first group of micro-operations have been executed a first number of times. Therefore, instead of having to jump to a new location in the code, the same input register set may be used and the transaction simply looped through again. As stated above, this is accomplished by biasing the input register set from receiving new data during the pendency of the transaction, as well as spilling and refilling an input register's contents upon re-use of the input register. On again in block 720, the shared memory locations associated with each load and each store micro-operation are locked and the first group of micro-operations are re-executed.

**[0076]** Transactional execution as described above avoids the false contention that potentially occurs in locking architectures and limits contention to actual contention by tracking invalidating accesses to memory locations during execution of a transaction. Furthermore, if the transaction is re-executed a predetermined number of times, because actual contention continues to occur, then the transaction is non-speculatively executed utilizing locks/semaphores to ensure the transaction is executed and committed after trying to speculatively execute the transaction the predetermined number of times. Alternatively, a software non-blocking mechanism might be employed instead of a non-

speculative execution method. As noted above, speculative register state updates/commits can be supported in software by ensuring that the “live-in” data of the transaction is preserved, either in the original input registers, or by copying the input data values to a save location, which may be either other registers or memory, from which they can be restored if the transaction must be retried. A processor may also contain hardware mechanisms to buffer the register state, possibly using a mechanism typically used to support out-of-order execution

[0077] In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative sense rather than a restrictive sense.

## CLAIMS

What is claimed is:

1. An apparatus comprising:  
a shared memory to be shared by a first agent and a remote agent;  
execution logic to execute a transaction, the transaction comprising a plurality of instructions;  
transaction tracking logic to track invalidating accesses made by the remote agent to each address loaded from and each address to be written to the shared memory during execution of the plurality of instructions; and  
transaction retirement logic to (1) retire the transaction, if an invalidating access to each address loaded from and each address to be written to the shared memory has not been tracked by the transaction tracking logic during execution of the transaction, and (2) initiate a re-execution of the transaction, if an invalidating access to any address loaded from or any address to be written to the shared memory has been tracked by the transaction tracking logic during execution of the transaction.
2. The microprocessor of claim 1, further comprising a lockout mechanism to deny the remote agent access to each address loaded from and to be written to the shared memory during execution of the transaction, if the transaction is re-executed a first number of times without retiring the transaction.
3. The microprocessor of claim 2, wherein the lockout mechanism comprises a lockout circuit operable to set a lockout bit to deny at least one remote agent

access to each address loaded from and to be written to the shared memory during execution of the transaction, if the transaction is re-executed a first number of times without retiring the transaction.

4. The microprocessor of claim 2, wherein the lockout mechanism comprises logic operable to execute code to invoke a semaphore to deny at least one remote agent access to each address loaded from and to be written to the shared memory during execution of the transaction, if the transaction is re-executed a first number of times without retiring the transaction.

5. The microprocessor of claim 1, wherein the transaction tracking logic comprises logic operable to store a load table to track each address loaded from the shared memory and a write buffer to track each address to be written to the shared memory during execution of the plurality of macro-operations.

6. The microprocessor of claim 1, wherein logic operable to store a load table includes an Advanced Load Address Table (ALAT).

7. The microprocessor of claim 5, wherein the load table is operable to store a load entry for each address loaded from the shared memory, each load entry comprising a representation of the address loaded from the shared memory and an invalidating access field, and wherein the write buffer is operable to store a write entry for each address to be written to the shared memory, each write entry comprising the address to be written to, a data line to write, and an invalidating access field.



8. The microprocessor of claim 5, wherein the load table further comprises a lock mechanism load entry, the lock mechanism load entry to track remote agent accesses to the software implemented lockout mechanism
9. The microprocessor of claim 1, wherein an invalidating access comprises (1) the remote agent writing to an address loaded from the shared memory during execution of the plurality of instructions or (2) the remote agent reading from or writing to an address to be written to the shared memory during execution of the plurality of micro-operations.
10. The microprocessor of claim 9, wherein the remote agent is selected from a group consisting of a core on an integrated circuit including the agent, a thread on an integrated circuit including the agent, a logical processor on an integrated circuit including the agent, a physical processor, a processor coupled to an integrated circuit including the agent.
11. The microprocessor of claim 1, wherein shared memory is a cache, and wherein the agent and remote agents are logical processors that share the cache.
12. A system comprising:  
software demarcating a transaction with a transaction declaration, the transaction comprising a critical section with a plurality of micro-operations to be executed, and the transaction declaration comprising an identifier to identify the bounds of the transaction, a count variable to represent the number of times the critical section has been executed, and a check instruction;

a first microprocessor to execute the transaction, wherein the first microprocessor comprises, logic to store a load tracking table for tracking invalidating accesses to addresses associated with load micro-operations within the plurality of micro-operations, logic to store a write-tracking table for tracking invalidating accesses to addresses associated with store micro-operations within the plurality of micro-operations, check logic to execute the check instruction for probing the load and store tracking tables for invalidating accesses, retirement logic to (1) retire the transaction if execution of the check instruction returns no invalidating accesses and (2) initiate re-execution of the transaction and change the count variable, if execution of the check instruction returns at least one invalidating access.

13. The system of claim 12, wherein the transaction declaration further comprises a locking predicate, when set, to execute the transaction using a lockout mechanism, and wherein the locking predicate is set, if the count variable represents the transaction has been re-executed a predetermined number of times.

14. The system of claim 12, further comprising a storage medium coupled to the first microprocessor for storing the software, a system memory for storing lines of data, and a cache in the first microprocessor for storing recently accessed lines of data from the system memory.

15. The system of claim 14, wherein invalidating accesses to addresses associated with load micro-operations comprise a first remote agent writing to an address loaded from the cache during execution of the transaction, and wherein

invalidating accesses to addresses associated with store micro-operations comprise a second remote agent reading or writing to an address to be written to the cache during execution of the transaction.

16. The system of claim 15, wherein the first microprocessor further comprises a plurality of cores, each core having a plurality of logical processors, and wherein the first and second remote agents are any one of the plurality of cores or plurality of logical processors that are not scheduled to execute the transaction.

17. A method comprising:

tracking invalidating accesses to a plurality of lines in a shared memory referenced by a first transaction during speculative execution of the first transaction;

speculatively re-executing the first transaction each time an invalidating access to the plurality lines in the shared memory is tracked during execution of the first transaction;

locking out access to the plurality of lines in the shared memory referenced by the first transaction after a first number of times speculatively re-executing the first transaction; and

non-speculatively re-executing the first transaction after locking out access to the plurality of lines in the shared memory.

18. The method of claim 17, wherein an invalidating access to the plurality of lines in the shared memory comprises (1) a remote agent writing to one of the plurality of lines in the shared memory that was loaded during speculative execution of the first transaction or (2) a remote agent writing to or reading from

one of the plurality of lines in the shared memory that is to be written to upon commitment of the first transaction.

19. The method of claim 17, wherein tracking invalidating accesses to lines in a shared memory comprises: storing a load entry in a load table for each line in the shared memory loaded during execution of the first transaction, each load entry comprising a representation of an address associated with the line loaded and an invalidating access field to (1) store a first value, upon storing the load entry in the load table to represent that no invalidating access has occurred during execution of the first transaction and (2) store a second value, if an invalidating access occurred during execution of the first transaction.

20. The method of claim 19, wherein tracking invalidating accesses to lines in a shared memory further comprises: storing a write entry in a write table for each line in the shared memory that is to be written to at the end of executing the first transaction, each write entry comprising a representation of a physical address associated with the line to be written to, a data field, and an invalidating access field to (1) store a first value, upon storing the load entry in the load table to represent that no invalidating access has occurred during execution of the first transaction and (2) store a second value, if an invalidating access occurred during execution of the first transaction.

21. The method of claim 20, wherein each write entry and each load entry further comprises a length field for storing the length of the line loaded or the line to be written.

22. The method of claim 20, wherein the length of each line loaded and each line to be written to is implicit in the design of the processor.

23. The method of claim 17, further comprising biasing input registers used during execution of the first transaction from receiving new data.

24. The method of claim 23, further comprising spilling a first input register's contents to a second register, if the first input register is re-used during execution of the first transaction.

25. The method of claim 24, further comprising refilling the first input register with the contents stored in the second register upon speculatively re-executing the transaction.




FIG. 1

2/7

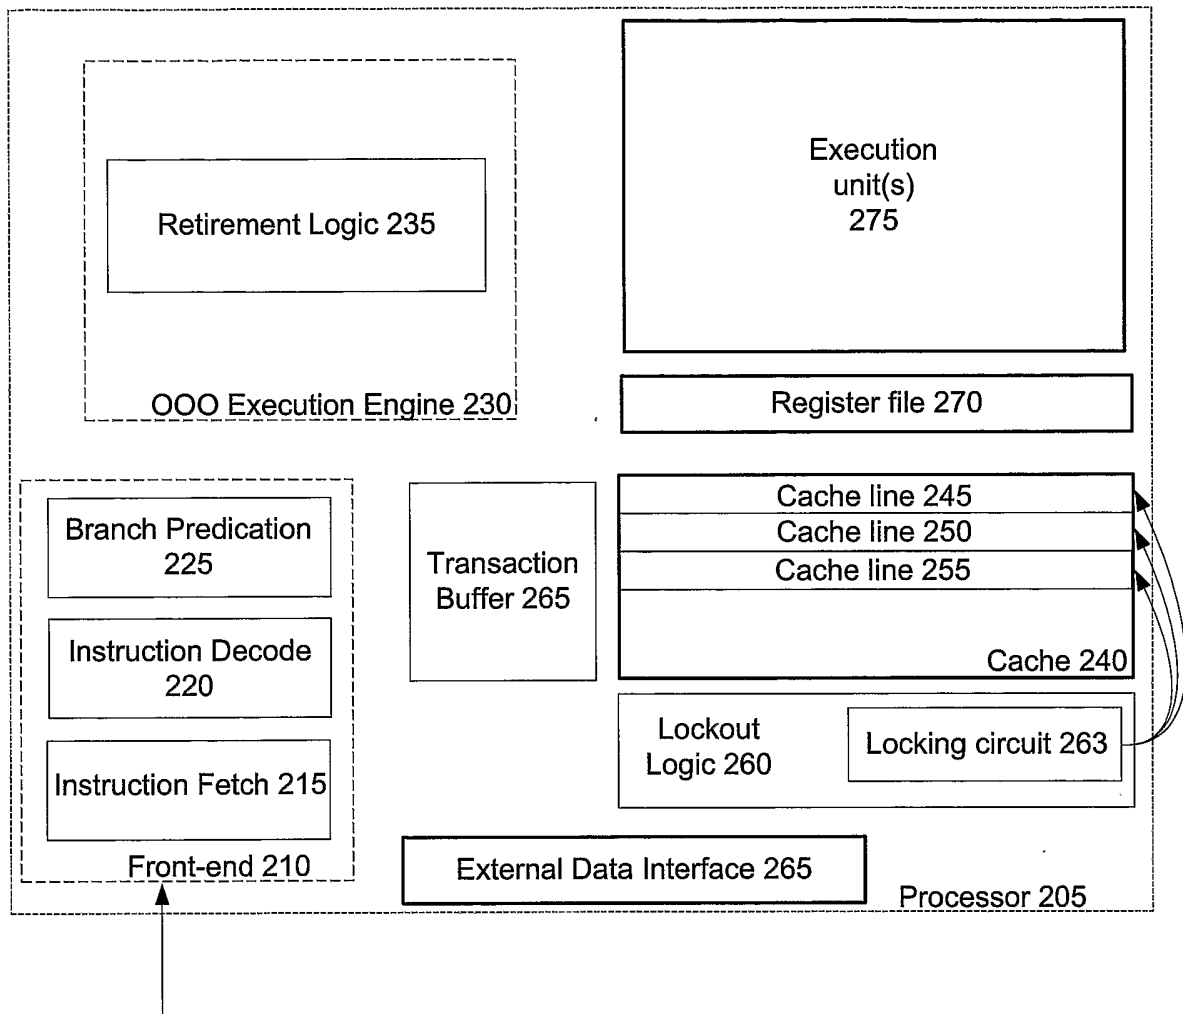


FIG. 2




FIG. 3



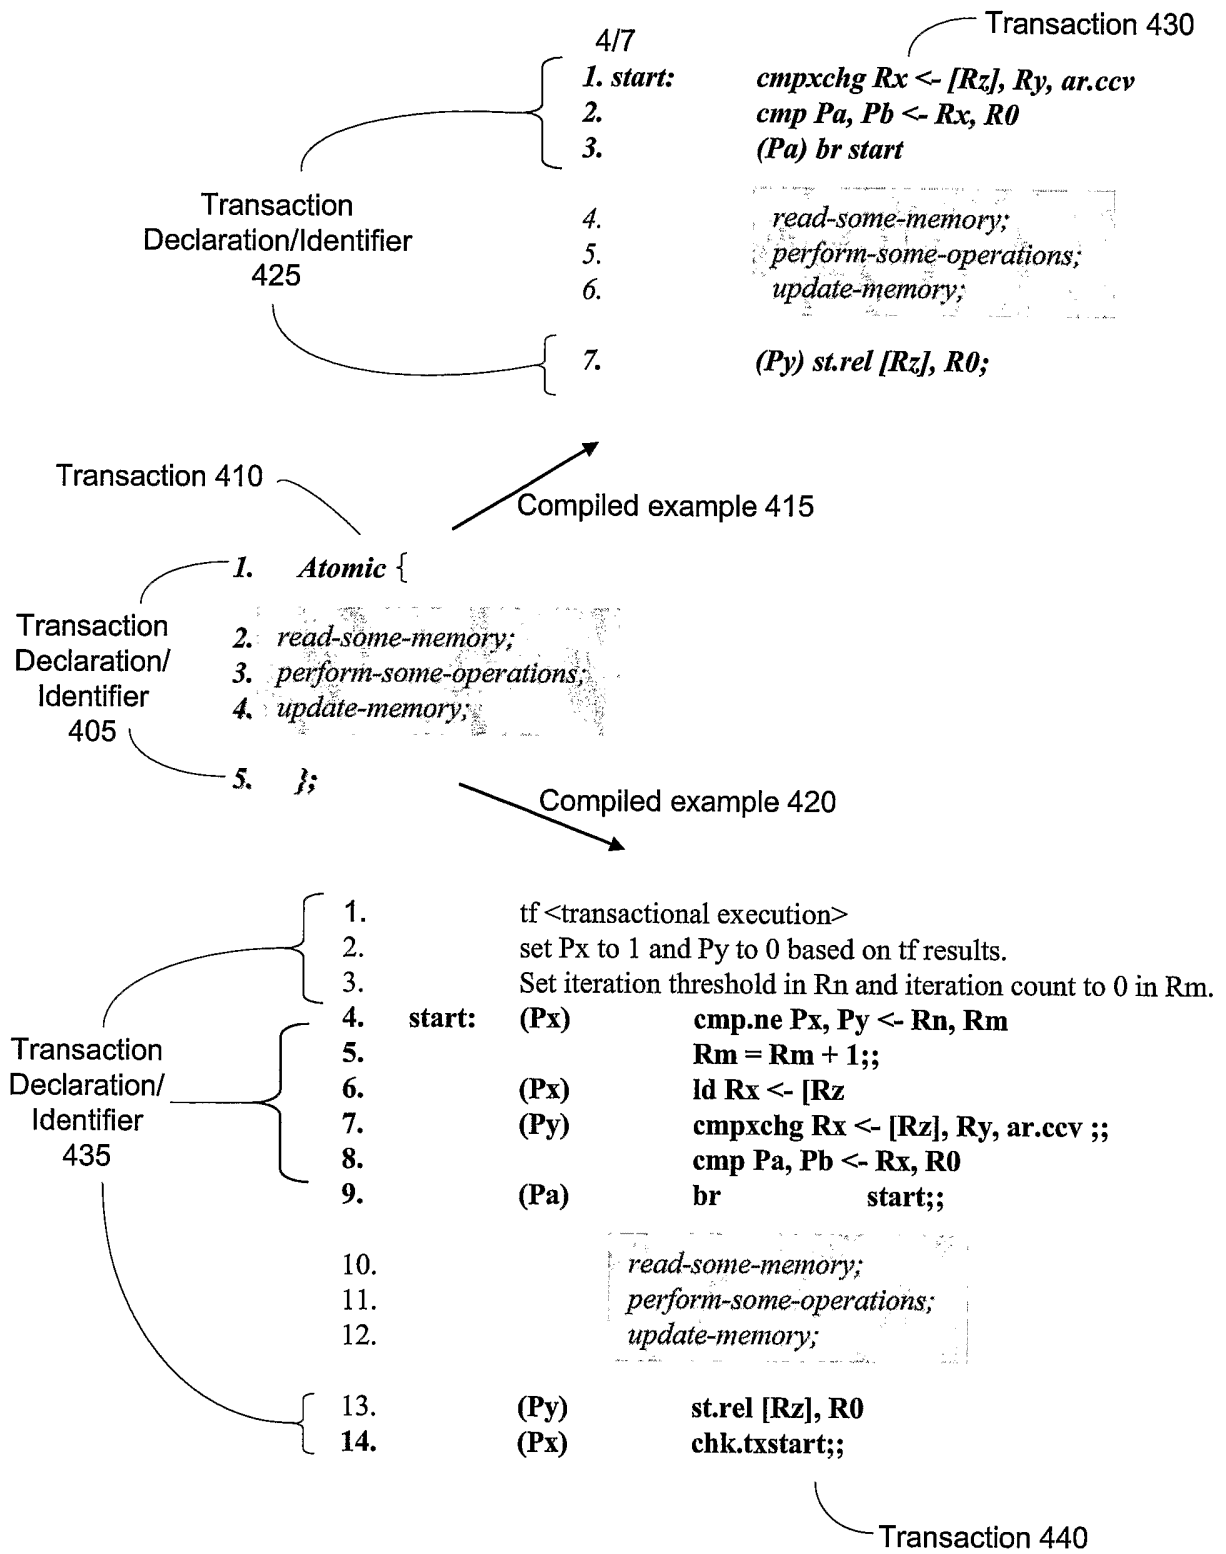


FIG. 4

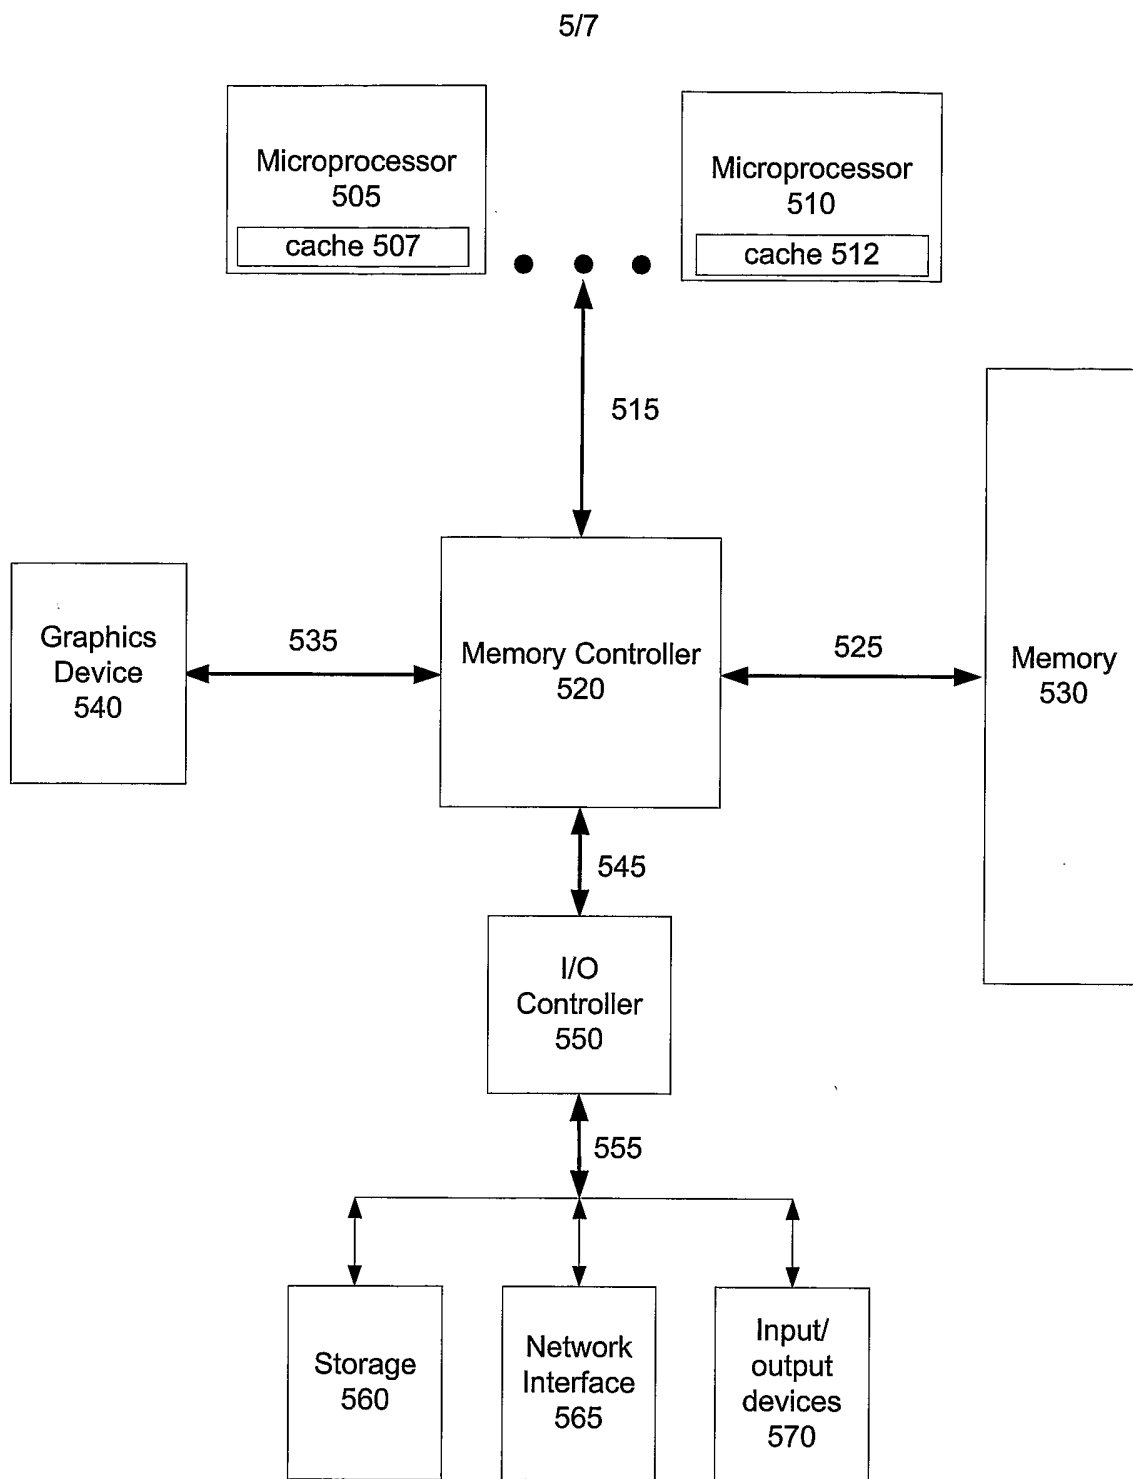


FIG. 5

6/7

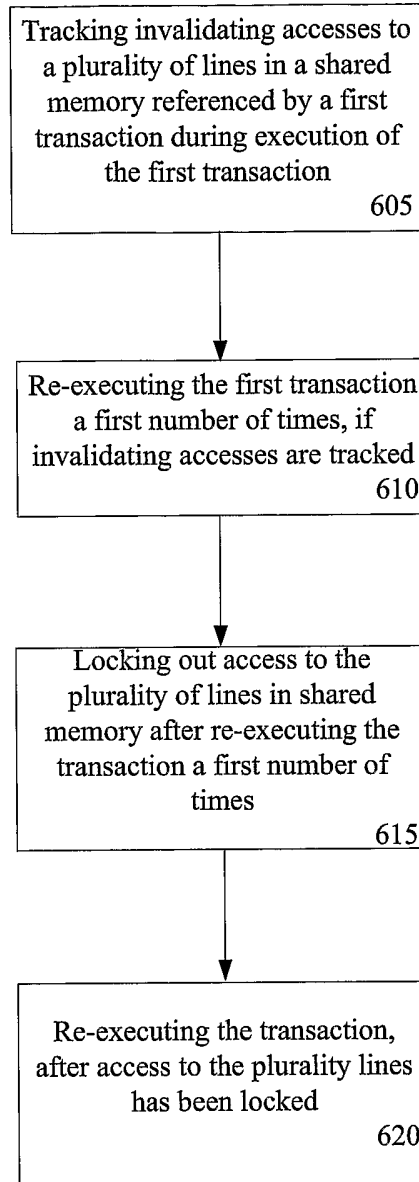


FIG. 6

7/7

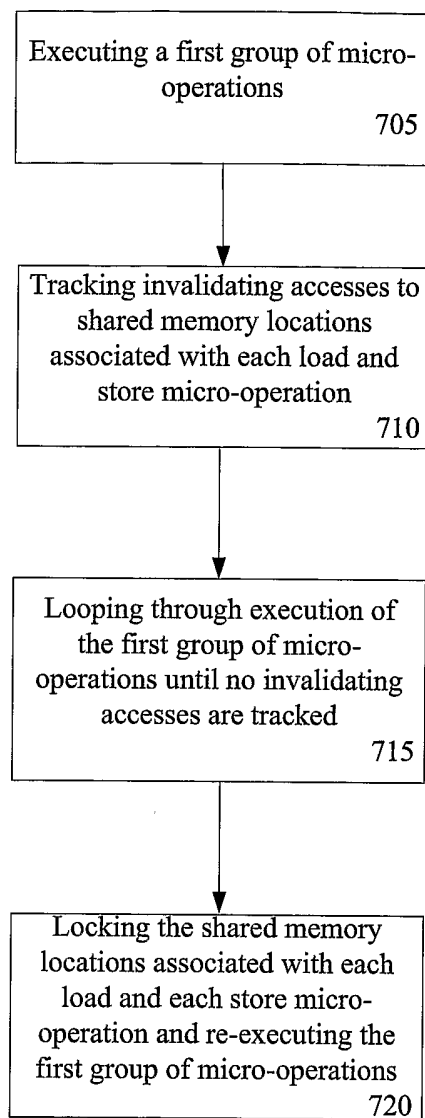


FIG. 7

**INTERNATIONAL SEARCH REPORT**

International application No  
**PCT/US2005/047376**

**A. CLASSIFICATION OF SUBJECT MATTER**  
INV. G06F9/52 G06F9/46

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 practical, search terms used)

**EPO-Internal, PAJ, IBM-TDB, INSPEC, WPI Data**

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2004/075044 A (SUN MICROSYSTEMS INC; TREMBLAY, MARC; JACOBSON, QUINN, A; CHAUDHRY, SH) 2 September 2004 (2004-09-02) abstract paragraph [0014] - paragraph [0033] paragraph [0048] - paragraph [0097]; claims 1-29	1-25
X	WO 2004/075045 A (SUN MICROSYSTEMS INC; TREMBLAY, MARC; JACOBSON, QUINN, A; CHAUDHRY, SH) 2 September 2004 (2004-09-02) abstract paragraph [0014] - paragraph [0033] paragraph [0047] - paragraph [0097]; claims 1-27	1-25
	----- -/--	

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 document 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

**13 April 2006**

Date of mailing of the international search report

**25/04/2006**

Name and mailing address of the ISA/  
European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

**Wierzejewski, P**

INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2005/047376

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 809 503 A (AOSHIMA ET AL) 15 September 1998 (1998-09-15) the whole document -----	1-25
A	OPLINGER J ET AL: "Enhancing software reliability with speculative threads" ACM SIGPLAN NOTICES, ACM, ASSOCIATION FOR COMPUTING MACHINERY, NEW YORK, NY, US, vol. 37, no. 10, October 2002 (2002-10), pages 184-196, XP002285168 ISSN: 0362-1340 the whole document -----	1-25
A	HERLIHY M ET AL: "Transactional Memory: Architectural Support For Lock-free Data Structures" PROCEEDINGS OF THE ANNUAL INTERNATIONAL SYMPOSIUM ON COMPUTER ARCHITECTURE. SAN DIEGO, MAY 16 - 19, 1993, LOS ALAMITOS, IEEE. COMP. SOC. PRESS, US, vol. SYMP. 20, 16 May 1993 (1993-05-16), pages 289-300, XP010095799 ISBN: 0-8186-3810-9 the whole document -----	1-25
A	SHAVIT N ET AL: "SOFTWARE TRANSACTIONAL MEMORY" PROCEEDINGS OF THE ANNUAL ACM SYMPOSIUM ON PRINCIPLES OF DISTRIBUTED COMPUTING. OTTAWA, AUG. 20 - 23, 1995, PROCEEDINGS OF THE ANNUAL ACM SYMPOSIUM ON PRINCIPLES OF DISTRIBUTED COMPUTING.(PODC), NEW YORK, ACM, US, vol. SYMP. 14, 20 August 1995 (1995-08-20), pages 204-213, XP000643146 ISBN: 0-89791-710-3 the whole document -----	1-25

**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No

PCT/US2005/047376

Patent document cited in search report		Publication date		Patent family member(s)	Publication date
WO 2004075044	A	02-09-2004	US	2004187115 A1	23-09-2004
WO 2004075045	A	02-09-2004	US	2004187116 A1	23-09-2004
US 5809503	A	15-09-1998	NONE		