

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization

International Bureau



(10) International Publication Number

WO 2013/048826 A1

(43) International Publication Date  
4 April 2013 (04.04.2013)

(51) International Patent Classification:  
*G06F 15/167* (2006.01)

(21) International Application Number:

PCT/US2012/055942

(22) International Filing Date:

18 September 2012 (18.09.2012)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/541,051 29 September 2011 (29.09.2011) US  
13/414,593 7 March 2012 (07.03.2012) US

(71) Applicant: ORACLE INTERNATIONAL CORPORATION [US/US]; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US).

(72) Inventors: SHEN, Xugang; 11-3-401, No. 30 Ma Jia Pu Xi Lu, Fengtai District, Beijing (CN). LI, Xiangdong; 9-5-1, Changchunjiexili, Xuanwu District, Beijing (CN).

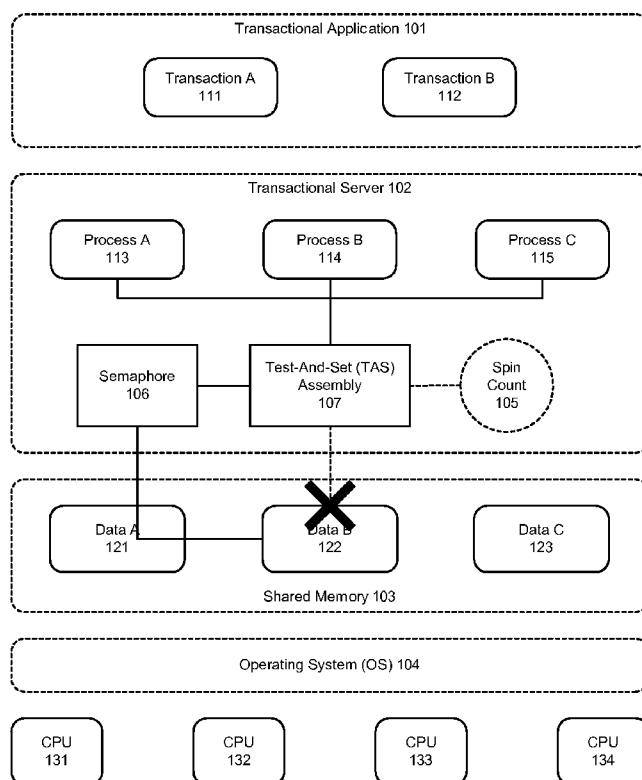
(74) Agents: MEYER, Sheldon, R. et al.; Fliesler Meyer LLP, 650 California Street, Fourteenth Floor, San Francisco, California 94108 (US).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR SUPPORTING A SELF-TUNING LOCKING MECHANISM IN A TRANSACTIONAL MIDDLEWARE MACHINE ENVIRONMENT



(57) Abstract: A lock mechanism can be supported in a transactional middleware system to protect transaction data in a shared memory when there are concurrent transactions. The transactional middleware machine environment comprises a semaphore provided by an operating system running on a plurality of processors. The plurality of processors operates to access data in the shared memory. The transactional middleware machine environment also comprises a test-and-set (TAS) assembly component that is associated with one or more processes. Each said process operates to use the TAS assembly component to perform one or more TAS operations in order to obtain a lock for data in the shared memory. Additionally, a process operates to be blocked on the semaphore and waits for a release of a lock on data in the shared memory, after the TAS component has performed a number of TAS operations and failed to obtain the lock.

WO 2013/048826 A1

FIGURE 1



MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, **Published:**  
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, — *with international search report (Art. 21(3))*  
GW, ML, MR, NE, SN, TD, TG).

## SYSTEM AND METHOD FOR SUPPORTING A SELF-TUNING LOCKING MECHANISM IN A TRANSACTIONAL MIDDLEWARE MACHINE ENVIRONMENT

### Copyright Notice:

5 [0001] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

10 **Field of Invention:**

[0003] The present invention is generally related to computer systems and software such as middleware, and is particularly related to supporting a transactional middleware machine environment.

15 **Background:**

[0004] A transactional middleware system, or a transaction oriented middleware, includes enterprise application servers that can process various transactions within an organization. With the developments in new technologies such as high performance network and multiprocessor computers, there is a need to further improve the performance of the transactional middleware.

20 These are the generally areas that embodiments of the invention are intended to address.

### Summary:

[0005] Described herein is a lock mechanism that can be supported in a transactional middleware system for protecting transaction data in a shared memory when there are concurrent transactions. The transactional middleware machine environment includes an operating system running on a plurality of processors, each of which operates to access data in a shared memory. The transactional middleware machine environment also comprises a semaphore that is provided by the operating system, and a test-and-set (TAS) assembly component that is associated with one or more processes. Each process can use the TAS assembly component to perform one or more TAS operations in order to obtain a lock for data in the shared memory. Furthermore, a process can block the semaphore and wait for a release of a lock on data in the shared memory, after the TAS component has failed to obtain the lock after performing a specified number of TAS operations.

35 **Brief Description of the Figures:**

[0006] **Figure 1** shows an illustration of a transactional middleware machine environment that supports a self-tuning lock mechanism, in accordance with an embodiment of the invention.

[0007] **Figure 2** illustrates an exemplary flow chart for supporting a self-tuning lock

mechanism in a transactional middleware machine environment, in accordance with an embodiment of the invention.

**[0008]** **Figure 3** is a functional block diagram of a system for supporting a locking mechanism in a transactional middleware machine environment, in accordance with an embodiment of the invention.

5

**Detailed Description:**

**[0009]** Described herein is a system and method for supporting a transactional middleware system, such as Tuxedo, that can take advantage of fast machines with multiple processors, and 10 a high performance network connection. A lock mechanism can be supported in a transactional middleware system to protect transaction data in a shared memory when there are concurrent transactions. The transactional middleware machine environment comprises a semaphore provided by an operating system running on a plurality of processors. The plurality of processors can access data in the shared memory. The transactional middleware machine environment also 15 comprises a test-and-set (TAS) assembly component that is associated with one or more processes. Each said process operates to use the TAS assembly component to perform one or more TAS operations in order to obtain a lock for data in the shared memory. Additionally, a process operates to be blocked on the semaphore and wait for a release of a lock on data in the shared memory, after the TAS component has performed a number of TAS operations and failed 20 to obtain the lock.

**[00010]** In accordance with an embodiment of the invention, the system comprises a combination of high performance hardware, e.g. 64-bit processor technology, high performance large memory, and redundant InfiniBand and Ethernet networking, together with an application server or middleware environment, such as WebLogic Suite, to provide a complete Java EE 25 application server complex which includes a massively parallel in-memory grid, that can be provisioned quickly, and can scale on demand. In accordance with an embodiment, the system can be deployed as a full, half, or quarter rack, or other configuration, that provides an application server grid, storage area network, and InfiniBand (IB) network. The middleware machine software can provide application server, middleware and other functionality such as, for 30 example, WebLogic Server, JRockit or Hotspot JVM, Oracle Linux or Solaris, and Oracle VM. In accordance with an embodiment, the system can include a plurality of compute nodes, IB switch gateway, and storage nodes or units, communicating with one another via an IB network. When implemented as a rack configuration, unused portions of the rack can be left empty or occupied by fillers.

**[00011]** In accordance with an embodiment of the invention, referred to herein as "Sun Oracle Exalogic" or "Exalogic", the system is an easy-to-deploy solution for hosting middleware or application server software, such as the Oracle Middleware SW suite, or Weblogic. As 35

described herein, in accordance with an embodiment the system is a “grid in a box” that comprises one or more servers, storage units, an IB fabric for storage networking, and all the other components required to host a middleware application. Significant performance can be delivered for all types of middleware applications by leveraging a massively parallel grid architecture using, e.g. Real Application Clusters and Exalogic Open storage. The system 5 delivers improved performance with linear I/O scalability, is simple to use and manage, and delivers mission-critical availability and reliability.

**[00012]** In accordance with an embodiment of the invention, Tuxedo is a set of software modules that enables the construction, execution, and administration of high performance, 10 distributed business applications and has been used as transactional middleware by a number of multi-tier application development tools. Tuxedo is a middleware platform that can be used to manage distributed transaction processing in distributed computing environments. It is a proven platform for unlocking enterprise legacy applications and extending them to a services oriented architecture, while delivering unlimited scalability and standards-based interoperability.

**[00013]** In accordance with an embodiment of the invention, a transactional middleware system, such as a Tuxedo system, can take advantage of fast machines with multiple processors, such as an Exalogic middleware machine, and a high performance network connection, such as an Infiniband (IB) network.

## 20 **A Self-tuning lock mechanism**

**[00014]** In accordance with an embodiment of the invention, a self-tuning lock mechanism can be supported in a transactional middleware system to protect transaction data in a shared memory when there are concurrent transactions. Using the self-tuning lock mechanism, the transactional middleware machine environment can achieve significant throughput improvement 25 in transactional application scenarios such as applications with massive concurrent transactions.

**[00015]** **Figure 1** shows an illustration of a transactional middleware machine environment that supports a self-tuning lock mechanism, in accordance with an embodiment of the invention. As shown in Figure 1, a transactional middleware machine comprises multiple CPUs 131-134 that support an operating system (OS) 104, and a shared memory 103 that includes various 30 transactional data 121-123. A transactional application 101 with multiple concurrent transactions 111-112 can run on a plurality of processes 113-115 in a transactional server 102, each of which can use an atomic TAS (Test-And-Set) assembly 107 to implement an effective locking mechanism. The locking mechanism can protect the transaction data in the shared memory when there are concurrent transactions. In addition, a process in the transactional application 35 can use a semaphore mechanism 106 provided by the OS to obtain a lock on data 122 if necessary.

**[00016]** In accordance with one embodiment, when a process wants to get a lock on data

122, the process can perform a TAS operation for a number of rounds. The system can specify a target spin count, which is the number of rounds of TAS operation that are allowed. The target spin count can be either preconfigured or dynamically determined.

**[00017]** If the lock becomes available before the target spin count is reached, the process can

5 obtain the lock with much less cost than the semaphore mechanism provided by the OS. On the other hand, if the lock is not available during this period, then the process can be configured to block the semaphore, and wait until the lock owner wakes up and releases the lock.

**[00018]** In accordance with an embodiment of the invention, the target spin count value can be decided in the context of the hardware configuration and the application scenario. Users can

10 fine tune the spin count value manually in order to find an optimized value. The decision may not be obvious in some situations, since there is a trade-off between the CPU usage and the time to get a lock. For example, users may have to spend more CPU power to perform more TAS operations in order to get the lock in a shorter time frame. Hence, there may not be an optimized target spin count value that is obvious for every case.

15 **[00019]** One drawback of user level semaphore implementation is that the transactional application may not be able to dynamically adjust the target spin count in real time depending on the specific machine type. In general, the user level semaphore only uses a statically configured target spin count value, and users can adjust the target spin count value manually only by experimental practices. Since the optimal spin count value is machine dependent, and there is 20 no one-fit-in-all value for all platforms, a more suitable approach is to employ a mechanism to calculate the target spin count value dynamically and in real time.

**[00020]** In accordance with an embodiment of the invention, the target spin count value can be stored in a shared memory. A special process, such as a Tuxedo daemon process, can periodically change the spin count value according to operation information collected in the 25 previous period. For example, the Tuxedo daemon can update the target spin count value once per 5 seconds by default.

**[00021]** In accordance with one embodiment, an algorithm can be used to configure the target spin count value. The algorithm can increase the target spin count value, if the CPU idle ratio is low, or too many TAS operations have failed to obtain the lock and the system switched to the 30 semaphore. Furthermore, the algorithm can decrease the target spin count value if the CPU idle ratio is too high.

**[00022]** **Figure 2** illustrates an exemplary flow chart for supporting a self- tuning lock mechanism in a transactional middleware machine environment, in accordance with an embodiment of the invention. As shown in Figure 2, at step 201, the system can provide a 35 semaphore associated with an operating system running on a plurality of processors, wherein the plurality of processors operate to access data in a shared memory . Then, at step 202, a process of one or more processes can use a test-and-set (TAS) assembly component to perform

one or more TAS operations in order to obtain a lock for data in the shared memory. Finally, at step 203, the process can be blocked on the semaphore and wait for a release of a lock on data in the shared memory, after the TAS component has performed a specified number of TAS operations and failed to obtain the lock.

5

### Configuring the SPINCOUNT value in Tuxedo

**[00023]** In accordance with an embodiment of the invention, metadata, such as a SPINCOUNT parameter in the Tuxedo configuration file, can be used to specify the target spin count. SPINCOUNT can be either a static configured value, or a dynamic one.

10 **[00024]** For example, Tuxedo can use a statically set value of SPINCOUNT to determine how many times the bulletin board lock waits are spinning before getting blocked on system level semaphore. The drawback of this algorithm is that the value set by user is not the optimal value of SPINCOUNT on the specific platform, because the optimal value of SPINCOUNT is dependent upon many dynamic factors such as the CPU amounts, workload, number of waits for  
15 the Bulletin Board (BB) lock etc..

20 **[00025]** Additionally, Tuxedo can dynamically tune the value of SPINCOUNT while taking the runtime environment into consideration. An algorithm can be used to determine an appropriate value for the SPINCOUNT parameter. The system can increase the SPINCOUNT, if too many TAS operations have failed in the previous period and the system switched to the semaphore, and there was an enough CPU idle ratio. On the other hand, the system can decrease the SPINCOUNT, if the CPU idle ratio was too high.

25 **[00026]** The above algorithm is based on the CPU usage, e.g. CPU idle rate, and the ratio of SPIN failure, e.g. a ratio of SPIN failure in every 10000 operations to obtain a lock. The ratio of SPIN failure can indicate how many times locks are obtained via the semaphore instead of via the TAS operation.

30 **[00027]** A process can be in a SPIN mode, when the process is in an active status. The process can be blocked on the semaphore, if the process fails to obtain the lock after trying to perform the TAS operations for a number of times, which is referred to as a SPIN failure. The process can try to perform a TAS operation again and again in order to get the lock before a SPIN failure happens. A configurable parameter, e.g. SPINCOUNT, can be used to specify the number of rounds of TAS operation to be invoked and performed.

35 **[00028]** In accordance with an embodiment of the invention, the minimum idle CPU rate and the SPIN failed rate can be defined using metadata in a configuration file. For example, in the Tuxedo configuration file, there can be a MINIDLECPU parameter for defining the minimum idle CPU rate, the value of which has a range of 1-100 with 20 as default. Additionally, there can be a FACTOR parameter for defining the SPIN failed rate, the value of which has a range of 1-10000 with 1000 as default.

[00029] When a minimum idle CPU rate and SPIN failed rate are given, the system can tune the SPINCOUNT in each scan unit. Then, the system can increase the SPINCOUNT if the SPIN failure rate is too high (e.g. the SPIN failure rate is set to be greater than (1/FACTOR\*1.1)), or the idle CPU time is enough (e.g., idle rate > MINIDLECPU% +0.05), with the new SPINCOUNT

5 = old SPINCOUNT + old SPINCOUNT \* (cpu\_idletime/cpu\_usertime) and the maximum SPINCOUNT can be set as 10,000,000. Otherwise, the system can decrease the SPINCOUNT when the idle CPU ratio is too low (e.g., idle rate < MINIDLECPU% - 0.05), with the new SPINCOUNT=old SPINCOUNT/4 and the minimum SPINCOUNT to be set as 50000.

[00030] The following Listing 1 is a Tuxedo example for SPINTUNING configuration.

```

10      *RESOURCES
11          IPCKEY      123456
12          DOMAINID   simpapp
13          MASTER       ALLEN
14
15          MAXACCESSERS 10
16          MAXSERVERS   5
17          MAXSERVICES  10
18          MODEL        SHM
19          LDBAL        N
20          OPTIONS      EXALOGIC,SPINTUNING

25      *MACHINES
26          ALLENHOST    LMID="ALLEN"
27          APPDIR="/home/allen/Workspace/Tuxedo11gR1PS2/simpdir"
28          TUXCONFIG="/home/allen/Workspace/Tuxedo11gR1PS2/simpdir/tuxconfig"
29          TUXDIR="/home/allen/Software/OraHome/tuxedo11gR1PS2"
30          SPINTUNING_TARGET=1000
31          SPINTUNING_MINIDLECPU=20
32          *GROUPS
33          GROUP1
34          LMID=ALLEN GRPNO=1 OPENINFO=NONE

35          *SOURCES
36          DEFAULT:
37          CLOPT="-A"
38          simperv  SRVGRP=GROUP1 SRVID=1

39          *SERVICES
40          TOUPPER

```

Listing 1

[00031] As shown in the above example, the \*MACHINES section of the configuration file includes an attribute "SPINTUNING\_TARGET" that is used for configure tuning target. The value of "SPINTUNING\_TARGET" is numeric, which can be greater than or equal to "0" and less than or equal to e.g. "10000". A value of 0 indicates that the value built into the binary is used. The built-in value is 1000. The default value of "SPINTUNING\_TARGET" is 0.

[00032] In the above example, the value of SPINTUNING\_TARGET means that there is at most one time for the bulletin board to be locked via system semaphore per every thousand locks. The system can increase the value of SPINCOUNT with more CPU spent to meet a larger value of SPINTUNING\_TARGET. The attribute can be set with a nonzero value only if the option 5 SPINTUNING is specified.

[00033] As shown in the above example, the \*MACHINES section of the configuration file includes an attribute "SPINTUNING\_MINIDLECPU" that is used for specifying idle CPU rate. The idle CPU rate can be used by the system to find a proper SPINCOUNT dynamically. The larger SPINCOUNT is, the more CPU that the system uses. User can set the minimum idle CPU 10 rate via "SPINTUNING\_MINIDLECPU" to avoid spending too much CPU.

[00034] The value of "SPINTUNING\_MINIDLECPU" is numeric and in percentages. It can be greater than or equal to "0" and less than or equal to "100". A value of 0 indicates that the value built into the delivered binary should be used. For example, in Tuxedo, the built-in value can be set as 20, and the default value of "SPINTUNING\_MINIDLECPU" is 0. The attribute can be set 15 with a nonzero value only if the option SPINTUNING is specified.

[00035] In accordance with some embodiments, Figure 3 shows a functional block diagram of a system 1000 for supporting a locking mechanism in a transactional middleware machine environment configured in accordance with the principles of the invention as described above. The functional blocks of the system 1000 may be implemented by hardware, software, or a 20 combination of hardware and software to carry out the principles of the invention. It is understood by persons of skill in the art that the functional blocks described in Figure 3 may be combined or separated into sub-blocks to implement the principles of the invention as described above. Therefore, the description herein may support any possible combination or separation or further definition of the functional blocks described herein.

[00036] As shown in Figure 3, the system 1000 for supporting a locking mechanism in a transactional middleware machine environment comprises a providing unit 1100 and a TAS assembly component 1200. The providing unit 1100 provides a semaphore associated with an operating system running on a plurality of processors, wherein the plurality of processors operate to access data in a shared memory. The TAS assembly component 1200 is associated 30 with one or more processes, wherein each said process operates to use the TAS assembly component 1200 to perform one or more TAS operations in order to obtain a lock for data in the shared memory. A process operates to be blocked on the semaphore and waits for a release of a lock on data in the shared memory, after the TAS component 1200 has performed a specified number of TAS operations and failed to obtain the lock.

[00037] In some embodiments, the system 1000 further comprises a protecting unit 1300 for protecting transaction data in the shared memory when there are multiple concurrent transactions.

**[00038]** In some embodiments, the system 1000 further comprises an allowing unit 1400 for allowing the lock mechanism to use a spin count, which is the specified number of maximum rounds of TAS operations allowed.

**[00039]** In some embodiments, the system 1000 further comprises a preconfiguring unit 1500

5 for preconfiguring the spin count in metadata.

**[00040]** In some embodiments, the system 1000 further comprises a dynamically determining unit 1600 for dynamically determining the spin count based on both a hardware configuration and an application scenario.

**[00041]** In some embodiments, the system 1000 further comprises a periodically determining

10 unit 1700 for periodically determining the spin count using a special process.

**[00042]** In some embodiments, the spin count is dynamically determined using an algorithm, wherein the algorithm specifies that the spin count is increased from the spin count of a previous period, if a number of spin failures in the previous period exceeds a spin failure limit and a CPU idle ratio in the previous period is below a CPU idle ratio limit, and the spin count is decreased from the spin count of a previous period, if the CPU idle ratio exceeds a CPU idle ratio limit.

**[00043]** In some embodiments, a spin failure happens when a process fails to obtain a lock on data after attempting the TAS operation for a specified number of times.

**[00044]** In some embodiments, the system 1000 further comprises a using unit 1800 for using the semaphore to obtain the lock when a lock owner wakes up and releases the lock.

**[00045]** In some embodiments, the system 1000 further comprises a fine tuning unit 1900 for fine tuning the spin count manually to find an optimized value.

**[00046]** Other embodiments include an apparatus for supporting a locking mechanism in a transactional middleware machine environment, comprising a means for providing a semaphore associated with an operating system running on a plurality of processors, wherein the plurality of processors operate to access data in a shared memory; a means for using, via a process of one or more processes, a test-and-set (TAS) assembly component to perform one or more TAS operations in order to obtain a lock for data in the shared memory; a means for blocking on the semaphore via the process and waiting for a release of a lock on data in the shared memory, after the TAS component has performed a specified number of TAS operations and failed to obtain the lock.

**[00047]** In another embodiment, the apparatus further comprises means for protecting transaction data in the shared memory when there are multiple concurrent transactions.

**[00048]** In another embodiment, the apparatus further comprising means for allowing the lock mechanism to use a spin count, which is the specified number of maximum rounds of TAS operations allowed.

**[00049]** In another embodiment, the apparatus further comprises means for preconfiguring the spin count in metadata.

**[00050]** In another embodiment, the apparatus further comprises means for dynamically determining the spin count based on both a hardware configuration and an application scenario.

**[00051]** In another embodiment, the apparatus further comprises means for periodically determining the spin count using a special process.

5 **[00052]** In an embodiment further comprises an apparatus wherein the spin count is dynamically determined using an algorithm, wherein the algorithm specifies that the spin count is increased from the spin count of a previous period, if a number of spin failures in the previous period exceeds a spin failure limit and a CPU idle ratio in the previous period is below a CPU idle ratio limit, and the spin count is decreased from the spin count of a previous period, if the CPU 10 idle ratio exceeds a CPU idle ratio limit.

**[00053]** In another embodiment, the apparatus wherein a spin failure happens when a process fails to obtain a lock on data after attempting the TAS operation for a specified number of times.

15 **[00054]** In another embodiment, the apparatus further comprising means for using the semaphore to obtain the lock when a lock owner wakes up and releases the lock.

**[00055]** In another embodiment, the apparatus further comprises means for fine tuning the spin count manually to find an optimized value.

20 **[00056]** And yet another embodiment includes a system for supporting a locking mechanism in a transactional middleware machine environment, comprising a unit for providing a semaphore associated with an operating system running on a plurality of processors, wherein the plurality of processors operate to access data in a shared memory; a test-and-set (TAS) assembly component that is associated with one or more processes, wherein each said process operates to use the TAS assembly component to perform one or more TAS operations in order to obtain a lock for data in the shared memory, wherein a process operates to be blocked on the semaphore 25 and waits for a release of a lock on data in the shared memory, after the TAS component has performed a specified number of TAS operations and failed to obtain the lock;

**[00057]** In another embodiment the system further comprises a unit for protecting transaction data in the shared memory when there are multiple concurrent transactions.

30 **[00058]** In another embodiment the system further comprises a unit for allowing the lock mechanism to use a spin count, which is the specified number of maximum rounds of TAS operations allowed.

**[00059]** In another embodiment the system further comprises a unit for preconfiguring the spin count in metadata.

35 **[00060]** In another embodiment the system further comprises a unit for dynamically determining the spin count based on both a hardware configuration and an application scenario.

**[00061]** In another embodiment the system further comprises a unit for periodically determining the spin count using a special process.

**[00062]** Another embodiment includes a system wherein the spin count is dynamically determined using an algorithm, wherein the algorithm specifies that the spin count is increased from the spin count of a previous period, if a number of spin failures in the previous period exceeds a spin failure limit and a CPU idle ratio in the previous period is below a CPU idle ratio limit, and the spin count is decreased from the spin count of a previous period, if the CPU idle ratio exceeds a CPU idle ratio limit.

**[00063]** Another embodiment includes a system wherein a spin failure happens when a process fails to obtain a lock on data after attempting the TAS operation for a specified number of times.

10 **[00064]** Another embodiment includes a system further comprising a unit for using the semaphore to obtain the lock when a lock owner wakes up and releases the lock.

**[00065]** And yet another embodiment includes a system further comprising a unit for fine tuning the spin count manually to find an optimized value.

15 **[00066]** The present invention may be conveniently implemented using one or more conventional general purpose or specialized digital computer, computing device, machine, or microprocessor, including one or more processors, memory and/or computer readable storage media programmed according to the teachings of the present disclosure. Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art.

20 **[00067]** In some embodiments, the present invention includes a computer program product which is a storage medium or computer readable medium (media) having instructions stored thereon/in which can be used to program a computer to perform any of the processes of the present invention. The storage medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, 25 ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems (including molecular memory ICs), or any type of media or device suitable for storing instructions and/or data.

30 **[00068]** The foregoing description of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations will be apparent to the practitioner skilled in the art. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby enabling others skilled in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope of the invention 35 be defined by the following claims and their equivalence.

**Claims:**

What is claimed is:

5 1. A system for supporting a locking mechanism in a transactional middleware machine environment, comprising:

    a semaphore provided by an operating system running on a plurality of processors, wherein the plurality of processors operate to access data in a shared memory;

10    a test-and-set (TAS) assembly component that is associated with one or more processes, wherein each said process operates to use the TAS assembly component to perform one or more TAS operations in order to obtain a lock for data in the shared memory,

    wherein a process operates to be blocked on the semaphore and waits for a release of a lock on data in the shared memory, after the TAS component has performed a specified number of TAS operations and failed to obtain the lock.

15

2. The system of claim 1, wherein the lock mechanism protects transaction data in the shared memory when there are multiple concurrent transactions.

3. The system of claim 1, wherein the lock mechanism use a spin count, which is the 20 specified number of maximum rounds of TAS operations allowed.

4. The system of claim 3, wherein the spin count is preconfigured in metadata.

5. The system of claim 3, wherein the spin count is dynamically determined based on both a 25 hardware configuration and an application scenario.

6. The system of claim 3, wherein the spin count is periodically determined using a special process.

30 7. The system of claim 6, wherein the spin count is dynamically determined using an algorithm, wherein the algorithm specifies that

        the spin count is increased from the spin count of a previous period, if a number of spin failures in the previous period exceeds a spin failure limit and a CPU idle ratio in the previous period is below a CPU idle ratio limit, and

35        the spin count is decreased from the spin count of a previous period, if the CPU idle ratio exceeds a CPU idle ratio limit.

8. The system of claim 7, wherein a spin failure happens when a process fails to obtain a lock on data after attempting the TAS operation for a specified number of times.

9. The system of claim 1, wherein, using the semaphore, the process operates to obtain the

5 lock when a lock owner wakes up and releases the lock.

10. The system of claim 1, wherein the spin count is fine tuned manually to find an optimized value.

10 11. A method for supporting a locking mechanism in a transactional middleware machine environment, comprising:

providing a semaphore associated with an operating system running on a plurality of processors, wherein the plurality of processors operate to access data in a shared memory;

15 using, via a process of one or more processes, a test-and-set (TAS) assembly component to perform one or more TAS operations in order to obtain a lock for data in the shared memory;

blocking on the semaphore via the process and waiting for a release of a lock on data in the shared memory, after the TAS component has performed a specified number of TAS operations and failed to obtain the lock.

20

12. The method of claim 11, further comprising protecting transaction data in the shared memory when there are multiple concurrent transactions.

25

13. The method of claim 11, further comprising allowing the lock mechanism to use a spin count, which is the specified number of maximum rounds of TAS operations allowed.

14. The method of claim 13, further comprising preconfiguring the spin count in metadata.

30

15. The method of claim 13, further comprising dynamically determining the spin count based on both a hardware configuration and an application scenario.

16. The method of claim 13, further comprising periodically determining the spin count using a special process.

35

17. The method of claim 16, wherein the spin count is dynamically determined using an algorithm, wherein the algorithm specifies that

the spin count is increased from the spin count of a previous period, if a number

of spin failures in the previous period exceeds a spin failure limit and a CPU idle ratio in the previous period is below a CPU idle ratio limit, and

the spin count is decreased from the spin count of a previous period, if the CPU idle ratio exceeds a CPU idle ratio limit.

5

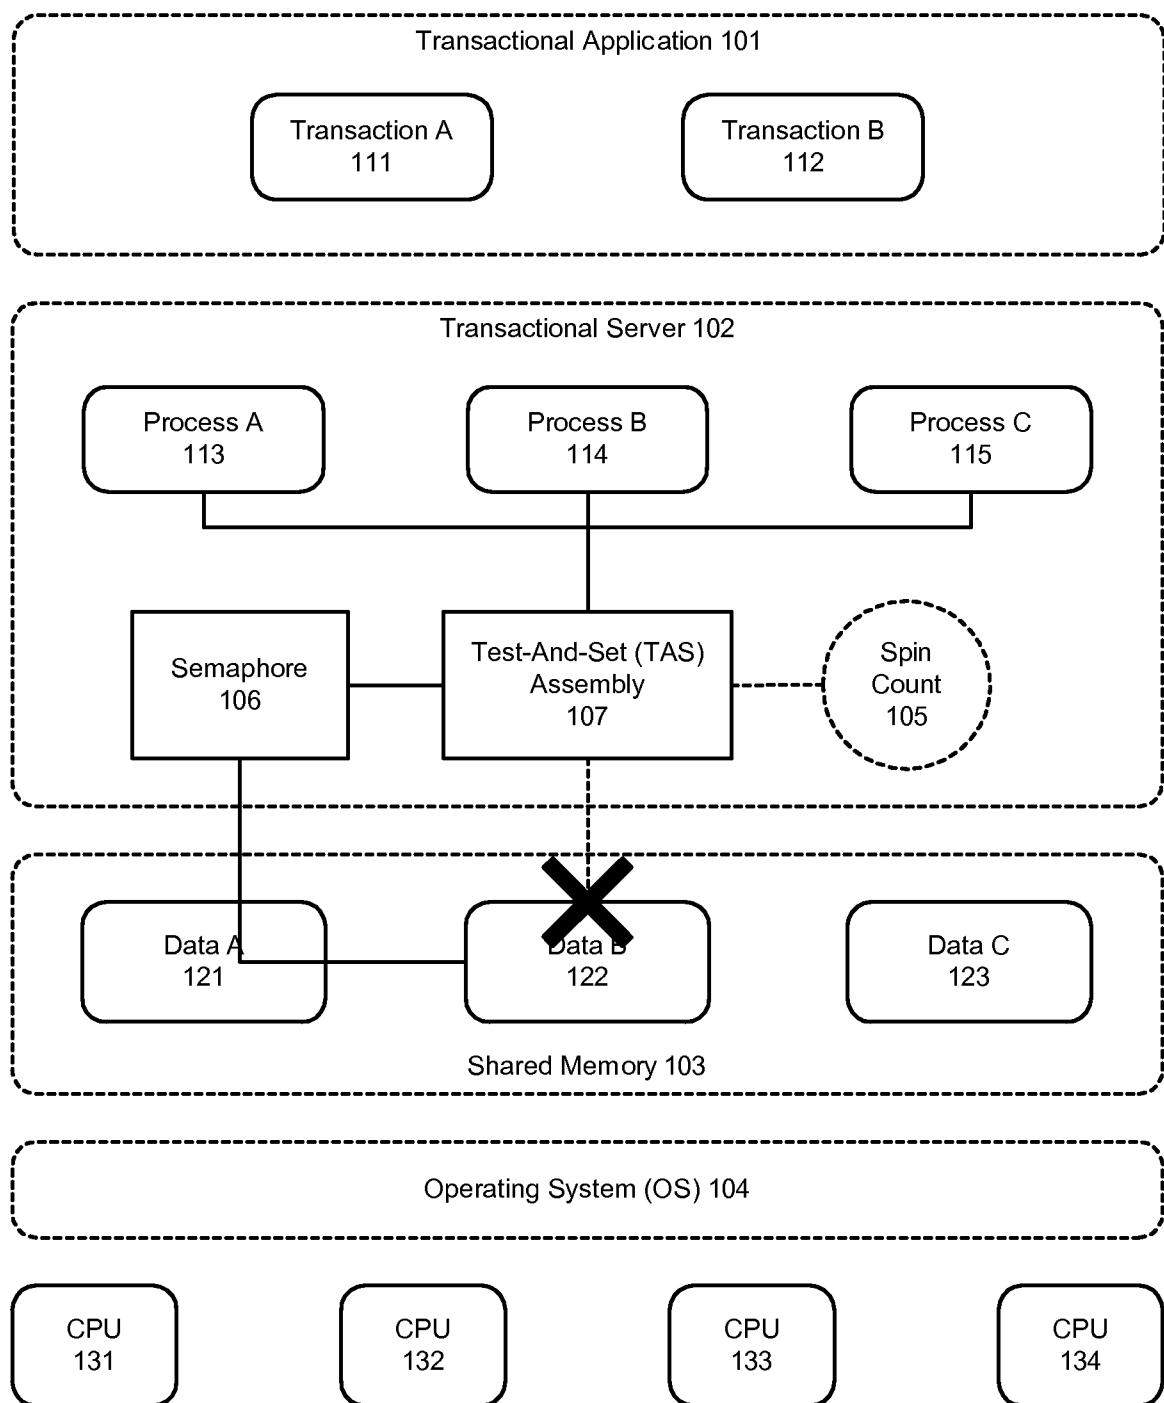
18. The method of claim 17, wherein a spin failure happens when a process fails to obtain a lock on data after attempting the TAS operation for a specified number of times.

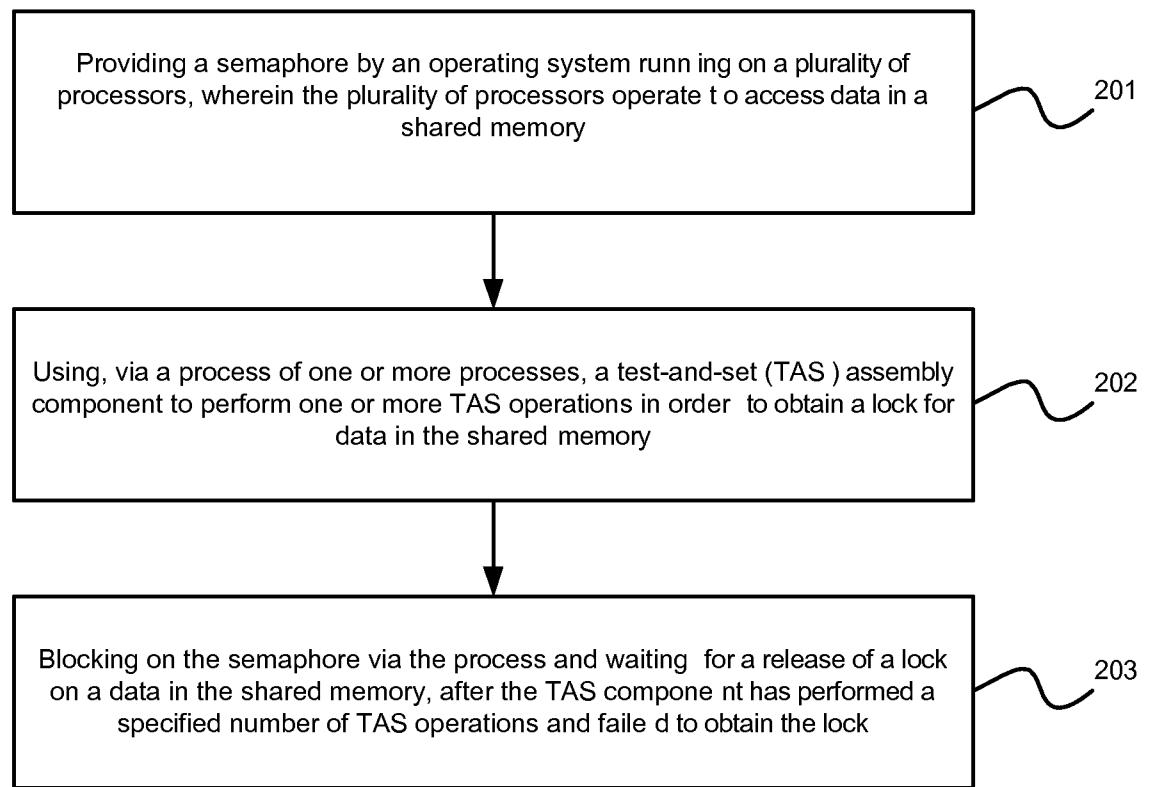
19. The method of claim 11, further comprising using the semaphore to obtain the lock when 10 a lock owner wakes up and releases the lock.

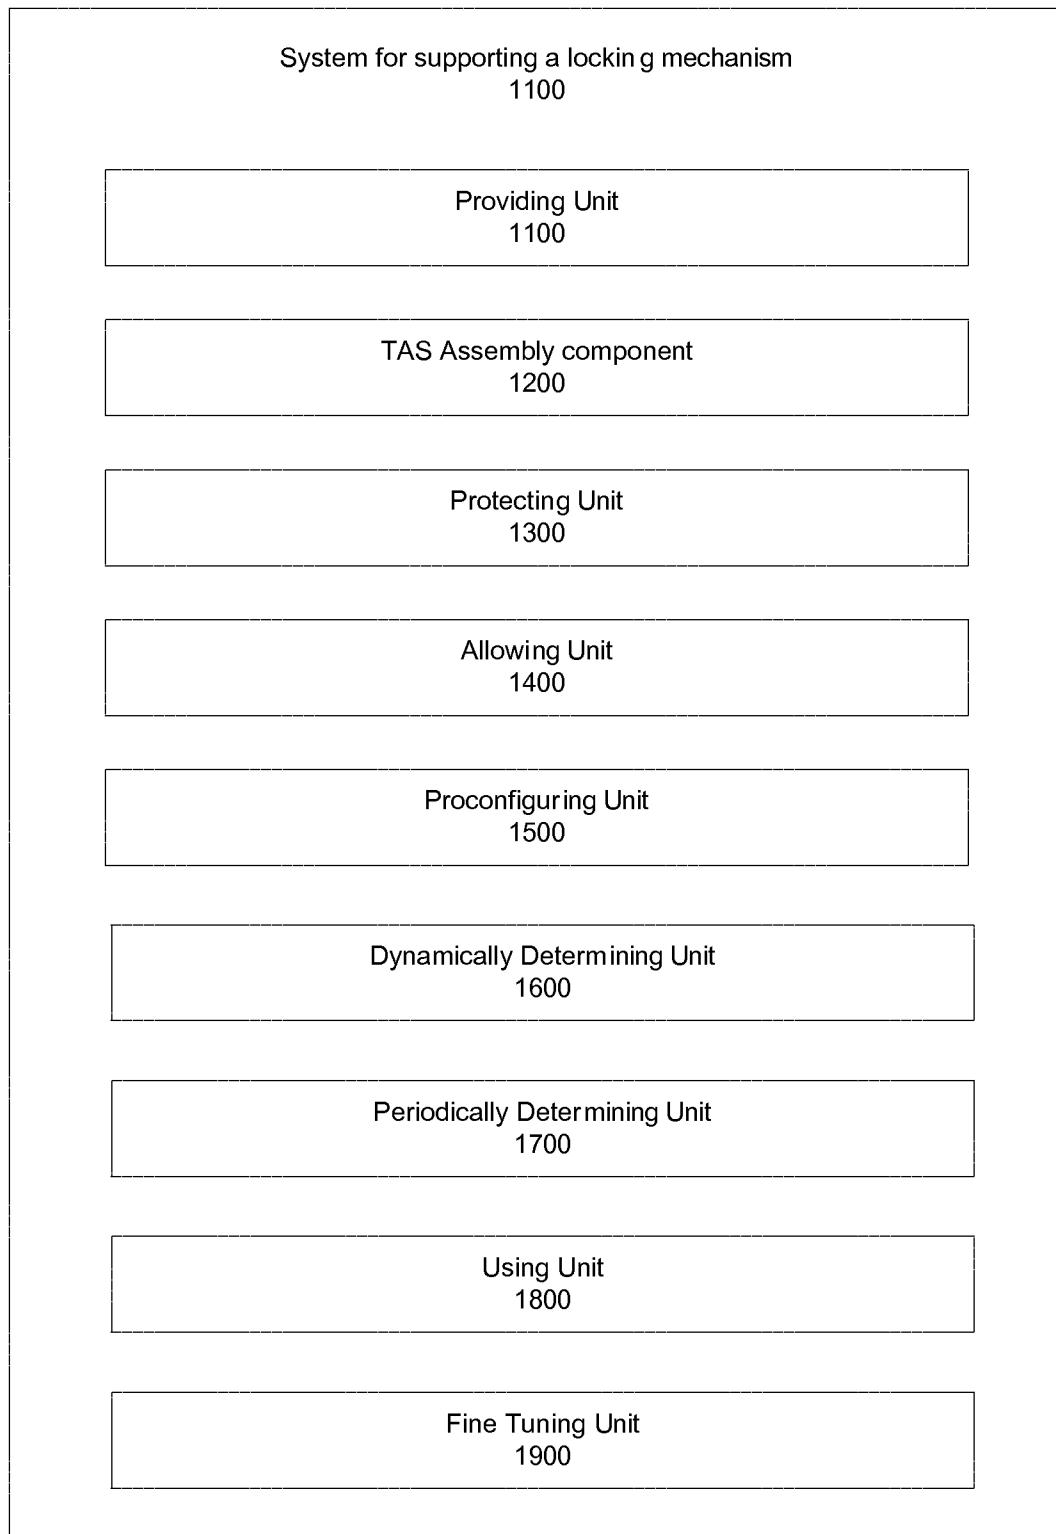
20. The method of claim 11, further comprising fine tuning the spin count manually to find an optimized value.

15 21. A program that causes a computer to perform the method of any of claim 11 to 20.

22. A computer readable non-volatile media that stores the program of claim 21.

**FIGURE 1**

**FIGURE 2**

**FIGURE 3**

# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 12/55942

## A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 15/167 (2012.01)

USPC - 709/214

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

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8): G06F 15/167 (2012.01)

USPC: 709/214

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched  
IPC(8): G06F 15/167 (2012.01); USPC: 709/201, 213, 214; 711/147, 149, 150, 151, 152; 718/1, 100, 104; 726/1 (keyword limited, see terms below)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
Google Scholar; Google Patents; PatBase(All). Search terms: semaphore, TAS, test, set, test-and-set, spin, failure, exceed, limit, number, iteration, count, lock, success, obtain, attempt, cycle, dynamic, adjust, optimize, change, modify, tune, compare, swap, CPU, processor, idle, load, measure, increase, decrease, tally, tries, etc.

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 7,594,234 B1 (Dice) 22 September 2009 (22.09.2009), entire document, especially fig. 2-3; col 1, ln 7-15; col 2, ln 51-59; col 4, ln 40-46; col 5, ln 57-60; col 6, ln 7-9, 24-28, 39-49, 61-67; col 7, ln 22-32; col 8, ln 3-8; col 12, ln 27-32	1-22
Y	US 6,549,961 B1 (Kloth) 15 April 2003 (15.04.2003), entire document, especially col 3, ln 18-20; abstract	1-22
Y	US 2007/0239943 A1 (Dice et al.) 11 October 2007 (11.10.2007), entire document, especially para [0035];[0010];[0090]	2, 12, 21-22
A	US 7,747,805 B2 (McKenney) 29 June 2010 (29.06.2010), entire document	1-22
A	US 5,050,072 A (Earnshaw et al.) 17 September 1991 (17.09.1991), entire document	1-22
A	US 2006/0143511 A1 (Huemiller) 29 June 2006 (29.06.2006), entire document	1-22

Further documents are listed in the continuation of Box C.

\* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

08 November 2012 (08.11.2012)

Date of mailing of the international search report

06 DEC 2012

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents  
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774