

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
12 July 2007 (12.07.2007)


PCT

(10) International Publication Number  
**WO 2007/078891 A1**

- (51) International Patent Classification:  
*G06F 11/14* (2006.01)
  - (21) International Application Number:  
PCT/US2006/048074
  - (22) International Filing Date:  
15 December 2006 (15.12.2006)
  - (25) Filing Language: English
  - (26) Publication Language: English
  - (30) Priority Data:  
11/323,092 30 December 2005 (30.12.2005) 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): **PETERSEN, Leaf** [US/US]; 101 E. San Fernando Street #540, San Jose, CA 95112 (US). **SAHA, Bratin** [IN/US]; 4301 Norwalk Drive U 206, San Jose, CA 95129 (US). **ADL-TABATABAI, Ali-Reza** [US/US]; 2125 Quinn Avenue, Santa Clara, CA 95051 (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, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
  - (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, 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, 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).
- Published:**  
— with international search report

[Continued on next page]

(54) Title: SOFTWARE ASSISTED NESTED HARDWARE TRANSACTIONS



(57) Abstract: A method and apparatus for efficiently executing nested transactions is herein described. Hardware support for execution of transactions is provided. Additionally, through the use of logging previous values immediately before a current nested transaction in a local memory and storage of a stack of handlers associated with a hierarchy of transactions, nested transactions are potentially efficiently executed. Upon a failure, abort, or invalidating event/access within a nested transaction, the state of variables or memory locations written to during execution of the nested transaction are rolled-back to immediately before the nested transaction, instead of all the way back to an original state of the variables or memory locations before an enclosing transaction. As a result, nested transactions may be re-executed within enclosing transactions, without flattening the enclosing and nested transactions to re-execute everything.

WO 2007/078891 A1



- 
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

## SOFTWARE ASSISTED NESTED HARDWARE TRANSACTIONS

### FIELD

This invention relates to the field of processor execution and, in particular, to providing for efficient execution of nested transactions.

### BACKGROUND

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.

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 at least some dedicated execution resources. 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

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 has created problems with synchronizing data shared among the software threads. 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. However, the ever increasing ability to execute multiple

software threads potentially results in false contention and a serialization of execution.

Another solution is using transactional execution to access shared memory to execute instructions and operate on data. Often transactional execution includes speculatively executing a grouping of a plurality of micro-operations, operations, or instructions. During speculative execution of a transaction by a processor, core, or thread, the memory locations read from and written to are tracked to see if another processor, core, or thread accesses those locations. If another thread does alter those locations, the transaction is restarted and it is re-executed from the beginning. Currently, values of memory locations to be changed in a transaction are saved elsewhere, so if the transaction needs to be re-executed the original state of all memory/registers may be restored.

However, as transactional execution has progressed, software programmers have begun to use nested transactions, i.e. a grouping of instructions/operations to be executed within and part of another outer/enclosing group of instructions/operations. As a consequence, current hardware for support for nested transactions has resulted in inefficient execution of nested transactions.

For example, assume an outer transaction and an inner transaction nested within the outer transaction is to be executed. Current hardware support typically saves values of memory locations to be changed before entering the outer transaction. Yet, when executing in the inner transaction, if an abort or invalidating event occurs, the state of memory locations is usually rolled-back to the original state of the memory locations requiring a re-execution of both the outer and the inner transaction. This simple example is magnified where more nested transactions exist within each other. Specifically, if an abort occurs within a nested transaction deep in a hierarchy of transactions, numerous nested transactions that were not associated with the abort would have to be re-executed for no reason.

## **BRIEF DESCRIPTION OF THE DRAWINGS**

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

Figure 1 illustrates an embodiment of a processor to efficiently execute nested transactions.

Figure 2 illustrates an embodiment of a logging module at different points in pseudo-code including nested transactions.

Figure 3 illustrates another embodiment of a processor to efficiently execute nested transactions.

Figure 4 illustrates an embodiment of a flow diagram for executing a current level nested transaction within a previous level transaction.

Figure 5 illustrates an embodiment of a flow diagram for executing a nested transaction.

#### **DETAILED DESCRIPTION**

In the following description, numerous specific details are set forth such as examples of specific hardware support for transactional execution, specific types of local memory in processors, and specific types of memory locations, etc. 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 as coding of transactions in software, demarcation of transactions in hardware, architectures of multi-core and multi-threaded processors, and specific operational details of microprocessors, have not been described in detail in order to avoid unnecessarily obscuring the present invention.

The method and apparatus described herein are for providing software assistance for nested hardware transactions in processors. However, the methods and apparatus for providing software assistance for nested hardware transactions are not so limited, as they may be implemented on or in association with any integrated circuit device or system, such as cell phones, personal digital assistants, embedded controllers, mobile platforms, desktop platforms, and server platforms.

#### **PROCESSORS**

Referring to **Figure 1**, an embodiment of processing element **100**, which is capable of transactional execution, is illustrated. Processing element **100** may be any element for executing instructions or operating on data. Examples of processing

element *100* include a processor, microprocessor, multi-resource host processor, a microprocessor, a processing core, a logical processor, and an embedded processor, a multi-threaded processor, and a multi-core processor.

In one of the examples, reference to a multi-resource processor is made. Often a resource is referred to as a processor core, logical processor, or threaded processor. Consequently, a multi-resource processor includes a processor with multiple cores, logical processors, threads, or any combination thereof. 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. In contrast, a logical processor typically refers to 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

Processor *100* may include any combination of cores or threads, such as a multi-core processor where each core supports execution of multiple software threads. Note that processor *100* is capable of individual execution within a system or may be combined with other processing elements in a multiple physical processor system. In one embodiment, to support speculative execution of transactions, processor *100* is capable of speculative execution. Other potential execution capabilities of processor *100* include in-order execution, out-of-order execution, serial execution, parallel execution, fixed point execution, floating-point execution, or other well-known types of execution. Specific examples of execution logic and resources are discussed below in reference to the execution resources section.

#### TRANSACTIONS/INVALIDATING EVENTS

Transactional execution usually includes grouping a plurality of instructions or operations into a transaction or a critical section of code. In some cases, use of the word instruction refers to a macro-instruction which is made up of a plurality of micro-operations. There are commonly two ways to identify transactions. The first example, includes demarcating the transaction in software. Here, some software demarcation is included in code to be identified during execution. In another embodiment, which may be implemented in conjunction with the foregoing software

demarcation, transactions are grouped by hardware or recognized by instructions indicating a beginning of a transaction and an end of a transaction. Begin transaction and end transaction instructions are discussed in more detail in reference to **Figures 4 and 5** below.

In a processor, a transaction is either executed speculatively or non-speculatively. In the second case, a grouping of instructions is executed with some form of lock or guaranteed valid access to shared memory locations to be accessed. In the alternative, speculative execution of a transaction is more common, where a transaction is speculatively executed and only committed upon verification that no invalidating events occurred during execution of the transaction.

An invalidating event refers to an event generated by the system, a user, or another resource, not associated with the execution of the current transaction, that requires re-execution of the transaction or invalidation of the transaction. For example, assume processor *100* is a two-core processor. The first core is executing a transaction accessing shared memory, which includes a read from the shared memory and a write-back to the shared memory. During speculative execution of a current transaction by the first core, if the second core writes to the shared memory location that the first core read from an invalidating event has occurred. The current transaction is invalid, because a shared memory location read from was written to by the second core; therefore, the location read from is potentially bad/outdated data.

Another example includes the write to a shared memory location, where the second core reads from the same memory location. Here, the write is potentially invalid as it invalidates the read the second core just completed. In these examples, the second core was “another resource” not associated with the transaction, such as not assigned to execute the transaction or not executing the transaction, that made an invalidating access.

However, other examples of invalidating events include system generated interrupts, user-generated interrupts, an access by another resource to a memory location to be read from or written to during pendency of a current transaction, or any other event that cause a manipulation of a shared memory location that a current transaction is reading from or writing to. A current transaction is considered

“pending” when it is being executed, i.e. a begin transaction instruction has been executed, but an end transaction has not been executed and the data from the transaction has not been committed or retired. In addition, another resource or other resource is not limited to a core, but may also include a logical processor, a hardware thread, a software thread, as well as another remote agent/resource.

### EXECUTION RESOURCES

Figure 1 also illustrates execution resources *105*, which are to execute transactions. Execution resources *105* may also refer to hardware, logic, or modules to support transactional execution. As an example, execution resources *105* are to execute a first transaction and a second transaction nested in the first transaction. A transaction is nested within another transaction, when either in software or hardware, a begin transaction demarcation for an inner transaction is within a transaction demarcation for an outer transaction.

Quickly turning to **Figure 2**, an embodiment of pseudo code nested transactions is illustrated. A `begin_transaction` call identifies either a software construct or a hardware execution of a begin transaction instruction, while an `end_transaction` call identifies also in software or hardware an end transaction instruction. Here first `begin_transaction` instruction *212* and first `end_transaction` instruction *232* illustrates the bounds of a first transaction. Second `begin_transaction` instruction *217* and second `end_transaction` instruction *227* outline the bounds a second transaction that is included in the bounds of the first transaction. In other words, the second transaction is nested within the first transaction. **Figure 2** is discussed in more detail below in the logging module section.

Note that a reference to a first and second transaction herein is arbitrary, in that, a first transaction references only a starting point within a hierarchy of nested transactions to view the interplay of operations between an outer transaction and inner transaction nested within an outer transaction. Reference to a first and second transaction does not necessarily refer to a numerical first and second transaction, but rather, may refer to the order/nesting of transactions, i.e. the second within the first.



In fact, the first transaction may be a nested transaction, as well as the 100<sup>th</sup> transaction executed in an execution flow. For example, a first transaction may refer to a nested transaction nested three levels down from a top-level transaction, and the second transaction may refer to the fourth level transaction nested within the first transaction.

Referring back to **Figure 1**, in one embodiment execution resource **105** includes logic to execute a begin transaction instruction at the start of a transaction, as well as execute an end transaction instruction at the end of a transaction, such as the instructions shown in **Figure 2**. Execution resources **105** may also include logic to retire transactions or instructions associated with transactions, as well as logic to buffer writes and commit writes to memory. Other well-known transactional execution resources and logic may also be included. A few examples include logic to detect data conflicts, to buffer speculative updates, to make updates atomically visible at transaction commit, to execute instructions, and to operate on data.

Other common modules, logic, and functional units not illustrated in **Figure 1** may also be included, but are not required to be included, in processor **100**, such as any one or any combination of the following: 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, a fetch unit, a decode unit, a cache, an instruction retirement unit, an instruction re-order unit, and any other logic that is be used for fetching, executing, or operating on instructions and/or data.

#### MEMORY

Memory **125** includes any memory location to be accessed by execution resources **105**, processing resources, such as cores, logical processors, or threads. In one embodiment, memory **125** is a shared memory shared by at least one other processing resource, such as a core, thread, logical processor, or remote agent. Examples of memory **125** include a cache, a plurality of registers, a register file, a static random access memory (SRAM), a plurality of latches, or other storage

element. Note that processor *100* or any processing resources on processor *100* may be addressing a system memory location, a virtual memory address, a physical address, or other address when reading from or writing to a memory location. Memory *125* will be discussed in more detail in reference to the exemplary embodiments below, when discussing memory location to be written to or read from, states of memory locations, or rolling-back/restoring states/values of memory locations.

As a specific illustrative example, assume that memory *125* is a cache memory, such as a trace cache, a first-level cache, a second-level cache, or a higher-level cache. Cache *125* includes cache lines *126*, *127*, and *128*, which may also be referred to as memory locations within memory *125*. As another example, assume memory *125* is a plurality of registers used by a processing element or resource as execution space or scratch pad to store variables, instructions, or data. Memory locations in grouping of registers *125* are registers *126*, *127*, and *128*, in this example.

#### MONITOR MODULE

In **Figure 1**, monitoring module *110* is illustrated coupled to memory *125*; however, monitor module *110* may be wholly or partially implemented in memory *125*, as well as implemented elsewhere in processor *100*. A module may be implemented in hardware, software, firmware, or any combination thereof. Commonly, module boundaries vary and functions are implemented together, as well as separately in different embodiments. As an example, which is discussed in more detail below, a cache coherency protocol is used to monitor/track invalidating accesses to memory *125*, where bits in the cache are used to represent invalidating accesses and other hardware, software, or firmware determines that an invalidating event has occurred based on the value of those bits. Here, monitoring module varies across boundaries, as parts are implemented in memory *125*, as well as in other hardware or firmware in processor *100*, as well as software executed on processor *100*.

Monitoring module *110* is to detect an invalidating event during execution of a transaction. In one embodiment, monitor module *110* is part of the hardware support for executing transactions, as discussed above. Here, monitor module *110* is

implemented in hardware to detect/monitor invalidating accesses to locations memory *125* to be changed/written to during execution of a transaction. As stated above an invalidating event refers to an event generated by the system, a user, or another resource, not associated with the execution of the current transaction, that requires re-execution of the transaction or invalidation of the transaction. For example, assume a current transaction writes to memory locations *126* and *127*. A write to memory location *126* or *127* by another resource is monitored/detected by monitoring module *110* as an invalidating access.

As a first example, monitoring module *110* includes a cache coherency protocol to detect invalidating accesses to memory *125*. In one embodiment of a cache coherency protocol there is a bit/flag for every line of the cache/memory location, such as lines *126 – 128*, to represent if an invalidating access occurred to a corresponding memory location during execution of a transaction.

To simplify the example, assume that a transaction is executing and is to write to line *127* of cache *125*. If a predetermined flag/bit in cache line *127* or a bit/flag associated with cache line *127* represents a first logical value, such as a high logical value, then there has been no invalidating access. However, upon an access to cache line *127* that would invalidate the current transaction, the flag/bit is flipped to represent that an invalidating access has occurred. Essentially, monitoring module *110* has detected an invalidating access and represented that in cache *125*. Detecting of an invalidating access may include any known-method, such as comparing a resource ID, such as a thread ID or core ID, with that of the ID executing the current transaction.

Note that specific reference to a high logical value for one implementation was used in the example above; however, conversely, a low logic level may represent that there was no invalidating access. Often, the use of logic levels or logical values is also referred to as 1's and 0's, which simply represents binary logic states. For example, a 1 refers to a high logic level and 0 refers to a low logic level. In many older systems a high voltage level was represented by a voltage, e.g. 5V, and a low voltage level, e.g. 0V. As another specific example, a high logic level is at 1.2V and a low logic level is at .3V. However, a high logic/voltage level may refer to any voltage

level above a threshold value, and inversely, a low logic level may refer to any voltage level below the threshold value. In addition, there may be more than two logical levels in a cell or waveform. As an example, a single waveform may represent four different logical values at different voltage levels.

Other methods of using hardware or software to monitor/track invalidating accesses to a memory, such as memory *125* may be used. One example of tracking invalidating accesses using at least in part tracking logic, such as an advanced load address table (ALAT), is discussed in, "Transaction based shared data operations in a multiprocessor environment," with serial no. 11/027,623. Other standard hardware support may also be used. Two examples of standard hardware support for transactions are discussed in the following articles: (1) "Transactional Memory: Architectural support for lock-free data structures," by Maurice Herlihy and J. Eliot B. Mossand in "Proceedings of the 20th Annual International Symposium on Computer Architecture," May 1993, pp. 289—300; (2) "Virtualizing Transactional Memory," by Ravi Rajwar and Maurice Herlihy and Konrad Lai in "Proceedings of the 32nd Annual International Symposium on Computer Architecture," June 2005, pp. 494--505.

#### LOGGING MODULE

Also illustrated in **Figure 1** is logging module *115* and rollback module *120*, which may be implemented together, separately, or vary across each other's boundaries. As stated above, a module may be implemented in hardware, software, firmware, or any combination thereof. Commonly, module boundaries vary and functions are implemented together, as well as separately in different embodiments. Logging module is to take a snapshot of values in memory, record states of memory locations, store handlers, store stacks of values in memory, store stacks of abort handlers, and/or perform other functions to enable roll-back/restore of memory locations to values immediately preceding a nested transaction.

Turning to **Figure 2**, an embodiment of a log at different points in pseudo-code is illustrated. **Figure 2** illustrates pseudo code on the left, and at checkpoints through the pseudo code, such as checkpoints *210*, *215*, *220*, *225*, *230*, and *235*, log *200*'s operation is illustrated. At the highest level, instruction *205* is shown, i.e. X =

0, which represents a store, write, or other representation of setting a memory location X or a memory location referenced by address X to a simple numerical value, 0. In addition, further calls to variable, address, or memory location X operate in a similar manner. As stated above, X may be any memory location such as a cache location, register, variable space, system memory, etc.

First, it is worth noting that a first transaction, which is a top-level transaction as it is not nested within another transaction, begins at *begi\_transaction\_1* instruction **212** and ends at *end\_transaction\_1* instruction **232**. Consequently, a second transaction, which is a nested transaction being within the bounds of the first transaction, starts at *begin\_transaction\_2* **217** and ends at *end\_transaction\_2* **227**. Finally, a third transaction, which is nested within the second transaction, starts at *begin\_transaction\_3* **222** instruction and ends at *end\_transaction\_3* instruction **224**. Labeling of begin and end of transactions is shown as separate instructions with different reference numerals to illustrate that matching of the beginning and ending of transactions. However, *begin\_transaction\_1* instruction **212**, *begin\_transaction\_2* instruction **217**, *begin\_transaction\_3* instruction **232** may be the same instruction to be executed at three different periods of time during execution. The same applies to the *end\_transaction* instructions.

Starting at line **205** memory location X gets a value 0. At checkpoint **210**, before entering the first transaction, there are two options to begin logging. In a first option, where all addresses written to/changed are enumerated, i.e. known, then log **200** stores a first set of values/takes a snapshot of memory locations, such as X, to be changed in the first transaction. Here, before entering the first transaction or at the start of the transaction before executing a first instruction in the transaction, the current value/snapshot of memory location X is taken and stored in log **200**, as represented by entry **211** including  $X = 0$ .

As a second option, where not all addresses to be written to are known, executing of the first transaction begins, and during the transaction but immediately before changing/writing to a memory location, the snapshot or storing of a state of the memory location to be changed/written to is then stored in entry **211**.

Therefore, upon execution of the first transaction, if an abort, failure, or invalidating access occurs, the memory location X that was changed is able to be rolled-back/restored to its original state immediately before the first transaction from entry 211 in log 200. However, assuming execution continues to checkpoint 215 memory location X now has a value of 1 from operation/execution of instruction 213. Now, in one embodiment a logging module, which includes log 200, is to take a second-level snapshot of a second-level value in a memory location to be changed during execution of the second transaction.

To illustrate, assume that the program flow is at checkpoint 215, before entering the second transaction, which is at a second-level and nested within the first transaction. At this point there the same two options for logging exist: (1) log before any locations are written to or (2) log as execution goes along immediately before any writes. Here, the second-level value is the value written to X at line/instruction 213, as it is the value at the same level as the second transaction. The second transaction is at a second-level as it is nested within another transaction, which is deemed the first-level. As a result, a second-level snapshot of a second-level value, i.e. 1 from line 213, in a memory location to be changed during execution of the second transaction, such as X which is changed/written to in line 218. Here, the second level-snapshot includes the storing of the second-level value, 1, in log 200, which is seen at entry 216.

It is worth noting at this point, that log 200 may be implemented in any manner to track values, such as regular storage in a stack configuration. In this example, the second-level value of 1 is stacked on the first-level value of 0. To illustrate this functionality, assume an abort occurs at line 218. Instead of previous transactional execution that would flatten nested transactions and restore X to its first-level original state of 0, the state of X is able to be rolled-back to just checkpoint 215 where X has a value of 1, from the assistance of log 200. Additionally, if execution had to be rolled back again to checkpoint 210 and the value 1 had been popped off a stack configuration to roll-back previously, then the state of X may further be rolled-back to 0 at checkpoint 210 based on log 200.

It should also be apparent that reference to a first transaction and a second transaction nested within a first transaction does not necessarily refer to a top-level transaction, such as the first transaction in **Figure 2**, and the second transaction does not necessarily refer to the first nested transaction, as in **Figure 2**. For example, a first transaction may be the second transaction in **Figure 2**, which is a nested transaction already, while the second transaction nested within the first transaction may be the third transaction that is nested in the second transaction of **Figure 2**. Basically, the first and second transaction may be viewed at any level in a hierarchy of nested transactions.

Continuing the example, from above, assume that execution has continued to checkpoint **220**. At this point, the third transaction is a current-level transaction or an inner transaction, while the second transaction is the outer transaction, enclosing transaction, or previous-level transaction. Once again, X is to be changed in the inner-level transaction, so a state of the previous-level or second-level memory location to be changed during execution of the transaction is stored in log **200**. Here, X has a value of 2, so the value 2 is pushed onto a stack in entry **221** including 0, 1, and 2. Similarly, if an abort occurs at this point, memory location X is restored to the previous-level value of 2 and the third transaction may be re-executed without re-executing both the first and second transactions.

Upon end-transaction\_3 instruction **224** when the third transaction is to be ended, retired, and/or committed a write log is merged into the write log of the enclosing or previous level transaction. In the embodiment above, values were appended or popped onto the entries store in log **200**. As a result upon executing an end\_transaction instruction, the nested transaction is to be committed, so the log pops, merges, or removes the value appended for the current-level transaction that is being ended. Therefore, upon an abort at this point the second transaction is able to be re-executed and the original state of the memory location X is able to be restored from the log to the value before entering the second transaction of 1. This embodiment is illustrated in entry **226**.

However, any other implementation of storing the values may be used. In fact, an implementation is tailored to an abort handler that rolls-back or restores the values, as discussed below. Similar merging operations are done at checkpoints **230** and **235**, as illustrated in entries **231** and **236**, respectively.

Returning to **Figure 1** briefly, roll-back module **120** is illustrated and is to roll-back the memory location to the second-level snapshot, upon the monitor module detecting an invalidating event during execution of the second transaction. In the example, above it is apparent, that roll-back module **120** restores values of memory locations to be changed during execution of a nested transaction to their states immediately before executing the nested transaction. Examples of roll-back module **120** are discussed below in reference to default handlers and abort handlers.

Turning over to **Figure 3**, an embodiment of a processor to efficiently execute nested transactions is illustrated. Similar to **Figure 1**, processor **300** includes execution logic **305**, memory **315** including memory locations **316**, and monitoring logic **310**. In one embodiment, processor **300** is a multi-resource processor, such as a multi-core, multi-threaded, or multi-logical processor with monitoring logic **310** including cache coherency logic. However, any of the foregoing examples of a processor and its contents may be included in processor **300**.

In addition, processor **300** includes local memory **320**. Local memory **320** includes any memory associated with processor **320**. In one embodiment, local memory consists of memory that is only accessible by one processor or one processing resource; however, local memory is not so limited. In another embodiment local memory includes a globally shared memory that is partitioned among processors or resources. In this case, a thread Id, core Id, or other way of identifying a resource is used to allow access to a section of memory, and inversely, to deny access to other resources not assigned to the section of memory. Specific examples of local memory **320** include a Random Access Memory (RAM), a Static RAM (SRAM), a Dynamic RAM (DRAM), a plurality of registers, a plurality of latches, a cache, or any storage element in processor **300**.

In one embodiment, log **325** is illustrated stored in local memory **320**. As an example, log **325** operates in a similar manner to log **200** illustrated in **Figure 2**. In



this case, notice that log 325 illustrates the state of log 200 at entry 220 in a stack format. Assuming that operation is at checkpoint 220 in **Figure 2**, the third nested transaction is the current-level transaction and the second transaction is the previous-level transaction. Therefore, upon an invalidating event occurring during execution of the current-level transaction, the memory location to be changed during execution of the current-level transaction, i.e. memory location X, is restored to the previous-level value/state of 2, which is on the top of the stack in log 325 at entry 326.

Moreover, handler module 330, which may also be referred to as a roll-back module, is shown stored in local memory 320, but as illustrated in Figure 1 handler module 330 may be stored in other memory or implemented in other logic/modules of processor 300. Handler module 330 is to restore the memory locations to be written to during execution of a second transaction nested within a first transaction to a second set of values or a plurality of previous-level values logged before writing to the memory locations to be written to, if an invalidating access occurs during execution of the second transaction.

A specific simplified embodiment of handler module 330 is discussed to illustrate its potential operation. Upon an invalidating access/event occurring during execution of a nested transaction, a default handler is initially invoked/executed. A default handler, although not specifically depicted, may be stored in local memory 320, memory 315, system memory 340, or other memory in processor 300 not depicted. A default handler is to transfer control of processor 300 to an abort handler associated with the current-level transaction, if an invalidating access occurs during execution of the current level-transaction.

For example, assume that upon execution of the first transaction in **Figure 2**, an abort handler associated with the first transaction, such as abort handler 331 in **Figure 3**, is pushed onto a stack of abort handlers, as in handler module 330. In addition, a first set of values of memory locations to be written to during execution of the first transaction are stored in log 325, which in this case includes only memory location X getting a value 0 in entry 328 of log 325. Next, upon entering the second transaction in **Figure 2**, second abort handler 332 associated with the second

transaction is pushed onto handler module **330**'s stack and a second set of values is logged in entry **327**, i.e. second/previous-level value 1, for X.

Now assume, that an invalidating access or event occurs during execution of the second transaction. The default handler, when executed, is to initiate execution and/or control to the top abort handler on the stack, which is abort handler **332** that is associated with the second transaction. Abort handler **332**, when executed by the execution unit, is to restore memory location X to the previous-level, i.e. the second set of values, logged in log **325**. At this point the value 1 is at the top of the stack and is loaded into memory location X by abort handler **332** to allow for re-execution of the second nested transaction.

However, if normal execution were to continue, then upon entering the third transaction in **Figure 2**, abort handler **333** associated with the third transaction is pushed onto the handler stack and previous-level value **326** is pushed onto log **325**'s stack.

#### AN EMBODIMENT OF A SYSTEM

Referring still to **Figure 3**, an embodiment of a system is illustrated. Processor **300** is coupled to system memory **340**. Although not shown, system memory **340** may be coupled to processor **300** through other components or devices, such as a memory controller hub. System memory includes any memory for storage in a system such as a SRAM, DRAM, double data rate (DDR) RAM, non-volatile (NV) RAM, EDO RAM, or other memory device. System memory **340** is to store elements, such as instructions and data operands to be executed by processor **300**. In one embodiment, system memory stores a plurality of instructions that are to be grouped into transactions, such as a first transaction and a second transaction to be nested in the first transaction.

#### AN EMBODIMENT OF A METHOD FOR EXECUTING A NESTED TRANSACTION

Turning to **Figure 4**, an embodiment of a flow diagram for a method of executing a current level transaction nested within a previous level transaction is illustrated. In flow **405**, a current level transaction nested within a previous-level transaction is executed. Here, as above, a begin transaction instruction or

demarcation may establish the starting points of both the current level and previous level transactions. In one embodiment, the transaction is executed with hardware support for determining bounds of transaction, detecting data conflicts, and other hardware support functions mentioned above.

In flow **410** a log including a first state of a memory location at the previous level, which is to be written to during execution of the current-level transaction, is maintained. Although, the flow diagram in **Figure 4**, illustrates an arrowed flow, the flows shown, such as flow **410** and **405** may occur in any order. As stated above, storing to the log may be done either before entering and executing the transaction, as well as during execution of the transaction and immediately before the memory location to be written to is changed, written to, or a write to is buffered. As illustrated, in the foregoing examples, the log may be stored in a local memory.

In one embodiment, local memory also maintains a stack of a first abort handler associated with the current level transaction stacked on a second abort handler associated with the previous level transaction. Similar in operation, to handler module **330** in **Figure 3**, a stack or other general configuration of abort handlers is stored in local memory. In another embodiment, the stack of handlers is stored in a different memory at random locations, not in a stack configuration. In this case, a group of pointers or other association device associates the handlers with their respective transaction.

In decision box **415**, it is determined if an invalidating event occurs during execution of the current level transaction. Examples of invalidating events are discussed above, in the transactions/invalidating events sections. In one embodiment, a cache coherency protocol or other hardware support is used to determine if the invalidating event occurs during execution of the current-level transaction.

In flow **420**, the memory location is restored to the first state of the memory location at the previous level, if an invalidating event is detected during execution of the current level transaction. As an example restoring includes popping the first abort handler off the stack, transferring control to the first abort handler, executing the first abort handler, wherein upon execution, the first abort handler restores the memory location with the first state of the memory location maintained in the log.

Finally, in flow **425**, the log is merged with a second state of the memory location at the current level, if no invalidating event occurs during execution of the current level transaction. In a first embodiment, merging includes appending the second state on the end of the log. In another example, merging may include discarding the last log entry to ensure a potential abort within the previous-level transaction is correctly rolled-back to an original state before its execution. In addition, upon committing or ending the previous-level transaction the log may be discarded completely, if the previous-level transaction is a top-level transaction, which is not nested within another transaction.

ANOTHER EMBODIMENT OF A METHOD FOR EXECUTING A  
NESTED TRANSACTION

Turning finally to **Figure 5**, another embodiment of a flow diagram for a method of executing a nested transaction is illustrated. In flow **505**, an abort handler associated with a current nested transaction is pushed onto a stack of abort handlers. In one embodiment, the stack of abort handlers is stored in a local memory.

In flow **510**, a first value of a first memory location, which is to be changed during execution of the current nested transaction, is logged in a local memory before changing the value of the first memory location. As a first example, logging a first value of the first memory location includes storing/logging the first value of the first memory location in a local memory before executing a first instruction in the current nested transaction. In another embodiment, logging a first value of the first memory locations includes storing/logging the first value of the first memory location in a local memory during execution of the current execution and immediately before the first memory location is changed/written to.

In decision box **515**, it is determined with hardware if an invalidating event occurs during execution of the current nested transaction. In an illustrative example, a functional unit of a microprocessor such as a cache, an advanced load address table (ALAT), and invalidating access tracking unit is at least partially used to determine if an invalidating event occurs during execution of the current nested transaction. A cache coherency protocol or other hardware based protocol may also be used.

In flow **520** and **525**, control is transferred to the abort handler associated with the current nested transaction and the first memory location is restored with the first value, if an invalidating event occurs during execution of the current nested transaction. In one embodiment, control is transferred by execution of a default handler to pop the abort handler associated with the current nested transaction off the stack of handlers to be executed by a processor. Restoring the location includes writing the first value to the first memory location before re-execution of the current transaction. The abort handler associated with the current transaction may accomplish this, when executed.

In decision box, **530** it is determined if an enclosing transaction is a top-level transaction, if no invalidating event occurs during execution of the current nested transaction. Essentially, here, the transaction is ended successfully with no invalidating access or event, such as a write or interrupt. If the transaction is not an enclosing transaction, but rather a nested transaction, a log associated with the current nested transaction including the first value is merged with an enclosing transaction log. In contrast, if the transaction is an enclosing transaction then a log associated with the current nested transaction including the first value is discarded.

As illustrated above, nested transaction may be executed efficiently. Previous nested transaction execution, typically, flattened the nested transaction. As a result, in a nested transaction hierarchy, like the transactions shown in **Figure 2**, if there was a failure, abort, or invalidating access during execution of the third transaction, then both the first and the second transaction would also have to be re-executed. However, by providing hardware support of the execution of transactions, and assistance through modules such as a logging module and handler module, only the failed third transaction is potentially re-executed, which saves precious execution time and allows for efficient execution of nested transactions.

The embodiments of methods, software, firmware or code set forth above may be implemented via instructions or code stored on a machine-accessible or machine readable medium which are executable by a processing element. A machine-accessible/readable medium includes any mechanism that provides (i.e., stores and/or transmits) information in a form readable by a machine, such as a computer or electronic system. For example, a machine-accessible medium includes random-access memory (RAM), such as static RAM (SRAM) or dynamic RAM (DRAM); ROM; magnetic or optical storage medium; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals); etc.

In the foregoing specification, a detailed description has been given with reference to specific exemplary embodiments. 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. Furthermore, the foregoing use of embodiment and other exemplarily language does not necessarily refer to the same embodiment or the same example, but may refer to different and distinct embodiments, as well as potentially the same embodiment.

**CLAIMS**

What is claimed is:

1. An apparatus comprising:  
execution resources to execute first transaction and a second transaction nested in the first transaction;  
a monitoring module to detect an invalidating event during execution of the second transaction;  
a logging module to take a second-level snapshot of a second-level value in a memory location to be changed during execution of the second transaction; and  
a roll-back module to roll-back the memory location to the second-level snapshot, upon the monitor module detecting an invalidating event during execution of the second transaction.
2. The apparatus of claim 1, wherein the execution resources includes logic to execute a begin transaction instruction at the start of the first transaction and the second transaction, and wherein the execution module also includes logic to execute an end transaction at the end of the first and the second transaction.
3. The apparatus of claim 1, wherein the monitoring module includes a cache coherency module to detect an invalidating event during execution of the second transaction.
4. The apparatus of claim 3, wherein the cache coherency module includes a bit associated with memory location, and wherein the bit is flipped to a first logical value upon detecting an invalidating event to represent an invalidating event occurred.
5. The apparatus of claim 1, wherein the logging module is also to store an abort handler stack including a second-level abort handler associated with the



second transaction stacked on a first-level abort handler associated with the first transaction.

6. The apparatus of claim 5, wherein the roll-back module includes a control transfer module to transfer control to the second-level abort handler, if the monitoring module detects an invalidating event during execution of the second transaction.
7. The apparatus of claim 1, wherein the logging module includes a local memory to store the second-level snapshot.
8. The apparatus of claim 7, wherein the local memory is to store the second-level snapshot is stored in a stacked snapshot of the second-level snapshot stacked on a first-level snapshot taken before executing the first transaction.
9. The apparatus of claim 7, wherein the local memory is selected from a group consisting of a Random Access Memory (RAM), a Static RAM (SRAM), a Dynamic RAM (DRAM), a plurality of registers, and a plurality of latches.
10. An apparatus comprising:
  - a processor including
  - execution logic to execute a first transaction and a second transaction nested within the first transaction;
  - monitoring logic to determine if an invalidating access occurs during execution of the first or second transactions;
  - a local memory to store a first set of values from memory locations to be written to during execution of the first transaction and, a second set of values from memory locations to be written to during execution of the second transaction;
  - a handler module to restore the memory locations to be written to during execution of the second transaction to the second set of values, if an invalidating access occurs during execution of the second transaction.

11. The apparatus of claim 10, wherein the processor is selected from a group consisting of a host processor, a microprocessor, a processing core, a logical processor, and an embedded processor, a multi-threaded processor, and a multi-core processor.
12. The apparatus of claim 10, wherein execution logic includes logic to detect data conflicts, to buffer speculative updates, to make updates atomically visible at transaction commit, to execute instructions, and to operate on data.
13. The apparatus of claim 10, wherein a first transaction is demarcated by a first begin transaction instruction and a first end transaction instruction, and wherein the second transaction is demarcated by a second start transaction instruction and a second end transaction instruction between the first begin and end transaction instruction.
14. The apparatus of claim 10, wherein monitoring logic includes a flag for every line of memory in a cache of the processor to represent if an invalidating access occurred to a corresponding memory location during execution of the first or second transactions.
15. The apparatus of claim 14, wherein an invalidating access occurs during execution of the first transaction when an access by another resource to a memory location to be read from or written to occurs during pendency of the first transaction, and wherein an invalidating access occurs during execution of the second transaction when an access by another resource to a memory location to be read from or written to occurs during pendency of the second transaction.
16. The apparatus of claim 10, wherein the another resource is a resource, not executing the transaction, the resource being selected from a group consisting

of a logical processor, a core, a thread, a software thread, a hardware thread, and a remote resource.

17. The apparatus of claim 10, wherein the first set of values from memory locations to be written to during execution of the first transaction are stored in the local memory before entering the first transaction, and wherein the second set of values from memory locations to be written to are stored in the local memory before entering the second transaction.
18. The apparatus of claim 10, wherein the first set of values from memory locations to be written to are stored in the local memory during execution of the first transaction immediately before the memory locations are written to, and wherein the second set of values from memory locations to be written to are stored in the local memory during execution of the second transaction immediately before the memory locations are written to.
19. The apparatus of claim 10, wherein the local memory is also to store a first abort handler associated with the first transaction and a second abort handler associated with the second transaction.
20. The apparatus of claim 19, wherein the handler module includes a default handler to transfer control of the processor to the first abort handler, if an invalidating access occurs during execution of the first transaction, and to transfer control of the processor to the second abort handler, if an invalidating access occurs during execution of the second transaction.
21. The apparatus of claim 20, wherein the second abort handler restores the memory locations to be written to during execution of the second transaction to the second set of values, after control is transferred to the second abort handler, if an invalidating access occurs during execution of the second transaction.

22. The apparatus of claim 21, wherein restoring the memory locations to be written to during execution of the second transaction to the second set of values comprises writing the second set of values, stored in the local memory, back into the memory locations to be written to during execution of the second transaction.
23. A system comprising:  
a multi-resource microprocessor including  
an execution unit to execute an outer transaction and an inner transaction nested within the outer transaction,  
a local memory to store  
a first set of values of a first plurality of locations, which are to be written to during execution of the inner transaction,  
a first abort handler associated with the inner transaction, when executed, to restore the first plurality of locations to the first set of values;  
a default handler, when executed by the execution unit, to initiate execution of the second abort handler, if an invalidating event occurs during execution of the second transaction; and  
a system memory coupled to the multi-resource microprocessor to store a first plurality of instructions to be grouped as the outer instruction and a second plurality of instructions to be grouped as the inner transaction.
24. The system of claim 23, wherein the local memory is also to store  
a second abort handler associated with the outer transaction, and  
a second set of values of a second plurality of locations, which are to be written to during execution of the outer transaction.
25. The system of claim 23, wherein the default handler, when executed, is also to  
initial execution of the second abort handler, if an invalidating event occurs during execution of the outer transaction, and wherein the second abort

- handler, when executed, is to restore the second plurality of locations to the second set of values.
26. The system of claim 23, wherein the multi-resource processor includes a plurality of resources, and wherein the plurality of resources are a plurality of execution threads.
27. The system of claim 26, wherein an invalidating event occurring during execution of the second transaction includes an event initiated by one of the plurality of resources, not associated with the inner transaction, accessing one of the first plurality of locations during pendency of the second transaction.
28. The system of claim 27, wherein a cache coherency protocol is used to determine if an invalidating event occurs during execution of the second transaction.
29. A method comprising:  
executing a current level transaction nested within a previous level transaction; and  
maintaining a log including a first state of a memory location at the previous level,  
which is to be written to during execution of the current level;  
determining if an invalidating event occurs during execution of the current level transaction;  
restoring the memory location to the first state of the memory location at the previous level, if an invalidating event is detected during execution of the current level transaction; and  
merging the log with a second state of the memory location at the current level, if no invalidating event occurs during execution of the current level transaction.
30. The method of claim 29, wherein the log is stored in a local memory in a microprocessor.

31. The method of claim 29, wherein determining if an invalidating event occurs during execution of the current level transaction includes using a cache coherency protocol to determine if the invalidating event occurs during execution of the current level transaction.
32. The method of claim 29, further comprising maintaining a stack of a first abort handler associated with the current level transaction stacked on a second abort handler associated with the previous level transaction.
33. The method of claim 32, wherein restoring the memory location to the first state of the memory location at the previous level comprises:  
popping the first abort handler off the stack;  
transferring control to the first abort handler;  
executing the first abort handler, wherein upon execution, the first abort handler restores the memory location with the first state of the memory location maintained in the log.
34. The method of claim 29, wherein merging the log with a second state of the memory location at a current level comprises: appending the second state on the end of the log.
35. The method of claim 29, further comprising discarding the log upon committing the previous-level transaction, if the previous-level transaction is a top-level transaction.
36. A method comprising:  
pushing an abort handler associated with a current nested transaction onto a stack of abort handlers;  
logging a first value of a first memory location, which is to be changed during execution of the current nested transaction, in a local memory, before changing the value of the first memory location; and

determining with hardware if an invalidating event occurs during execution of the current nested transaction;  
if an invalidating event occurs during execution of the current nested transaction transferring control to the abort handler associated with the current nested transaction,  
and  
restoring the first memory location with the first value.

37. The method of claim 36, wherein the stack of abort handlers is stored in the local memory.
38. The method of claim 36, wherein logging a first value of a first memory location in a local memory before changing the value of the first memory location comprises: logging the first value of the first memory location in a local memory before executing a first instruction in the current nested transaction.
39. The method of claim 36, wherein logging a first value of a first memory location in a local memory before changing the value of the first memory location comprises: logging the first value of the first memory location in a local memory during execution of the current execution and immediately before the first memory location is changed.
40. The method of claim 36, wherein the hardware used to determine if an invalidating event occurs during execution of the current nested transaction includes a functional unit of a microprocessor selected from a group consisting of a cache, an advanced load address table (ALAT), and invalidating access tracking unit.
41. The method of claim 36, further comprising merging a log associated with the current nested transaction including the first value with an enclosing

transaction log, if the current nested transaction is committed and the enclosing transaction is a nested transaction.

42. The method of claim 36, further comprising discarding a log associated with the current nested transaction including the first value, if the current nested transaction is committed and the enclosing transaction is a top-level transaction.






FIG. 1

2/5




FIG. 2

3/5




FIG. 3

4/5

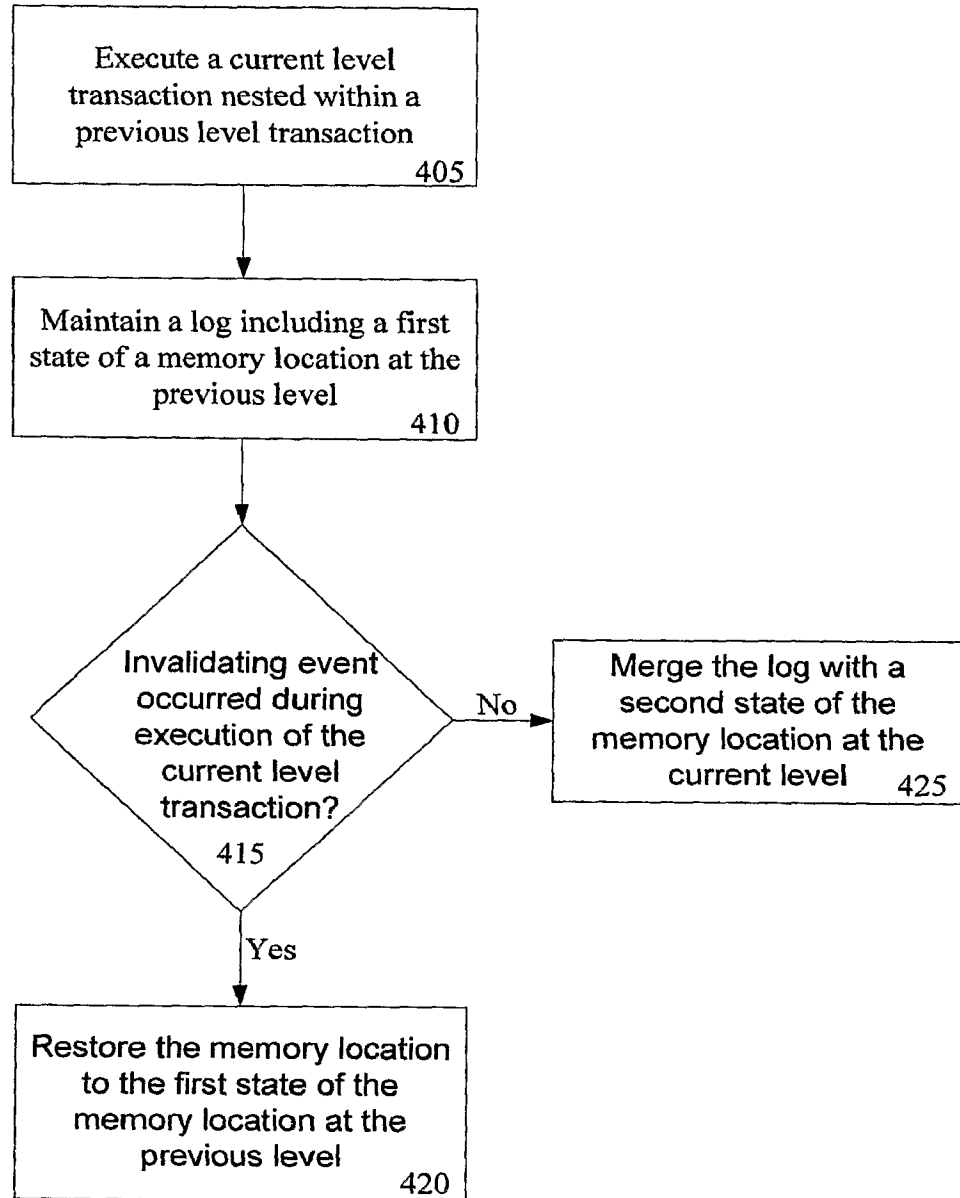


FIG. 4

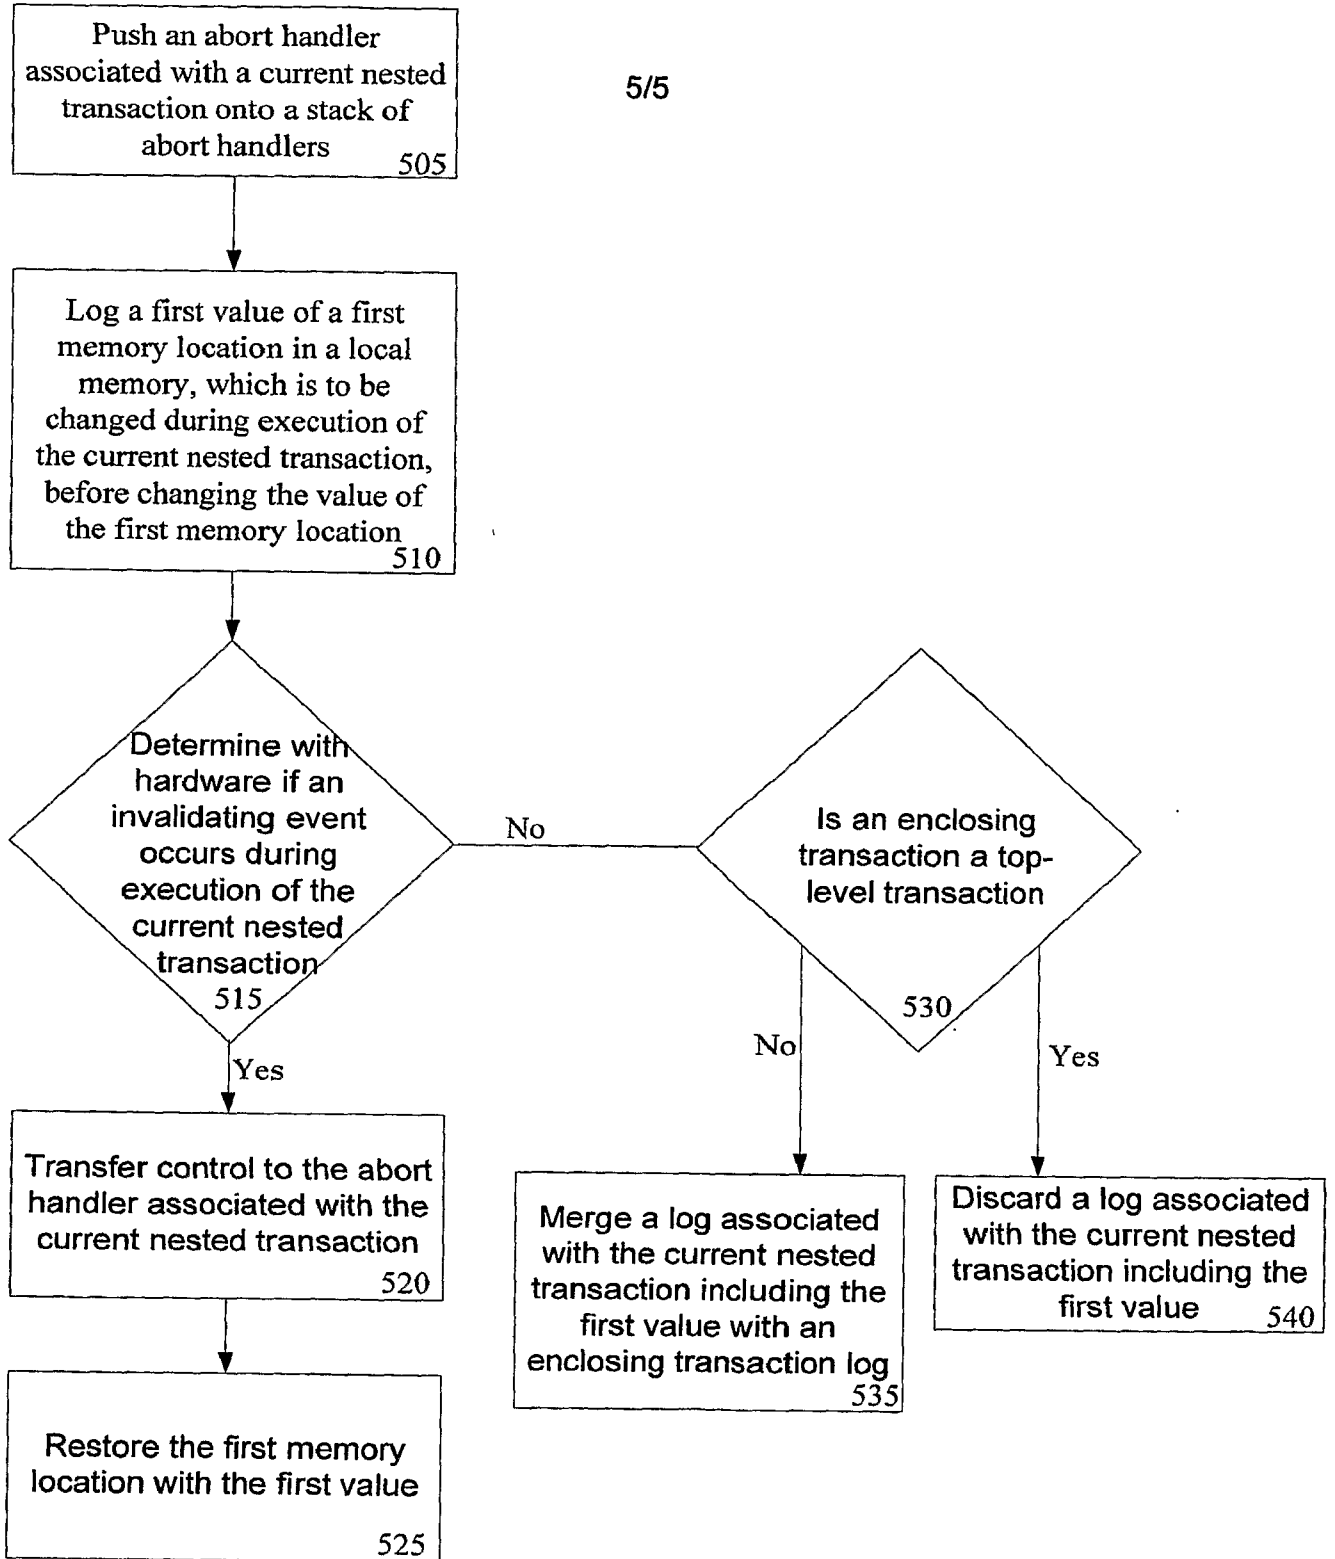


FIG. 5

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2006/048074

<b>A. CLASSIFICATION OF SUBJECT MATTER</b> INV. G06F11/14				
According to International Patent Classification (IPC) or to both national classification and IPC				
<b>B. FIELDS SEARCHED</b>				
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, WPI Data, IBM-TDB, INSPEC				
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>				
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.		
	-/--			
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C. <span style="margin-left: 200px;"><input checked="" type="checkbox"/> See patent family annex.</span>				
* Special categories of cited documents :				
<table style="width: 100%; border: none;"> <tr> <td style="width: 50%; border: none; vertical-align: top;"> <ul style="list-style-type: none"> <li>*A* document defining the general state of the art which is not considered to be of particular relevance</li> <li>*E* earlier document but published on or after the international filing date</li> <li>*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)</li> <li>*O* document referring to an oral disclosure, use, exhibition or other means</li> <li>*P* document published prior to the international filing date but later than the priority date claimed</li> </ul> </td> <td style="width: 50%; border: none; vertical-align: top;"> <ul style="list-style-type: none"> <li>*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</li> <li>*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</li> <li>*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.</li> <li>*Z* document member of the same patent family</li> </ul> </td> </tr> </table>			<ul style="list-style-type: none"> <li>*A* document defining the general state of the art which is not considered to be of particular relevance</li> <li>*E* earlier document but published on or after the international filing date</li> <li>*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)</li> <li>*O* document referring to an oral disclosure, use, exhibition or other means</li> <li>*P* document published prior to the international filing date but later than the priority date claimed</li> </ul>	<ul style="list-style-type: none"> <li>*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</li> <li>*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</li> <li>*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.</li> <li>*Z* document member of the same patent family</li> </ul>
<ul style="list-style-type: none"> <li>*A* document defining the general state of the art which is not considered to be of particular relevance</li> <li>*E* earlier document but published on or after the international filing date</li> <li>*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)</li> <li>*O* document referring to an oral disclosure, use, exhibition or other means</li> <li>*P* document published prior to the international filing date but later than the priority date claimed</li> </ul>	<ul style="list-style-type: none"> <li>*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</li> <li>*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</li> <li>*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.</li> <li>*Z* document member of the same patent family</li> </ul>			
Date of the actual completion of the international search  <div style="text-align: center; font-size: 1.2em;">5 June 2007</div>		Date of mailing of the international search report  <div style="text-align: center; font-size: 1.2em;">13/06/2007</div>		
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  <div style="text-align: center; font-size: 1.2em;">Johansson, Ulf</div>		

INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2006/048074

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X  A	<p>US 6 219 666 B1 (KRISHNASWAMY JANAKI [US] ET AL) 17 April 2001 (2001-04-17)</p> <p>column 2, line 1 - line 30 column 3, line 11 - line 22 column 3, line 44 - column 4, line 6 column 5, line 4 - line 25 figures 2,3a,3b column 7, line 45 - column 8, line 19 column 9, line 1 - line 14 column 9, line 65 - column 10, line 30 column 16, line 14 - column 17, line 23 column 20, line 47 - line 59 column 24, line 26 - line 36 claims 12,13</p> <p>-----</p>	<p>1,5-12, 16-27, 36-39, 41,42 2-4, 13-15, 28-35,40</p>
X  A	<p>US 4 164 017 A (HORNING JAMES J [CA] ET AL) 7 August 1979 (1979-08-07)</p> <p>column 1, line 9 - line 68 column 2, line 25 - column 3, line 32 column 5, line 5 - line 36 column 7, line 35 - column 8, line 17 column 11, line 27 - line 50 column 12, line 25 - line 64 column 17, line 46 - line 60 claims 1-3,6,8 figures 1,10</p> <p>-----</p>	<p>1,5-12, 16-27, 36-39, 41,42 2-4, 13-15, 28-35,40</p>
X  A	<p>WO 01/50264 A (ORACLE CORP [US]) 12 July 2001 (2001-07-12)</p> <p>page 6, line 4 - line 10 page 8, column 4 - page 9, column 3 page 11, line 4 - page 12, line 22 page 14, line 1 - page 17, line 23 figures 1,4,6</p> <p>-----</p>	<p>1,7, 10-12, 16-27, 36-39, 41,42 2-6,8,9, 13-15, 28-35,40</p>

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2006/048074
---

Patent document cited in search report	Publication date	Publication date	Patent family member(s)	Publication date
US 6219666	B1	17-04-2001	NONE	
US 4164017	A	07-08-1979	GB 1509193 A	04-05-1978
WO 0150264	A	12-07-2001	AT 242896 T	15-06-2003
			AU 771514 B2	25-03-2004
			AU 2291701 A	16-07-2001
			CA 2395282 A1	12-07-2001
			DE 60003339 D1	17-07-2003
			DE 60003339 T2	29-04-2004
			EP 1247182 A1	09-10-2002
			HK 1046455 A1	11-09-2003
			US 6574750 B1	03-06-2003