



(19) **United States**

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

(10) **Pub. No.: US 2017/0262227 A1**
(43) **Pub. Date: Sep. 14, 2017**

(54) **HARDWARE TRANSACTION TRANSIENT CONFLICT RESOLUTION**

Publication Classification

(71) Applicant: **International Business Machines Corporation**, Armonk, NY (US)

(51) **Int. Cl.**
G06F 3/06 (2006.01)
(52) **U.S. Cl.**
CPC **G06F 3/0659** (2013.01); **G06F 3/061** (2013.01); **G06F 3/0673** (2013.01)

(72) Inventors: **Jonathan D. Bradbury**, Poughkeepsie, NY (US); **Michael Karl Gschwind**, Chappaqua, NY (US); **Maged M. Michael**, Danbury, CT (US); **Valentina Salapura**, Chappaqua, NY (US); **Eric M. Schwarz**, Gardiner, NY (US); **Chung-Lung K. Shum**, Wappingers Falls, NY (US); **Timothy J. Slegel**, Staatsburg, NY (US)

(57) **ABSTRACT**

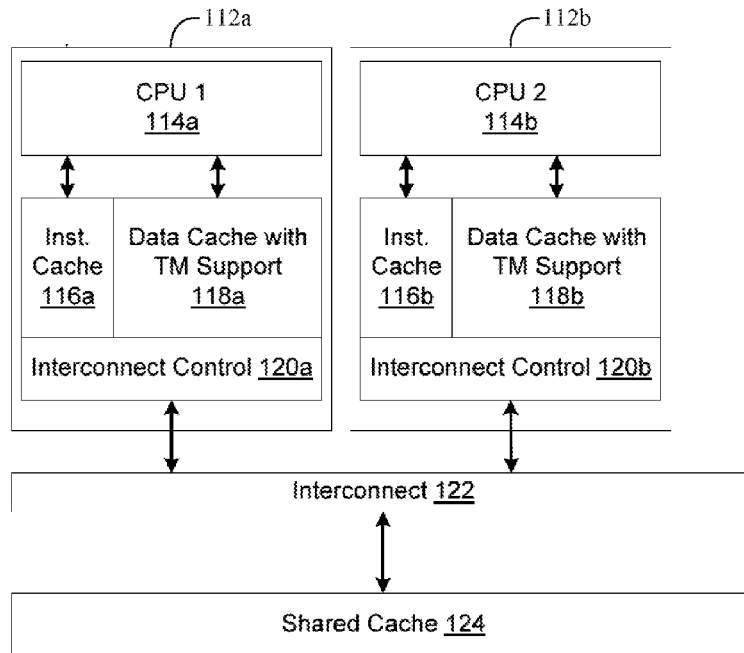
In an approach for resolving terminated transactions in a transactional memory environment, a processor initiates a hardware transaction in a computing environment, wherein the hardware transaction accesses a memory location, and wherein the hardware transaction includes a transaction begin indicator and a transaction end indicator. A processor detects a conflicting access of the memory location while executing the hardware transaction. A processor aborts the hardware transaction based on the conflicting access of the memory location. Hardware determines that the conflicting access of the memory location is a transient condition. A processor reinitiates the hardware transaction.

(21) Appl. No.: **15/180,428**

(22) Filed: **Jun. 13, 2016**

Related U.S. Application Data

(63) Continuation of application No. 15/064,029, filed on Mar. 8, 2016.



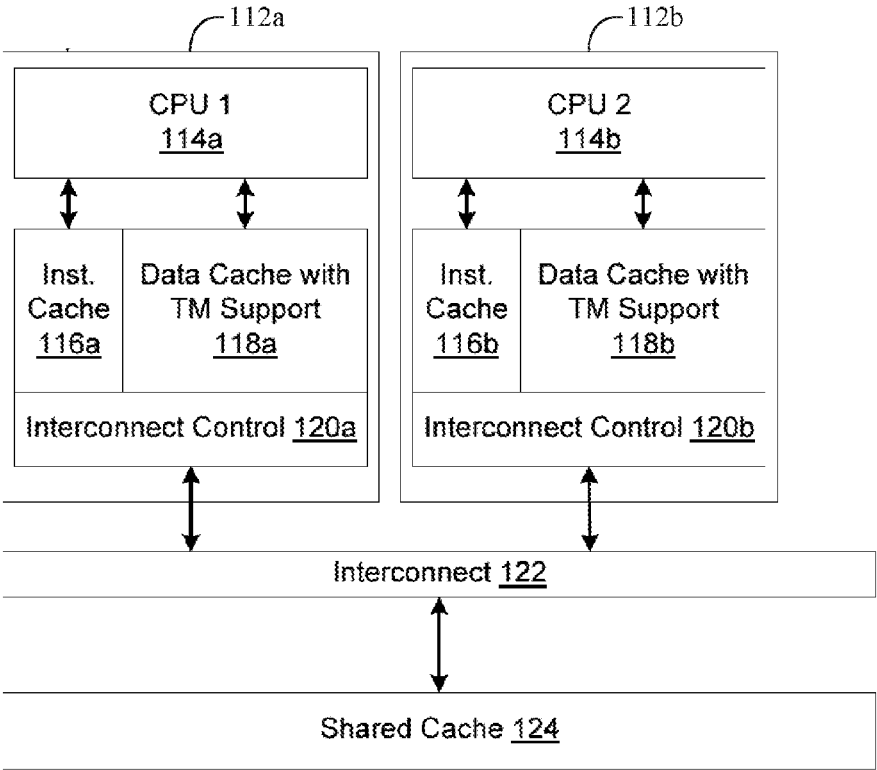



FIG. 1

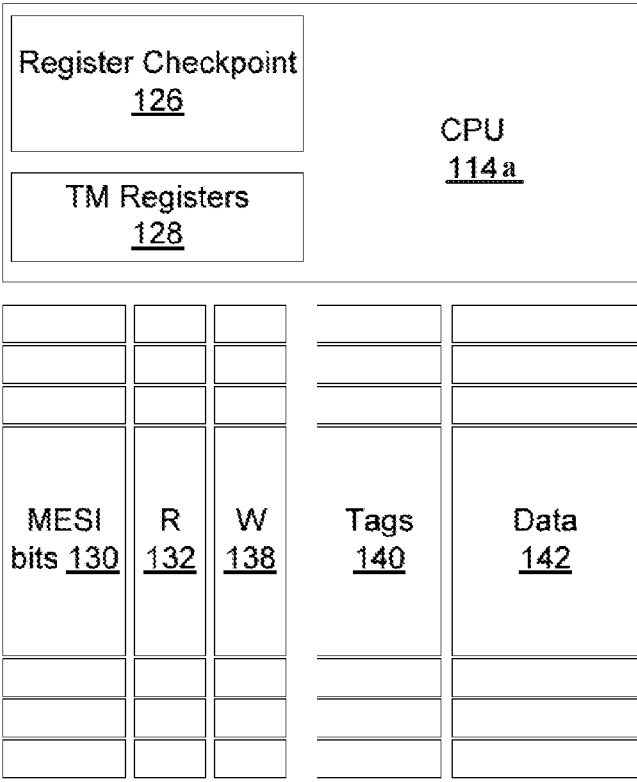



FIG. 2




FIG. 3




FIG. 4

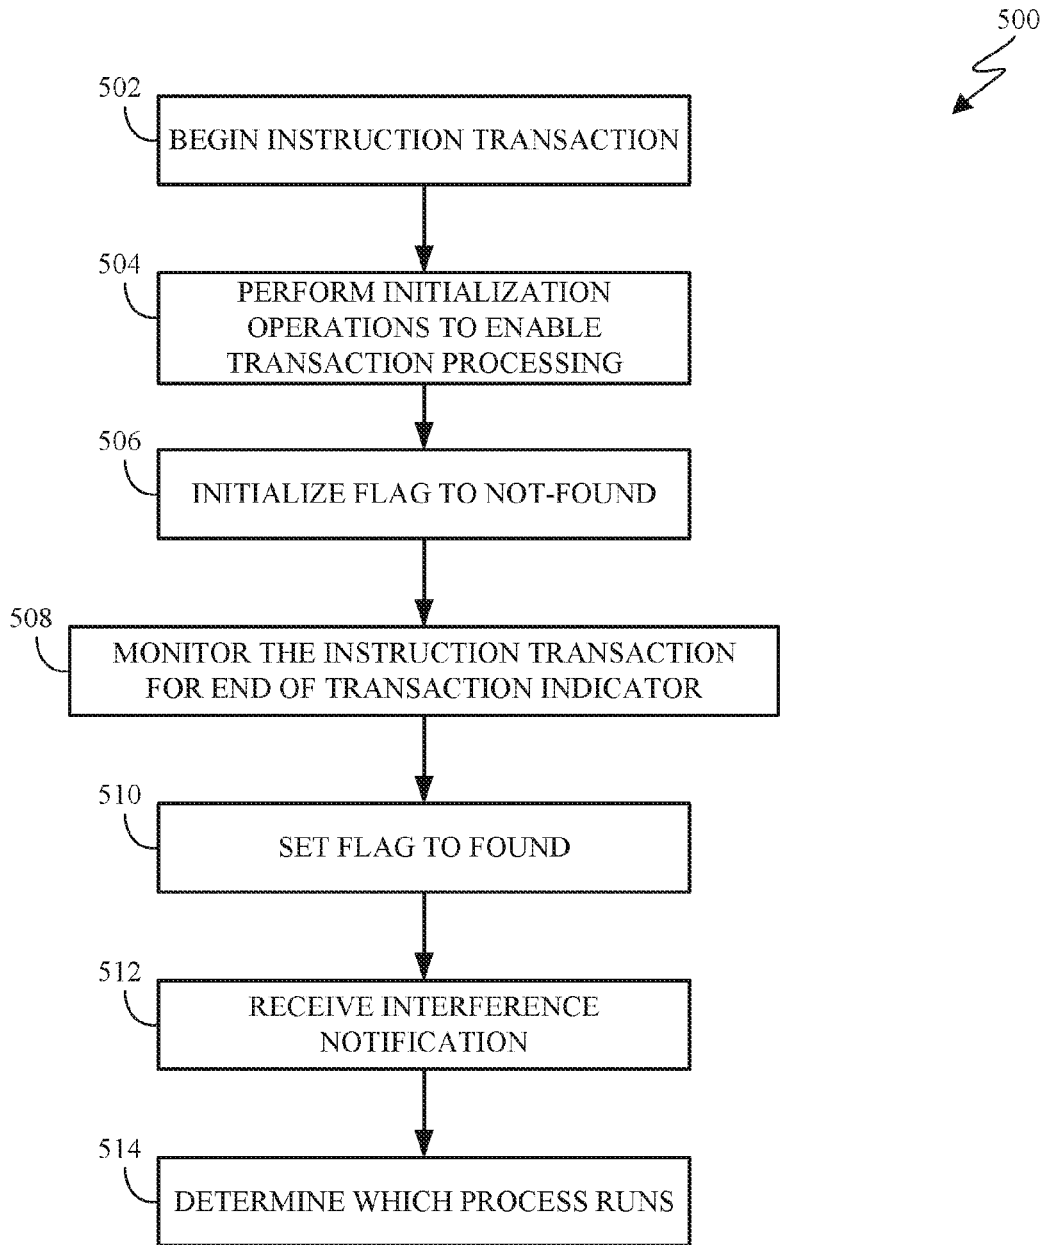


FIG. 5

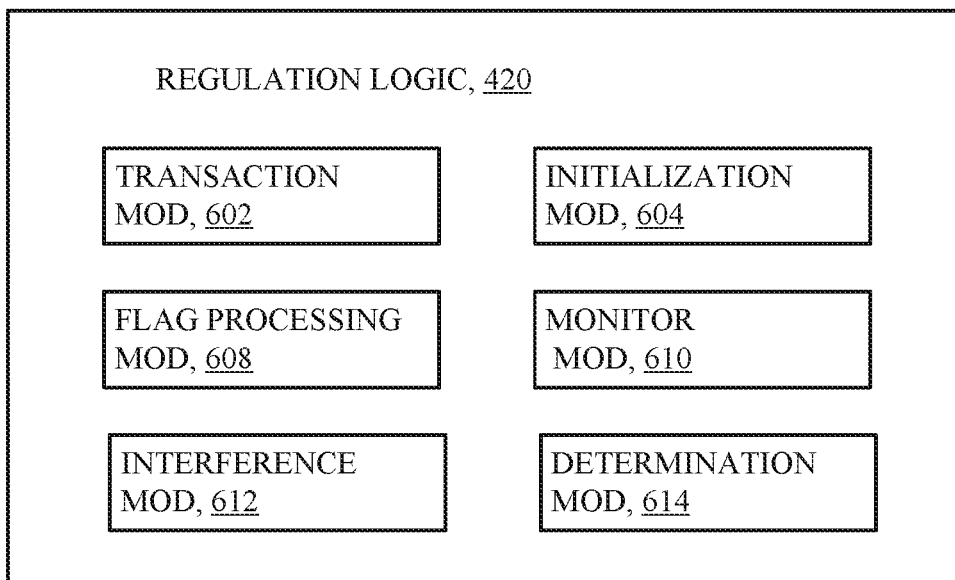


FIG. 6

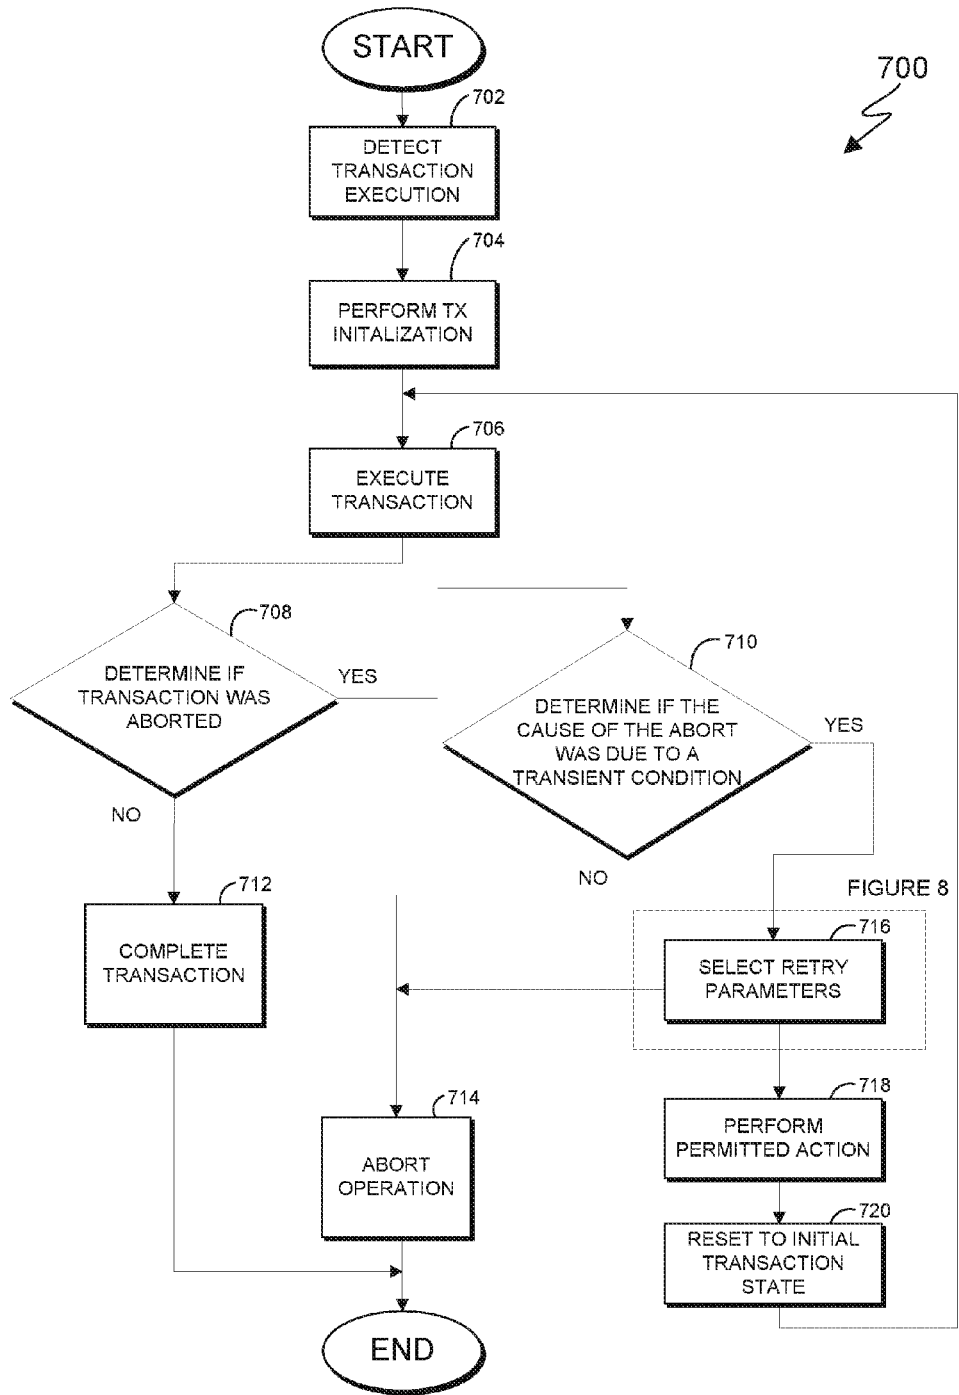


FIG. 7

800
↘

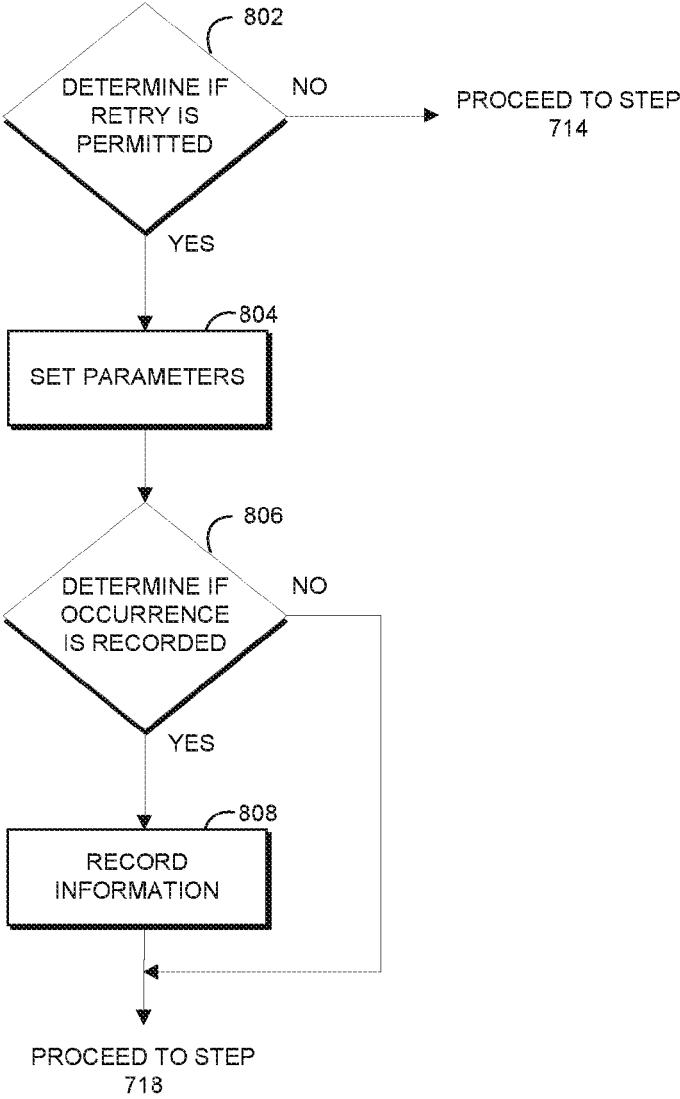


FIG. 8

HARDWARE TRANSACTION TRANSIENT CONFLICT RESOLUTION

BACKGROUND

[0001] This disclosure relates generally to transactional memory systems and more specifically to a method, computer program and computer system for improving the efficiency of transactional instruction processing.

[0002] The number of central processing unit (CPU) cores on a chip and the number of CPU cores connected to a shared memory continues to grow significantly to support growing workload capacity demand. The increasing number of CPUs cooperating to process the same workloads puts a significant burden on software scalability, for example, shared queues or data structures protected by traditional semaphores become hot spots and lead to sub-linear n-way scaling curves. Traditionally this has been countered by implementing finer-grained locking in software, and with lower latency/higher bandwidth interconnects in hardware. Implementing fine-grained locking to improve software scalability can be very complicated and error-prone, and at today's CPU frequencies, the latencies of hardware interconnects are limited by the physical dimension of the chips and systems, and by the speed of light.

[0003] Implementations of hardware Transactional Memory (HTM, or in this discussion, simply TM) have been introduced, wherein a group of instructions—called a transaction—operate in an atomic manner on a data structure in memory, as viewed by other central processing units (CPUs) and the I/O subsystem (atomic operation is also known as “block concurrent” or “serialized” in other literature). The transaction executes optimistically without obtaining a lock, but may need to abort and retry the transaction execution if an operation, of the executing transaction, on a memory location conflicts with another operation on the same memory location. Previously, software transactional memory implementations have been proposed to support software Transactional Memory (TM). However, hardware TM can provide improved performance aspects and ease of use over software TM.

[0004] U.S. Patent Application Publication US20080244354 A1 titled “Apparatus and method for redundant multi-threading with recovery” filed 2007 Mar. 28 and incorporated by reference herein teaches a method and apparatus for reducing the effect of soft errors in a computer system is provided. Soft errors are detected by combining software redundant threading and instruction duplication. Upon detection of a soft error, errors are recovered through the use of software check pointing/rollback technology. Reliable regions are identified by vulnerability profiling and redundant multi-threading is applied to the identified reliable regions.

[0005] U.S. Patent Application Publication US20120210162 A1 titled “State recovery and lockstep execution restart in a system with multiprocessor pairing” filed 2011 Feb. 15 and incorporated by reference herein teaches a system, method and computer program product for a multiprocessing system to offer selective pairing of processor cores for increased processing reliability. A selective pairing facility is provided that selectively connects, i.e., pairs, multiple microprocessor or processor cores to provide one highly reliable thread (or thread group). Each paired microprocessor or processor cores that provide one highly reliable thread for high-reliability connect with a system

components such as a memory “nest” (or memory hierarchy), an optional system controller, and optional interrupt controller, optional I/O or peripheral devices, etc. The memory nest is attached to a selective pairing facility via a switch or a bus. Each selectively paired processor core includes a transactional execution facility, wherein the system is configured to enable processor rollback to a previous state and reinitialize lockstep execution in order to recover from an incorrect execution when an incorrect execution has been detected by the selective pairing facility.

SUMMARY

[0006] Aspects of an embodiment of the present invention disclose a method, computer program product, and computing system for resolving terminated transactions in a transactional memory environment. A processor initiates a hardware transaction in a computing environment, wherein the hardware transaction accesses a memory location, and wherein the hardware transaction includes a transaction begin indicator and a transaction end indicator. A processor detects a conflicting access of the memory location while executing the hardware transaction. A processor aborts the hardware transaction based on the conflicting access of the memory location. Hardware determines that the conflicting access of the memory location is a transient condition. A processor reinitiates the hardware transaction.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 depicts an example multicore transactional memory environment, in accordance with an illustrative embodiment.

[0008] FIG. 2 depicts an example multicore transactional memory environment, in accordance with an illustrative embodiment.

[0009] FIG. 3 depicts example components of an example CPU, in accordance with an illustrative embodiment.

[0010] FIG. 4 is a block diagram of internal and external components of a computing device, in accordance with one embodiment of the present invention.

[0011] FIG. 5 is a flowchart of the steps of an approach for processing a hardware transaction and identifying that an interference has occurred, in accordance with one embodiment of the present invention.

[0012] FIG. 6 is a block diagram of regulation logic 420 and associated modules, in accordance with one embodiment of the present invention.

[0013] FIG. 7 depicts a flowchart of the steps of regulation logic for detecting whether a transient condition was the cause of an aborted hardware transaction and determining whether to retry the hardware transaction, in accordance with one embodiment of the present invention.

[0014] FIG. 8 depicts a flowchart of the steps of regulation logic for selecting retry parameters, in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION

[0015] A transaction within a computer program or computer application comprises program instructions performing multiple store operations that appear to run and complete as a single, atomic operation. The program instructions forming a current transaction comprise a transaction begin indicator, a plurality of instructions (e.g., arithmetic, load, branch or store operations), and a transaction end indicator.

A near-end of transaction indicator is triggered based on a speculative look ahead operation, and enabling near-end-transaction processing mode, such that an interfering operation may be delayed to allow the current transaction to complete. A halt operation, also referred to as an abort operation, as used herein refers to an operation responsive to a condition where two transactions have been detected to interfere where at least one transaction must be aborted and the state of the processor is reset to the state at the beginning of the aborted transaction by performing a rollback. This Detailed Description section is divided into the following subsections: (i) The Hardware and Software Environment; (ii) Example Embodiment; (iii) Further Comments and/or Embodiments; and (iv) Definitions.

I. The Hardware and Software Environments

[0016] The present invention may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

A. Transaction Execution Environment

[0017] Historically, a computer system or processor had only a single processor (aka processing unit or central processing unit). The processor included an instruction processing unit (IPU), a branch unit, a memory control unit and the like. Such processors were capable of executing a single thread of a program at a time. Operating systems were developed that could time-share a processor by dispatching a program to be executed on the processor for a period of time, and then dispatching another program to be executed on the processor for another period of time. As technology evolved, memory subsystem caches were often added to the processor as well as complex dynamic address translation including translation lookaside buffers (TLBs). The IPU itself was often referred to as a processor. As technology continued to evolve, an entire processor could be packaged on a single semiconductor chip or die, such a processor was referred to as a microprocessor. Then processors were developed that incorporated multiple IPUs, such processors were often referred to as multi-processors. Each such processor of a multi-processor computer system (processor) may include individual or shared caches, memory interfaces, system bus, address translation mechanism and the like. Virtual machine and instruction set architecture (ISA) emulators added a layer of software to a processor, that provided the virtual machine with multiple “virtual processors” (aka processors) by time-slice usage of a single IPU in a single hardware processor. As technology further evolved, multi-threaded processors were developed, enabling a single hardware processor having a single multi-thread IPU to provide a capability of simultaneously executing threads of different programs, thus each thread of a multi-threaded processor appeared to the operating system as a processor. As technology further evolved, it was possible to put multiple processors (each having an IPU) on a single semiconductor chip or die. These processors were referred to processor cores or just cores. Thus the terms such as processor, central processing unit, processing unit, microprocessor, core, processor core, processor thread, and thread, for example, are often used interchangeably. Aspects of embodiments herein

may be practiced by any or all processors including those shown supra, without departing from the teachings herein. Wherein the term “thread” or “processor thread” is used herein, it is expected that particular advantage of the embodiment may be had in a processor thread implementation.

Transaction Execution in Intel Based Embodiments

[0018] In “Intel Architecture Instruction Set Extensions Programming Reference” 319433-012A, February 2012, incorporated herein by reference in its entirety, Chapter 8 teaches, in part, that multithreaded applications may take advantage of increasing numbers of CPU cores to achieve higher performance. (Note: the term “Intel” may be subject to trademark rights in various jurisdictions throughout the world and are used here only in reference to the products or services properly denominated by the marks to the extent that such trademark rights may exist.) However, the writing of multi-threaded applications requires programmers to understand and take into account data sharing among the multiple threads. Access to shared data typically requires synchronization mechanisms. These synchronization mechanisms are used to ensure that multiple threads update shared data by serializing operations that are applied to the shared data, often through the use of a critical section that is protected by a lock. Since serialization limits concurrency, programmers try to limit the overhead due to synchronization.

[0019] Intel Transactional Synchronization Extensions (Intel TSX) allow a processor to dynamically determine whether threads need to be serialized through lock-protected critical sections, and to perform that serialization only when required. (Note: the term(s) “Intel,” “TSX,” and/or “Intel TSX” may be subject to trademark rights in various jurisdictions throughout the world and are used here only in reference to the products or services properly denominated by the marks to the extent that such trademark rights may exist.) This allows the processor to expose and exploit concurrency that is hidden in an application because of dynamically unnecessary synchronization.

[0020] With Intel TSX, programmer-specified code regions (also referred to as “transactional regions” or just “transactions”) are executed transactionally. If the transactional execution completes successfully, then all memory operations performed within the transactional region will appear to have occurred instantaneously when viewed from other processors. A processor makes the memory operations of the executed transaction, performed within the transactional region, visible to other processors only when a successful commit occurs, i.e., when the transaction successfully completes execution. This process is often referred to as an atomic commit.

[0021] Intel TSX provides two software interfaces to specify regions of code for transactional execution. Hardware Lock Elision (HLE) is a legacy compatible instruction set extension (comprising the XACQUIRE and XRELEASE prefixes) to specify transactional regions. Restricted Transactional Memory (RTM) is a new instruction set interface (comprising the XBEGIN, XEND, and XABORT instructions) for programmers to define transactional regions in a more flexible manner than that possible with HLE. HLE is for programmers who prefer the backward compatibility of the conventional mutual exclusion programming model and would like to run HLE-enabled software on legacy hardware

but would also like to take advantage of the new lock elision capabilities on hardware with HLE support. RTM is for programmers who prefer a flexible interface to the transactional execution hardware. In addition, Intel TSX also provides an XTEST instruction. This instruction allows software to query whether the logical processor is transactionally executing in a transactional region identified by either HLE or RTM.

[0022] Since a successful transactional execution ensures an atomic commit, the processor executes the code region optimistically without explicit synchronization. If synchronization was unnecessary for that specific execution, execution can commit without any cross-thread serialization. If the processor cannot commit atomically, then the optimistic execution fails. When this happens, the processor will roll back the execution, a process referred to as a transactional abort. On a transactional abort, the processor will discard all updates performed in the memory region used by the transaction, restore architectural state to appear as if the optimistic execution never occurred, and resume execution non-transactionally.

[0023] A processor can perform a transactional abort for numerous reasons. A primary reason to abort a transaction is due to conflicting memory accesses between the transactionally executing logical processor and another logical processor. Such conflicting memory accesses may prevent a successful transactional execution. Memory addresses read from within a transactional region constitute the read-set of the transactional region and addresses written to within the transactional region constitute the write-set of the transactional region. Intel TSX maintains the read- and write-sets at the granularity of a cache line. A conflicting memory access occurs if another logical processor either reads a location that is part of the transactional region's write-set or writes a location that is a part of either the read- or write-set of the transactional region. A conflicting access typically means that serialization is required for this code region. Since Intel TSX detects data conflicts at the granularity of a cache line, unrelated data locations placed in the same cache line will be detected as conflicts that result in transactional aborts. Transactional aborts may also occur due to limited transactional resources. For example, the amount of data accessed in the region may exceed an implementation-specific capacity. Additionally, some instructions and system events may cause transactional aborts. Frequent transactional aborts result in wasted cycles and increased inefficiency.

Hardware Lock Elision

[0024] Hardware Lock Elision (HLE) provides a legacy compatible instruction set interface for programmers to use transactional execution. HLE provides two new instruction prefix hints: XACQUIRE and XRELEASE.

[0025] With HLE, a programmer adds the XACQUIRE prefix to the front of the instruction that is used to acquire the lock that is protecting the critical section. The processor treats the prefix as a hint to elide the write associated with the lock acquire operation. Even though the lock acquire has an associated write operation to the lock, the processor does not add the address of the lock to the transactional region's write-set nor does it issue any write requests to the lock. Instead, the address of the lock is added to the read-set. The logical processor enters transactional execution. If the lock was available before the XACQUIRE prefixed instruction, then all other processors will continue to see the lock as

available afterwards. Since the transactionally executing logical processor neither added the address of the lock to its write-set nor performed externally visible write operations to the lock, other logical processors can read the lock without causing a data conflict. This allows other logical processors to also enter and concurrently execute the critical section protected by the lock. The processor automatically detects any data conflicts that occur during the transactional execution and will perform a transactional abort if necessary.

[0026] Even though the eliding processor did not perform any external write operations to the lock, the hardware ensures program order of operations on the lock. If the eliding processor itself reads the value of the lock in the critical section, it will appear as if the processor had acquired the lock, i.e. the read will return the non-elided value. This behavior allows an HLE execution to be functionally equivalent to an execution without the HLE prefixes.

[0027] An XRELEASE prefix can be added in front of an instruction that is used to release the lock protecting a critical section. Releasing the lock involves a write to the lock. If the instruction is to restore the value of the lock to the value the lock had prior to the XACQUIRE prefixed lock acquire operation on the same lock, then the processor elides the external write request associated with the release of the lock and does not add the address of the lock to the write-set. The processor then attempts to commit the transactional execution.

[0028] With HLE, if multiple threads execute critical sections protected by the same lock but they do not perform any conflicting operations on each other's data, then the threads can execute concurrently and without serialization. Even though the software uses lock acquisition operations on a common lock, the hardware recognizes this, elides the lock, and executes the critical sections on the two threads without requiring any communication through the lock—if such communication was dynamically unnecessary.

[0029] If the processor is unable to execute the region transactionally, then the processor will execute the region non-transactionally and without elision. HLE enabled software has the same forward progress guarantees as the underlying non-HLE lock-based execution. For successful HLE execution, the lock and the critical section code must follow certain guidelines. These guidelines only affect performance; and failure to follow these guidelines will not result in a functional failure. Hardware without HLE support will ignore the XACQUIRE and XRELEASE prefix hints and will not perform any elision since these prefixes correspond to the REPNE/REPE IA-32 prefixes which are ignored on the instructions where XACQUIRE and XRELEASE are valid. Importantly, HLE is compatible with the existing lock-based programming model. Improper use of hints will not cause functional bugs though it may expose latent bugs already in the code.

[0030] Restricted Transactional Memory (RTM) provides a flexible software interface for transactional execution. RTM provides three new instructions—XBEGIN, XEND, and XABORT—for programmers to start, commit, and abort a transactional execution.

[0031] The programmer uses the XBEGIN instruction to specify the start of a transactional code region and the XEND instruction to specify the end of the transactional code region. If the RTM region could not be successfully

executed transactionally, then the XBEGIN instruction takes an operand that provides a relative offset to the fallback instruction address.

[0032] A processor may abort RTM transactional execution for many reasons. In many instances, the hardware automatically detects transactional abort conditions and restarts execution from the fallback instruction address with the architectural state corresponding to that present at the start of the XBEGIN instruction and the EAX register updated to describe the abort status.

[0033] The XABORT instruction allows programmers to abort the execution of an RTM region explicitly. The XABORT instruction takes an 8-bit immediate argument that is loaded into the EAX register and will thus be available to software following an RTM abort. RTM instructions do not have any data memory location associated with them. While the hardware provides no guarantees as to whether an RTM region will ever successfully commit transactionally, most transactions that follow the recommended guidelines are expected to successfully commit transactionally. However, programmers must always provide an alternative code sequence in the fallback path to guarantee forward progress. This may be as simple as acquiring a lock and executing the specified code region non-transactionally. Further, a transaction that always aborts on a given implementation may complete transactionally on a future implementation. Therefore, programmers must ensure the code paths for the transactional region and the alternative code sequence are functionally tested.

Detection of HLE Support

[0034] A processor supports HLE execution if CPUID.07H.EBX.HLE [bit 4]=1. However, an application can use the HLE prefixes (XACQUIRE and XRELEASE) without checking whether the processor supports HLE. Processors without HLE support ignore these prefixes and will execute the code without entering transactional execution.

Detection of RTM Support

[0035] A processor supports RTM execution if CPUID.07H.EBX.RTM [bit 11]=1. An application must check if the processor supports RTM before it uses the RTM instructions (XBEGIN, XEND, XABORT). These instructions will generate a #UD exception when used on a processor that does not support RTM.

Detection of XTEST Instruction

[0036] A processor supports the XTEST instruction if it supports either HLE or RTM. An application must check either of these feature flags before using the XTEST instruction. This instruction will generate a #UD exception when used on a processor that does not support either HLE or RTM.

Querying Transactional Execution Status

[0037] The XTEST instruction can be used to determine the transactional status of a transactional region specified by HLE or RTM. Note, while the HLE prefixes are ignored on processors that do not support HLE, the XTEST instruction will generate a #UD exception when used on processors that do not support either HLE or RTM.

Requirements for HLE Locks

[0038] For HLE execution to successfully commit transactionally, the lock must satisfy certain properties and access to the lock must follow certain guidelines.

[0039] An XRELEASE prefixed instruction must restore the value of the elided lock to the value it had before the lock acquisition. This allows hardware to safely elide locks by not adding them to the write-set. The data size and data address of the lock release (XRELEASE prefixed) instruction must match that of the lock acquire (XACQUIRE prefixed) and the lock must not cross a cache line boundary.

[0040] Software should not write to the elided lock inside a transactional HLE region with any instruction other than an XRELEASE prefixed instruction, otherwise such a write may cause a transactional abort. In addition, recursive locks (where a thread acquires the same lock multiple times without first releasing the lock) may also cause a transactional abort. Note that software can observe the result of the elided lock acquire inside the critical section. Such a read operation will return the value of the write to the lock.

[0041] The processor automatically detects violations to these guidelines, and safely transitions to a non-transactional execution without elision. Since Intel TSX detects conflicts at the granularity of a cache line, writes to data collocated on the same cache line as the elided lock may be detected as data conflicts by other logical processors eliding the same lock.

Transactional Nesting

[0042] Both HLE and RTM support nested transactional regions. However, a transactional abort restores state to the operation that started transactional execution: either the outermost XACQUIRE prefixed HLE eligible instruction or the outermost XBEGIN instruction. The processor treats all nested transactions as one transaction.

HLE Nesting and Elision

[0043] Programmers can nest HLE regions up to an implementation specific depth of MAX_HLE_NEST_COUNT. Each logical processor tracks the nesting count internally but this count is not available to software. An XACQUIRE prefixed HLE-eligible instruction increments the nesting count, and an XRELEASE prefixed HLE-eligible instruction decrements it. The logical processor enters transactional execution when the nesting count goes from zero to one. The logical processor attempts to commit only when the nesting count becomes zero. A transactional abort may occur if the nesting count exceeds MAX_HLE_NEST_COUNT.

[0044] In addition to supporting nested HLE regions, the processor can also elide multiple nested locks. The processor tracks a lock for elision beginning with the XACQUIRE prefixed HLE eligible instruction for that lock and ending with the XRELEASE prefixed HLE eligible instruction for that same lock. The processor can, at any one time, track up to a MAX_HLE_ELIDED_LOCKS number of locks. For example, if the implementation supports a MAX_HLE_ELIDED_LOCKS value of two and if the programmer nests three HLE identified critical sections (by performing XACQUIRE prefixed HLE eligible instructions on three distinct locks without performing an intervening XRELEASE prefixed HLE eligible instruction on any one of the locks), then the first two locks will be elided, but the third won't be elided (but will be added to the transaction's write-set).

However, the execution will still continue transactionally. Once an XRELEASE for one of the two elided locks is encountered, a subsequent lock acquired through the XACQUIRE prefixed HLE eligible instruction will be elided.

[0045] The processor attempts to commit the HLE execution when all elided XACQUIRE and XRELEASE pairs have been matched, the nesting count goes to zero, and the locks have satisfied requirements. If execution cannot commit atomically, then execution transitions to a non-transactional execution without elision as if the first instruction did not have an XACQUIRE prefix.

RTM Nesting

[0046] Programmers can nest RTM regions up to an implementation specific MAX_RTM_NEST_COUNT. The logical processor tracks the nesting count internally but this count is not available to software. An XBEGIN instruction increments the nesting count, and an XEND instruction decrements the nesting count. The logical processor attempts to commit only if the nesting count becomes zero. A transactional abort occurs if the nesting count exceeds MAX_RTM_NEST_COUNT.

Nesting HLE and RTM

[0047] HLE and RTM provide two alternative software interfaces to a common transactional execution capability. Transactional processing behavior is implementation specific when HLE and RTM are nested together, e.g., HLE is inside RTM or RTM is inside HLE. However, in all cases, the implementation will maintain HLE and RTM semantics. An implementation may choose to ignore HLE hints when used inside RTM regions, and may cause a transactional abort when RTM instructions are used inside HLE regions. In the latter case, the transition from transactional to non-transactional execution occurs seamlessly since the processor will re-execute the HLE region without actually doing elision, and then execute the RTM instructions.

Abort Status Definition

[0048] RTM uses the EAX register to communicate abort status to software. Following an RTM abort the EAX register has the following definition.

TABLE 1

RTM Abort Status Definition	
EAX Register Bit Position	Meaning
0	Set if abort caused by XABORT instruction
1	If set, the transaction may succeed on retry, this bit is always clear if bit 0 is set
2	Set if another logical processor conflicted with a memory address that was part of the transaction that aborted
3	Set if an internal buffer overflowed
4	Set if a debug breakpoint was hit
5	Set if an abort occurred during execution of a nested transaction
23:6	Reserved
31-24	XABORT argument (only valid if bit 0 set, otherwise reserved)

[0049] The EAX abort status for RTM only provides causes for aborts. It does not by itself encode whether an abort or commit occurred for the RTM region. The value of

EAX can be 0 following an RTM abort. For example, a CPUID instruction when used inside an RTM region causes a transactional abort and may not satisfy the requirements for setting any of the EAX bits. This may result in an EAX value of 0.

RTM Memory Ordering

[0050] A successful RTM commit causes all memory operations in the RTM region to appear to execute atomically. A successfully committed RTM region consisting of an XBEGIN followed by an XEND, even with no memory operations in the RTM region, has the same ordering semantics as a LOCK prefixed instruction.

[0051] The XBEGIN instruction does not have fencing semantics. However, if an RTM execution aborts, then all memory updates from within the RTM region are discarded and are not made visible to any other logical processor.

RTM-Enabled Debugger Support

[0052] By default, any debug exception inside an RTM region will cause a transactional abort and will redirect control flow to the fallback instruction address with architectural state recovered and bit 4 in EAX set. However, to allow software debuggers to intercept execution on debug exceptions, the RTM architecture provides additional capability.

[0053] If bit 11 of DR7 and bit 15 of the IA32_DEBUGCTL_MSR are both 1, any RTM abort due to a debug exception (#DB) or breakpoint exception (#BP) causes execution to roll back and restart from the XBEGIN instruction instead of the fallback address. In this scenario, the EAX register will also be restored back to the point of the XBEGIN instruction.

Programming Considerations

[0054] Typical programmer-identified regions are expected to transactionally execute and commit successfully. However, Intel TSX does not provide any such guarantee. A transactional execution may abort for many reasons. To take full advantage of the transactional capabilities, programmers should follow certain guidelines to increase the probability of their transactional execution committing successfully.

[0055] This section discusses various events that may cause transactional aborts. The architecture ensures that updates performed within a transaction that subsequently aborts execution will never become visible. Only committed transactional executions initiate an update to the architectural state. Transactional aborts never cause functional failures and only affect performance.

Instruction Based Considerations

[0056] Programmers can use any instruction safely inside a transaction (HLE or RTM) and can use transactions at any privilege level. However, some instructions will always abort the transactional execution and cause execution to seamlessly and safely transition to a non-transactional path.

[0057] Intel TSX allows for most common instructions to be used inside transactions without causing aborts. The following operations inside a transaction do not typically cause an abort:

[0058] Operations on the instruction pointer register, general purpose registers (GPRs) and the status flags (CF, OF, SF, PF, AF, and ZF); and

[0059] Operations on XMM and YMM registers and the MXCSR register.

[0060] However, programmers must be careful when intermixing SSE and AVX operations inside a transactional region. Intermixing SSE instructions accessing XMM registers and AVX instructions accessing YMM registers may cause transactions to abort. Programmers may use REP/REPNE prefixed string operations inside transactions. However, long strings may cause aborts. Further, the use of CLD and STD instructions may cause aborts if they change the value of the DF flag. However, if DF is 1, the STD instruction will not cause an abort. Similarly, if DF is 0, then the CLD instruction will not cause an abort.

[0061] Instructions not enumerated here as causing abort when used inside a transaction will typically not cause a transaction to abort (examples include but are not limited to MFENCE, LFENCE, SFENCE, RDTSC, RDTSCP, etc.).

[0062] The following instructions will abort transactional execution on any implementation:

[0063] XABORT

[0064] CPUID

[0065] PAUSE

[0066] In addition, in some implementations, the following instructions may always cause transactional aborts. These instructions are not expected to be commonly used inside typical transactional regions. However, programmers must not rely on these instructions to force a transactional abort, since whether they cause transactional aborts is implementation dependent.

[0067] Operations on X87 and MMX architecture state. This includes all MMX and X87 instructions, including the FXRSTOR and FXSAVE instructions.

[0068] Update to non-status portion of EFLAGS: CLI, STI, POPFD, POPFQ, CLTS.

[0069] Instructions that update segment registers, debug registers and/or control registers:

[0070] MOV to DS/ES/FS/GS/SS, POP DS/ES/FS/GS/SS, LDS, LES, LFS, LGS, LSS, SWAPGS, WRFSBASE, WRGSBASE, LGDT, SGDT, LIDT, SIDT, LLDT, SLDT, LTR, STR, Far CALL, Far JMP, Far RET, IRET, MOV to DRx, MOV to CR0/CR2/CR3/CR4/CR8 and LMSW.

[0071] Ring transitions: SYSENTER, SYSCALL, SYSEXIT, and SYSRET.

[0072] TLB and Cacheability control: CLFLUSH, INVD, WBINVD, INVLPG, INVPCID, and memory instructions with a non-temporal hint (MOVNTDQA, MOVNTDQ, MOVNTI, MOVNTPD, MOVNTPS, and MOVNTQ).

[0073] Processor state save: XSAVE, XSAVEOPT, and XRSTOR.

[0074] Interrupts: INTn, INTO.

[0075] IO: IN, INS, REP INS, OUT, OUTS, REP OUTS and their variants.

[0076] VMX: VMPTRLD, VMPTRST, VMCLEAR, VMREAD, VMWRITE, VMCALL, VMLAUNCH, VMRESUME, VMXOFF, VMXON, INVEPT, and INVVPID.

[0077] SMX: GETSEC.

[0078] UD2, RSM, RDMSR, WRMSR, HLT, MONITOR, MWAIT, XSETBV, VZEROUPPER, MASKMOVQ, and V/MASKMOVDQU.

Runtime Considerations

[0079] In addition to the instruction-based considerations, runtime events may cause transactional execution to abort. These may be due to data access patterns or micro-architectural implementation features. The following list is not a comprehensive discussion of all abort causes.

[0080] Any fault or trap in a transaction that must be exposed to software will be suppressed. Transactional execution will abort and execution will transition to a non-transactional execution, as if the fault or trap had never occurred. If an exception is not masked, then that un-masked exception will result in a transactional abort and the state will appear as if the exception had never occurred.

[0081] Synchronous exception events (#DE, #OF, #NP, #SS, #GP, #BR, #UD, #AC, #XF, #PF, #NM, #TS, #ME, #DB, #BP/INT3) that occur during transactional execution may cause an execution not to commit transactionally, and require a non-transactional execution. These events are suppressed as if they had never occurred. With HLE, since the non-transactional code path is identical to the transactional code path, these events will typically reappear when the instruction that caused the exception is re-executed non-transactionally, causing the associated synchronous events to be delivered appropriately in the non-transactional execution. Asynchronous events (NMI, SMI, INTR, IPI, PMI, etc.) occurring during transactional execution may cause the transactional execution to abort and transition to a non-transactional execution. The asynchronous events will be pending and handled after the transactional abort is processed.

[0082] Transactions only support write-back cacheable memory type operations. A transaction may always abort if the transaction includes operations on any other memory type. This includes instruction fetches to UC memory type.

[0083] Memory accesses within a transactional region may require the processor to set the Accessed and Dirty flags of the referenced page table entry. The behavior of how the processor handles this is implementation specific. Some implementations may allow the updates to these flags to become externally visible even if the transactional region subsequently aborts. Some Intel TSX implementations may choose to abort the transactional execution if these flags need to be updated. Further, a processor's page-table walk may generate accesses to its own transactionally written but uncommitted state. Some Intel TSX implementations may choose to abort the execution of a transactional region in such situations. Regardless, the architecture ensures that, if the transactional region aborts, then the transactionally written state will not be made architecturally visible through the behavior of structures such as TLBs.

[0084] Executing self-modifying code transactionally may also cause transactional aborts. Programmers must continue to follow the Intel recommended guidelines for writing self-modifying and cross-modifying code even when employing HLE and RTM. While an implementation of RTM and HLE will typically provide sufficient resources for executing common transactional regions, implementation constraints and excessive sizes for transactional regions may cause a transactional execution to abort and transition to a non-transactional execution. The architecture provides no

guarantee of the amount of resources available to do transactional execution and does not guarantee that a transactional execution will ever succeed.

[0085] Conflicting requests to a cache line accessed within a transactional region may prevent the transaction from executing successfully. For example, if logical processor P0 reads line A in a transactional region and another logical processor P1 writes line A (either inside or outside a transactional region) then logical processor P0 may abort if logical processor P1's write interferes with processor P0's ability to execute transactionally.

[0086] Similarly, if P0 writes line A in a transactional region and P1 reads or writes line A (either inside or outside a transactional region), then P0 may abort if P1's access to line A interferes with P0's ability to execute transactionally. In addition, other coherence traffic may at times appear as conflicting requests and may cause aborts. While these false conflicts may happen, they are expected to be uncommon. The conflict resolution policy to determine whether P0 or P1 aborts in the above scenarios is implementation specific.

Generic Transaction Execution embodiments:

[0087] According to "ARCHITECTURES FOR TRANSACTIONAL MEMORY", a dissertation submitted to the Department of Computer Science and the Committee on Graduate Studies of Stanford University in partial fulfillment of the requirements for the Degree of Doctor of Philosophy, by Austen McDonald, June 2009, incorporated by reference herein in its entirety, fundamentally, there are three mechanisms needed to implement an atomic and isolated transactional region: versioning, conflict detection, and contention management.

[0088] To make a transactional code region appear atomic, all the modifications performed by that transactional code region must be stored and kept isolated from other transactions until commit time. The system does this by implementing a versioning policy. Two versioning paradigms exist: eager and lazy. An eager versioning system stores newly generated transactional values in place and stores previous memory values on the side, in what is called an undo-log. A lazy versioning system stores new values temporarily in what is called a write buffer, copying them to memory only on commit. In either system, the cache is used to optimize storage of new versions.

[0089] To ensure that transactions appear to be performed atomically, conflicts must be detected and resolved. The two systems, i.e., the eager and lazy versioning systems, detect conflicts by implementing a conflict detection policy, either optimistic or pessimistic. An optimistic system executes transactions in parallel, checking for conflicts only when a transaction commits. A pessimistic system checks for conflicts at each load and store. Similar to versioning, conflict detection also uses the cache, marking each line as either part of the read-set, part of the write-set, or both. The two systems resolve conflicts by implementing a contention management policy. Many contention management policies exist, some are more appropriate for optimistic conflict detection and some are more appropriate for pessimistic. Described below are some example policies.

[0090] Since each transactional memory (TM) system needs both versioning detection and conflict detection, these options give rise to four distinct TM designs: Eager-Pessimistic (EP), Eager-Optimistic (EO), Lazy-Pessimistic (LP), and Lazy-Optimistic (LO). Table 2 briefly describes all four distinct TM designs.

[0091] FIGS. 1 and 2 depict an example of a multicore TM environment. FIG. 1 shows many TM-enabled CPUs (CPU1 114a, CPU2 114b, and other CPUs not shown) on one die 100, connected with an interconnect 122, under management of an interconnect control 120a, 120b. Each CPU 114a and 114b (also known as a Processor) may have a split cache consisting of Instruction Cache 116a and 116b for caching instructions from memory to be executed and Data Cache with TM support 118a and 118b for caching data (operands) of memory locations to be operated on by CPU 114a and 114b (in FIG. 1, each CPU 114a, 114b and its associated caches are referenced as 112a and 112b). In an implementation, caches of multiple dies 100 are interconnected to support cache coherency between the caches of multiple dies 100. In an implementation, a single cache, rather than the split cache is employed holding both instructions and data. In implementations, the CPU caches are one level of caching in a hierarchical cache structure. For example each die 100 may employ shared cache 124 to be shared amongst all the CPUs on die 100. In another implementation, each die may have access to shared cache 124, shared amongst all the processors of all dies 100.

[0092] FIG. 2 shows the details of transactional CPU environment 112a, having CPU 114a, including additions to support TM. Transactional CPU (processor) 114a may include hardware for supporting Register Checkpoints 126 and special TM Registers 128. The transactional CPU cache may have MESI bits 130, Tags 140 and Data 142 of a conventional cache but also, for example, R bits 132 showing a line has been read by CPU 114a while executing a transaction and W bits 138 showing a line has been written to by CPU 114a while executing a transaction.

[0093] A key detail for programmers in any TM system is how non-transactional accesses interact with transactions. By design, transactional accesses are screened from each other using the mechanisms above. However, the interaction between a regular, non-transactional load with a transaction containing a new value for that address must still be considered. In addition, the interaction between a non-transactional store with a transaction that has read that address must also be explored. These are issues of the database concept isolation.

[0094] A TM system is said to implement strong isolation, sometimes called strong atomicity, when every non-transactional load and store acts like an atomic transaction. Therefore, non-transactional loads cannot see uncommitted data and non-transactional stores cause atomicity violations in any transactions that have read that address. A system where this is not the case is said to implement weak isolation, sometimes called weak atomicity.

[0095] Strong isolation is often more desirable than weak isolation due to the relative ease of conceptualization and implementation of strong isolation. Additionally, if a programmer has forgotten to surround some shared memory references with transactions, causing bugs, then with strong isolation, the programmer will often detect that oversight using a simple debug interface because the programmer will see a non-transactional region causing atomicity violations. Also, programs written in one model may work differently on another model. Further, strong isolation is often easier to support in hardware TM than weak isolation. With strong isolation, since the coherence protocol already manages load and store communication between processors, transactions can detect non-transactional loads and stores and act appro-

privately. To implement strong isolation in software Transactional Memory (TM), non-transactional code must be modified to include read- and write-barriers; potentially crippling performance. Although great effort has been expended to remove many unneeded barriers, such techniques are often complex and performance is typically far lower than that of HTMs.

needed line. If the other processors have the needed line non-speculatively or have the line R **132** (Read), they downgrade that line to S, and in certain cases issue a cache-to-cache transfer if they have the line in MESI's **130** M or E state. However, if the cache has the line W **138**, then a conflict is detected between the two transactions and additional action(s) must be taken.

TABLE 2

Transactional Memory Design Space			
VERSIONING			
		Lazy	Eager
CONFLICT DETECTION	Optimistic	Storing updates in a write buffer; detecting conflicts at commit time.	Not practical: waiting to update memory until commit time but detecting conflicts at access time guarantees wasted work and provides no advantage.
	Pessimistic	Storing updates in a write buffer; detecting conflicts at access time.	Updating memory, keeping old values in undo log; detecting conflicts at access time.

[0096] Table 2 illustrates the fundamental design space of transactional memory (versioning and conflict detection).

Eager-Pessimistic (EP)

[0097] This first TM design described below is known as Eager-Pessimistic. An EP system stores its write-set “in place” (hence the name “eager”) and, to support rollback, stores the old values of overwritten lines in an “undo log”. Processors use the W **138** and R **132** cache bits to track read and write-sets and detect conflicts when receiving snooped load requests. Perhaps the most notable examples of EP systems in known literature are LogTM and UTM.

[0098] Beginning a transaction in an EP system is much like beginning a transaction in other systems: `tm_begin()` takes a register checkpoint, and initializes any status registers. An EP system also requires initializing the undo log, the details of which are dependent on the log format, but often involve initializing a log base pointer to a region of pre-allocated, thread-private memory, and clearing a log bounds register.

[0099] Versioning: In EP, due to the way eager versioning is designed to function, the MESI **130** state transitions (cache line indicators corresponding to Modified, Exclusive, Shared, and Invalid code states) are left mostly unchanged. Outside of a transaction, the MESI **130** state transitions are left completely unchanged. When reading a line inside a transaction, the standard coherence transitions apply (S (Shared)→S, I (Invalid)→S, or I→E (Exclusive)), issuing a load miss as needed, but the R **132** bit is also set. Likewise, writing a line applies the standard transitions (S→M, E→I, I→M), issuing a miss as needed, but also sets the W **138** (Written) bit. The first time a line is written, the old version of the entire line is loaded then written to the undo log to preserve it in case the current transaction aborts. The newly written data is then stored “in-place,” over the old data.

[0100] Conflict Detection: Pessimistic conflict detection uses coherence messages exchanged on misses, or upgrades, to look for conflicts between transactions. When a read miss occurs within a transaction, other processors receive a load request; but they ignore the request if they do not have the

[0101] Similarly, when a transaction seeks to upgrade a line from shared to modified (on a first write), the transaction issues an exclusive load request, which is also used to detect conflicts. If a receiving cache has the line non-speculatively, then the line is invalidated, and in certain cases a cache-to-cache transfer (M or E states) is issued. But, if the line is R **132** or W **138**, a conflict is detected.

[0102] Validation: Because conflict detection is performed on every load, a transaction always has exclusive access to its own write-set. Therefore, validation does not require any additional work.

[0103] Commit: Since eager versioning stores the new version of data items in place, the commit process simply clears the W **138** and R **132** bits and discards the undo log.

[0104] Abort: When a transaction rolls back, the original version of each cache line in the undo log must be restored, a process called “unrolling” or “applying” the log. This is done during `tm_discard()` and must be atomic with regard to other transactions. Specifically, the write-set must still be used to detect conflicts: this transaction has the only correct version of lines in its undo log, and requesting transactions must wait for the correct version to be restored from that log. Such a log can be applied using a hardware state machine or software abort handler.

[0105] Eager-Pessimistic has the characteristics of: Commit is simple and since it is in-place, very fast. Similarly, validation is a no-op. Pessimistic conflict detection detects conflicts early, thereby reducing the number of “doomed” transactions. For example, if two transactions are involved in a Write-After-Read dependency, then that dependency is detected immediately in pessimistic conflict detection. However, in optimistic conflict detection such conflicts are not detected until the writer commits.

[0106] Eager-Pessimistic also has the characteristics of: As described above, the first time a cache line is written, the old value must be written to the log, incurring extra cache accesses. Aborts are expensive as they require undoing the log. For each cache line in the log, a load must be issued, perhaps going as far as main memory before continuing to the next line. Pessimistic conflict detection also prevents certain serializable schedules from existing.

[0107] Additionally, because conflicts are handled as they occur, there is a potential for livelock and careful contention management mechanisms must be employed to guarantee forward progress.

Lazy-Optimistic (LO)

[0108] Another popular TM design is Lazy-Optimistic (LO), which stores its write-set in a “write buffer” or “redo log” and detects conflicts at commit time (still using the R 132 and W 138 bits).

[0109] Versioning: Just as in the EP system, the MESI protocol of the LO design is enforced outside of the transactions. Once inside a transaction, reading a line incurs the standard MESI transitions but also sets the R 132 bit. Likewise, writing a line sets the W 138 bit of the line, but handling the MESI transitions of the LO design is different from that of the EP design. First, with lazy versioning, the new versions of written data are stored in the cache hierarchy until commit while other transactions have access to old versions available in memory or other caches. To make available the old versions, dirty lines (M lines) must be evicted when first written by a transaction. Second, no upgrade misses are needed because of the optimistic conflict detection feature: if a transaction has a line in the S state, it can simply write to it and upgrade that line to an M state without communicating the changes with other transactions because conflict detection is done at commit time.

[0110] Conflict Detection and Validation: To validate a transaction and detect conflicts, LO communicates the addresses of speculatively modified lines to other transactions only when it is preparing to commit. On validation, the processor sends one, potentially large, network packet containing all the addresses in the write-set. Data is not sent, but left in the cache of the committer and marked dirty (M). To build this packet without searching the cache for lines marked W, a simple bit vector is used, called a “store buffer,” with one bit per cache line to track these speculatively modified lines. Other transactions use this address packet to detect conflicts: if an address is found in the cache and the R 132 and/or W 138 bits are set, then a conflict is initiated. If the line is found but neither R 132 nor W 138 is set, then the line is simply invalidated, which is similar to processing an exclusive load.

[0111] To support transaction atomicity, these address packets must be handled atomically, i.e., no two address packets may exist at once with the same addresses. In an LO system, this can be achieved by simply acquiring a global commit token before sending the address packet. However, a two-phase commit scheme could be employed by first sending out the address packet, collecting responses, enforcing an ordering protocol (perhaps oldest transaction first), and committing once all responses are satisfactory.

[0112] Commit: Once validation has occurred, commit needs no special treatment: simply clear W 138 and R 132 bits and the store buffer. The transaction’s writes are already marked dirty in the cache and other caches’ copies of these lines have been invalidated via the address packet. Other processors can then access the committed data through the regular coherence protocol.

[0113] Abort: Rollback is equally easy: because the write-set is contained within the local caches, these lines can be invalidated, then clear W 138 and R 132 bits and the store buffer. The store buffer allows W lines to be found to invalidate without the need to search the cache.

[0114] Lazy-Optimistic has the characteristics of: Aborts are very fast, requiring no additional loads or stores and making only local changes. More serializable schedules can exist than found in EP, which allows an LO system to more aggressively speculate that transactions are independent, which can yield higher performance. Finally, the late detection of conflicts can increase the likelihood of forward progress.

[0115] Lazy-Optimistic also has the characteristics of: Validation takes global communication time proportional to size of write set. Doomed transactions can waste work since conflicts are detected only at commit time.

Lazy-Pessimistic (LP)

[0116] Lazy-Pessimistic (LP) represents a third TM design option, sitting somewhere between EP and LO: storing newly written lines in a write buffer but detecting conflicts on a per access basis.

[0117] Versioning: Versioning is similar but not identical to that of LO: reading a line sets its R bit 132, writing a line sets its W bit 138, and a store buffer is used to track W lines in the cache. Also, dirty (M) lines must be evicted when first written by a transaction, just as in LO. However, since conflict detection is pessimistic, load exclusives must be performed when upgrading a transactional line from I, S→M, which is unlike LO.

[0118] Conflict Detection: LP’s conflict detection operates the same as EP’s: using coherence messages to look for conflicts between transactions.

[0119] Validation: Like in EP, pessimistic conflict detection ensures that at any point, a running transaction has no conflicts with any other running transaction, so validation is a no-op.

[0120] Commit: Commit needs no special treatment: simply clear W 138 and R 132 bits and the store buffer, like in LO.

[0121] Abort: Rollback is also like that of LO: simply invalidate the write-set using the store buffer and clear the W and R bits and the store buffer.

Eager-Optimistic (EO)

[0122] The LP has the characteristics of: Like LO, aborts are very fast. Like EP, the use of pessimistic conflict detection reduces the number of “doomed” transactions. Like EP, some serializable schedules are not allowed and conflict detection must be performed on each cache miss.

[0123] The final combination of versioning and conflict detection is Eager-Optimistic (EO). EO may be a less than optimal choice for HTM systems: since new transactional versions are written in-place, other transactions have no choice but to notice conflicts as they occur (i.e., as cache misses occur). But since EO waits until commit time to detect conflicts, those transactions become “zombies,” continuing to execute, wasting resources, yet are “doomed” to abort.

[0124] EO has proven to be useful in STMs and is implemented by Bartok-STM and McRT. A lazy versioning STM needs to check its write buffer on each read to ensure that it is reading the most recent value. Since the write buffer is not a hardware structure, this is expensive, hence the preference for write-in-place eager versioning. Additionally, since

checking for conflicts is also expensive in an STM, optimistic conflict detection offers the advantage of performing this operation in bulk.

Contention Management

[0125] How a transaction rolls back once the system has decided to abort that transaction has been described above, but, since a conflict involves two transactions, the topics of which transaction should abort, how that abort should be initiated, and when should the aborted transaction be retried need to be explored. These are topics that are addressed by Contention Management (CM), a key component of transactional memory. Described below are policies regarding how the systems initiate aborts and the various established methods of managing which transactions should abort in a conflict.

Contention Management Policies

[0126] A Contention Management (CM) Policy is a mechanism that determines which transaction involved in a conflict should abort and when the aborted transaction should be retried. For example, it is often the case that retrying an aborted transaction immediately does not lead to the best performance. Conversely, employing a back-off mechanism, which delays the retrying of an aborted transaction, can yield better performance. STMs first grappled with finding the best contention management policies and many of the policies outlined below were originally developed for STMs.

[0127] CM Policies draw on a number of measures to make decisions, including ages of the transactions, size of read- and write-sets, the number of previous aborts, etc. The combinations of measures to make such decisions are endless, but certain combinations are described below, roughly in order of increasing complexity.

[0128] To establish some nomenclature, first note that in a conflict there are two sides: the attacker and the defender. The attacker is the transaction requesting access to a shared memory location. In pessimistic conflict detection, the attacker is the transaction issuing the load or load exclusive. In optimistic, the attacker is the transaction attempting to validate. The defender in both cases is the transaction receiving the attacker's request.

[0129] An Aggressive CM Policy immediately and always retries either the attacker or the defender. In LO, Aggressive means that the attacker always wins, and so Aggressive is sometimes called committer wins. Such a policy was used for the earliest LO systems. In the case of EP, Aggressive can be either defender wins or attacker wins.

[0130] Restarting a conflicting transaction that will immediately experience another conflict is bound to waste work—namely interconnect bandwidth refilling cache misses. A Polite CM Policy employs exponential back-off (but linear could also be used) before restarting conflicts. To prevent starvation, a situation where a process does not have resources allocated to it by the scheduler, the exponential back-off greatly increases the odds of transaction success after some n retries.

[0131] Another approach to conflict resolution is to randomly abort the attacker or defender (a policy called Randomized). Such a policy may be combined with a randomized back-off scheme to avoid unneeded contention.

[0132] However, making random choices, when selecting a transaction to abort, can result in aborting transactions that have completed “a lot of work,” which can waste resources. To avoid such waste, the amount of work completed on the transaction can be taken into account when determining which transaction to abort. One measure of work could be a transaction's age. Other methods include Oldest, Bulk TM, Size Matters, Karma, and Polka. Oldest is a simple time-stamp method that aborts the younger transaction in a conflict. Bulk TM uses this scheme. Size Matters is like Oldest but instead of transaction age, the number of read/written words is used as the priority, reverting to Oldest after a fixed number of aborts. Karma is similar, using the size of the write-set as priority. Rollback then proceeds after backing off a fixed amount of time. Aborted transactions keep their priorities after being aborted (hence the name Karma). Polka works like Karma but instead of backing off a pre-defined amount of time, it backs off exponentially more each time.

[0133] Since aborting wastes work, it is logical to argue that stalling an attacker until the defender has finished their transaction would lead to better performance. Unfortunately, such a simple scheme easily leads to deadlock.

[0134] Deadlock avoidance techniques can be used to solve this problem. Greedy uses two rules to avoid deadlock. The first rule is, if a first transaction, T1, has lower priority than a second transaction, T0, or if T1 is waiting for another transaction, then T1 aborts when conflicting with T0. The second rule is, if T1 has higher priority than T0 and is not waiting, then T0 waits until T1 commits, aborts, or starts waiting (in which case the first rule is applied). Greedy provides some guarantees about time bounds for executing a set of transactions. One EP design (LogTM) uses a CM policy similar to Greedy to achieve stalling with conservative deadlock avoidance.

[0135] Example MESI coherency rules provide for four possible states in which a cache line of a multiprocessor cache system may reside, M, E, S, and I, defined as follows:

[0136] Modified (M): The cache line is present only in the current cache, and is dirty; it has been modified from the value in main memory. The cache is required to write the data back to main memory at some time in the future, before permitting any other read of the (no longer valid) main memory state. The write-back changes the line to the Exclusive state.

[0137] Exclusive (E): The cache line is present only in the current cache, but is clean; it matches main memory. It may be changed to the Shared state at any time, in response to a read request. Alternatively, it may be changed to the Modified state when writing to it.

[0138] Shared (S): Indicates that this cache line may be stored in other caches of the machine and is “clean”; it matches the main memory. The line may be discarded (changed to the Invalid state) at any time.

[0139] Invalid (I): Indicates that this cache line is invalid (unused).

[0140] TM coherency status indicators (R 132, W 138) may be provided for each cache line, in addition to, or encoded in the MESI coherency bits. An R 132 indicator indicates the current transaction has read from the data of the cache line, and a W 138 indicator indicates the current transaction has written to the data of the cache line.

[0141] In another aspect of TM design, a system is designed using transactional store buffers. U.S. Pat. No.

6,349,361 titled “Methods and Apparatus for Reordering and Renaming Memory References in a Multiprocessor Computer System,” filed Mar. 31, 2000 and incorporated by reference herein in its entirety, teaches a method for reordering and renaming memory references in a multiprocessor computer system having at least a first and a second processor. The first processor has a first private cache and a first buffer, and the second processor has a second private cache and a second buffer. The method includes the steps of, for each of a plurality of gated store requests received by the first processor to store a datum, exclusively acquiring a cache line that contains the datum by the first private cache, and storing the datum in the first buffer. Upon the first buffer receiving a load request from the first processor to load a particular datum, the particular datum is provided to the first processor from among the data stored in the first buffer based on an in-order sequence of load and store operations. Upon the first cache receiving a load request from the second cache for a given datum, an error condition is indicated and a current state of at least one of the processors is reset to an earlier state when the load request for the given datum corresponds to the data stored in the first buffer.

[0142] The main implementation components of one such transactional memory facility are a transaction-backup register file for holding pre-transaction GR (general register) content, a cache directory to track the cache lines accessed during the transaction, a store cache to buffer stores until the transaction ends, and firmware routines to perform various complex functions. In this section a detailed implementation is described.

IBM zEnterprise EC12 Enterprise Server Embodiment

[0143] The IBM zEnterprise EC12 enterprise server introduces transactional execution (TX) in transactional memory, and is described in part in a paper, “Transactional Memory Architecture and Implementation for IBM System z” of Proceedings Pages 25-36 presented at MICRO-45, 1-5 Dec. 2012, Vancouver, British Columbia, Canada, available from IEEE Computer Society Conference Publishing Services (CPS), which is incorporated by reference herein in its entirety. “IBM,” “zEnterprise,” “System z,” “EC12,” and/or “MICRO-45” may be subject to trademark rights in various jurisdictions throughout the world and are used here only in reference to the products or services properly denominated by the marks to the extent that such trademark rights may exist.)

[0144] Table 3 shows an example transaction. Transactions started with TBEGIN are not assured to ever successfully complete with TEND, since they can experience an aborting condition at every attempted execution, e.g., due to repeating conflicts with other CPUs. This requires that the program support a fallback path to perform the same operation non-transactionally, e.g., by using traditional locking schemes. This puts significant burden on the programming and software verification teams, especially where the fallback path is not automatically generated by a reliable compiler.

TABLE 3-continued

Example Transaction Code			
	JNZ	lckbzy	*branch if lock busy
	... perform operation ...		
	TEND		*end transaction
		
lckbzy	TABORT		*abort if lock busy; this *resumes after TBEGIN
abort	JO	fallback	*no retry if CC=3
	AHI	R0, 1	*increment retry count
	CIJNL	R0,6, fallback	*give up after 6 attempts
	PPA	R0, TX	*random delay based on retry count
	... potentially wait for lock to become free ...		
	J	loop	*jump back to retry fallback
	OBTAIN	lock	*using Compare&Swap
	... perform operation ...		
	RELEASE	lock	
		

[0145] The requirement of providing a fallback path for aborted Transaction Execution (TX) transactions can be onerous. Many transactions operating on shared data structures are expected to be short, touch only a few distinct memory locations, and use simple instructions only. For those transactions, the IBM zEnterprise EC12 introduces the concept of constrained transactions; under normal conditions, the CPU 114a (FIG. 2) assures that constrained transactions eventually end successfully, albeit without giving a strict limit on the number of necessary retries. A constrained transaction starts with a TBEGINC instruction and ends with a regular TEND. Implementing a task as a constrained or non-constrained transaction typically results in very comparable performance, but constrained transactions simplify software development by removing the need for a fallback path. IBM’s Transactional Execution architecture is further described in z/Architecture, Principles of Operation, Tenth Edition, SA22-7832-09 published September 2012 from IBM, incorporated by reference herein in its entirety.

[0146] A constrained transaction starts with the TBEGINC instruction. A transaction initiated with TBEGINC must follow a list of programming constraints; otherwise the program takes a non-filterable constraint-violation interruption. Exemplary constraints may include, but not be limited to: the transaction can execute a maximum of 32 instructions, all instruction text must be within 256 consecutive bytes of memory; the transaction contains only forward-pointing relative branches (i.e., no loops or subroutine calls); the transaction can access a maximum of 4 aligned octowords (an octoword is 32 bytes) of memory; and restriction of the instruction-set to exclude complex instructions like decimal or floating-point operations. The constraints are chosen such that many common operations like doubly linked list-insert/delete operations can be performed, including the very powerful concept of atomic compare-and-swap targeting up to 4 aligned octowords. At the same time, the constraints were chosen conservatively such that future CPU implementations can assure transaction success without needing to adjust the constraints, since that would otherwise lead to software incompatibility.

[0147] TBEGINC mostly behaves like XBEGIN in TSX or TBEGIN on IBM’s zEC12 servers, except that the floating-point register (FPR) control and the program interruption filtering fields do not exist and the controls are considered to be zero. On a transaction abort, the instruction

TABLE 3

Example Transaction Code			
	LHI	R0,0	*initialize retry count=0
loop	TBEGIN		*begin transaction
	JNZ	abort	*go to abort code if CC1=0
	LT	R1, lock	*load and test the fallback lock

address is set back directly to the TBEGINC instead of to the instruction after, reflecting the immediate retry and absence of an abort path for constrained transactions.

[0148] Nested transactions are not allowed within constrained transactions, but if a TBEGINC occurs within a non-constrained transaction it is treated as opening a new non-constrained nesting level just like TBEGIN would. This can occur, e.g., if a non-constrained transaction calls a subroutine that uses a constrained transaction internally.

[0149] Since interruption filtering is implicitly off, all exceptions during a constrained transaction lead to an interruption into the operating system (OS). Eventual successful finishing of the transaction relies on the capability of the OS to page-in the at most 4 pages touched by any constrained transaction. The OS must also ensure time-slices long enough to allow the transaction to complete.

TABLE 4

Transaction Code Example	
TBEGINC	*begin constrained transaction
...	... perform operation ...
TEND	*end transaction

[0150] Table 4 shows the constrained-transactional implementation of the code in Table 3, assuming that the constrained transactions do not interact with other locking-based code. No lock testing is shown therefore, but could be added if constrained transactions and lock-based code were mixed.

[0151] When failure occurs repeatedly, software emulation is performed using millicode as part of system firmware. Advantageously, constrained transactions have desirable properties because of the burden removed from programmers.

[0152] With reference to FIG. 3, the IBM zEnterprise EC12 processor introduced transactional execution facility 200. The processor can decode 3 instructions per clock cycle; simple instructions are dispatched as single micro-ops, and more complex instructions are cracked into multiple micro-ops. Micro-ops (Uops 234a, 234b, and 234c) are written into unified issue queue 216, from where they can be issued out-of-order. Up to two fixed-point, one floating-point, two load/store, and two branch instructions can execute every cycle. Global Completion Table (GCT) 230 holds every micro-op 234a, 234b, and 234c and transaction nesting depth (TND) 232. The GCT 230 is written in-order at decode time, tracks the execution status of each micro-op 234a, 234b, and 234c, and completes instructions when all micro-ops 234a, 234b, and 234c of the oldest instruction group have successfully executed.

[0153] Level 1 (L1) data cache 240 is a 96 KB (kilo-byte) 6-way associative cache with 256 byte cache-lines and 4 cycle use latency, coupled to a private 1 MB (mega-byte) 8-way associative 2nd-level (L2) data cache 268 with 7 cycles use-latency penalty for L1 240 misses. The L1 240 cache is the cache closest to a processor and Ln cache is a cache at the nth level of caching. Both L1 240 and L2 268 caches are store-through. Six cores on each central processor (CP) chip share a 48 MB 3rd-level store-in cache, and six CP chips are connected to an off-chip 384 MB 4th-level cache, packaged together on a glass ceramic multi-chip module (MCM). Up to 4 multi-chip modules (MCMs) can be

connected to a coherent symmetric multi-processor (SMP) system with up to 144 cores (not all cores are available to run customer workload).

[0154] Coherency is managed with a variant of the MESI protocol. Cache-lines can be owned read-only (shared) or exclusive; the L1 240 and L2 268 are store-through and thus do not contain dirty lines. The L3 272 and L4 caches (not shown) are store-in and track dirty states. Each cache is inclusive of all its connected lower level caches.

[0155] Coherency requests are called “cross interrogates” (XI) and are sent hierarchically from higher level to lower-level caches, and between the L4s. When one core misses the L1 240 and L2 268 and requests the cache line from its local L3 272, the L3 272 checks whether it owns the line, and if necessary sends an XI to the currently owning L2 268/L1 240 under that L3 272 to ensure coherency, before it returns the cache line to the requestor. If the request also misses the L3 272, the L3 272 sends a request to the L4 (not shown), which enforces coherency by sending XIs to all necessary L3s under that L4, and to the neighboring L4s. Then the L4 responds to the requesting L3 which forwards the response to the L2 268/L1 240.

[0156] Note that due to the inclusivity rule of the cache hierarchy, sometimes cache lines are XI’ed from lower-level caches due to evictions on higher-level caches caused by associativity overflows from requests to other cache lines. These XIs can be called “LRU XIs”, where LRU stands for least recently used.

[0157] Making reference to yet another type of XI requests, Demote-XIs transition cache-ownership from exclusive into read-only state, and Exclusive-XIs transition cache ownership from exclusive into invalid state. Demote-XIs and Exclusive-XIs need a response back to the XI sender. The target cache can “accept” the XI, or send a “reject” response if it first needs to evict dirty data before accepting the XI. The L1 240/L2 268 caches are store through, but may reject demote-XIs and exclusive XIs if they have stores in their store queues that need to be sent to L3 before downgrading the exclusive state. A rejected XI will be repeated by the sender. Read-only-XIs are sent to caches that own the line read-only; no response is needed for such XIs since they cannot be rejected. The details of the SMP protocol are similar to those described for the IBM z10 by P. Mak, C. Walters, and G. Strait, in “IBM System z10 processor cache subsystem microarchitecture”, IBM Journal of Research and Development, Vol 53:1, 2009, which is incorporated by reference herein in its entirety.

Transactional Instruction Execution

[0158] FIG. 3 depicts example components of an example transactional execution environment, including a CPU and caches/components with which it interacts (such as those depicted in FIGS. 1 and 2). The instruction decode unit 208 (IDU) keeps track of the current transaction nesting depth 212 (TND). When the IDU 208 receives a TBEGIN instruction, the nesting depth 212 is incremented, and conversely decremented on TEND instructions. The nesting depth 212 is written into the GCT 230 for every dispatched instruction. When a TBEGIN or TEND is decoded on a speculative path that later gets flushed, the IDU’s 208 nesting depth 212 is refreshed from the youngest GCT 230 entry that is not flushed. The transactional state is also written into the issue queue 216 for consumption by the execution units, mostly by the Load/Store Unit (LSU) 280, which also has an

effective address calculator **236** is included in the LSU **280**. The TBEGIN instruction may specify a transaction diagnostic block (TDB) for recording status information, should the transaction abort before reaching a TEND instruction. **[0159]** Similar to the nesting depth, the IDU **208**/GCT **230** collaboratively track the access register/floating-point register (AR/FPR) modification masks through the transaction nest; the IDU **208** can place an abort request into the GCT **230** when an AR/FPR-modifying instruction is decoded and the modification mask blocks that. When the instruction becomes next-to-complete, completion is blocked and the transaction aborts. Other restricted instructions are handled similarly, including TBEGIN if decoded while in a constrained transaction, or exceeding the maximum nesting depth.

[0160] An outermost TBEGIN is cracked into multiple micro-ops depending on the GR-Save-Mask; each micro-op **234a**, **234b**, and **234c** (including, for example uop 0, uop 1, and uop2) will be executed by one of the two fixed point units (FXUs) **220** to save a pair of GRs **228** into a special transaction-backup register file **224**, that is used to later restore the GR **228** content in case of a transaction abort. Also the TBEGIN spawns micro-ops **234a**, **234b**, and **234c** to perform an accessibility test for the TDB if one is specified; the address is saved in a special purpose register for later usage in the abort case. At the decoding of an outermost TBEGIN, the instruction address and the instruction text of the TBEGIN are also saved in special purpose registers for a potential abort processing later on.

[0161] TEND and NTSTG are single micro-op **234a**, **234b**, and **234c** instructions; NTSTG (non-transactional store) is handled like a normal store except that it is marked as non-transactional in the issue queue **216** so that the LSU **280** can treat it appropriately. TEND is a no-op at execution time, the ending of the transaction is performed when TEND completes.

[0162] As mentioned, instructions that are within a transaction are marked as such in the issue queue **216**, but otherwise execute mostly unchanged; the LSU **280** performs isolation tracking as described in the next section.

[0163] Since decoding is in-order, and since the IDU **208** keeps track of the current transactional state and writes it into the issue queue **216** along with every instruction from the transaction, execution of TBEGIN, TEND, and instructions before, within, and after the transaction can be performed out of order. It is even possible (though unlikely) that TEND is executed first, then the entire transaction, and lastly the TBEGIN executes. Program order is restored through the GCT **230** at completion time. The length of transactions is not limited by the size of the GCT **230**, since general purpose registers (GRs) **228** can be restored from the backup register file **224**.

[0164] During execution, the program event recording (PER) events are filtered based on the Event Suppression Control, and a PER TEND event is detected if enabled. Similarly, while in transactional mode, a pseudo-random generator may be causing the random aborts as enabled by the Transaction Diagnostics Control.

Tracking for Transactional Isolation

[0165] The Load/Store Unit **280** tracks cache lines that were accessed during transactional execution, and triggers an abort if an XI from another CPU (or an LRU-XI) conflicts with the footprint. If the conflicting XI is an exclusive or

demote XI, the LSU **280** rejects the XI back to the L3 **272** in the hope of finishing the transaction before the L3 **272** repeats the XI. This “stiff-arming” is very efficient in highly contended transactions. In order to prevent hangs when two CPUs stiff-arm each other, a XI-reject counter is implemented, which triggers a transaction abort when a threshold is met.

[0166] The L1 cache directory **240** is traditionally implemented with static random access memories (SRAMs). For the transactional memory implementation, the valid bits **244** (64 rows×6 ways) of the directory have been moved into normal logic latches, and are supplemented with two more bits per cache line: the TX-read **248** and TX-dirty **252** bits.

[0167] The TX-read **248** bits are reset when a new outermost TBEGIN is decoded (which is interlocked against a prior still pending transaction). The TX-read **248** bit is set at execution time by every load instruction that is marked “transactional” in the issue queue. Note that this can lead to over-marking if speculative loads are executed, for example on a mispredicted branch path. The alternative of setting the TX-read **248** bit at load completion time was too expensive for silicon area, since multiple loads can complete at the same time, requiring many read-ports on the load-queue.

[0168] Stores execute the same way as in non-transactional mode, but a transaction mark is placed in the store queue (STQ) **260** entry of the store instruction. At write-back time, when the data from the STQ **260** is written into the L1 **240**, the TX-dirty bit **252** in the L1-directory **256** is set for the written cache line. Store write-back into the L1 **240** occurs only after the store instruction has completed, and at most one store is written back per cycle. Before completion and write-back, loads can access the data from the STQ **260** by means of store-forwarding; after write-back, the CPU **114a** (FIG. 2) can access the speculatively updated data in the L1 **240**. If the transaction ends successfully, the TX-dirty bits **252** of all cache-lines are cleared, and also the TX-marks of not yet written stores are cleared in the STQ **260**, effectively turning the pending stores into normal stores.

[0169] On a transaction abort, all pending transactional stores are invalidated from the STQ **260**, even those already completed. All cache lines that were modified by the transaction in the L1 **240**, that is, have the TX-dirty bit **252** on, have their valid bits turned off, effectively removing them from the L1 **240** cache instantaneously.

[0170] The architecture requires that before completing a new instruction, the isolation of the transaction read- and write-set is maintained. This isolation is ensured by stalling instruction completion at appropriate times when XIs are pending; speculative out-of order execution is allowed, optimistically assuming that the pending XIs are to different addresses and not actually cause a transaction conflict. This design fits very naturally with the XI-vs-completion interlocks that are implemented on prior systems to ensure the strong memory ordering that the architecture requires.

[0171] When the L1 **240** receives an XI, L1 **240** accesses the directory to check validity of the XI’ed address in the L1 **240**, and if the TX-read bit **248** is active on the XI’ed line and the XI is not rejected, the LSU **280** triggers an abort. When a cache line with active TX-read bit **248** is LRU’ed from the L1 **240**, a special LRU-extension vector remembers for each of the 64 rows of the L1 **240** that a TX-read line existed on that row. Since no precise address tracking exists for the LRU extensions, any non-rejected XI that hits a valid

extension row the LSU 280 triggers an abort. Providing the LRU-extension effectively increases the read footprint capability from the L1-size to the L2-size and associativity, provided no conflicts with other CPUs 114a and 114b (FIGS. 1 and 2) against the non-precise LRU-extension tracking causes aborts.

[0172] The store footprint is limited by the store cache size (the store cache is discussed in more detail below) and thus implicitly by the L2 268 size and associativity. No LRU-extension action needs to be performed when a TX-dirty 252 cache line is LRU'ed from the L1 240.

Store Cache

[0173] In prior systems, since the L1 240 and L2 268 are store-through caches, every store instruction causes an L3 272 store access; with now 6 cores per L3 272 and further improved performance of each core, the store rate for the L3 272 (and to a lesser extent for the L2 268) becomes problematic for certain workloads. In order to avoid store queuing delays, a gathering store cache 264 had to be added, that combines stores to neighboring addresses before sending them to the L3 272.

[0174] For transactional memory performance, it is acceptable to invalidate every TX-dirty 252 cache line from the L1 240 on transaction aborts, because the L2 268 cache is very close (7 cycles L1 240 miss penalty) to bring back the clean lines. However, it would be unacceptable for performance (and silicon area for tracking) to have transactional stores write the L2 268 before the transaction ends and then invalidate all dirty L2 268 cache lines on abort (or even worse on the shared L3 272).

[0175] The two problems of store bandwidth and transactional memory store handling can both be addressed with the gathering store cache 264. The cache 264 is a circular queue of 64 entries, each entry holding 128 bytes of data with byte-precise valid bits. In non-transactional operation, when a store is received from the LSU 280, the store cache 264 checks whether an entry exists for the same address, and if so gathers the new store into the existing entry. If no entry exists, a new entry is written into the queue, and if the number of free entries falls under a threshold, the oldest entries are written back to the L2 268 and L3 272 caches.

[0176] When a new outermost transaction begins, all existing entries in the store cache are marked closed so that no new stores can be gathered into them, and eviction of those entries to L2 268 and L3 272 is started. From that point on, the transactional stores coming out of the LSU 280 STQ 260 allocate new entries, or gather into existing transactional entries. The write-back of those stores into L2 268 and L3 272 is blocked, until the transaction ends successfully; at that point subsequent (post-transaction) stores can continue to gather into existing entries, until the next transaction closes those entries again.

[0177] The store cache 264 is queried on every exclusive or demote XI, and causes an XI reject if the XI compares to any active entry. If the core is not completing further instructions while continuously rejecting XIs, the transaction is aborted at a certain threshold to avoid hangs.

[0178] The LSU 280 requests a transaction abort when the store cache 264 overflows. The LSU 280 detects this condition when it tries to send a new store that cannot merge into an existing entry, and the entire store cache 264 is filled with stores from the current transaction. The store cache 264 is managed as a subset of the L2 268: while transactionally

dirty lines can be evicted from the L1 240, they have to stay resident in the L2 268 throughout the transaction. The maximum store footprint is thus limited to the store cache size of 64×128 bytes, and it is also limited by the associativity of the L2 268. Since the L2 268 is 8-way associative and has 512 rows, it is typically large enough to not cause transaction aborts.

[0179] If a transaction aborts, the store cache 264 is notified and all entries holding transactional data are invalidated. The store cache 264 also has a mark per doubleword (8 bytes) whether the entry was written by a NTSTG instruction—those doublewords stay valid across transaction aborts.

Millicode-Implemented Functions

[0180] Traditionally, IBM mainframe server processors contain a layer of firmware called millicode which performs complex functions like certain CISC instruction executions, interruption handling, system synchronization, and RAS. Millicode includes machine dependent instructions as well as instructions of the instruction set architecture (ISA) that are fetched and executed from memory similarly to instructions of application programs and the operating system (OS). Firmware resides in a restricted area of main memory that customer programs cannot access. When hardware detects a situation that needs to invoke millicode, the instruction fetching unit 204 switches into “millicode mode” and starts fetching at the appropriate location in the millicode memory area. Millicode may be fetched and executed in the same way as instructions of the instruction set architecture (ISA), and may include ISA instructions.

[0181] For transactional memory, millicode is involved in various complex situations. Every transaction abort invokes a dedicated millicode subroutine to perform the necessary abort steps. The transaction-abort millicode starts by reading special-purpose registers (SPRs) holding the hardware internal abort reason, potential exception reasons, and the aborted instruction address, which millicode then uses to store a TDB if one is specified. The TBEGIN instruction text is loaded from an SPR to obtain the GR-save-mask, which is needed for millicode to know which GRs 238 to restore.

[0182] The CPU 114a (FIG. 2) supports a special millicode-only instruction to read out the backup-GRs 224 and copy them into the main GRs 228. The TBEGIN instruction address is also loaded from an SPR to set the new instruction address in the PSW to continue execution after the TBEGIN once the millicode abort subroutine finishes. That PSW may later be saved as program-old PSW in case the abort is caused by a non-filtered program interruption.

[0183] The TABORT instruction may be millicode implemented; when the IDU 208 decodes TABORT, it instructs the instruction fetch unit to branch into TABORT's millicode, from which millicode branches into the common abort subroutine.

[0184] The Extract Transaction Nesting Depth (ETND) instruction may also be millicoded, since it is not performance critical; millicode loads the current nesting depth out of a special hardware register and places it into a GR 228. The PPA instruction is millicoded; it performs the optimal delay based on the current abort count provided by software as an operand to PPA, and also based on other hardware internal state.

[0185] For constrained transactions, millicode may keep track of the number of aborts. The counter is reset to 0 on

successful TEND completion, or if an interruption into the OS occurs (since it is not known if or when the OS will return to the program). Depending on the current abort count, millicode can invoke certain mechanisms to improve the chance of success for the subsequent transaction retry. The mechanisms involve, for example, successively increasing random delays between retries, and reducing the amount of speculative execution to avoid encountering aborts caused by speculative accesses to data that the transaction is not actually using. As a last resort, millicode can broadcast to other CPUs other than 114a which is processing the local transaction, to stop all conflicting work and retry the local transaction before releasing the other CPUs to continue normal processing. Where multiple CPUs are enabled, their activity must be coordinated to not cause deadlocks, so some serialization between millicode instances on different CPUs 114 is required.

B. Computer Program Product Claim Support

[0186] A computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punchcards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0187] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0188] Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or

either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++ or the like, and conventional procedural programming languages, such as the “C” programming language or similar programming languages. The computer readable program instructions may execute entirely on the user’s computer, partly on the user’s computer, as a stand-alone software package, partly on the user’s computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user’s computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[0189] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0190] These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0191] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0192] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, seg-

ment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

C. Computer Program Product Claim Support

[0193] An embodiment of a possible hardware and software environment for software and/or methods according to the present invention will now be described in detail with reference to the Figures. FIG. 4 depicts a block diagram of components of a computing device 400, in accordance with an illustrative embodiment of the present invention. It should be appreciated that FIG. 4 provides only an illustration of one implementation and does not imply any limitations with regard to the environments in which different embodiments may be implemented. Many modifications to the depicted environment may be made. It should be appreciated FIG. 4 provides only an illustration of one implementation and does not imply any limitations with regard to the environments in which different embodiments may be implemented.

[0194] The computing environment of FIG. 4 is, in many respects, representative of the various computer subsystem (s) in the present invention. Accordingly, several portions of the computing environment will now be discussed in the following paragraphs.

[0195] Computing device 400 includes communications fabric 402, which provides communications between computer processor(s) 404, memory 406, persistent storage 408, communications unit 410, and input/output (I/O) interface(s) 412. Communications fabric 402 can be implemented with any architecture designed for passing data and/or control information between processors (such as microprocessors, communications and network processors, etc.), system memory, peripheral devices, and any additional hardware components within a system. For example, communications fabric 402 can be implemented with one or more buses.

[0196] Computing device 400 is capable of communicating with other computer subsystems via network 401. Network 401 can be, for example, a local area network (LAN), a wide area network (WAN) such as the Internet, or a combination of the two, and can include wired, wireless, or fiber optic connections. In general, network 401 can be any combination of connections and protocols that will support communications between computing device 400 and other computing devices.

[0197] Memory 406 and persistent storage 408 are computer-readable storage media. In one embodiment, memory 406 includes random access memory (RAM) and cache memory 414. In general, memory 406 can include any suitable volatile or non-volatile computer-readable storage media.

[0198] In some embodiments, regulation logic 420 may be stored for execution by one or more of the respective

computer processors 404 of computing device 400 via one or more memories of memory 406 of computing device 400. In the depicted embodiment, persistent storage 408 includes a magnetic hard disk drive. Alternatively, or in addition to a magnetic hard disk drive, persistent storage 408 can include a solid state hard drive, a semiconductor storage device, read-only memory (ROM), erasable programmable read-only memory (EPROM), flash memory, or any other computer-readable storage media that is capable of storing program instructions or digital information. In some embodiments, regulation logic 420 may be implemented using logic gates.

[0199] The media used by persistent storage 408 may also be removable. For example, a removable hard drive may be used for persistent storage 408. Other examples include optical and magnetic disks, thumb drives, and smart cards that are inserted into a drive for transfer onto another computer-readable storage medium that is also part of persistent storage 408.

[0200] Communications unit 410, in the examples, provides for communications with other data processing systems or devices, including computing device 400. In the examples, communications unit 410 includes one or more network interface cards. Communications unit 410 may provide communications through the use of either or both physical and wireless communications links.

[0201] I/O interface(s) 412 allows for input and output of data with other devices that may be connected to computing device 400. For example, I/O interface 412 may provide a connection to external devices 416 such as a keyboard, keypad, camera, a touch screen, and/or some other suitable input device. External devices 416 can also include portable computer-readable storage media such as, for example, thumb drives, portable optical or magnetic disks, and memory cards. Software and data used to practice embodiments of the present invention, e.g., regulation logic 420 can be stored on such portable computer-readable storage media and can be loaded onto persistent storage 408 of computing device 400 via I/O interface(s) 412 of computing device 400. It should be noted that, in some embodiments, regulation logic 420 is implemented as a hardware module.

[0202] Display 418 provides a mechanism to display data to a user and may be, for example, a computer monitor.

[0203] The logic described herein is identified based upon the application for which it is implemented in a specific embodiment of the invention. However, it should be appreciated that any particular logic nomenclature herein is used merely for convenience, and thus the invention should not be limited to use solely in any specific application identified and/or implied by such nomenclature.

[0204] Regulation logic 420 detects the cause of the termination of the transaction and selects the parameters for the retrying of the operation. A transaction is a group of instructions that operate in an atomic manner on a data structure in memory, as viewed by other CPUs and the I/O subsystem. For a transaction to be complete the changes need to be finalized and made permanent in their entirety. The processing of a transaction can either be successful or fail, it cannot be partially completed. In additional embodiments, a transaction is an individual, or indivisible, operation which is part of a larger operation. One event which can cause the termination of the transaction before completion is a transient condition. A transient condition, or transient property of an element of the system, is one which is

temporary. Transient conditions may be, for example asynchronous interruptions, another CPU trying to access memory used within the transaction, or another thread on the same CPU causing a cache line to be evicted using an algorithm of least recently used (LRU). Regulation logic 420 detects the transient condition or element, or another cause of the transaction being terminated. In additional embodiments, regulation logic 420 records the cause of the transaction being terminated in a repository, such as, for example, memory 406, persistent storage 408, or as an internal hardware logic state.

[0205] Regulation logic 420 also controls the procedure performed by computing device 400 once regulation logic 420 determines the cause of the premature termination of the transaction. Regulation logic 420 decides if the cause of the premature termination of the transaction can be fixed with a series of retries of the transaction, or other methods of allowing the transaction more attempts to be completed. In one embodiment, regulation logic 420 permits the transaction a predetermined number of retries to be completed. The predetermined number of retries does not guarantee a successful transaction, but can be used for assistance in future transactions to increase the speed of the transaction or to anticipate failures.

[0206] In additional embodiments, regulation logic 420 can permit the transaction to retry until the transaction is successful. In a portion of the additional embodiments, regulation logic 420 may record information related to the transaction, this information can be, for example, the cause of the premature termination, the number of retries, and the successful transaction. This information can potentially be used in future transactions which fail to find a known solution and decrease the time for future transaction to be successful. The programs described herein are identified based upon the application for which they are implemented in a specific embodiment of the invention. However, it should be appreciated that any particular logic nomenclature herein is used merely for convenience, and thus the invention should not be limited to use solely in any specific application identified and/or implied by such nomenclature.

II. Example Embodiment

[0207] FIG. 5 shows flowchart 500 depicting an approach according to the present invention. FIG. 6 shows regulation logic 420 for performing at least some of the steps of flowchart 500. This approach will now be discussed, over the course of the following paragraphs, with extensive reference to FIG. 5 and FIG. 6.

[0208] Processing begins at step 502, where transaction module (“mod”) 602 receives a begin-transaction instruction (e.g., T_BEGIN, TXBEGIN, or XBEGIN among others) indicating the beginning of a transaction. The transaction comprises the begin-transaction instruction and all the instructions immediately following the begin-transaction program instruction up to, and including an end-transaction instruction (e.g., T_END, TXEND, or XEND among others) that corresponds to the begin-transaction instruction.

[0209] Processing proceeds to step 504, where initialization mod 604 carries out initialization operations necessary to enable processing of the transaction to begin. Initialization operations may include, but are not limited to: (i) setting a transaction indicator to indicate a transaction is currently being processed; (ii) setting a metric counter to zero; (iii) preparing for a rollback situation by keeping a backup copy

of the processor state at the beginning of transaction processing; and/or (iv) invoking flag processing mod 608 to indicate the end-of-transaction has not been detected.

[0210] A transaction indicator may be implemented in software, hardware, or a combination of the two. In some embodiments of the present invention, the transaction indicator is a software implementation using a Boolean flag, wherein a zero value indicates that no transaction is being processed, and a one value indicates that a transaction is being processed. Additionally, in some embodiments of the present invention, the transaction indicator is a hardware implementation and a status register is used to indicate whether or not a transaction is being processed. Further, in some embodiments of the present invention, nested transactions are supported, and the transaction indicator is a counter that is incremented each time a begin-transaction is encountered, and decremented each time an end-transaction is encountered, thus, indicating all nested transactions are complete when the value in the transaction indicator reaches zero.

[0211] A metric counter may be used by the instruction processor to determine how far the current instruction is from the end of the transaction (e.g., how many instructions remain in the transaction). When a begin-transaction instruction is encountered, a metric counter corresponding to the current transaction is initialized to zero. In some embodiments of the present invention, the metric counter is incremented once for each instruction identified during a speculative look ahead operation, and the metric counter is frozen when an end-transaction instruction is encountered during the speculative look ahead operation (i.e., the metric counter will contain the total number of instructions included in the transaction corresponding to the metric counter). In some embodiments, the metric counter is decremented for each completed instruction within the transaction.

[0212] In the event that a transaction is unable to successfully complete, the processor shall be able to perform a rollback (e.g., restore) of the environment (e.g., transaction memory, registers, variables, and the like) to a state corresponding to the environment at the beginning of the transaction operation. In some embodiments of the present invention, a rollback is necessary when another process causes interference by attempting to access the transactional memory corresponding to the transaction. If the interference causes the transaction to halt (also referred to herein as an abort of the transaction) without reaching the end-transaction instruction, all processing performed during the transaction is discarded and a rollback operation is performed to restore the environment to the state equivalent to that at the start of the transaction.

[0213] When an environment utilizes speculative look ahead (i.e., when the environment fetches and decodes program instructions prior to execution), an indicator (e.g., TXEND_INSIGHT) may be maintained to determine whether an end-of-transaction has been detected during the speculative look ahead operation. When a begin-transaction instruction is processed, the indicator (e.g., TXEND_INSIGHT) may be initialized to indicate that no end-transaction has been detected.

[0214] Processing proceeds to step 506, where flag processing mod 608 receives a request to update an indicator (e.g., a flag) to contain a specified value. The received request may identify a flag and the operation that is to be performed on the flag. In some embodiments of the present

invention, the flag is TXEND_INSIGHT, and the operation is indicating no end-transaction instruction has been detected. Additionally, in some embodiments of the present invention, additional operations are carried out, such as, when TXEND_INSIGHT is changed to indicate that no end-transaction instruction has been detected, the metric counter is unfrozen to enable the metric counter to be incremented while the speculative look ahead operation proceeds.

[0215] Processing proceeds to step 508, where monitor mod 610 monitors the speculative look ahead operation and detect when an end-transaction instruction is encountered. Speculative look ahead is an operation that predicts an execution path that is likely to be followed in the future, however, it is possible that the prediction is incorrect and the predicted execution path is not actually followed. Predicting an execution path that is not actually followed may occur when the speculative look ahead operation encounters a branch instruction in the instruction stream. Since the information that determines the actual behavior of the branch instruction may not yet be available, the speculative look ahead operation may predict the behavior of the branch instruction based on previous behavior (e.g., including but not limited to local and global branch history, branch address history, branch history vectors and/or a return address stack), static branch prediction information encoded in the instruction, branch policies tangibly incorporated into the processor (e.g., backward branches, i.e., branches to lower addresses than the address of a branch instruction may be predicted as “taken”), meta data, or the like.

[0216] Instead of waiting for the information that determines the actual behavior of the branch instruction to become available, the speculative look ahead operation proceeds on a predicted execution path. In some embodiments of the present invention, a backup copy of the metric counter is maintained each time a branch instruction is encountered. If it is later determined that an incorrect execution path was predicted for the branch, the metric counter can be restored to the value it was when the branch instruction was being speculatively looked at, and the speculative look ahead operation can resume following a different execution path. In some embodiments of the present invention, instruction decode is also performed speculatively ahead of instruction execution and at least one decoupling queue (e.g., dispatch queue, issue queue, reorder queue, and the like) is available to store speculatively decoded instructions after the instruction decode operation and prior to instruction execution.

[0217] If the speculative look ahead operation encounters an end-transaction instruction, monitor mod 610 performs tasks to indicate that an end-transaction has been encountered (i.e., an end-transaction has been detected). Tasks to be performed include, but are not limited to: (i) freezing the metric counter; and/or (ii) notifying flag processing mod 608 that that an indicator (e.g., a flag) such as TXEND_INSIGHT is to be set to indicate an end-transaction instruction has been encountered.

[0218] If an end-transaction instruction has been identified, and it turns out that the predicted execution path is incorrect, then the effects of detecting the end-transaction instruction shall be rolled back to the point at which the incorrect branch prediction was determined. The rollback operation may include restoring the speculative look ahead operation to the point at which the incorrect branch predic-

tion was determined. The tasks included in the rollback operation may include, but are not limited to: (i) unfreezing the metric counter; (ii) restoring (e.g., rolling back) the metric counter to the state (i.e., value) of the metric counter when the speculative look ahead operation was looking at the branch instruction; and/or (iii) notifying flag processing mod 608 that that an indicator (e.g., a flag) such as TXEND_INSIGHT is to be set to indicate that no end-transaction instruction has been encountered.

[0219] Processing proceeds to step 510, where flag processing mod 608 receives a request to update an indicator (e.g., a flag) to contain a specified value. The received request may identify a flag and an operation that is to be performed on the flag. In some embodiments of the present invention, the flag is TXEND_INSIGHT, and the operation is indicating that an end-transaction instruction has been detected. Additionally, in some embodiments of the present invention, additional operations are carried out, such as when TXEND_INSIGHT is altered to indicate that an end-transaction instruction has been detected, the metric counter is frozen to allow calculations to determine how close the current instruction is to the end-transaction instruction.

[0220] Processing proceeds to step 512, where interference mod 612 receives a notification identifying another process that is attempting to access transactional memory corresponding to an active transaction. In accordance with one embodiment of the present invention, interference is detected in conjunction with the tracking of read and write sets of memory that has been the subject of accesses of the present transaction, and further explained with reference to FIGS. 1, 2, and 3 herein. Interference mod 612 may analyze the circumstances corresponding to the interference. The analysis may include, but is not limited to: (i) determining if delaying the requested halt of the transaction is an option (it should be noted that the terms “halt” and “abort” are used interchangeably herein and have the same meaning); (ii) determining if the transaction is near the end (i.e., an end-transaction instruction has been encountered by the speculative look ahead operation); and/or (iii) if both processes are transactions, determining which transaction is closer to completion.

[0221] Processing ends with step 514, where determination mod 614 determines an appropriate action to take with regard to the current transaction. Possible actions may include, but are not limited to: (i) delay the halt request, and continue processing the transaction; (ii) halt the transaction and perform a rollback operation; and/or (iii) request that the interfering transaction halt execution.

[0222] Delaying the halt request may include the interfering process waiting for the transaction to complete before the interfering process can obtain the requested data. In this example embodiment of the present invention: (i) a transaction is being processed; (ii) an interference occurs; and (iii) the speculative look ahead operation has encountered an end-transaction instruction. Determination mod 614 determines that the transaction will continue to be executed, however, it is later determined that the encountered end-transaction instruction was in an incorrect execution path, and therefore no end-transaction instruction has been encountered, resulting in the transaction being halted.

[0223] When a transaction halts, no additional instructions corresponding to the transaction are run, and a rollback operation may be performed so the processing environment appears as if the transaction had never begun processing. In

some embodiments of the present invention, whenever a transaction ends (e.g., halts or runs to completion), a transaction indicator is updated to indicate the transaction is no longer being processed, freeing the processor from any limitations in place during transaction processing.

[0224] In some embodiments of the current invention, the interfering transaction is of a lower priority, and the interfering transaction halts, allowing the current transaction to continue processing. In some embodiments of the present invention, the interfering transaction is operating with a higher priority, and therefore the current transaction halts. Additionally, in some embodiments of the present invention, no end-transaction instruction has been encountered by the speculative look ahead operation, and therefore the current transaction halts.

[0225] As noted above, interference mod 612 receives access requests from remote processes and determines interference. When no interference is detected, and a local processor includes the requested data, a response with the requested data is provided. In some embodiments of the present invention, the interference module provides the requested data as well as an indication that a present transaction of the present processor is to be aborted (action (ii), discussed above).

[0226] In some embodiments of the present invention, a decision is made as to whether to provide the data immediately (and cause a transaction abort responsive to such determination), or to defer a response (action (i), above). When a response is deferred, at a later time, responsive to completing the present transaction interference mod 612 provides the data corresponding to a deferred response. Further, when the possibility of a deadlock has been detected, and a present transaction is to be aborted, interference module is notified to provide the data in conjunction with the initiation of a transaction abort.

[0227] It should be noted as described above that when an abort operation is delayed, or held, both the instruction to abort as well as the data of the interfering transaction are withheld from processing.

III. Further Comments and/or Embodiments

[0228] Some embodiments of the present invention recognize the following facts, potential problems, and/or potential areas for improvement with respect to the current state of the art: (i) when deciding to make decisions about interference between at least two transactions, and deciding which instruction to abort, it is desirable to know which transaction will be ending shortly (e.g., whether holding off on responding to a transaction may prevent an interference); (ii) when an interference must cause one transaction to abort, it may be desirable to halt the transaction that is not close to completing (e.g., do not sacrifice a transaction that would only require a few instructions to be successfully completed); and/or (iii) it is desirable to offer a compatible way of indicating the impending end of a transaction that allows computer code and computer architecture to remain backwardly compatible.

[0229] FIG. 7 depicts a flowchart of the operational steps of regulation logic 420, within the computing environment of FIG. 4, in accordance with one embodiment of the present invention. Flowchart 700 depicts the steps taken by regulation logic 420 to control the steps taken by the hardware or software. It should be appreciated that FIG. 7 provides only an illustration of one implementation and does not imply any

limitations with regard to the environments in which different embodiments may be implemented. Many modifications to the depicted environment may be made.

[0230] In step 702, regulation logic 420 detects transaction execution. This step is carried out in the process described above in FIG. 5.

[0231] In step 704, regulation logic 420 performs TX initialization. This step is described above in reference to FIG. 2 and FIG. 5.

[0232] In step 706, regulation logic 420 executes the transaction. This step is carried out in the process described above in FIG. 5.

[0233] In decision 708, regulation logic 420 determines if the transaction was aborted. Regulation logic 420 detects the transaction has aborted as described above. An abort may be caused by, for example, a halt operation, also referred to as an abort operation, which, as used herein, refers to an operation responsive to a condition where two transactions have been detected to interfere and where at least one transaction must be aborted and the state of the processor is reset to the state at the beginning of the aborted transaction by performing a rollback. If regulation logic 420 determines the transaction was aborted (yes branch, proceed to decision 710), regulation logic 420 determines if the cause of the abort was due to a transient condition. If regulation logic 420 determines the transaction was not aborted (no branch, proceed to step 712), regulation logic 420 completes the transaction.

[0234] In decision 710, regulation logic 420 determines if the cause of the abort was due to a transient condition. A transient condition is a condition which may show up during one invocation of a transaction but not in a following transaction. Such conditions may be, for example, interference from another CPU, an asynchronous interruption, another thread LRUing an entry out of the L1 cache with TX_read or TX_dirty bits set. Non-transient conditions that may cause an abortion may include, for example, issuing an illegal instruction or, when running single threaded, exceeding a transaction footprint. If regulation logic 420 determines the cause of the termination is due to a transient condition (yes branch, proceed to step 716), regulation logic 420 selects the retry parameters. If regulation logic 420 determines the cause of the termination is not due to a transient condition (no branch, proceed to step 714), regulation logic 420 aborts the operation.

[0235] In step 712, regulation logic 420 completes the transaction. Regulation logic 420 completes the execution of the transaction till the transaction has been completed.

[0236] In step 714, regulation logic 420 aborts the operation. Regulation logic 420 aborts the operation because the cause of the error is beyond a transient condition or element, or the transaction cannot be completed for other reasons. This can be due to, for example, a loss of information necessary for the transaction, an invalid instruction, a footprint overflow, a transaction taking too long to execute, a hardware failure, or other mechanical or technical issues which can arise which would not be related to a transient condition or element.

[0237] In step 716, regulation logic 420 selects the retry parameters. This step selects what actions regulation logic 420 permits the hardware of FIG. 4 to perform to attempt to resolve the transient condition. FIG. 8 explains the step of selecting the retry parameters in greater detail.

[0238] In step 718, regulation logic 420 performs the permitted action. Regulation logic 420 performs the actions approved, and selected in the process performed in FIG. 7. Regulation logic 420 may, for example, automatically perform these actions once the permitted actions are selected, or regulation logic 420 may wait a predetermined time before starting the permitted action. In additional embodiments, regulation logic 420 may permit numerous permitted actions.

[0239] In step 720, regulation logic 420 resets to initial transaction state. In some embodiments, regulation logic 420 determines if the permitted action allowed the transaction to succeed. Regulation logic 420 determines if the transaction succeeded by detecting if the transaction was able to be performed fully to completion, without interruption or issues. If regulation logic 420 determines the permitted action allowed the transaction to succeed (yes branch, proceed to END), regulation logic 420 concludes the operation. If regulation logic 420 determines the permitted action causes the transaction to fail (no branch, proceed to 204), regulation logic 420 reselects the retry parameters. In an attempt to retry the transaction with a new set of parameters to attempt to allow the transaction to succeed.

[0240] FIG. 8 depicts a flowchart of the operational steps of regulation logic 420 operating within the computing environment of FIG. 1, in accordance with one embodiment of the present invention. Flowchart 800 depicts the steps taken by regulation logic 420 to structure the limitations of the actions performed by the computing device. It should be appreciated that FIG. 8 provides only an illustration of one implementation and does not imply any limitations with regard to the environments in which different embodiments may be implemented.

[0241] In decision 802, regulation logic 420 determines if retry is permitted. Regulation logic 420 may, in some instances, retry the transaction in order to resolve the conflict. The retry of the transaction can be, for example, because regulation logic 420 has solved the conflict before, the transient condition is down to regulation logic 420, or retrying the transaction can potentially fix the transient condition. If regulation logic 420 determines that the retry is permitted (yes branch, proceed to step 804), regulation logic 420 sets the parameters for the retry. If regulation logic 420 determines the retry is not permitted (no branch, proceed to step 714), regulation logic 420 aborts the operation.

[0242] In step 804, regulation logic 420 sets parameter. Regulation logic 420 sets the parameter which computing device 400 performs in an attempt to have the transaction succeed. These parameters can be, for example, number of retries which computing device 400 performs in an attempt to allow the transaction to succeed, time limit which computing device 400 is allowed to perform retries, or the action which is performed by each retry. In additional embodiments there is a predetermined delay in responding to coherency requests. In additional embodiments, regulation logic 420 retries the transaction in the same manner as the first instance which failed, or regulation logic 420 retries the transaction in a different manner than the first instance which failed. In additional embodiments, regulation logic 420 may set the parameters to record each retry, the actions performed by computing device 400, or if the retry was successful or a failure.

[0243] In decision 806, regulation logic 420 determines if occurrence is recorded. In certain instances where a trans-

action is terminated and it is known that the termination is caused by a transient condition, regulation logic 420 records information regarding the instance, which can be, for example, the number of retries, the transient condition, if the retries solved the transient condition, or the solution to the transient condition for future use. If regulation logic 420 determines the information regarding the occurrence is to be recorded (yes branch, proceed to step 718), regulation logic 420 records the information regarding the occurrence (see step 808). If regulation logic 420 determines the information regarding the occurrence is to not be recorded (no branch, proceed to step 714), regulation logic 420 proceeds to step 718 (see FIG. 7).

[0244] The present invention may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

[0245] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0246] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0247] Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or

either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++ or the like, and conventional procedural programming languages, such as the “C” programming language or similar programming languages. The computer readable program instructions may execute entirely on the user’s computer, partly on the user’s computer, as a stand-alone software package, partly on the user’s computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user’s computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, to perform aspects of the present invention.

[0248] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0249] These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0250] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0251] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, seg-

ment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

IV. Definitions

[0252] Present invention: should not be taken as an absolute indication that the subject matter described by the term “present invention” is covered by either the claims as they are filed, or by the claims that may eventually issue after patent prosecution; while the term “present invention” is used to help the reader to get a general feel for which disclosures herein that are believed as maybe being new, this understanding, as indicated by use of the term “present invention,” is tentative and provisional and subject to change over the course of patent prosecution as relevant information is developed and as the claims are potentially amended.

[0253] Embodiment: see definition of “present invention” above—similar cautions apply to the term “embodiment.”

[0254] and/or: inclusive or; for example, A, B “and/or” C means that at least one of A or B or C is true and applicable.

[0255] User/subscriber: includes, but is not necessarily limited to, the following: (i) a single individual human; (ii) an artificial intelligence entity with sufficient intelligence to act as a user or subscriber; and/or (iii) a group of related users or subscribers.

[0256] Module/Sub-Module: any set of hardware, firmware and/or software that operatively works to do some kind of function, without regard to whether the module is: (i) in a single local proximity; (ii) distributed over a wide area; (iii) in a single proximity within a larger piece of software code; (iv) located within a single piece of software code; (v) located in a single storage device, memory or medium; (vi) mechanically connected; (vii) electrically connected; and/or (viii) connected in data communication.

[0257] Computer: any device with significant data processing and/or machine readable instruction reading capabilities including, but not limited to: desktop computers, mainframe computers, laptop computers, field-programmable gate array (FPGA) based devices, smart phones, personal digital assistants (PDAs), body-mounted or inserted computers, embedded device style computers, application-specific integrated circuit (ASIC) based devices.

What is claimed is:

1. A method for resolving terminated transactions in a transactional memory environment, the method comprising:
 - initiating a hardware transaction in a computing environment, wherein the hardware transaction accesses a memory location, and wherein the hardware transaction includes a transaction begin indicator and a transaction end indicator;
 - detecting a conflicting access of the memory location while executing the hardware transaction;

aborting the hardware transaction based on the conflicting access of the memory location;
determining, by hardware, that the conflicting access of the memory location is a transient condition; and
reinitiating the hardware transaction.

2. The method of claim 1, further comprising:
generating, by software, a retry parameter based on, at least, the conflicting access of the memory, wherein the retry parameter indicates, to hardware, a maximum number of attempts to reinitiate the hardware transaction prior to aborting the hardware transaction.

3. The method of claim 2, further comprising:
determining, by hardware, that the retry parameter has not been met; and
wherein reinitiating the hardware transaction is based on the determination that the retry parameter has not been met.

4. The method of claim 1, wherein the transient condition is a condition which may cause conflicting access of the memory location during a first invocation of the hardware transaction, but which may not result in conflicting access of the memory location during a subsequent invocation of the hardware transaction.

5. The method of claim 1, wherein:
the conflicting access of the memory location is an intervening store to the memory location in a near-end transaction processing mode, wherein the near-end-transaction processing mode indicates that the conflicting access is near end of completion based on, at least, the transaction end indicator; and
aborting the hardware transaction is further based on the near-end transaction processing mode of the conflicting access of the memory location.

6. The method of claim 1, wherein:
the step of initiating the hardware transaction in the computing environment is performed by a first processor; and
the conflicting access of the memory location is an intervening store to the memory location by a second processor.

7. The method of claim 1, further comprising:
logging, by hardware, information about the reinitiated hardware transaction including, at least, the initiated hardware transaction attempt number and information about the conflicting access.

* * * * *