



Office de la Propriété

Intellectuelle
du Canada

Un organisme
d'Industrie Canada

Canadian
Intellectual Property
Office

An agency of
Industry Canada

CA 2664041 A1 2008/05/29

(21) **2 664 041**

**(12) DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION**

(13) A1

(86) Date de dépôt PCT/PCT Filing Date: 2007/11/17
(87) Date publication PCT/PCT Publication Date: 2008/05/29
(85) Entrée phase nationale/National Entry: 2009/03/19
(86) N° demande PCT/PCT Application No.: US 2007/085035
(87) N° publication PCT/PCT Publication No.: 2008/064139
(30) Priorité/Priority: 2006/11/17 (US11/601,541)

(51) Cl.Int./Int.Cl. *G06F 9/06* (2006.01)

(71) Demandeur/Applicant:
MICROSOFT CORPORATION, US

(72) Inventeurs/Inventors:
ZHANG, LINGLI, US;
GROVER, VINOD K., US;
MAGRUDER, MICHAEL M., US;
DETLEFS, DAVID, US;
DUFFY, JOHN JOSEPH, US;
GRAEFE, GOETZ, US

(74) Agent: SMART & BIGGAR

(54) Titre : ORDRE DE VALIDATION DE TRANSACTIONS BASEES SUR UN LOGICIEL ET GESTION DE CONFLITS
(54) Title: SOFTWARE TRANSACTION COMMIT ORDER AND CONFLICT MANAGEMENT

SOFTWARE TRANSACTIONAL MEMORY APPLICATION

200

PROGRAM LOGIC

204

LOGIC FOR PROVIDING A SOFTWARE TRANSACTIONAL MEMORY (STM) SYSTEM 206

LOGIC FOR PROVIDING A COMMIT ARBITRATOR THAT ALLOWS A PRE-DETERMINED COMMIT ORDER TO BE
SPECIFIED (STATICALLY OR DYNAMICALLY) FOR A PLURALITY OF TRANSACTIONS IN THE STM SYSTEM 208

LOGIC FOR ALLOWING THE COMMIT ARBITRATOR TO USE THE PRE-DETERMINED COMMIT ORDER AT
RUNTIME TO AID IN DETERMINING AN ORDER IN WHICH TO COMMIT THE PLURALITY OF TRANSACTIONS IN
THE SOFTWARE TRANSACTIONAL MEMORY SYSTEM 210

LOGIC FOR PROVIDING A CONTENTION MANAGEMENT PROCESS THAT IS INVOKED WHEN A CONFLICT
OCCURS BETWEEN A FIRST TRANSACTION AND A SECOND TRANSACTION 212

LOGIC FOR USING THE PRE-DETERMINED COMMIT ORDER IN THE CONTENTION MANAGEMENT PROCESS
TO AID IN DETERMINING WHETHER THE FIRST TRANSACTION OR THE SECOND TRANSACTION SHOULD WIN
THE CONFLICT AND BE ALLOWED TO PROCEED (E.G. DEPENDING ON WHICH ONE HAS THE LOWER COMMIT
ORDER NUMBER OF TWO TRANSACTIONS IN THE SAME TRANSACTION GROUP) 214

LOGIC FOR ALLOWING THE COMMIT ARBITRATOR TO BE OPERABLE TO USE THE PRE-DETERMINED
ORDERING TO TRACK ONE OR MORE ORDERING VALUES (E.G. IN TOTAL ORDERING – A NEXT-TO-COMMIT
FIELD THAT REPRESENTS A NEXT TRANSACTION OF THE PLURALITY OF TRANSACTIONS THAT SHOULD BE
ALLOWED TO COMMIT) AND FOR COMPARING THE ONE OR MORE ORDERING VALUES TO A PARTICULAR
COMMIT ORDER NUMBER OF A GIVEN TRANSACTION TO SEE IF THE COMMIT OF THE TRANSACTION IS
PROPER GIVEN THE ORDERING THAT SHOULD BE ENFORCED 216

OTHER LOGIC FOR OPERATING THE APPLICATION 220

(57) Abrégé/Abstract:

Various technologies and techniques are disclosed for applying ordering to transactions in a software transactional memory system. A software transactional memory system is provided with a feature to allow a pre-determined commit order to be specified

(57) Abrégé(suite)/Abstract(continued):

for a plurality of transactions. The pre-determined commit order is used at runtime to aid in determining an order in which to commit the transactions in the software transactional memory system. A contention management process is invoked when a conflict occurs between a first transaction and a second transaction. The pre-determined commit order is used in the contention management process to aid in determining whether the first transaction or the second transaction should win the conflict and be allowed to proceed.

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

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
29 May 2008 (29.05.2008)

PCT

(10) International Publication Number
WO 2008/064139 A1

(51) International Patent Classification:

G06F 9/06 (2006.01)

(21) International Application Number:

PCT/US2007/085035

(22) International Filing Date:

17 November 2007 (17.11.2007)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

11/601,541 17 November 2006 (17.11.2006) US

(71) Applicant (for all designated States except US): **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).(72) Inventors: **ZHANG, Lingli**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **GROVER, Vinod, K.**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **MAGRUDER, Michael, M.**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **DETLEFS, David**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **DUFFY, John, Joseph**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **GRAEFE, Goetz**; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

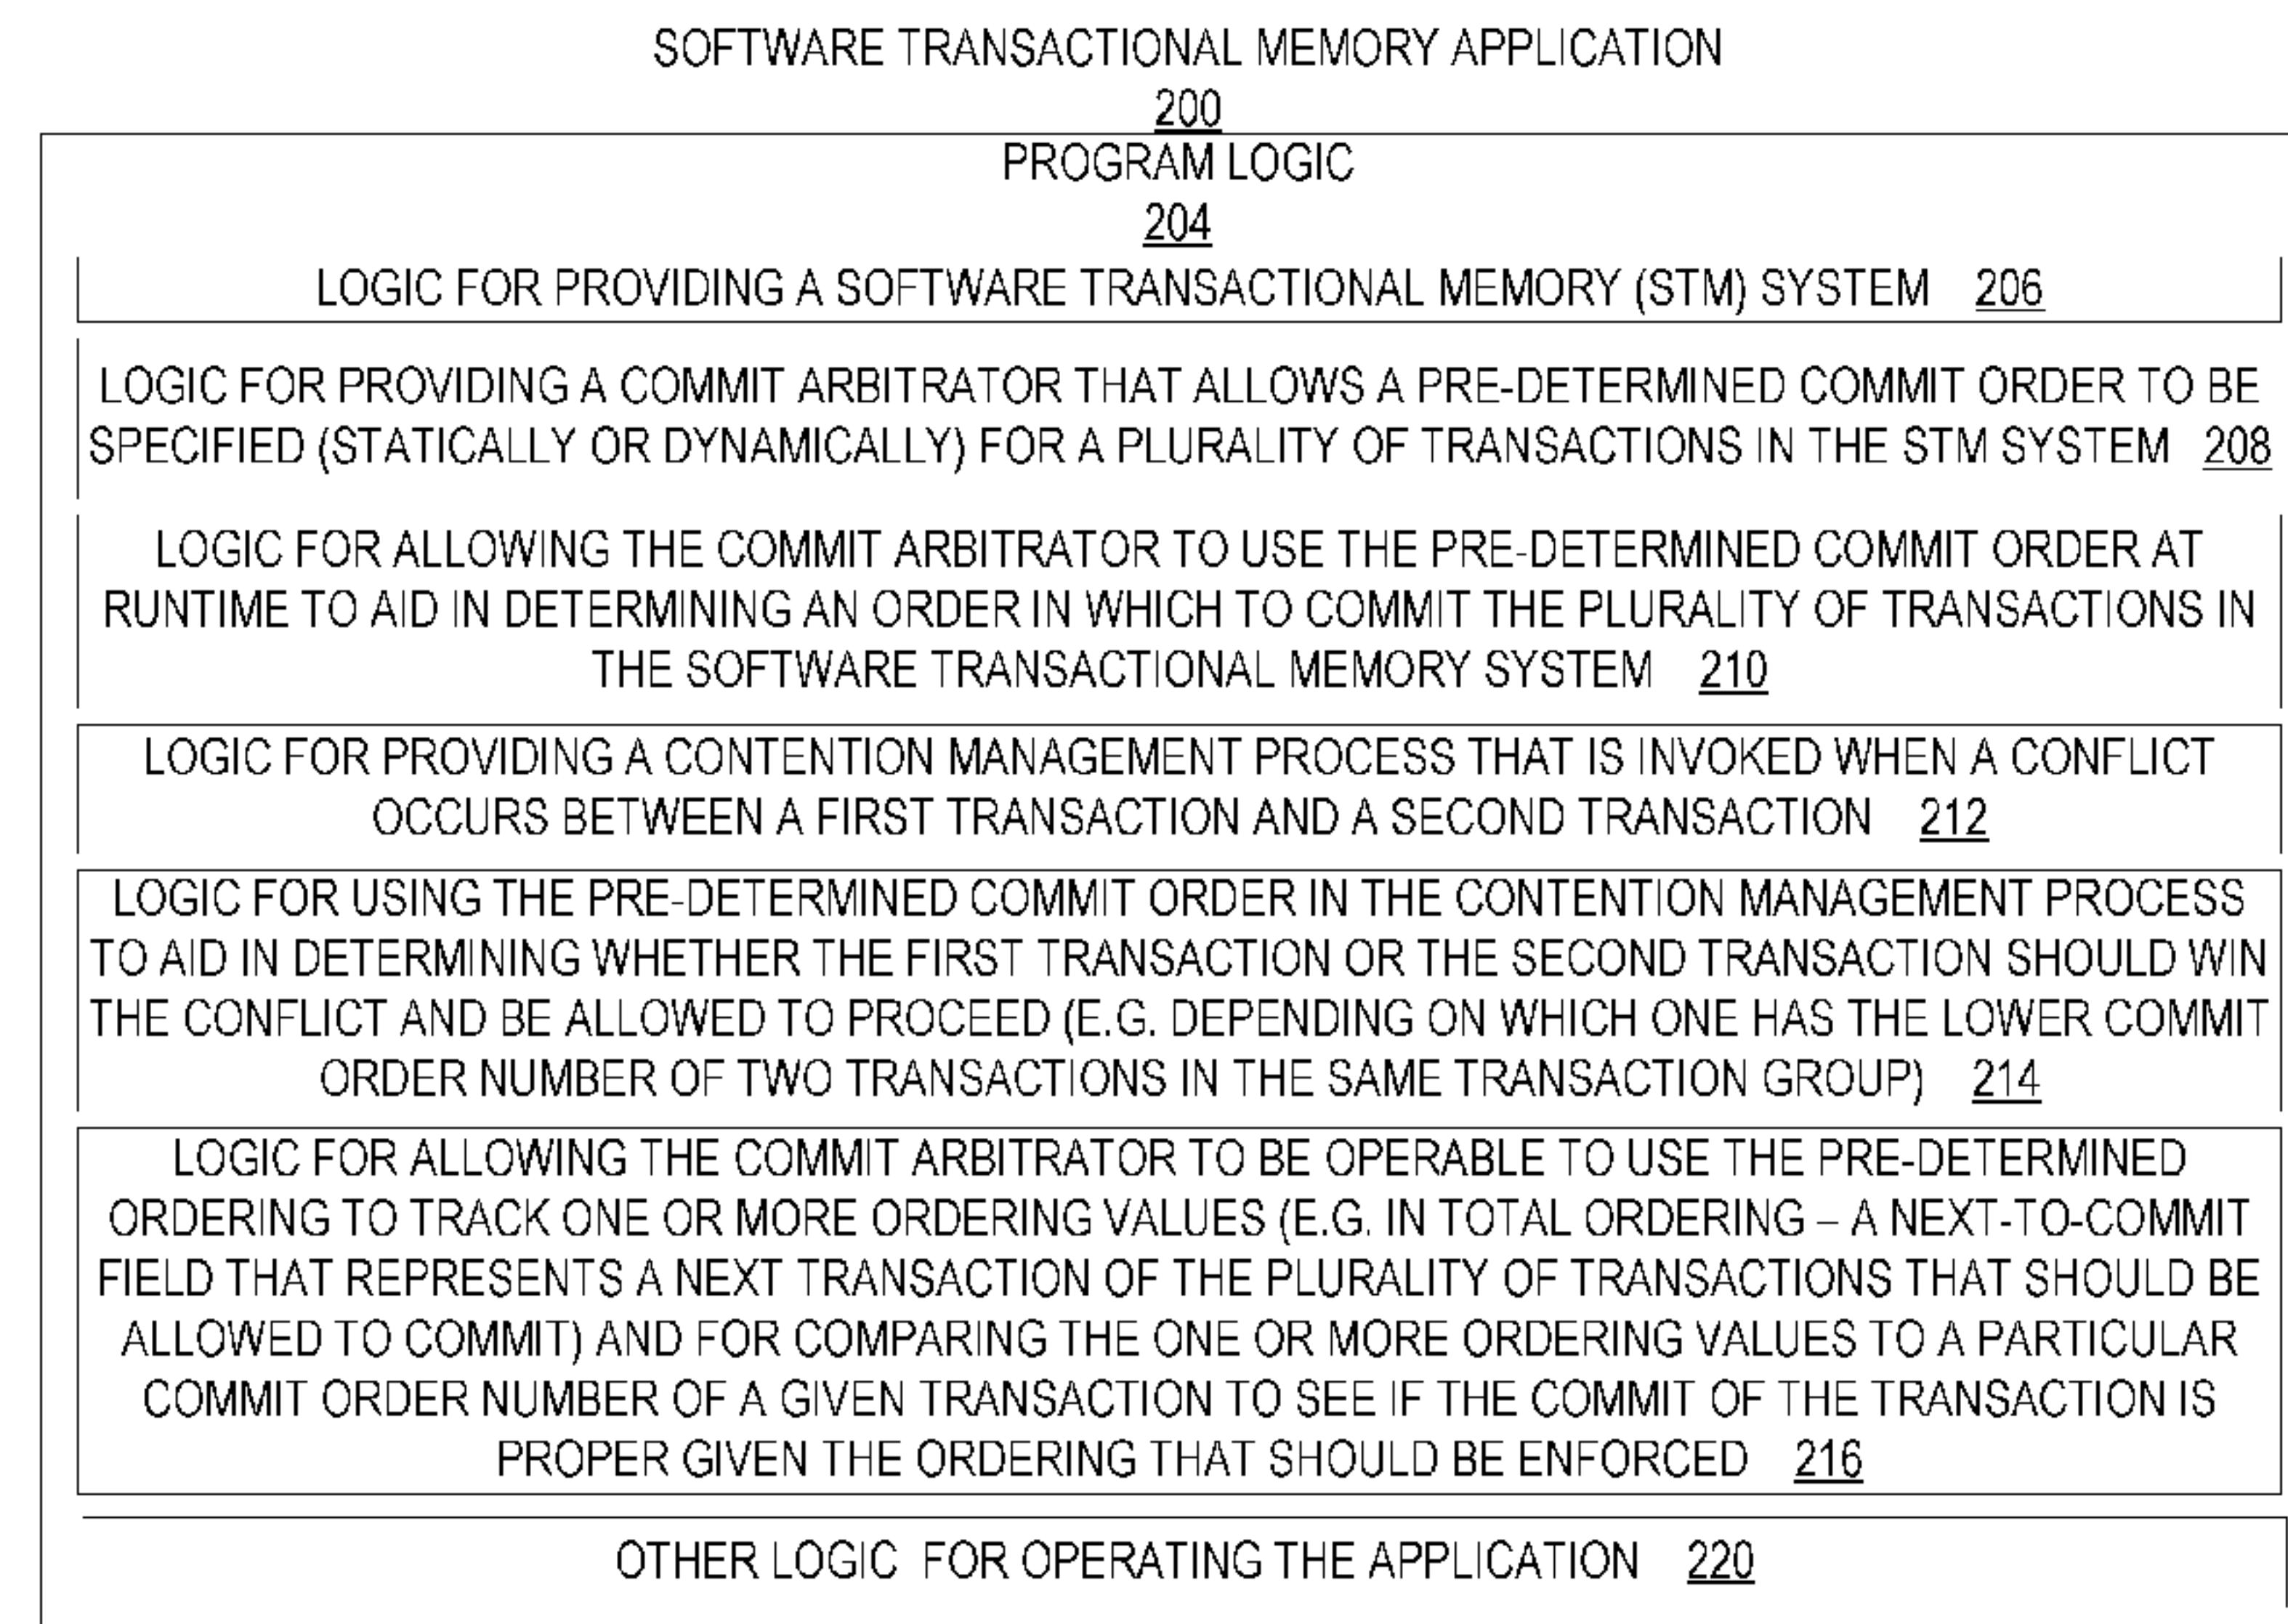
(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

[Continued on next page]

(54) Title: SOFTWARE TRANSACTION COMMIT ORDER AND CONFLICT MANAGEMENT



(57) **Abstract:** Various technologies and techniques are disclosed for applying ordering to transactions in a software transactional memory system. A software transactional memory system is provided with a feature to allow a pre-determined commit order to be specified for a plurality of transactions. The pre-determined commit order is used at runtime to aid in determining an order in which to commit the transactions in the software transactional memory system. A contention management process is invoked when a conflict occurs between a first transaction and a second transaction. The pre-determined commit order is used in the contention management process to aid in determining whether the first transaction or the second transaction should win the conflict and be allowed to proceed.

WO 2008/064139 A1

WO 2008/064139 A1



Published:

- *with international search report*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

SOFTWARE TRANSACTION COMMIT ORDER AND CONFLICT MANAGEMENT BACKGROUND

[001] Software transactional memory (STM) is a concurrency control mechanism analogous to database transactions for controlling access to shared memory in concurrent computing. A transaction in the context of transactional memory is a piece of code that executes a series of reads and writes to shared memory. STM is used as an alternative to traditional locking mechanisms. Programmers put a declarative annotation (e.g. atomic) around a code block to indicate safety properties they require and the system automatically guarantees that this block executes atomically with respect to other protected code regions. The software transactional memory programming model prevents lock-based priority-inversion and deadlock problems.

[002] While typical STM systems have many advantages, they still require the programmer to be careful in avoiding unintended memory access orderings. For example, the order in which transactions are committed (i.e. commit processing) in a typical STM environment is unconstrained. Transactions race with one another to commit, meaning that whether transaction 1 commits before transaction 2 or after is often a product of the dynamic scheduling of the program (and often by program-specific logic too). Moreover, if two transactions conflict, such as by trying to write to the same piece of memory, then their committing order can be arbitrarily decided based on one of many possible contention management policies. In both of these scenarios, no particular commit order is guaranteed; therefore the burden is on the programmer to make sure that his/her program works correctly with either order. This makes parallel programming very difficult.

[003] One approach to simplifying parallel programming is to automatically parallelize sequential programs, in a manner that guarantees that the semantics of the program are unchanged. In other words, if the sequential program works correctly, so does the parallelized version. Two (separate) variations of this concept to parallelize sequential programs have been termed, respectively, *safe futures* and *speculative loop parallelization*. In safe futures, the sequential version

of a program might perform “A; B” (that is, do A then do B). The programmer can add an annotation (a “future”) indicating that he or she thinks it might be possible to perform A and B in parallel without changing the program semantics – that A does not read any memory locations that B reads, nor vice-versa. But the system
5 treats this strictly as a “hint” whose validity must be checked. It executes A and B as transactions, and if they conflict, it prevents B from committing if it would be serialized before A. This is an “undesirable” aspect of undetermined commit order referred to above.

[004] Speculative loop parallelization is a similar idea, where the actions
10 performed in the sequential program are the successive iterations of a loop. The programmer (or some static analysis) indicates that it may be advantageous to execute the loop in parallel, and the system runs each iteration of the loop as a parallel transaction, requiring that these transactions commit in the order the iterations would have committed in the original program.

15 SUMMARY

[005] Various technologies and techniques are disclosed for applying ordering to transactions in a software transactional memory system. A software transactional memory system is provided with a feature to allow a pre-determined commit order to be specified for a plurality of transactions. The pre-determined commit order is
20 used at runtime to aid in determining an order in which to commit the transactions in the software transactional memory system. In one implementation, the pre-determined commit order can be either total ordering or partial ordering. In the case of total ordering, the transactions are forced to commit in a linear order. In the case of partial ordering, the transactions are allowed to commit in one of multiple
25 acceptable scenarios. In one implementation, a commit arbitrator keeps track of the next-to-commit value representing the transaction that should be allowed to commit next, and when a particular transaction is ready to commit, it is allowed to do so if its commit order number matches the next-to-commit value of the commit arbitrator.

[006] A contention management process is invoked when a conflict occurs between a first transaction and a second transaction. The pre-determined commit

order is used in the contention management process to aid in determining whether the first transaction or the second transaction should win the conflict and be allowed to proceed.

[007] This Summary was provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[008] Figure 1 is a diagrammatic view of a computer system of one implementation.

[009] Figure 2 is a diagrammatic view of a software transactional memory application of one implementation operating on the computer system of Figure 1.

[010] Figure 3 is a high-level process flow diagram for one implementation of the system of Figure 1.

[011] Figure 4 is a process flow diagram for one implementation of the system of Figure 1 illustrating the stages involved in using a commit arbitrator to enforce a pre-determined commit order.

[012] Figure 5 is a process flow diagram for one implementation of the system of Figure 1 illustrating the stages involved in using a commit arbitrator to enforce a total ordering of a plurality of transactions.

[013] Figure 6 is a process flow diagram for one implementation of the system of Figure 1 illustrating the stages involved in using a commit arbitrator to enforce a partial ordering of a plurality of transactions.

[014] Figure 7 is a process flow for one implementation of the system of Figure 1 that illustrates the stages involved in providing a contention management process that manages conflicts using the pre-determined commit order information.

[015] Figure 8 is a process flow for one implementation of the system of Figure 1 that illustrates the stages involved in providing a contention management process that manages conflicts with nested transactions using the pre-determined commit order information.

[016] Figure 9 is a logical diagram illustrating an exemplary ancestor tree with top level ancestors that have a common ancestor.

[017] Figure 10 is a logical diagram illustrating an exemplary ancestor tree with top level ancestors that do not have a common ancestor.

5 [018] Figure 11 is a process flow diagram for one implementation of the system of Figure 1 that illustrates the stages involved in reducing an amount of wasted work by using a commit arbitrator in a software transactional memory system.

[019] Figure 12 is a process flow diagram for one implementation of the system of Figure 1 that illustrates the stages involved in analyzing an entire ancestor chain in 10 a contention management process to determine the proper conflict resolution.

DETAILED DESCRIPTION

[020] For the purposes of promoting an understanding of the principles of the invention, reference will now be made to the embodiments illustrated in the drawings and specific language will be used to describe the same. It will nevertheless be understood that no limitation of the scope is thereby intended. Any alterations and further modifications in the described embodiments, and any further applications of the principles as described herein are contemplated as would normally occur to one skilled in the art.

[021] The system may be described in the general context as a software transactional memory system, but the system also serves other purposes in addition to these. In one implementation, one or more of the techniques described herein can be implemented as features within a framework program such as MICROSOFT® .NET Framework, or from any other type of program or service that provides platforms for developers to develop software applications. In another implementation, one or more of the techniques described herein are implemented as features with other applications that deal with developing applications that execute in concurrent environments.

[022] A feature is provided in the software transactional memory system to allow a pre-determined commit order to be specified for a plurality of transactions. The pre-determined commit order is used to aid in determining an order in which to commit the transactions. In one implementation, a contention management process

is invoked when a conflict occurs between a first transaction and a second transaction. The pre-determined commit order is then used in the contention management process to aid in determining whether the first transaction or the second transaction should win the conflict and be allowed to proceed.

5 [023] As shown in Figure 1, an exemplary computer system to use for implementing one or more parts of the system includes a computing device, such as computing device 100. In its most basic configuration, computing device 100 typically includes at least one processing unit 102 and memory 104. Depending on the exact configuration and type of computing device, memory 104 may be volatile 10 (such as RAM), non-volatile (such as ROM, flash memory, etc.) or some combination of the two. This most basic configuration is illustrated in Figure 1 by dashed line 106.

15 [024] Additionally, device 100 may also have additional features/functionality. For example, device 100 may also include additional storage (removable and/or non-removable) including, but not limited to, magnetic or optical disks or tape. Such additional storage is illustrated in Figure 1 by removable storage 108 and non-removable storage 110. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, 20 program modules or other data. Memory 104, removable storage 108 and non-removable storage 110 are all examples of computer storage media. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other 25 magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by device 100. Any such computer storage media may be part of device 100.

30 [025] Computing device 100 includes one or more communication connections 114 that allow computing device 100 to communicate with other computers/applications 115. Device 100 may also have input device(s) 112 such as keyboard, mouse, pen, voice input device, touch input device, etc. Output device(s)

111 such as a display, speakers, printer, etc. may also be included. These devices are well known in the art and need not be discussed at length here. In one implementation, computing device 100 includes software transactional memory application 200. Software transactional memory application 200 will be described 5 in further detail in Figure 2.

[026] Turning now to Figure 2 with continued reference to Figure 1, a software transactional memory application 200 operating on computing device 100 is illustrated. Software transactional memory application 200 is one of the application programs that reside on computing device 100. However, it will be understood that 10 software transactional memory application 200 can alternatively or additionally be embodied as computer-executable instructions on one or more computers and/or in different variations than shown on Figure 1. Alternatively or additionally, one or more parts of software transactional memory application 200 can be part of system memory 104, on other computers and/or applications 115, or other such variations 15 as would occur to one in the computer software art.

[027] Software transactional memory application 200 includes program logic 204, which is responsible for carrying out some or all of the techniques described herein. Program logic 204 includes logic for providing a software transactional memory (STM) system 206; logic for providing a commit arbitrator that allows a 20 pre-determined commit order to be specified, statically or dynamically, for a plurality of transactions in the STM system 208; logic for allowing the commit arbitrator to use the pre-determined commit order at runtime to aid in determining an order in which to commit the plurality of transactions in the software transactional memory system 210; logic for providing a contention management 25 process that is invoked when a conflict occurs between a first transaction and a second transaction 212; logic for using the pre-determined commit order in the contention management process to aid in determining whether the first transaction or the second transaction should win the conflict and be allowed to proceed (e.g. depending on which one has the lower commit order number of two transaction in 30 the same transaction group) 214; logic for allowing the commit arbitrator to be operable to use the pre-determined commit ordering to track one or more ordering

values (e.g. in total ordering – a next-to-commit field that represents a next transaction of the plurality of transaction that should be allowed to commit) and for comparing the one or more ordering values to a particular commit order number of a given transaction to see if the commit of the given transaction is proper given the ordering that should be enforced) 216; and other logic for operating the application 220. In one implementation, program logic 204 is operable to be called programmatically from another program, such as using a single call to a procedure in program logic 204.

[028] Turning now to Figures 3-10 with continued reference to Figures 1-2, the stages for implementing one or more implementations of software transactional memory application 200 are described in further detail. Figure 3 is a high level process flow diagram for software transactional memory application 200. In one form, the process of Figure 3 is at least partially implemented in the operating logic of computing device 100. The procedure begins at start point 240 with providing a software transactional memory system (stage 242). A feature is provided to allow a pre-determined commit order (e.g. a total ordering or partial ordering) to be specified for a plurality of transactions (e.g. assigned dynamically or statically) (stage 244). The term “pre-determined commit order” as used herein is meant to include a specific order in which a particular group of related transactions should be committed, as determined at any point in time before the transactions start running. The term “group” of transactions as used herein includes a particular set of (e.g. plurality of) transactions managed by the same commit arbitrator, as well as nested children of those transactions.

[029] The pre-determined commit order is used at runtime to aid in determining an order in which to commit the plurality of transactions in the software transactional memory system (stage 246). The pre-determined commit order is used to aid in resolving conflicts occurring between two or more of the plurality of transactions (stage 248). The process ends at end point 250.

[030] Figure 4 illustrates one implementation of the stages involved in using a commit arbitrator to enforce a pre-determined commit order. In one form, the process of Figure 4 is at least partially implemented in the operating logic of

computing device 100. The procedure begins at start point 270 with providing one or more commit arbitrators for a software transaction memory system, the commit arbitrator being operable to allow a pre-determined commit order to be specified for a plurality of transactions (stage 272). The term “commit arbitrator” as used herein is meant to include any type of program, feature, or process that is responsible for managing one or more groups of transactions that should be ordered with respect to one another. In one implementation, there can be one or more commit arbitrators active within a program at any given time. For example, as many commit arbitrators as are needed can be created to manage the different groups of 5 transactions. The commit arbitrator tracks and updates one or more ordering values that are used to determine the proper ordering of transactions with respect to one another (stage 274). In the case of total ordering, a next-to-commit field can be used to represent a next transaction of a plurality of transactions that should be committed next) (stage 274). In the case of partial ordering, a directed graph of 10 different possible orders is tracked using the ordering values. As appropriate, the commit arbitrator uses the pre-determined commit order to provide a commit order 15 number for each of the plurality of transactions (stage 276).

[031] When a particular transaction of the plurality of transactions prepares to commit, if the commit order number for the particular transaction when compared 20 to the one or more ordering values reveals that the commit is proper, then the commit arbitrator allows the transaction to commit (stage 278). In the case of total ordering, this scenario occurs when the next-to-commit field and the commit order number for the particular transaction have the same value. In such a scenario, the commit arbitrator allows the transaction to commit and then increments the next-to- 25 commit field to a next number in a sequence (e.g. next higher number) if the commit is successful (stage 278). When the particular transaction of the plurality of transactions prepares to commit, if the commit order number for the particular transaction when compared to the ordering values reveals that the commit is not proper, then the particular transaction is placed in a hold mode until it is awakened 30 at a later point in time after a predecessor transaction commits (stage 280). In the

case of total ordering, this hold mode is entered when the next-to-commit field and the order number for the particular transaction do not have the same value.

[032] In one implementation, the system may wake a transaction after its immediate predecessor has committed, in which case it may try to commit right away. Alternatively, the system may choose to wake a transaction after some non-immediate predecessor has committed, even though its immediate predecessor may not yet have committed. After being awakened, the system checks to see if it is appropriate for the transaction to really commit. If so, the transaction is committed. The process ends at end point 282.

[033] Figure 5 illustrates one implementation of the stages involved in using a commit arbitrator to enforce a total ordering of a plurality of transactions. In one form, the process of Figure 5 is at least partially implemented in the operating logic of computing device 100. The procedure begins at start point 290 with providing one or more commit arbitrators operable to allow a pre-determined total ordering to be specified for a plurality of transactions (e.g. one specifying an exact order in which the plurality of transactions should be committed) (stage 292). When a particular transaction of the plurality of transactions reaches its commit point, to enforce the commit order, the commit order of the particular transaction is compared with a next-to-commit field of the commit arbitrator (stage 296). In one implementation, if the system determines that enforcement of the total ordering is not necessary (e.g. such as because there is definitely no conflict), then the total ordering requirement can be broken as appropriate (stage 294), then the process ends at end point 302.

[034] If commit ordering is to be enforced, and if the commit order of the particular transaction has a same value as the next-to-commit field of the commit arbitrator (decision point 296), then the particular transaction is committed, and if the commit is successful, the next-to-commit field is incremented and the next successor is awakened, if any exist (stage 298). If the commit order of the particular transaction does not have the same value as the next-to-commit field of the commit arbitrator (decision point 296), then the particular transaction is put in a hold/sleep mode until it is awakened at a later point in time after a predecessor

transaction commits (stage 300). In one implementation, at that later point in time, if a conflict occurs with a predecessor, that particular transaction may be asked to abort and rollback such that a predecessor may make forward progress. Otherwise, if no such conflict has occurred, that particular transaction should be able to 5 commit once the commit order requirements described herein are met. The process then ends at end point 302.

[035] Figure 6 illustrates one implementation of the stages involved in using a commit arbitrator to enforce a partial ordering of a plurality of transactions. In one form, the process of Figure 6 is at least partially implemented in the operating logic 10 of computing device 100. The procedure begins at start point 310 with providing one or more commit arbitrators operable to allow a pre-determined partial ordering to be specified for a plurality of transactions (e.g. one specifying a plurality of acceptable orders in which the plurality of transactions should be committed – e.g. in the form of a directed graph) (stage 312). When a particular transaction of the 15 plurality of transactions reaches its commit point, to enforce the commit order, the state of the predecessor transactions (e.g. one or more ordering values) are consulted for the particular committing transaction (e.g. as tracked by the commit arbitrator) (stage 314). If all predecessors to the particular transaction have committed (decision point 316), then the particular transaction is committed (stage 20 318). If the commit is successful, one or more values tracked by the commit arbitrator are updated as appropriate, and all possible next successors are awakened, if any exist (stage 318).

[036] If all predecessors to the particular transaction have not committed (decision point 316), then the particular transaction is put in a hold/sleep mode until it is 25 awakened at a later point in time after a predecessor transaction commits (stage 320). The process ends at end point 322.

[037] Figure 7 illustrates one implementation of the stages involved in providing a contention management process that manages conflicts using the pre-determined commit order information. In one form, the process of Figure 7 is at least partially 30 implemented in the operating logic of computing device 100. The procedure begins at start point 340 with providing a software transactional memory system

that supports a pre-determined commit order for one or more groups of transactions (stage 342). A contention management process is provided that is invoked when a conflict occurs between a first transaction and a second transaction (stage 344).

The pre-determined commit order is used in the contention management process to aid in determining whether the first transaction or the second transaction should win the conflict and be allowed to proceed (stage 346). If the first transaction and second transaction are not part of the same transaction group (decision point 348), then a pre-determined commit order is not enforced between these two transactions (because none existed) (stage 350). In such a scenario, since the two transactions are not in a same transaction group, the ordering factor is not used to help resolve the conflict (stage 350).

[038] If the first transaction and the second transaction are part of the same transaction group (decision point 348), then the system compares the first order number of the first transaction and the second order number of the second transaction (stage 352). The transaction with the lower order number is allowed to proceed (or with another suitable priority ordering) (stage 354). The process ends at end point 356.

[039] Figure 8 illustrates one implementation of the stages involved in providing a contention management process that manages conflicts with nested transactions using the pre-determined commit order information. In one form, the process of Figure 8 is at least partially implemented in the operating logic of computing device 100. In one implementation, the entire ancestor chain is considered for each transaction before committing the particular transaction, so that any ordering present in that chain is enforced. The procedure begins at start point 370 with providing a contention management process that is invoked when a conflict occurs between a first transaction and a second transaction (stage 372). A pre-determined commit order is used in the contention management process to aid in determining whether the first transaction or the second transaction should win the conflict and be allowed to proceed (stage 372). If the first and second transactions are not part of the same transaction group (decision point 376), then a pre-determined commit order is not enforced between those two transactions (because none existed) (stage

378) and the process ends at end point 388. If the first and second transactions are part of the same transaction group (decision point 376), then the system checks to see if nested transactions are involved (decision point 380).

[040] If nested transactions are not involved (decision point 380), then the order number (or other ordering indicator) of the first transaction is compared with the order number (or other ordering indicator) of the second transaction (stage 384). The transaction with the lower order number is allowed to proceed (or the one determined to be next in order by using other suitable ordering criteria) (stage 386).

[041] If nested transactions are involved (decision point 380), then the order number (or other ordering indicator) of the top level ancestor of the first transaction is compared with the order number (or other ordering indicator) of the top level ancestor of the second transaction (stage 382). The term “top level ancestor” as used herein is meant to include the immediate children of common ancestors where common ancestors are involved, and the top level ancestor of each transaction where there is no common ancestor involved. These scenarios involving common and uncommon ancestors are illustrated in further detail in Figures 9 and 10. The transaction with the lower order number is allowed to proceed (e.g. the transaction related to the ancestor that had the lower order number or other suitable criteria) (stage 386). The process ends at end point 388.

[042] Figure 9 is a logical diagram illustrating an exemplary ancestor tree with top level ancestors that have a common ancestor. In the example shown, transaction A is a common ancestor of D and E. In conflicts occurring between D and E, the order number of transactions B and C (the immediate children of common ancestor A) are analyzed to determine which transaction D or E should be allowed to proceed (stage 382 in Figure 8).

[043] Figure 10 is a logical diagram illustrating an exemplary ancestor tree with top level ancestors that do not have common ancestors. In the example shown, transaction A is an ancestor of transaction C. Transaction D is an ancestor of transaction F. In conflicts occurring between transactions C and F, then the order number of transactions A and D (the top level ancestor of each) are compared to

determine which transaction C or F should be allowed to proceed (stage 382 in Figure 8).

[044] Figure 11 illustrates one implementation of the stages involved in reducing the amount of wasted work by using a commit arbitrator in a software transactional memory system. In one form, the process of Figure 11 is at least partially implemented in the operating logic of computing device 100. The procedure begins at start point 400 with providing one or more commit arbitrators for a software transactional memory system, the commit arbitrator being operable to allow a pre-determined commit order to be specified for a plurality of transactions (stage 402). The commit arbitrator is operable to put a transaction into sleep/hold mode to block that transaction from re-executing when a predecessor transaction is still executing (e.g. by analyzing the pre-determined commit order to determine the proper order (stage 404). The commit arbitrator is also operable to wake up transactions that were put on hold once the predecessor transaction(s) have finished (e.g. by again analyzing the pre-determined commit order to determine the proper order) (stage 406). By providing these blocking and waking mechanisms, the commit arbitrator helps reduce the amount of work that is wasted by keeping operations from being performed that would have to be undone later (stage 408). The process ends at end point 410.

[045] Figure 12 illustrates one implementation of the stages involved in analyzing an entire ancestor chain in a contention management process to determine the proper conflict resolution. In one form, the process of Figure 12 is at least partially implemented in the operating logic of computing device 100. The procedure begins at start point 430 with providing a contention management process that is invoked when a conflict occurs between a first transaction and a second transaction (stage 432). A pre-determined commit order is used in the contention management process to aid in determining whether the first transaction or the second transaction should win the conflict and be allowed to proceed (stage 434). An entire ancestor chain of a pre-determined commit order is analyzed to help determine the proper conflict management (stage 436). For example, if there are four transactions, two parents and two children, where B is nested within A and D is nested within C.

Suppose there is an ordering relationship between A and C where A should commit before C. If B and D conflict, the contention management process should favor B because favoring D is useless given that A must commit before C. (stage 436). The process ends at end point 438.

5 **[046]** Although the examples discussed herein talked about enforcing commit ordering using various technologies and techniques, it should be noted that a transaction may not have a commit arbitrator at all. In such a case that a transaction does not have a commit arbitrator at all, a normal unordered commit will occur.

10 **[047]** Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims. All equivalents, 15 changes, and modifications that come within the spirit of the implementations as described herein and/or by the following claims are desired to be protected.

16 **[048]** For example, a person of ordinary skill in the computer software art will recognize that the client and/or server arrangements, user interface screen content, and/or data layouts as described in the examples discussed herein could be 20 organized differently on one or more computers to include fewer or additional options or features than as portrayed in the examples.

What is claimed is:

1. A method for applying ordering to transactions in a software transactional memory system comprising the steps of:
 - providing a software transactional memory system (242);
 - 5 providing a feature to allow a pre-determined commit order to be specified for a plurality of transactions (244); and
 - using the pre-determined commit order at runtime to aid in determining an order in which to commit the plurality of transactions in the software transactional memory system (246).
- 10 2. The method of claim 1, further comprising:
 - using the pre-determined commit order to aid in resolving conflicts occurring between two or more of the plurality of transactions (248).
3. The method of claim 1, wherein the pre-determined commit order comprises a pre-determined total ordering of the plurality of transactions (244).
- 15 4. The method of claim 1, wherein the pre-determined commit order comprises a pre-determined partial ordering of the plurality of transactions (244).
5. The method of claim 1, wherein a commit arbitrator is provided that tracks one or more ordering values for the pre-determined commit order (272).
- 20 6. The method of claim 5, wherein when a particular transaction of the plurality of transactions prepares to commit, comparing the one or more ordering values to an ordering value for the particular transaction to determine if the particular transaction can commit (278).
7. The method of claim 1, wherein the pre-determined commit order is assigned dynamically (244).
- 25 8. The method of claim 1, wherein the pre-determined commit order is assigned statically (244).
9. The method of claim 1, wherein a commit arbitrator is provided that tracks a next-to-commit field that represents a next transaction of the plurality of transactions that should be allowed to commit (274).
- 30 10. The method of claim 9, wherein a commit order number is provided for each of the plurality of transactions (276).

11. The method of claim 10, wherein when a particular transaction of the plurality of transactions prepares to commit, determining if the commit order number for the particular transaction has a same value as the next-to-commit field tracked by the commit arbitrator (278).

5 12. The method of claim 11, wherein if the commit order number for the particular transaction and the next-to-commit field have the same value, allowing the commit to proceed (278).

10 13. The method of claim 12, wherein after the commit proceeds and is successful, the commit arbitrator increments the next-to-commit field to a next number in a sequence (278).

14. The method of claim 11, wherein if the commit order number for the particular transaction and the next-to-commit field do not have the same value, putting the particular transaction into a hold mode until it is awakened at a later point in time after a predecessor transaction commits (280).

15 15. A computer-readable medium having computer-executable instructions for causing a computer to perform the steps recited in claim 1 (200).

16. A method for providing contention management with ordering comprising the steps of:

20 providing a software transactional memory system that supports a pre-determined commit order for one or more groups of transactions (342);

providing a contention management process that is invoked when a conflict occurs between a first transaction and a second transaction (344); and

25 using the pre-determined commit order in the contention management process to aid in determining whether the first transaction or the second transaction should win the conflict and be allowed to proceed (346).

17. The method of claim 16, wherein if the contention management process determines that the first transaction and the second transaction are part of a same transaction group (348), then a first order number field representing the first transaction and a second order number field representing the second transaction are compared (352), and a particular transaction of the first transaction and the second transaction that has a lower respective order number is allowed to proceed (354).

18. A computer-readable medium having computer-executable instructions for causing a computer to perform the steps recited in claim 16 (200).

19. A computer-readable medium having computer-executable instructions for causing a computer to perform steps comprising:

5 provide a software transactional memory system (206);

provide a commit arbitrator that allows a pre-determined commit order to be specified for a plurality of transactions (208), the commit arbitrator being operable to use the pre-determined commit order at runtime to aid in determining an order in which to commit the plurality of transactions in the software transactional memory system (210);

provide a contention management process that is invoked when a conflict occurs between a first transaction and a second transaction (212); and

use the pre-determined commit order in the contention management process to aid in determining whether the first transaction or the second transaction should 15 win the conflict and be allowed to proceed (214).

20. The computer-readable medium of claim 19, wherein the commit arbitrator is operable to enforce commit ordering within nested transactions (380).

1 / 11

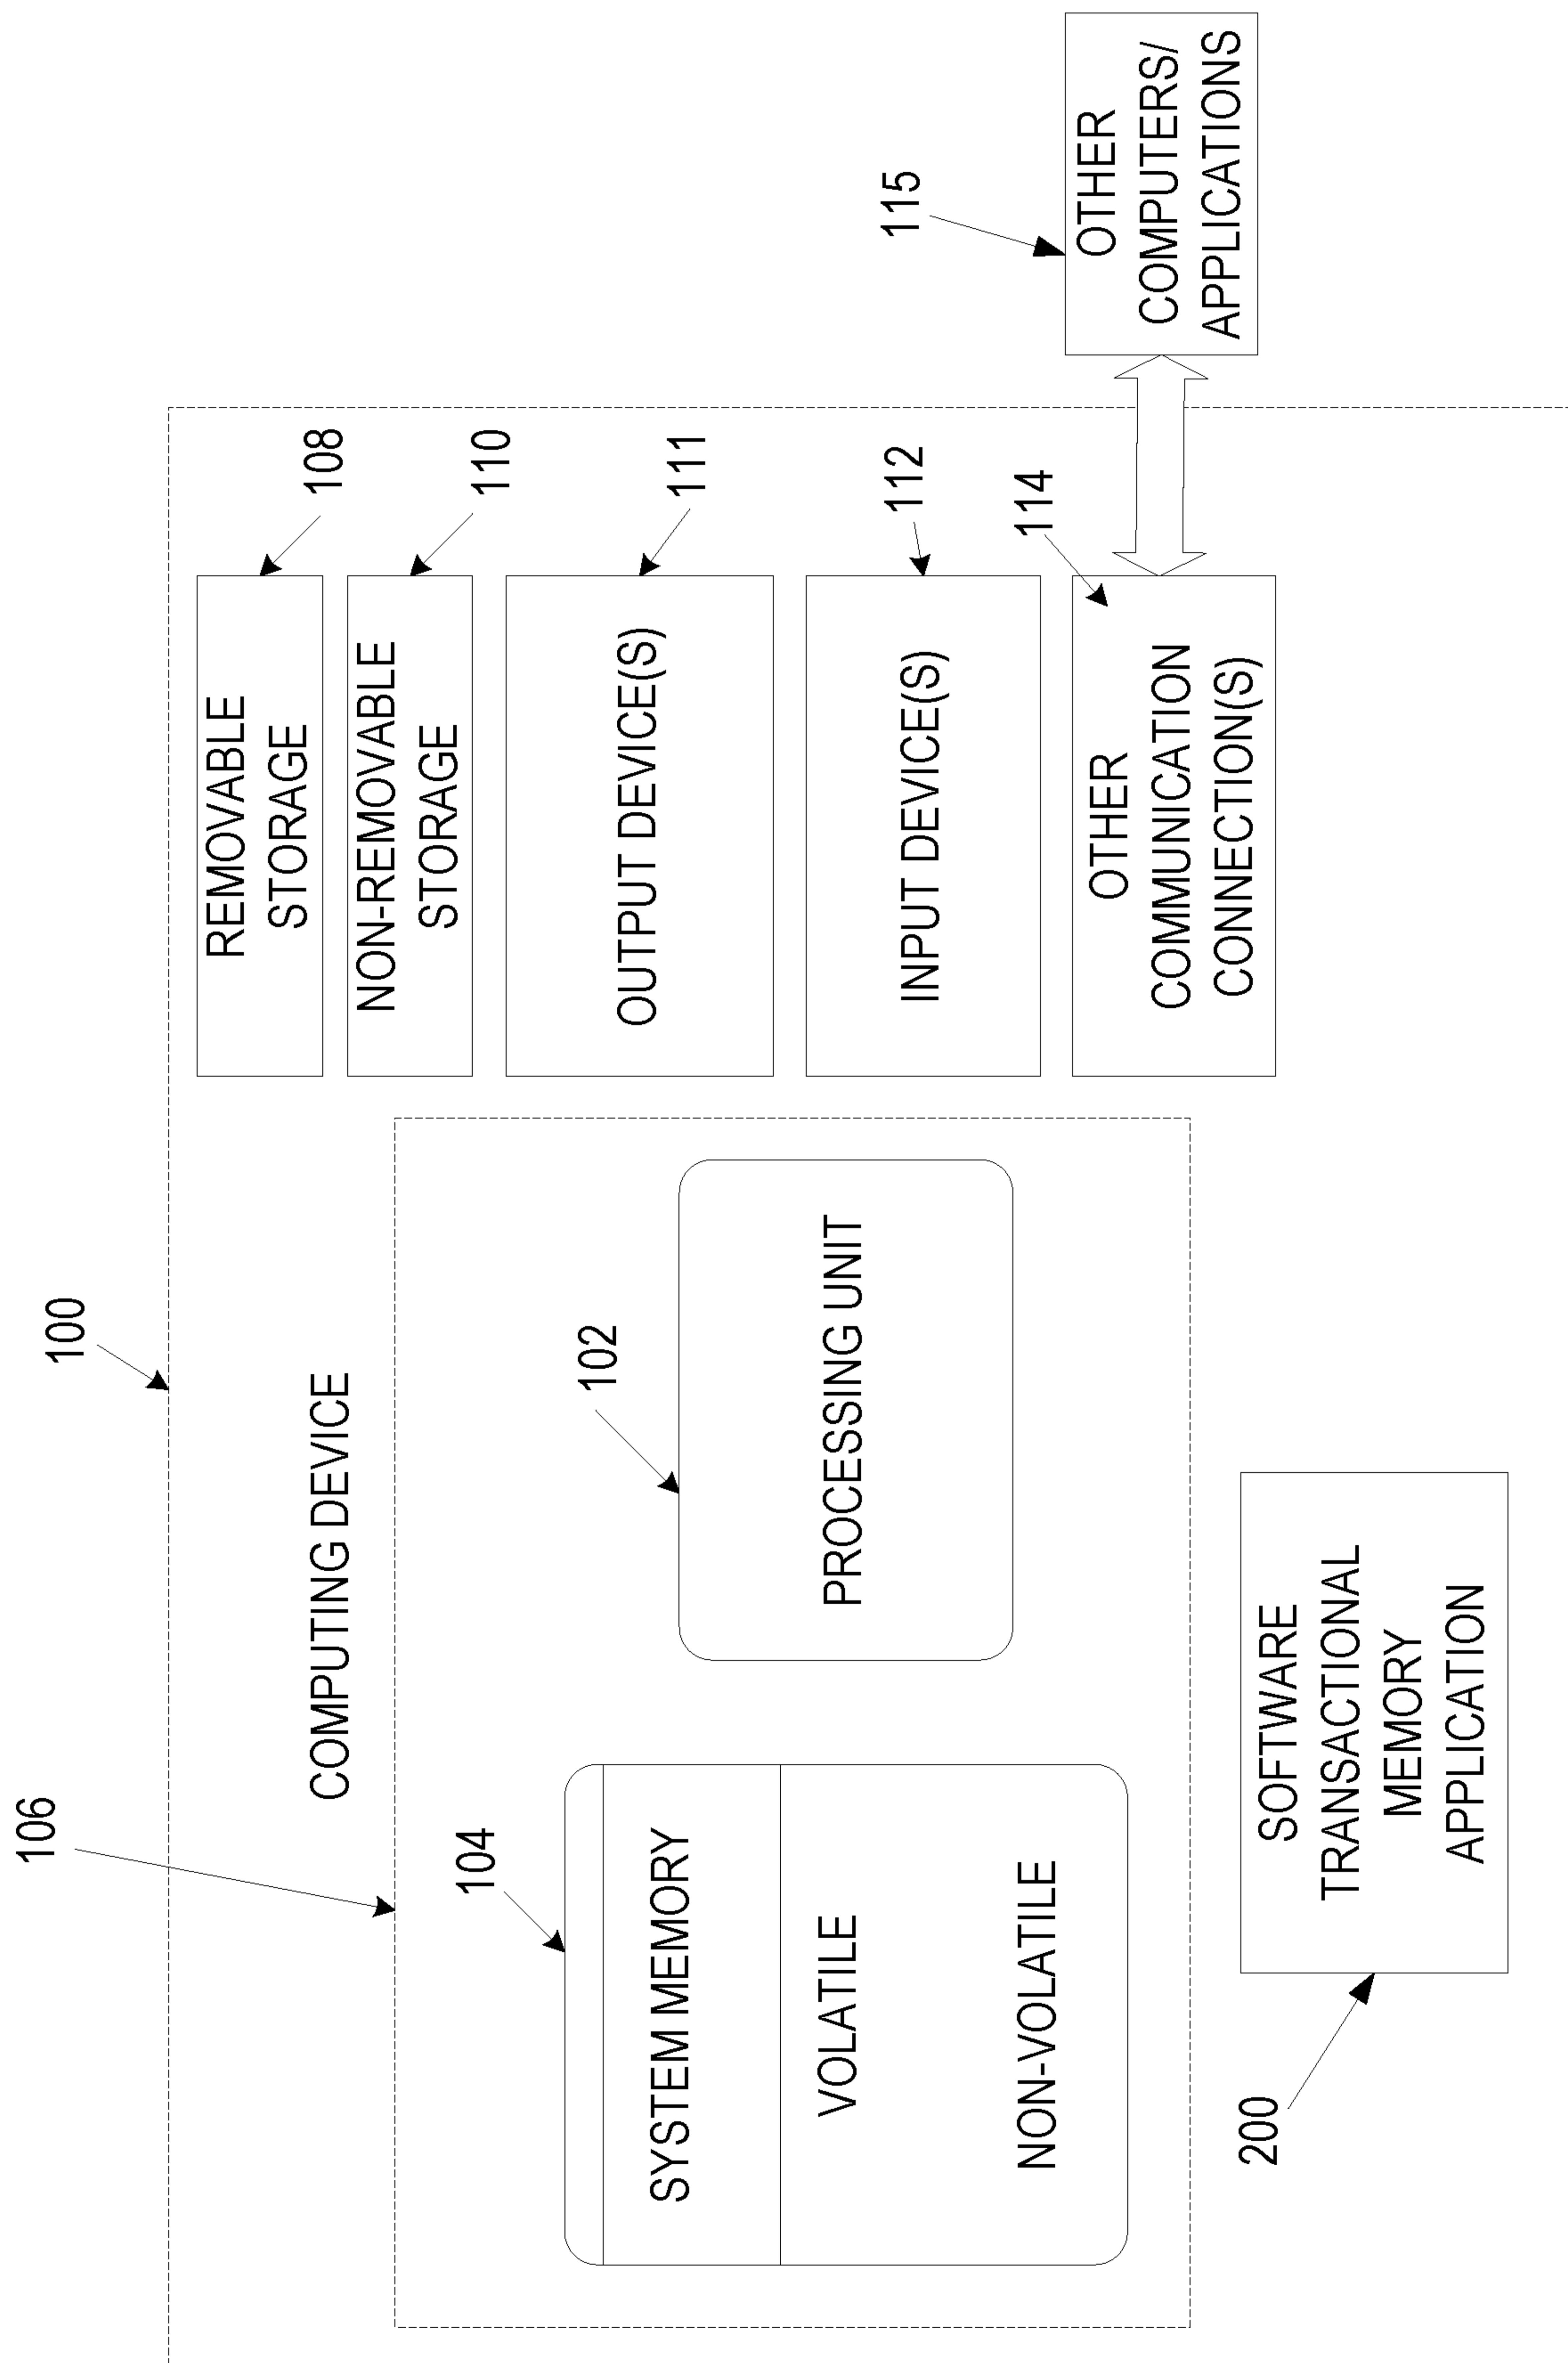


FIG. 1

SOFTWARE TRANSACTIONAL MEMORY APPLICATION

200

PROGRAM LOGIC

204LOGIC FOR PROVIDING A SOFTWARE TRANSACTIONAL MEMORY (STM) SYSTEM 206

LOGIC FOR PROVIDING A COMMIT ARBITRATOR THAT ALLOWS A PRE-DETERMINED COMMIT ORDER TO BE SPECIFIED (STATICALLY OR DYNAMICALLY) FOR A PLURALITY OF TRANSACTIONS IN THE STM SYSTEM 208

LOGIC FOR ALLOWING THE COMMIT ARBITRATOR TO USE THE PRE-DETERMINED COMMIT ORDER AT RUNTIME TO AID IN DETERMINING AN ORDER IN WHICH TO COMMIT THE PLURALITY OF TRANSACTIONS IN THE SOFTWARE TRANSACTIONAL MEMORY SYSTEM 210

LOGIC FOR PROVIDING A CONTENTION MANAGEMENT PROCESS THAT IS INVOKED WHEN A CONFLICT OCCURS BETWEEN A FIRST TRANSACTION AND A SECOND TRANSACTION 212

LOGIC FOR USING THE PRE-DETERMINED COMMIT ORDER IN THE CONTENTION MANAGEMENT PROCESS TO AID IN DETERMINING WHETHER THE FIRST TRANSACTION OR THE SECOND TRANSACTION SHOULD WIN THE CONFLICT AND BE ALLOWED TO PROCEED (E.G. DEPENDING ON WHICH ONE HAS THE LOWER COMMIT ORDER NUMBER OF TWO TRANSACTIONS IN THE SAME TRANSACTION GROUP) 214

LOGIC FOR ALLOWING THE COMMIT ARBITRATOR TO BE OPERABLE TO USE THE PRE-DETERMINED ORDERING TO TRACK ONE OR MORE ORDERING VALUES (E.G. IN TOTAL ORDERING - A NEXT-TO-COMMIT FIELD THAT REPRESENTS A NEXT TRANSACTION OF THE PLURALITY OF TRANSACTIONS THAT SHOULD BE ALLOWED TO COMMIT) AND FOR COMPARING THE ONE OR MORE ORDERING VALUES TO A PARTICULAR COMMIT ORDER NUMBER OF A GIVEN TRANSACTION TO SEE IF THE COMMIT OF THE TRANSACTION IS PROPER GIVEN THE ORDERING THAT SHOULD BE ENFORCED 216

OTHER LOGIC FOR OPERATING THE APPLICATION 220

FIG. 2

3 / 11

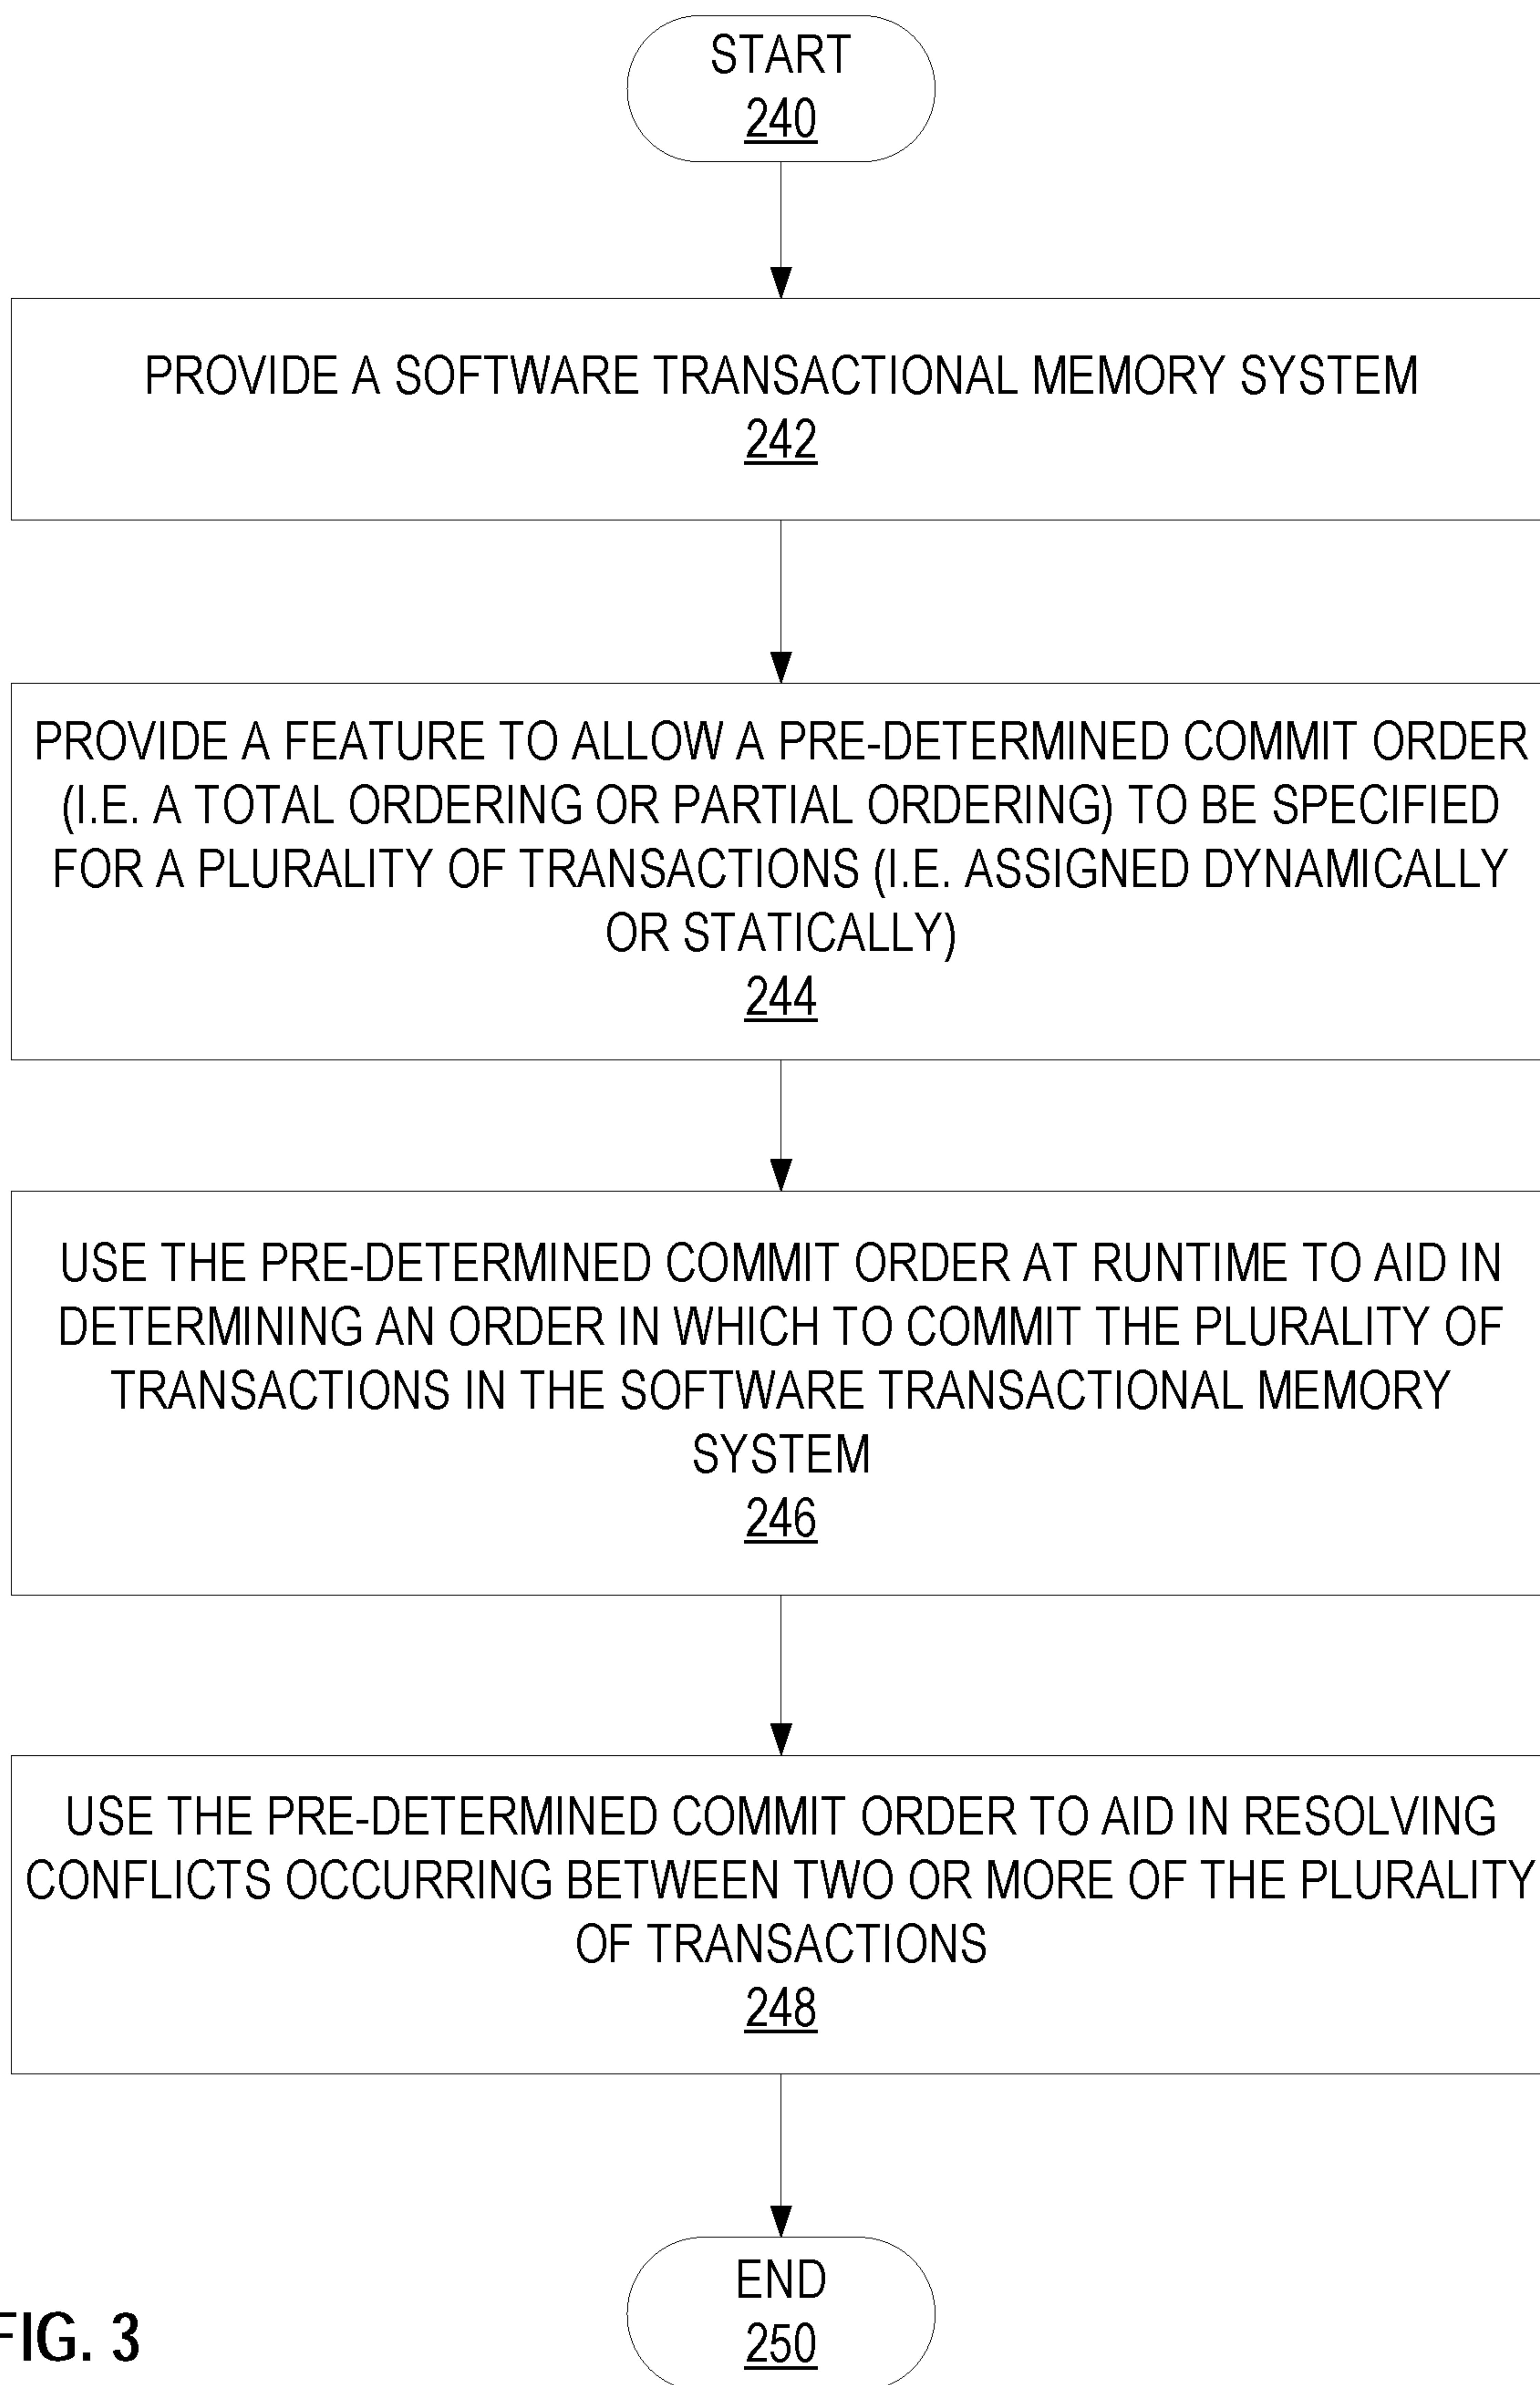


FIG. 3

4 / 11

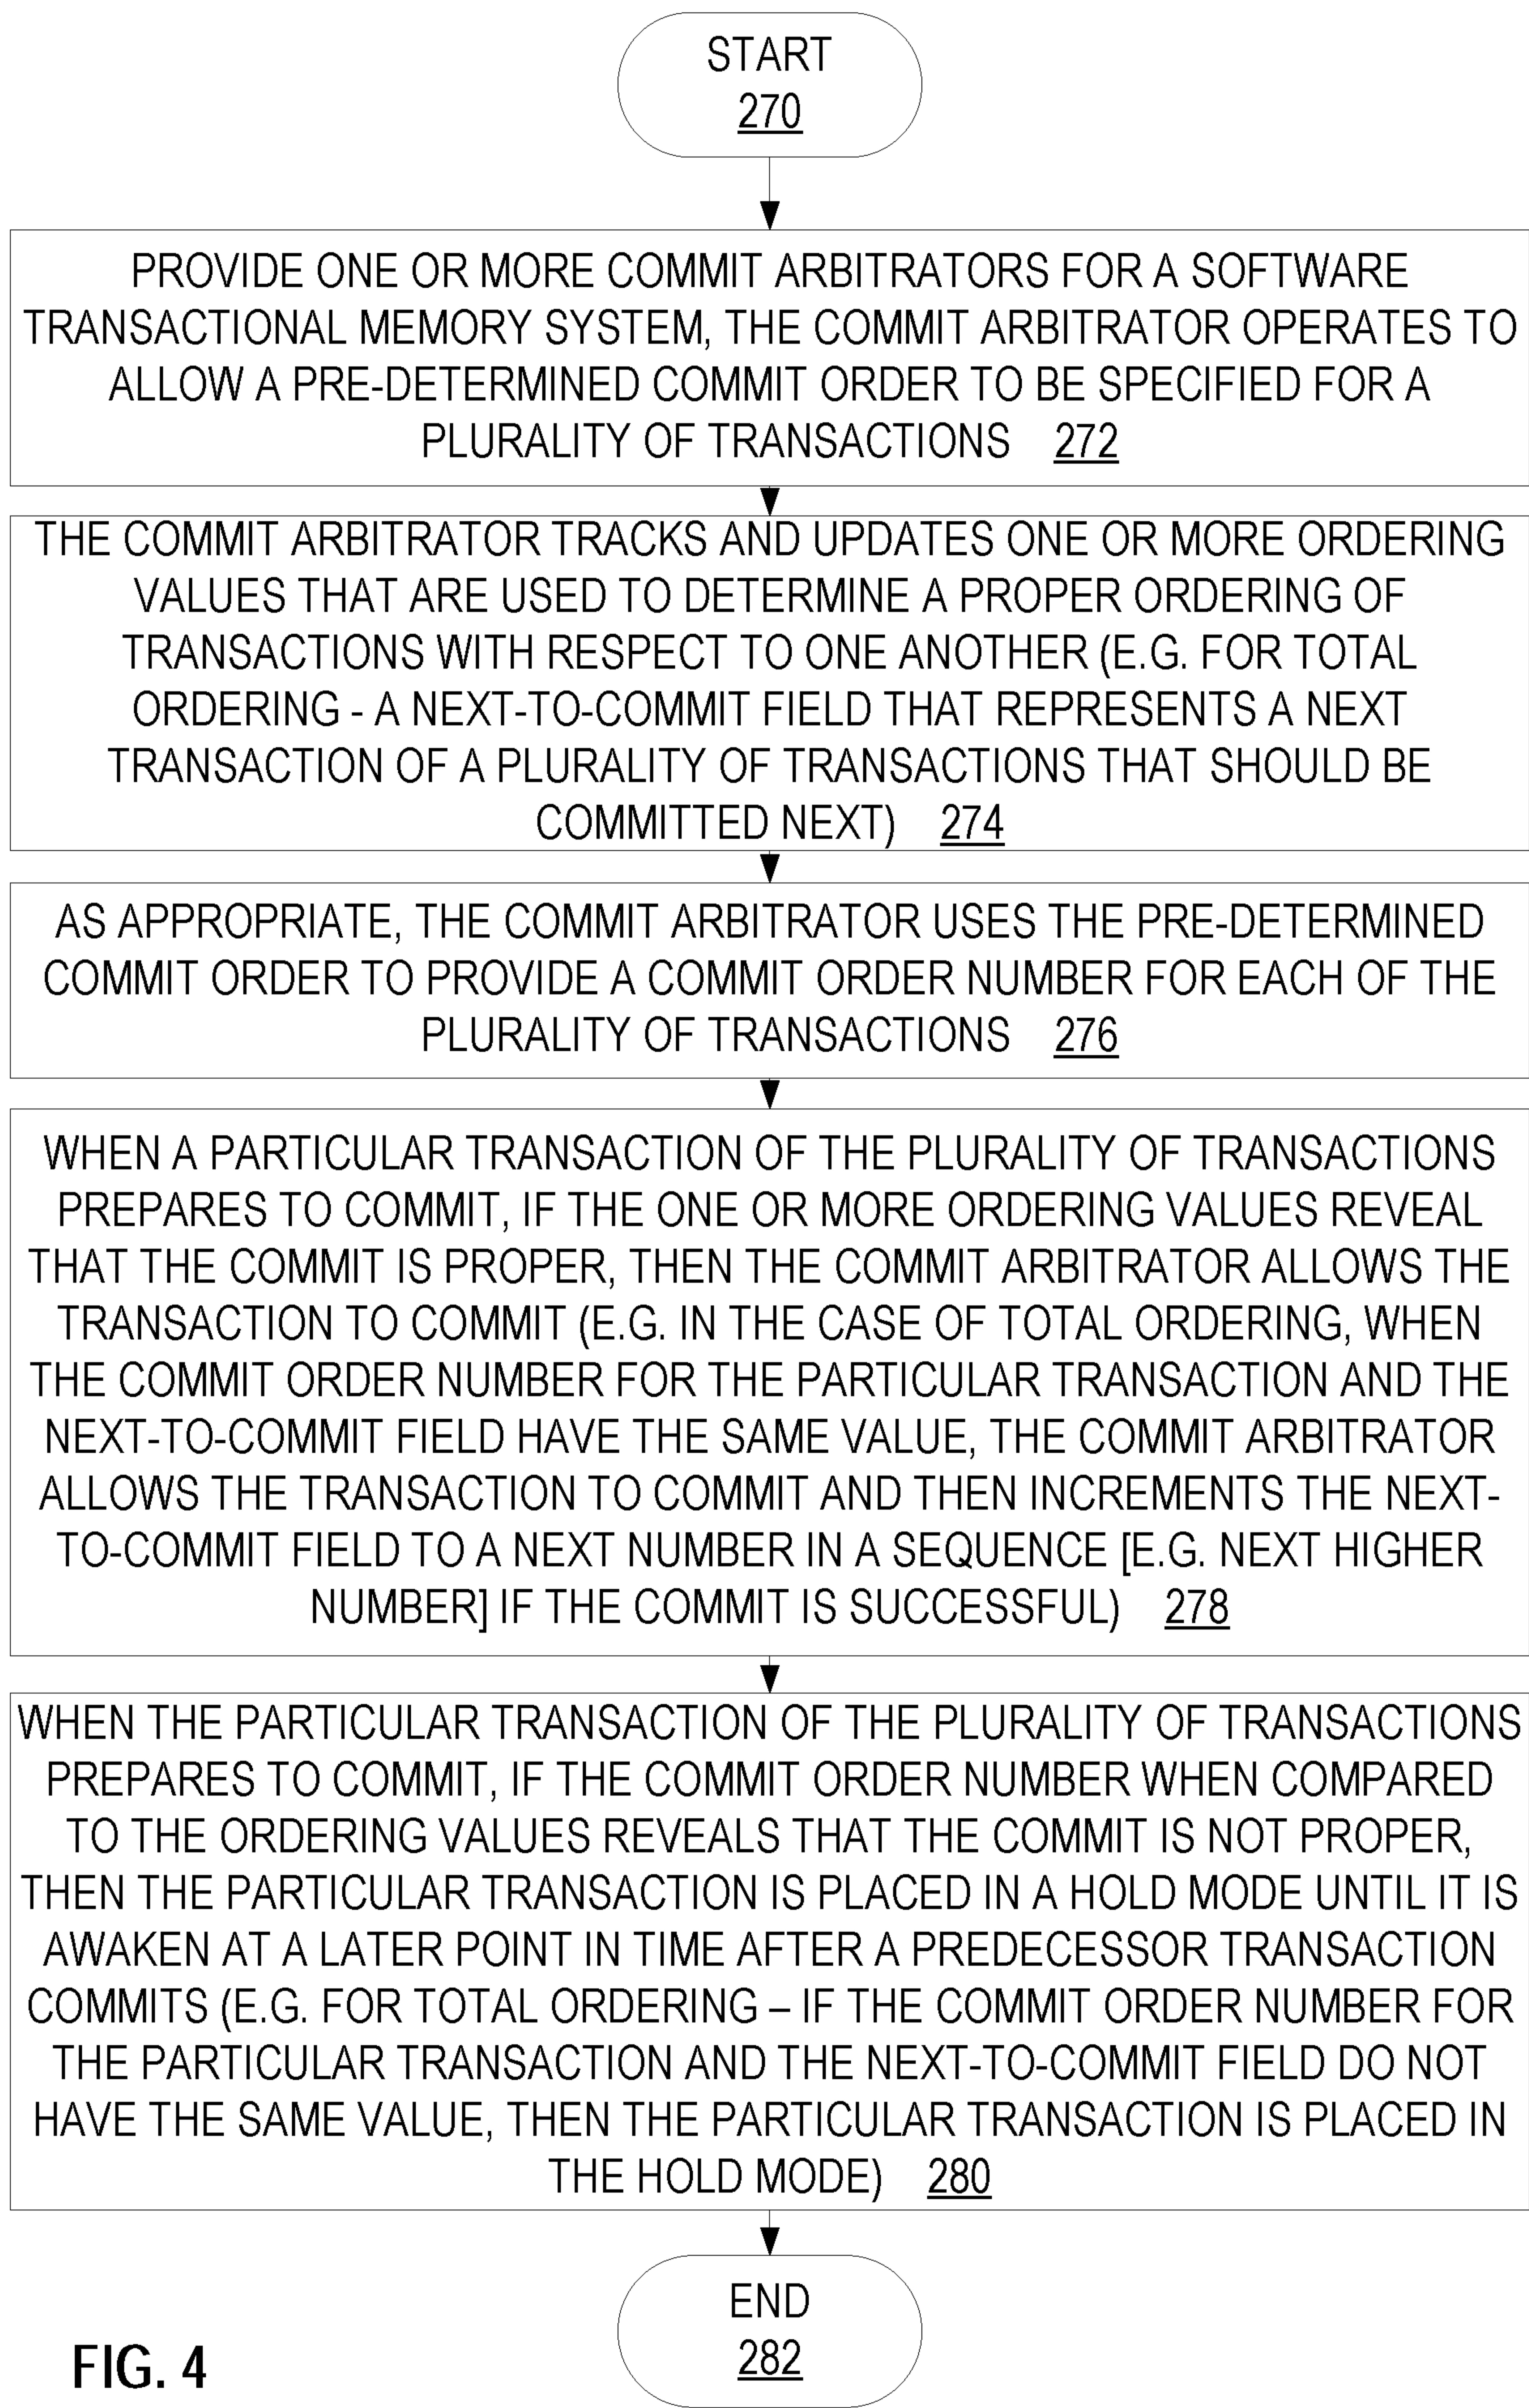


FIG. 4

5 / 11

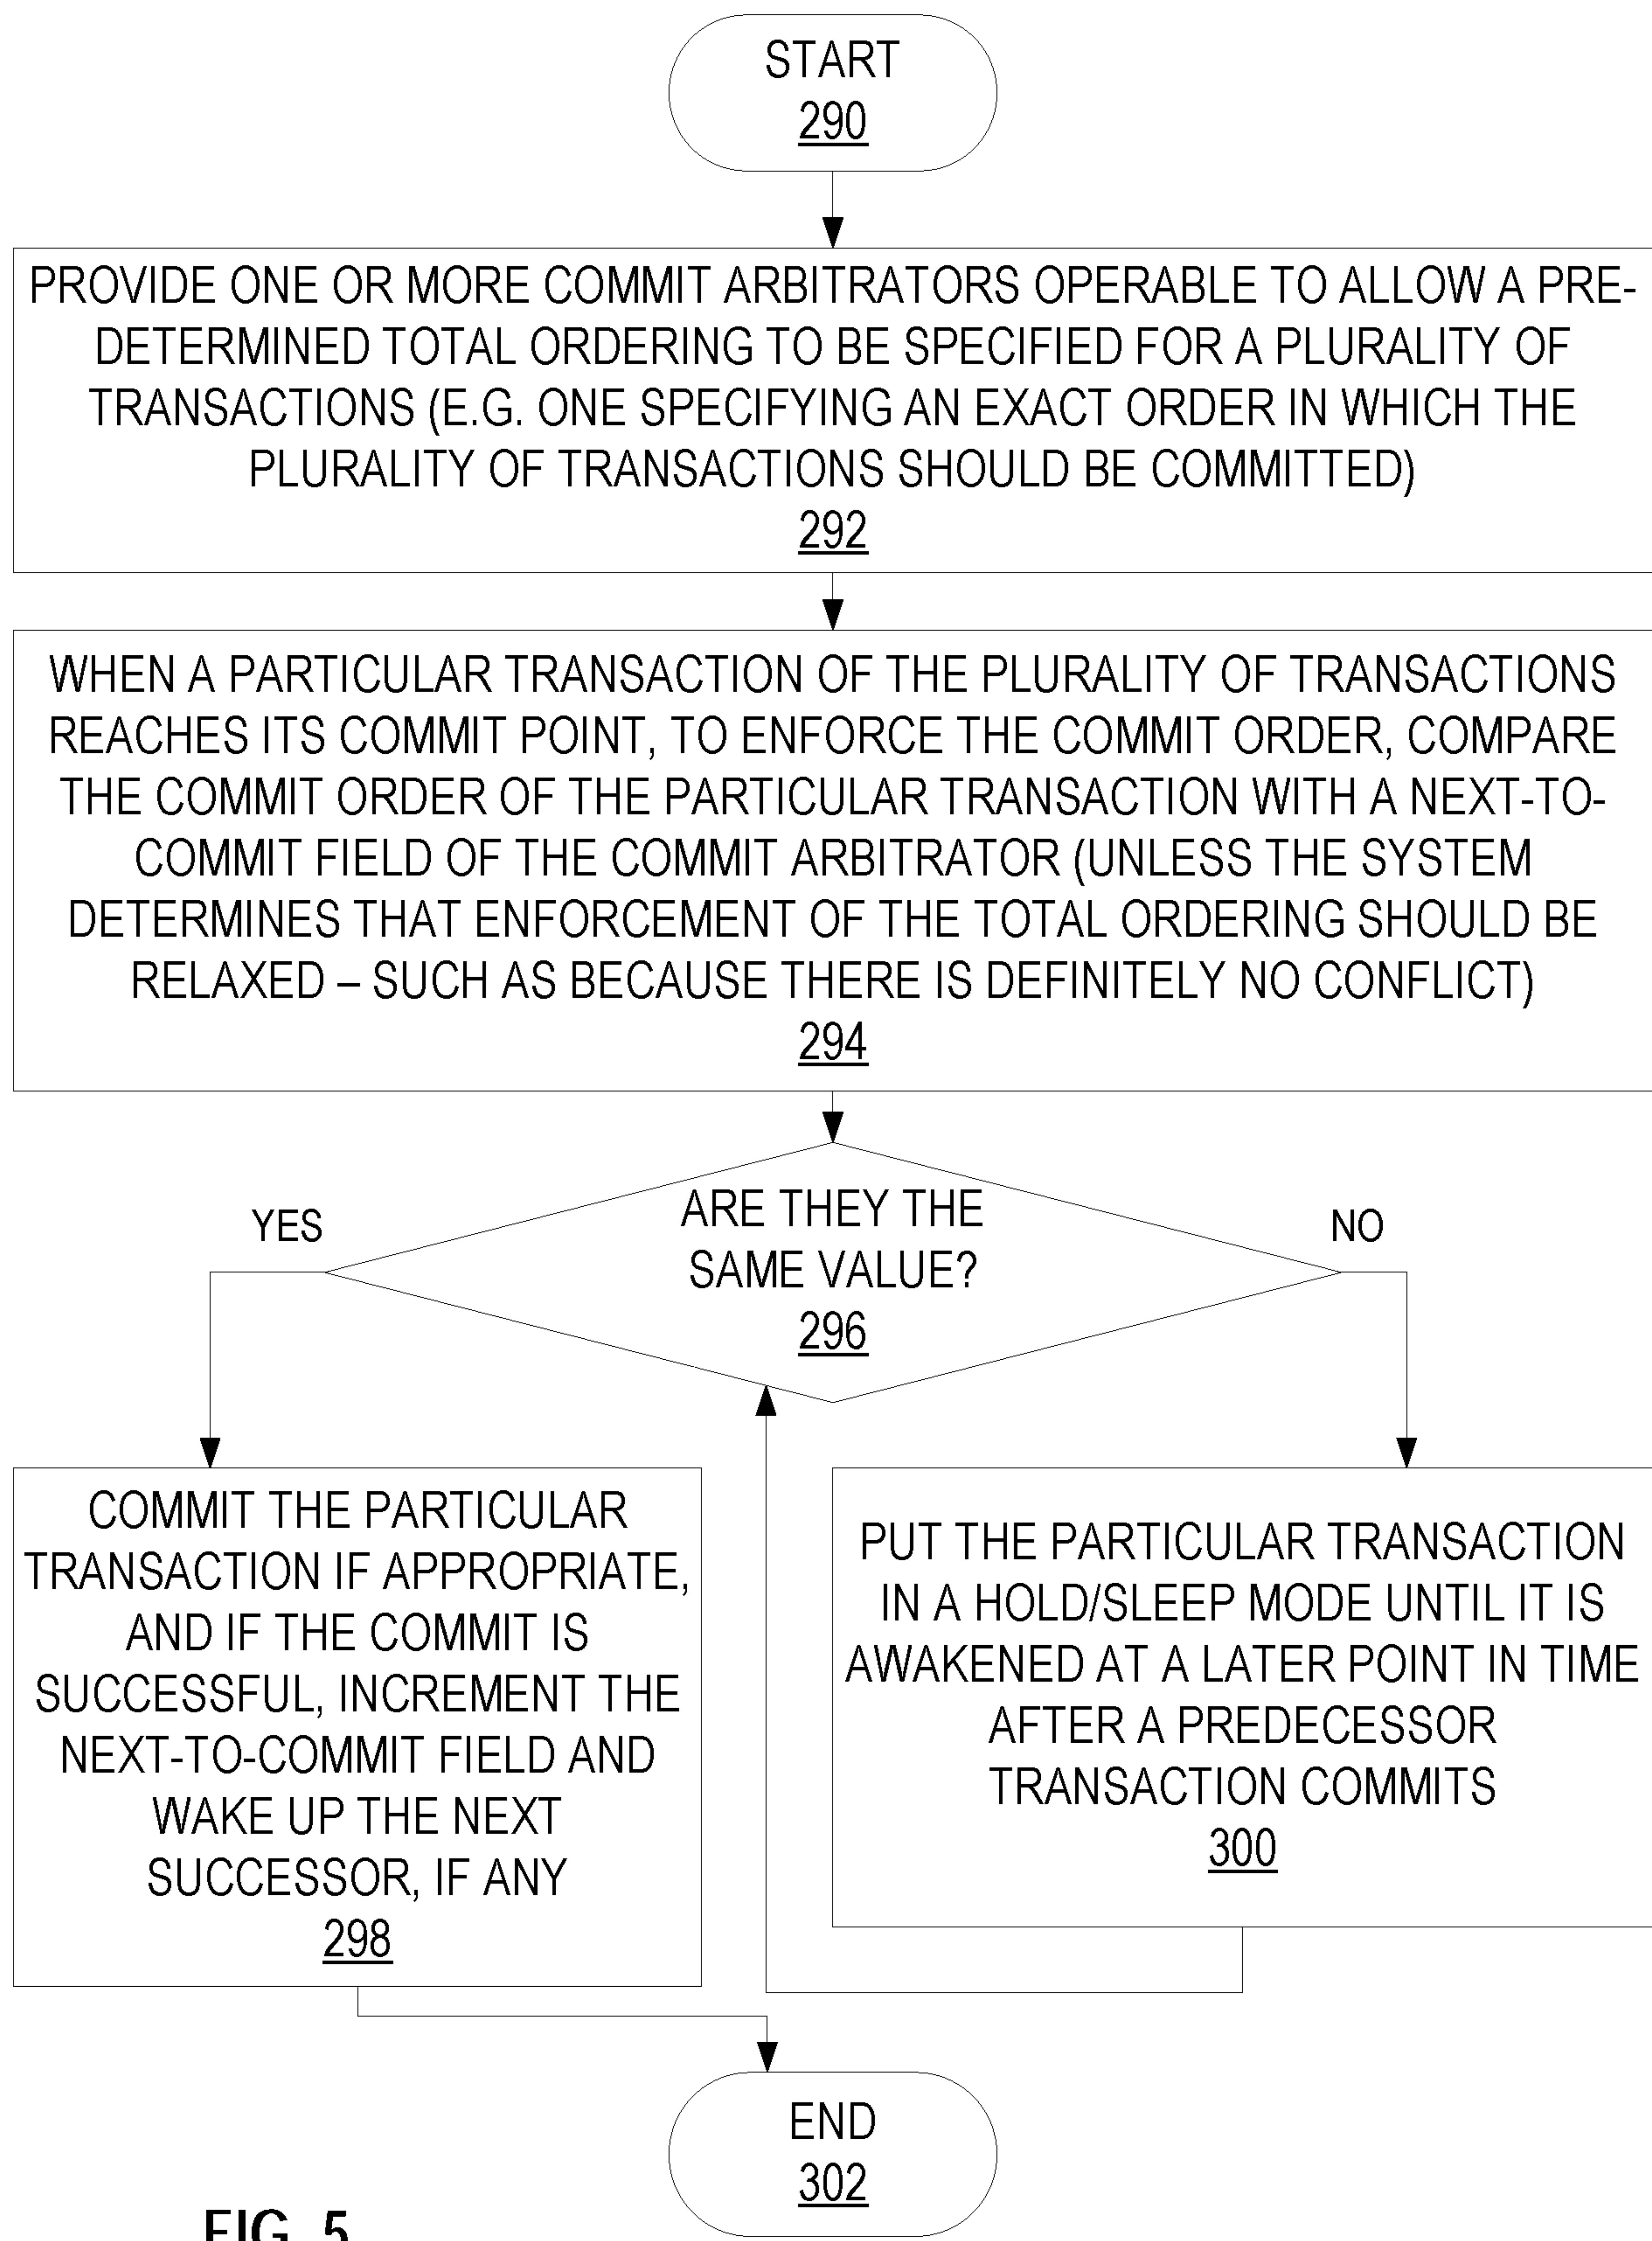


FIG. 5

6 / 11

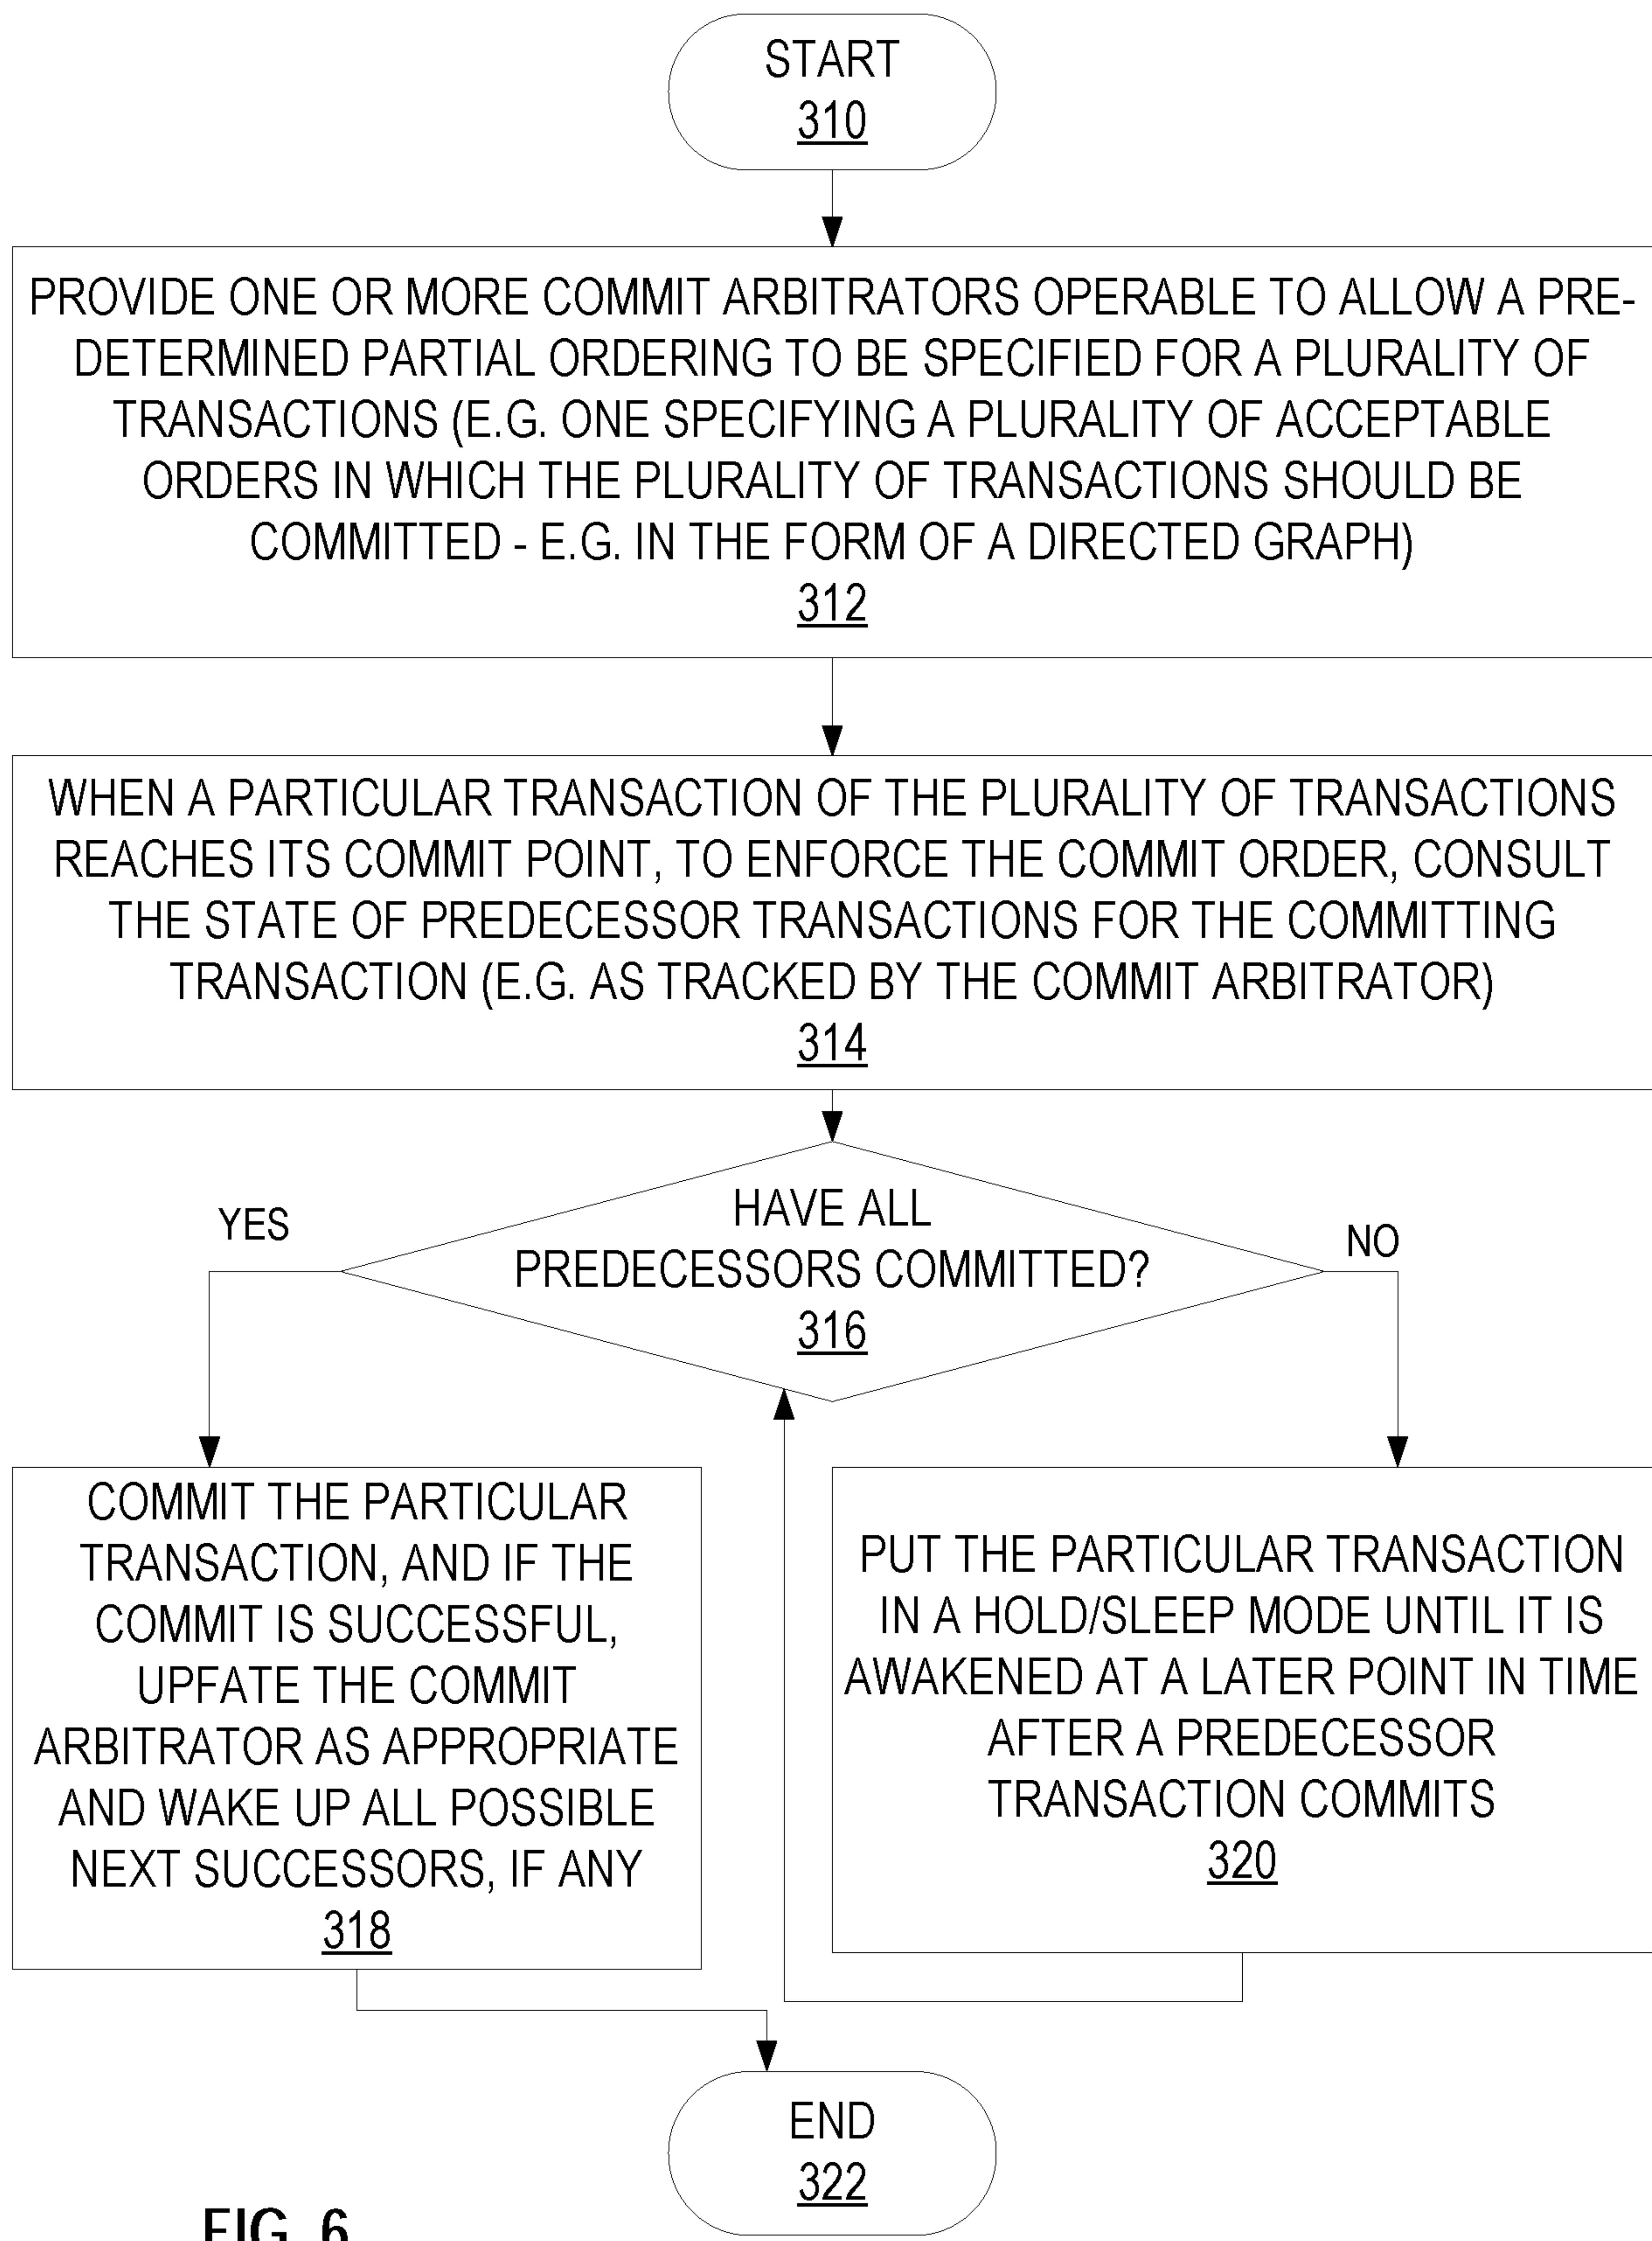


FIG. 6

7 / 11



START
340

PROVIDE A SOFTWARE TRANSACTIONAL MEMORY SYSTEM THAT SUPPORTS A PRE-DETERMINED COMMIT ORDER FOR ONE OR MORE GROUPS OF TRANSACTIONS

342

PROVIDE A CONTENTION MANAGEMENT PROCESS THAT IS INVOKED WHEN A CONFLICT OCCURS BETWEEN A FIRST TRANSACTION AND A SECOND TRANSACTION

344

USE THE PRE-DETERMINED COMMIT ORDER IN THE CONTENTION MANAGEMENT PROCESS TO AID IN DETERMINING WHETHER THE FIRST TRANSACTION OR THE SECOND TRANSACTION SHOULD WIN THE CONFLICT AND BE ALLOWED TO PROCEED

346

NO

ARE THEY PART OF THE SAME TRANSACTION GROUP?

YES

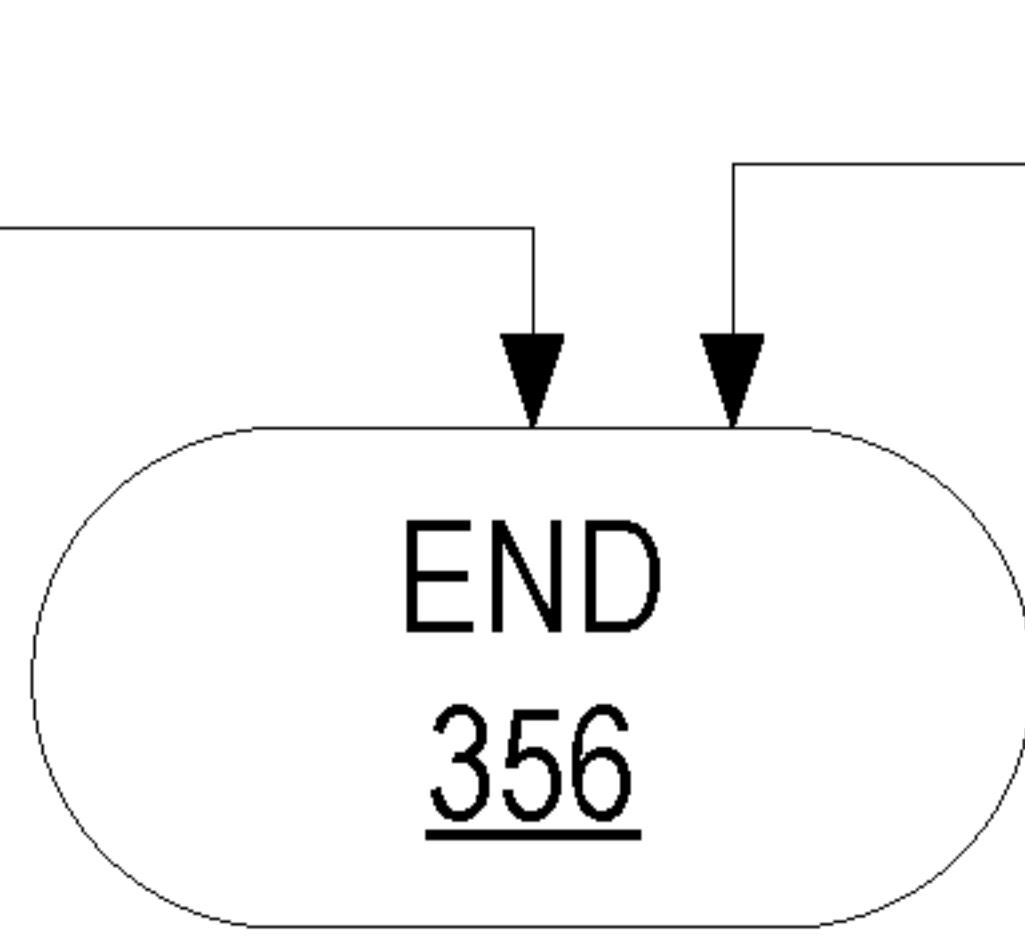
PRE-DETERMINED COMMIT ORDER IS NOT ENFORCED BETWEEN THESE TWO TRANSACTIONS (SINCE THEY ARE NOT IN A SAME TRANSACTION GROUP, NONE EXISTED)

350

COMPARE THE FIRST ORDER NUMBER OF THE FIRST TRANSACTION AND THE SECOND ORDER NUMBER OF THE SECOND TRANSACTION

352

ALLOW THE TRANSACTION WITH THE LOWER ORDER NUMBER TO PROCEED (OR OTHER SUITABLE PRIORITY ORDERING)

354


END
356

FIG. 7

8 / 11

START

370

PROVIDE A CONTENTION MANAGEMENT PROCESS THAT IS INVOKED WHEN A CONFLICT OCCURS BETWEEN A FIRST TRANSACTION AND A SECOND TRANSACTION

372

USE A PRE-DETERMINED COMMIT ORDER IN THE CONTENTION MANAGEMENT PROCESS TO AID IN DETERMINING WHETHER THE FIRST TRANSACTION OR THE SECOND TRANSACTION SHOULD WIN THE CONFLICT AND BE ALLOWED TO PROCEED

374

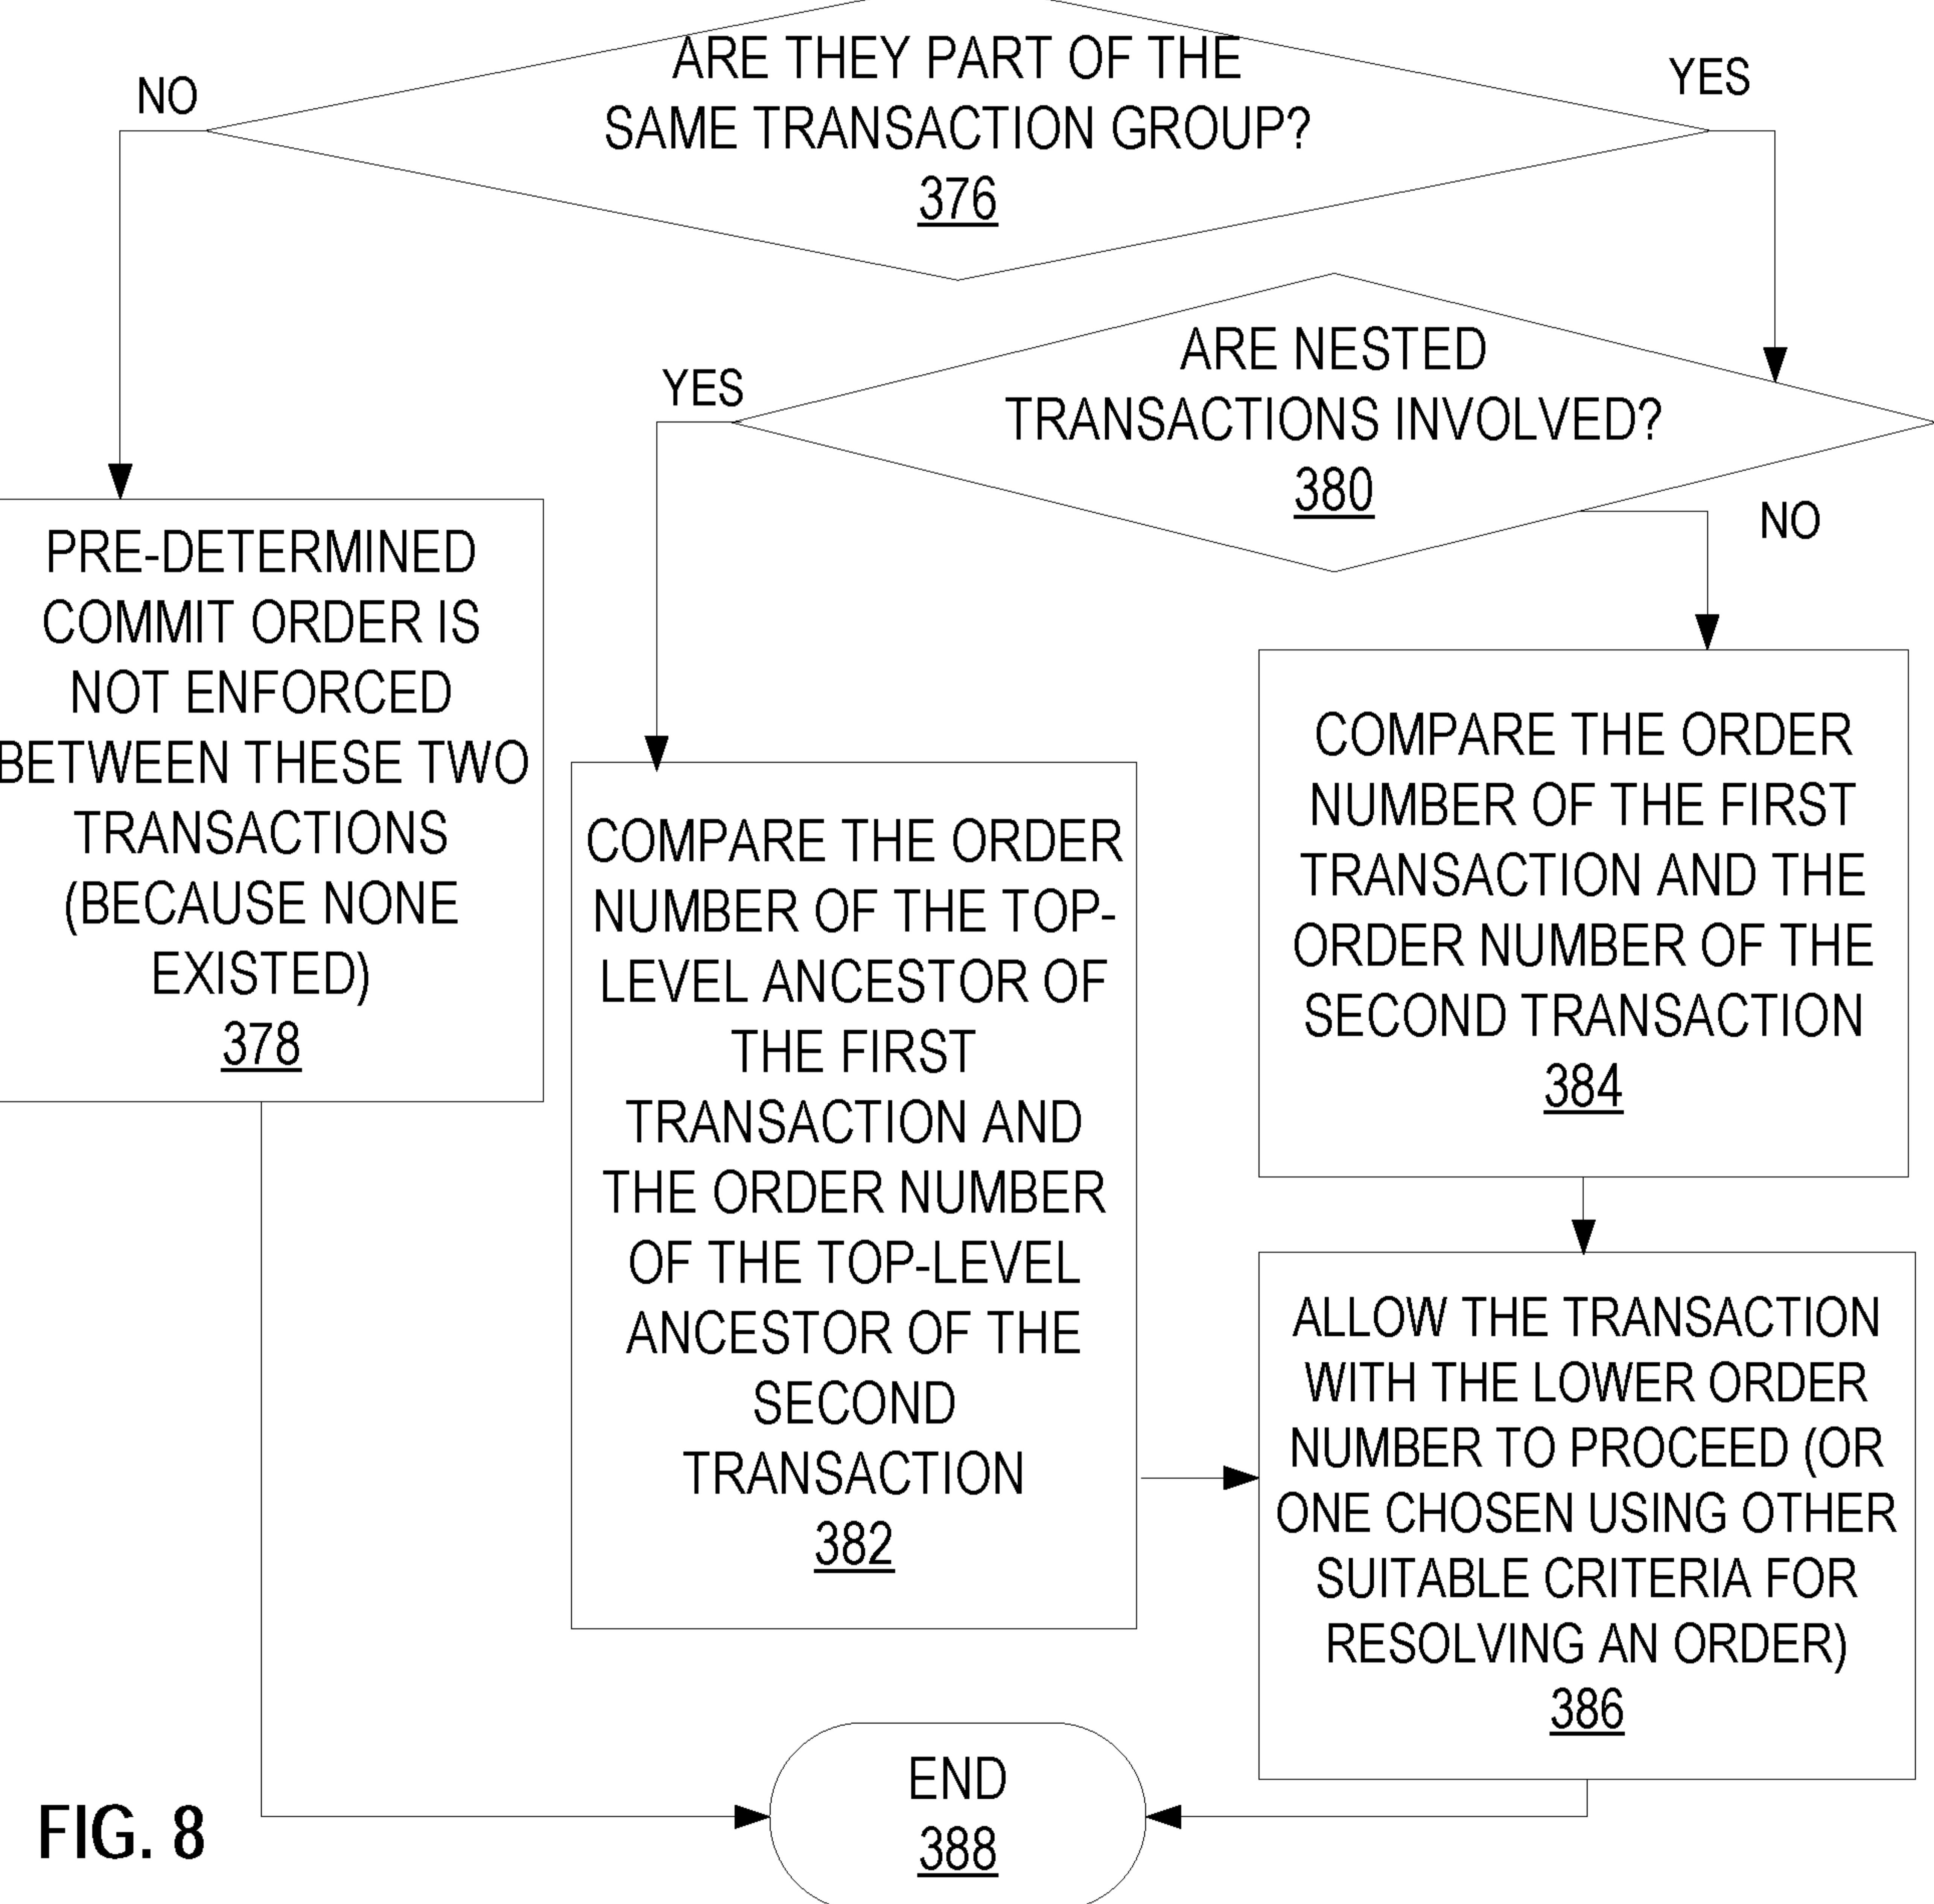


FIG. 8

9 / 11

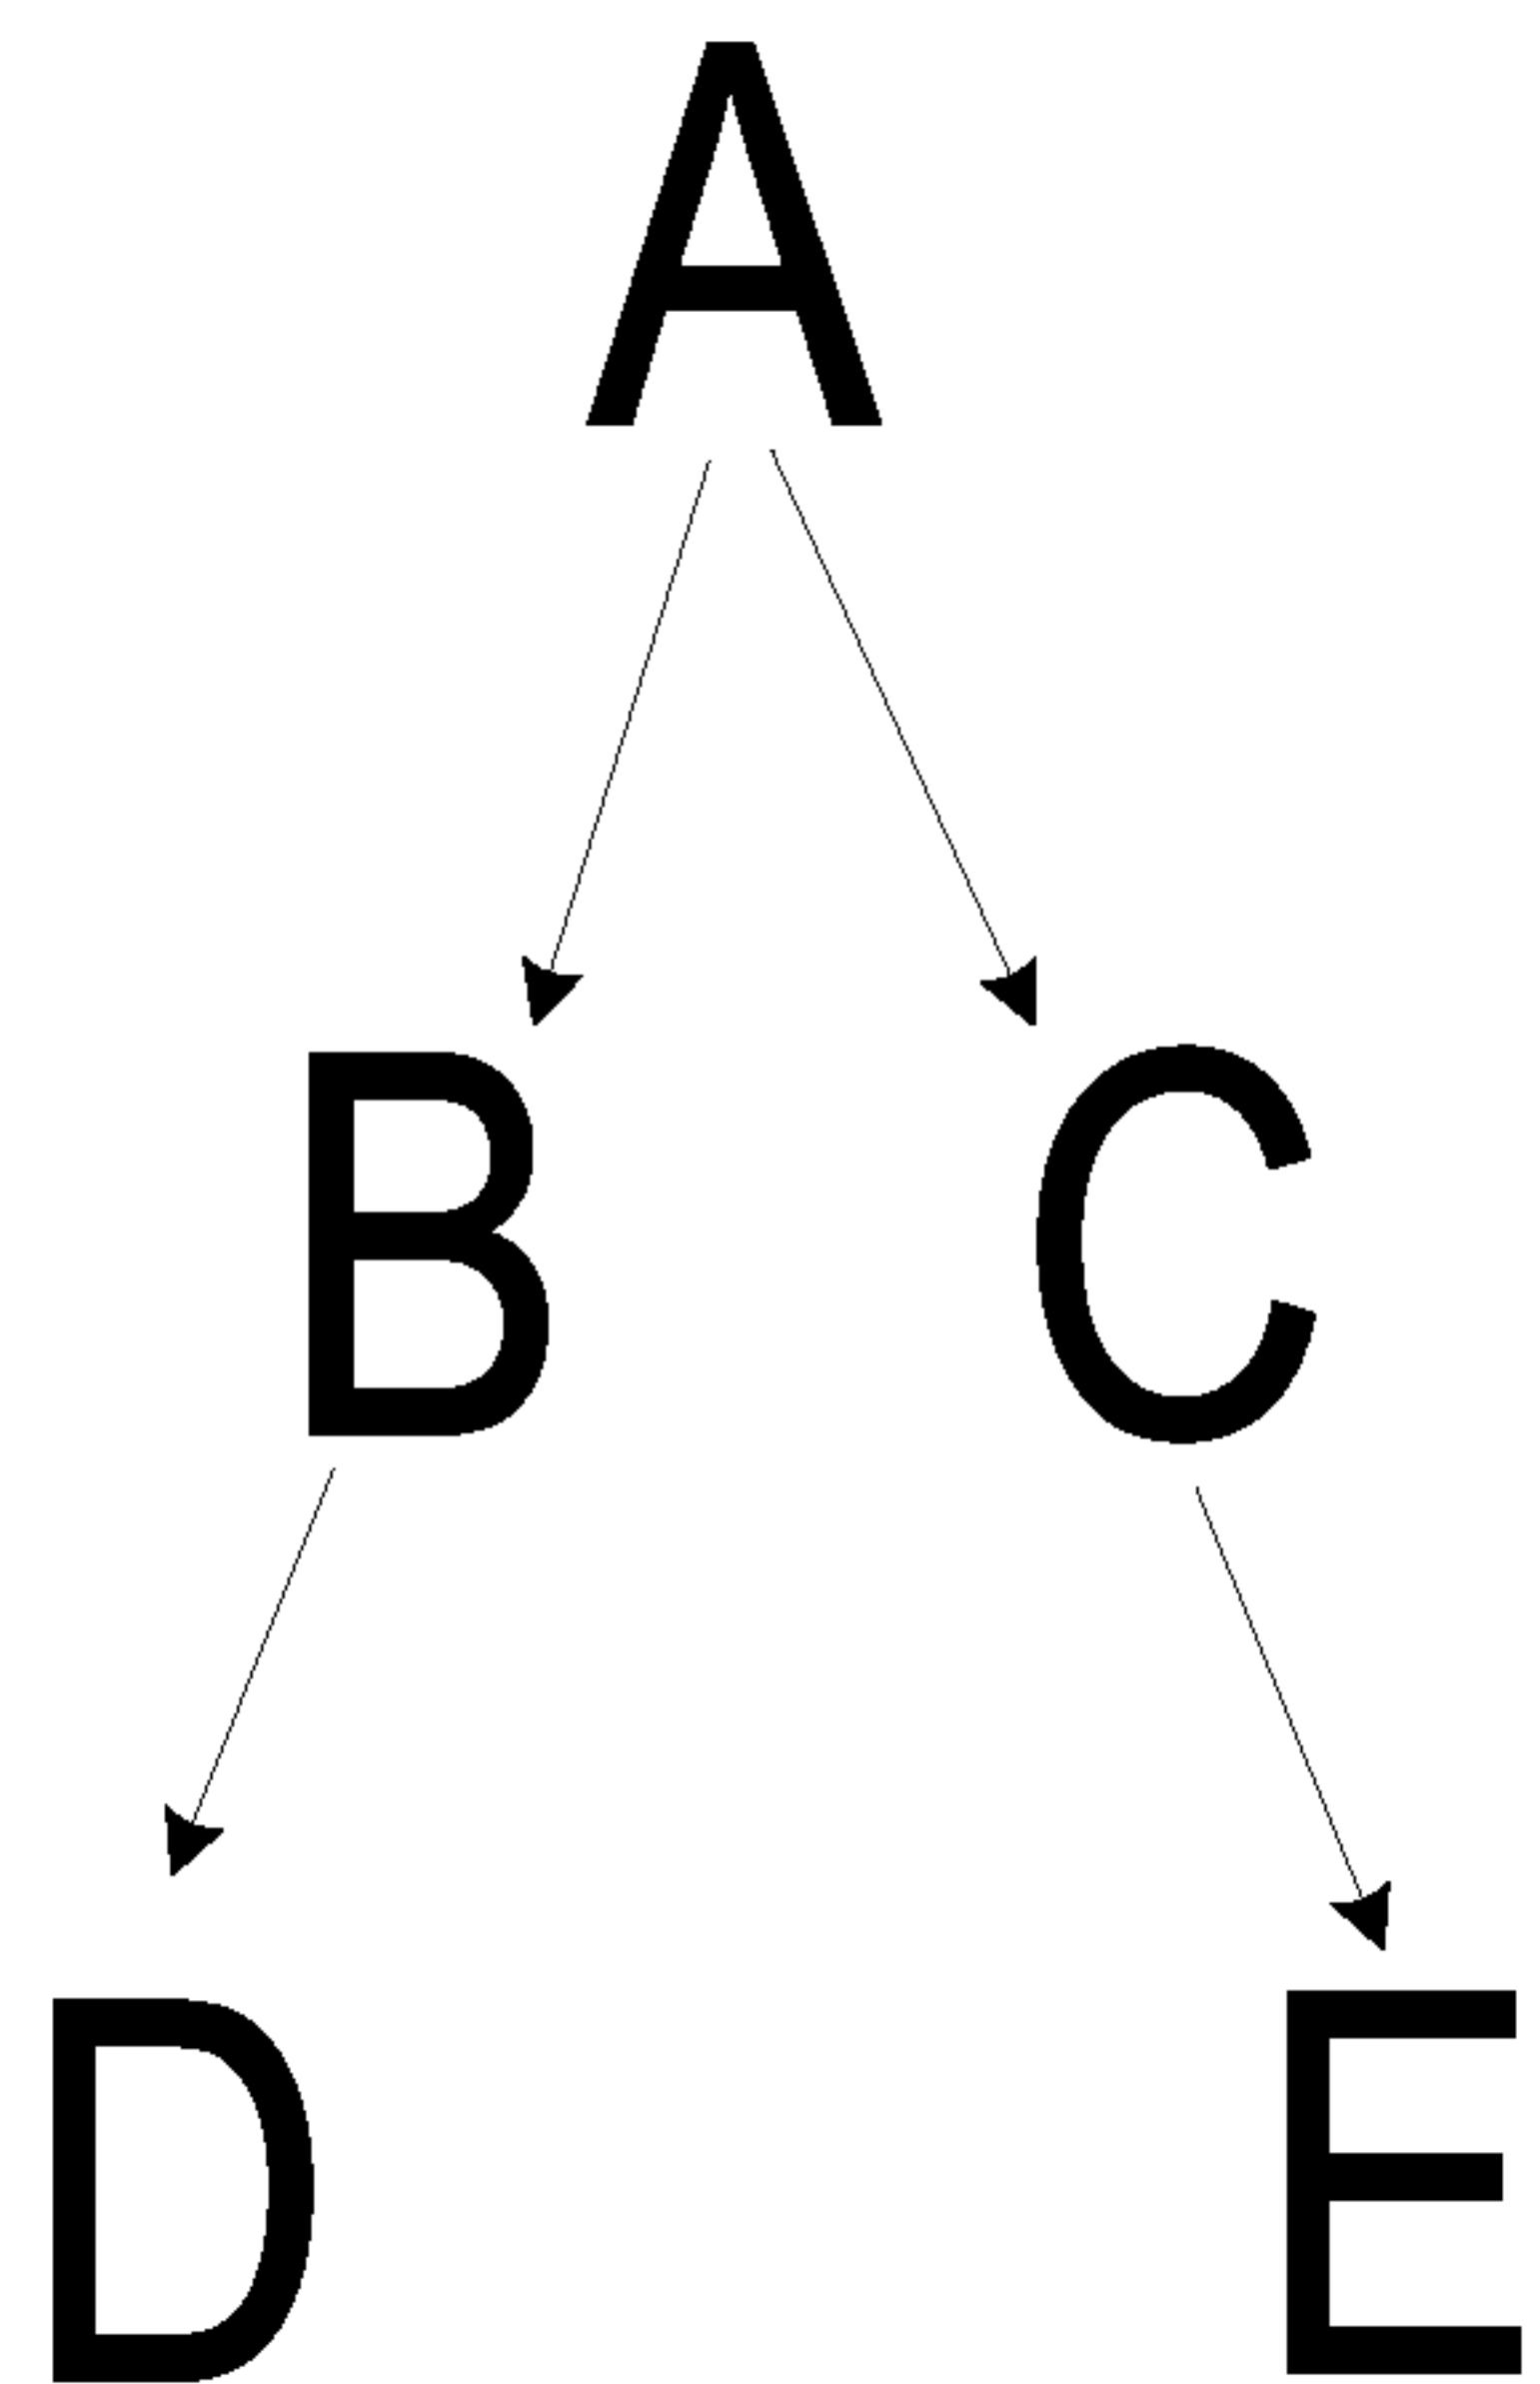


FIG. 9

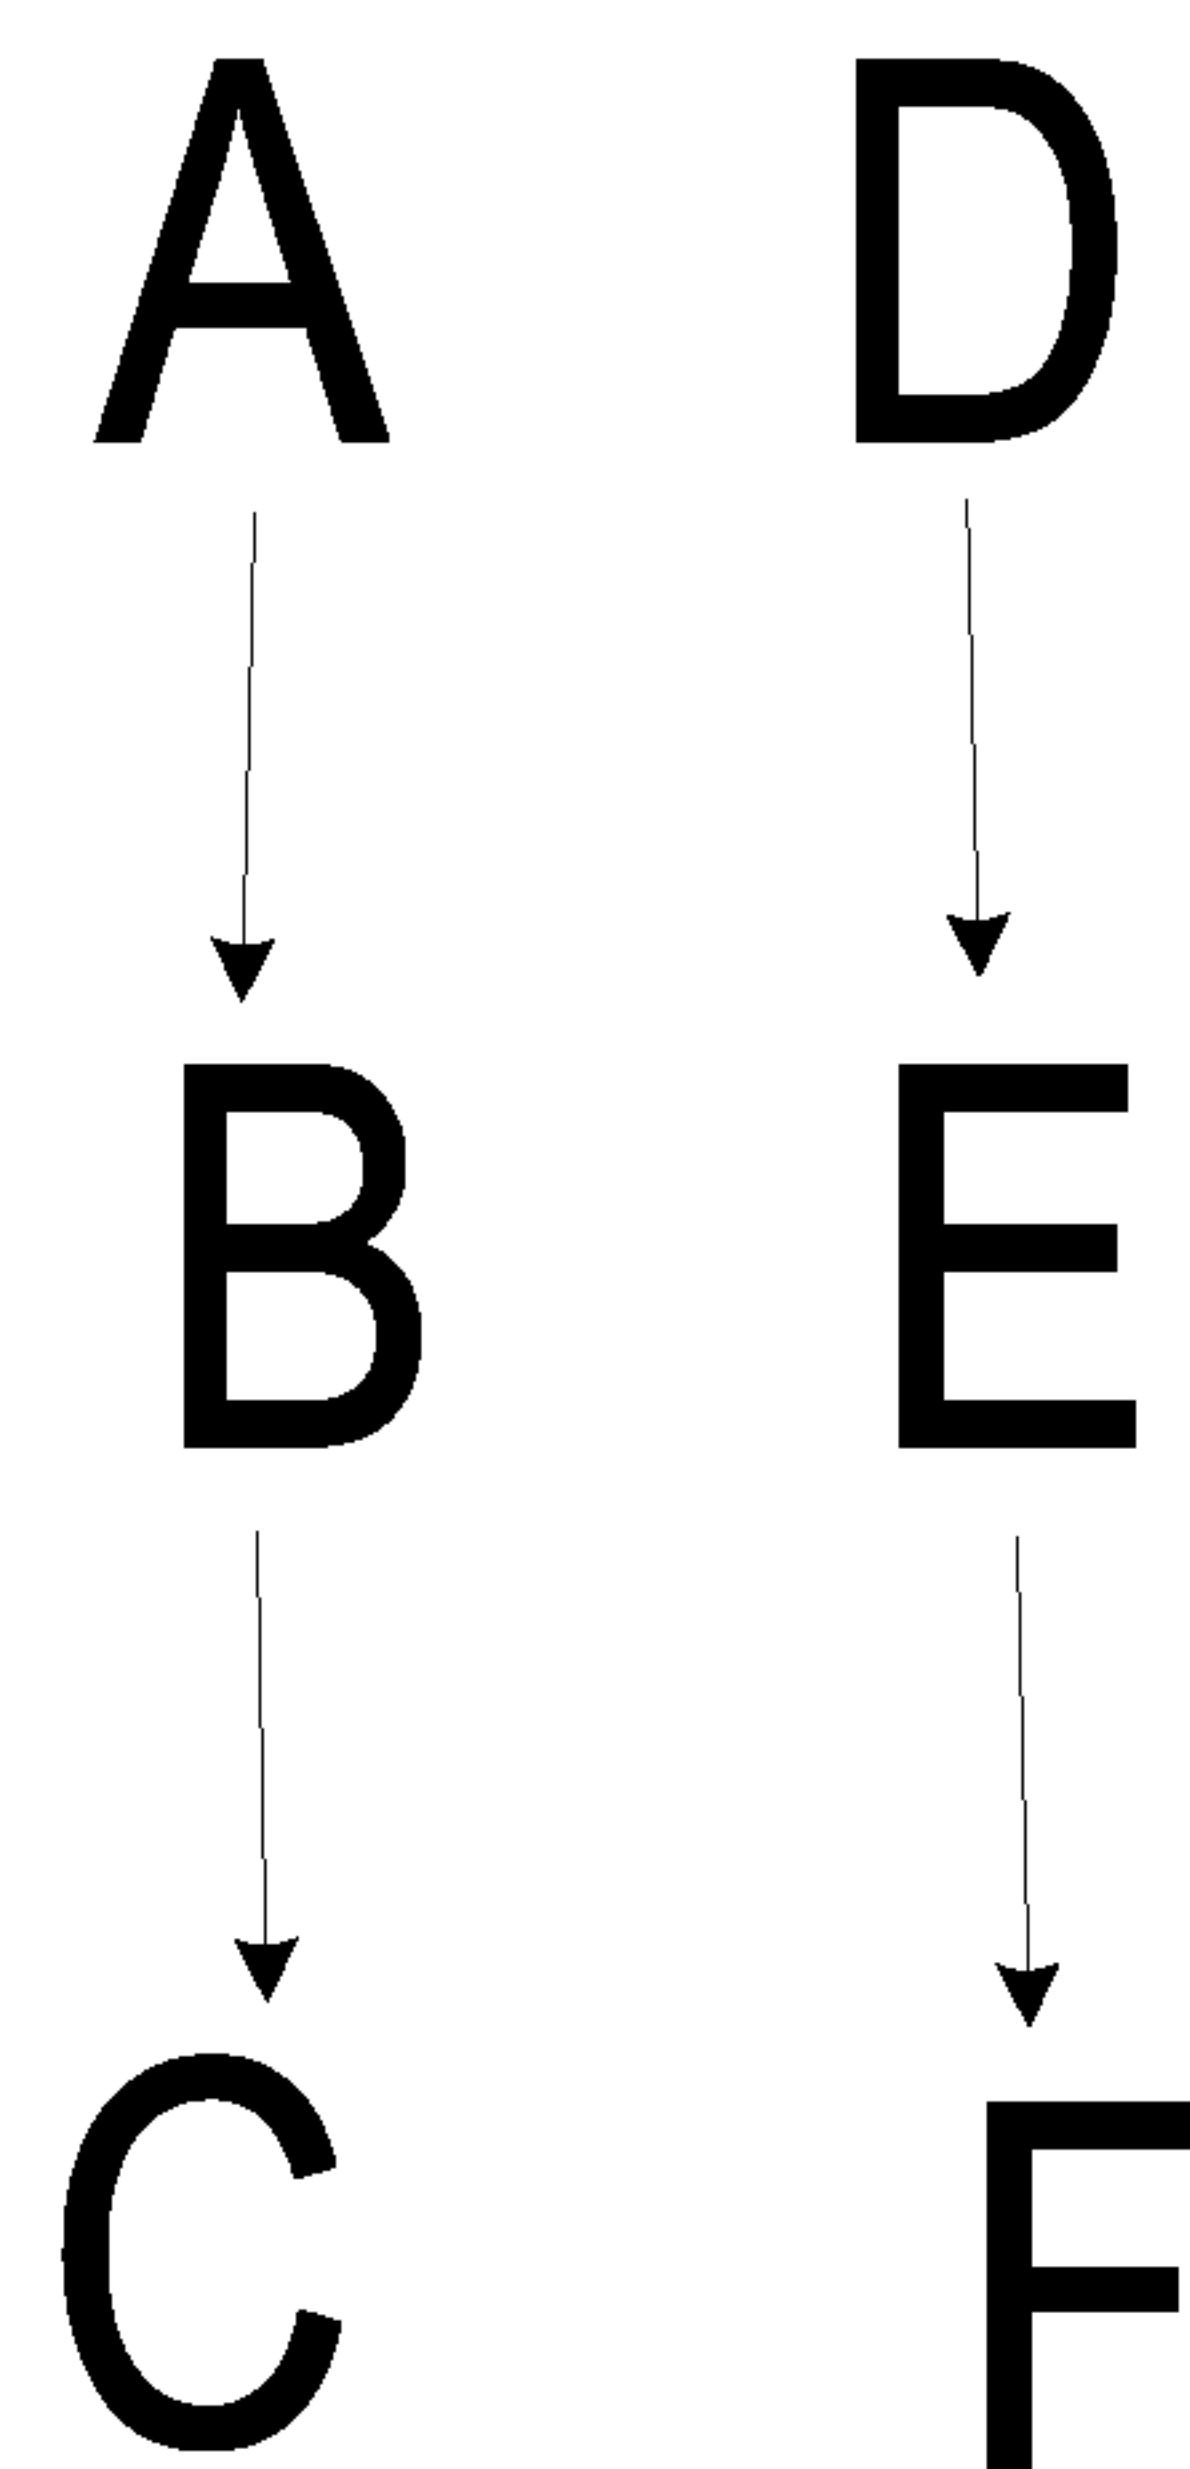


FIG. 10

10 / 11

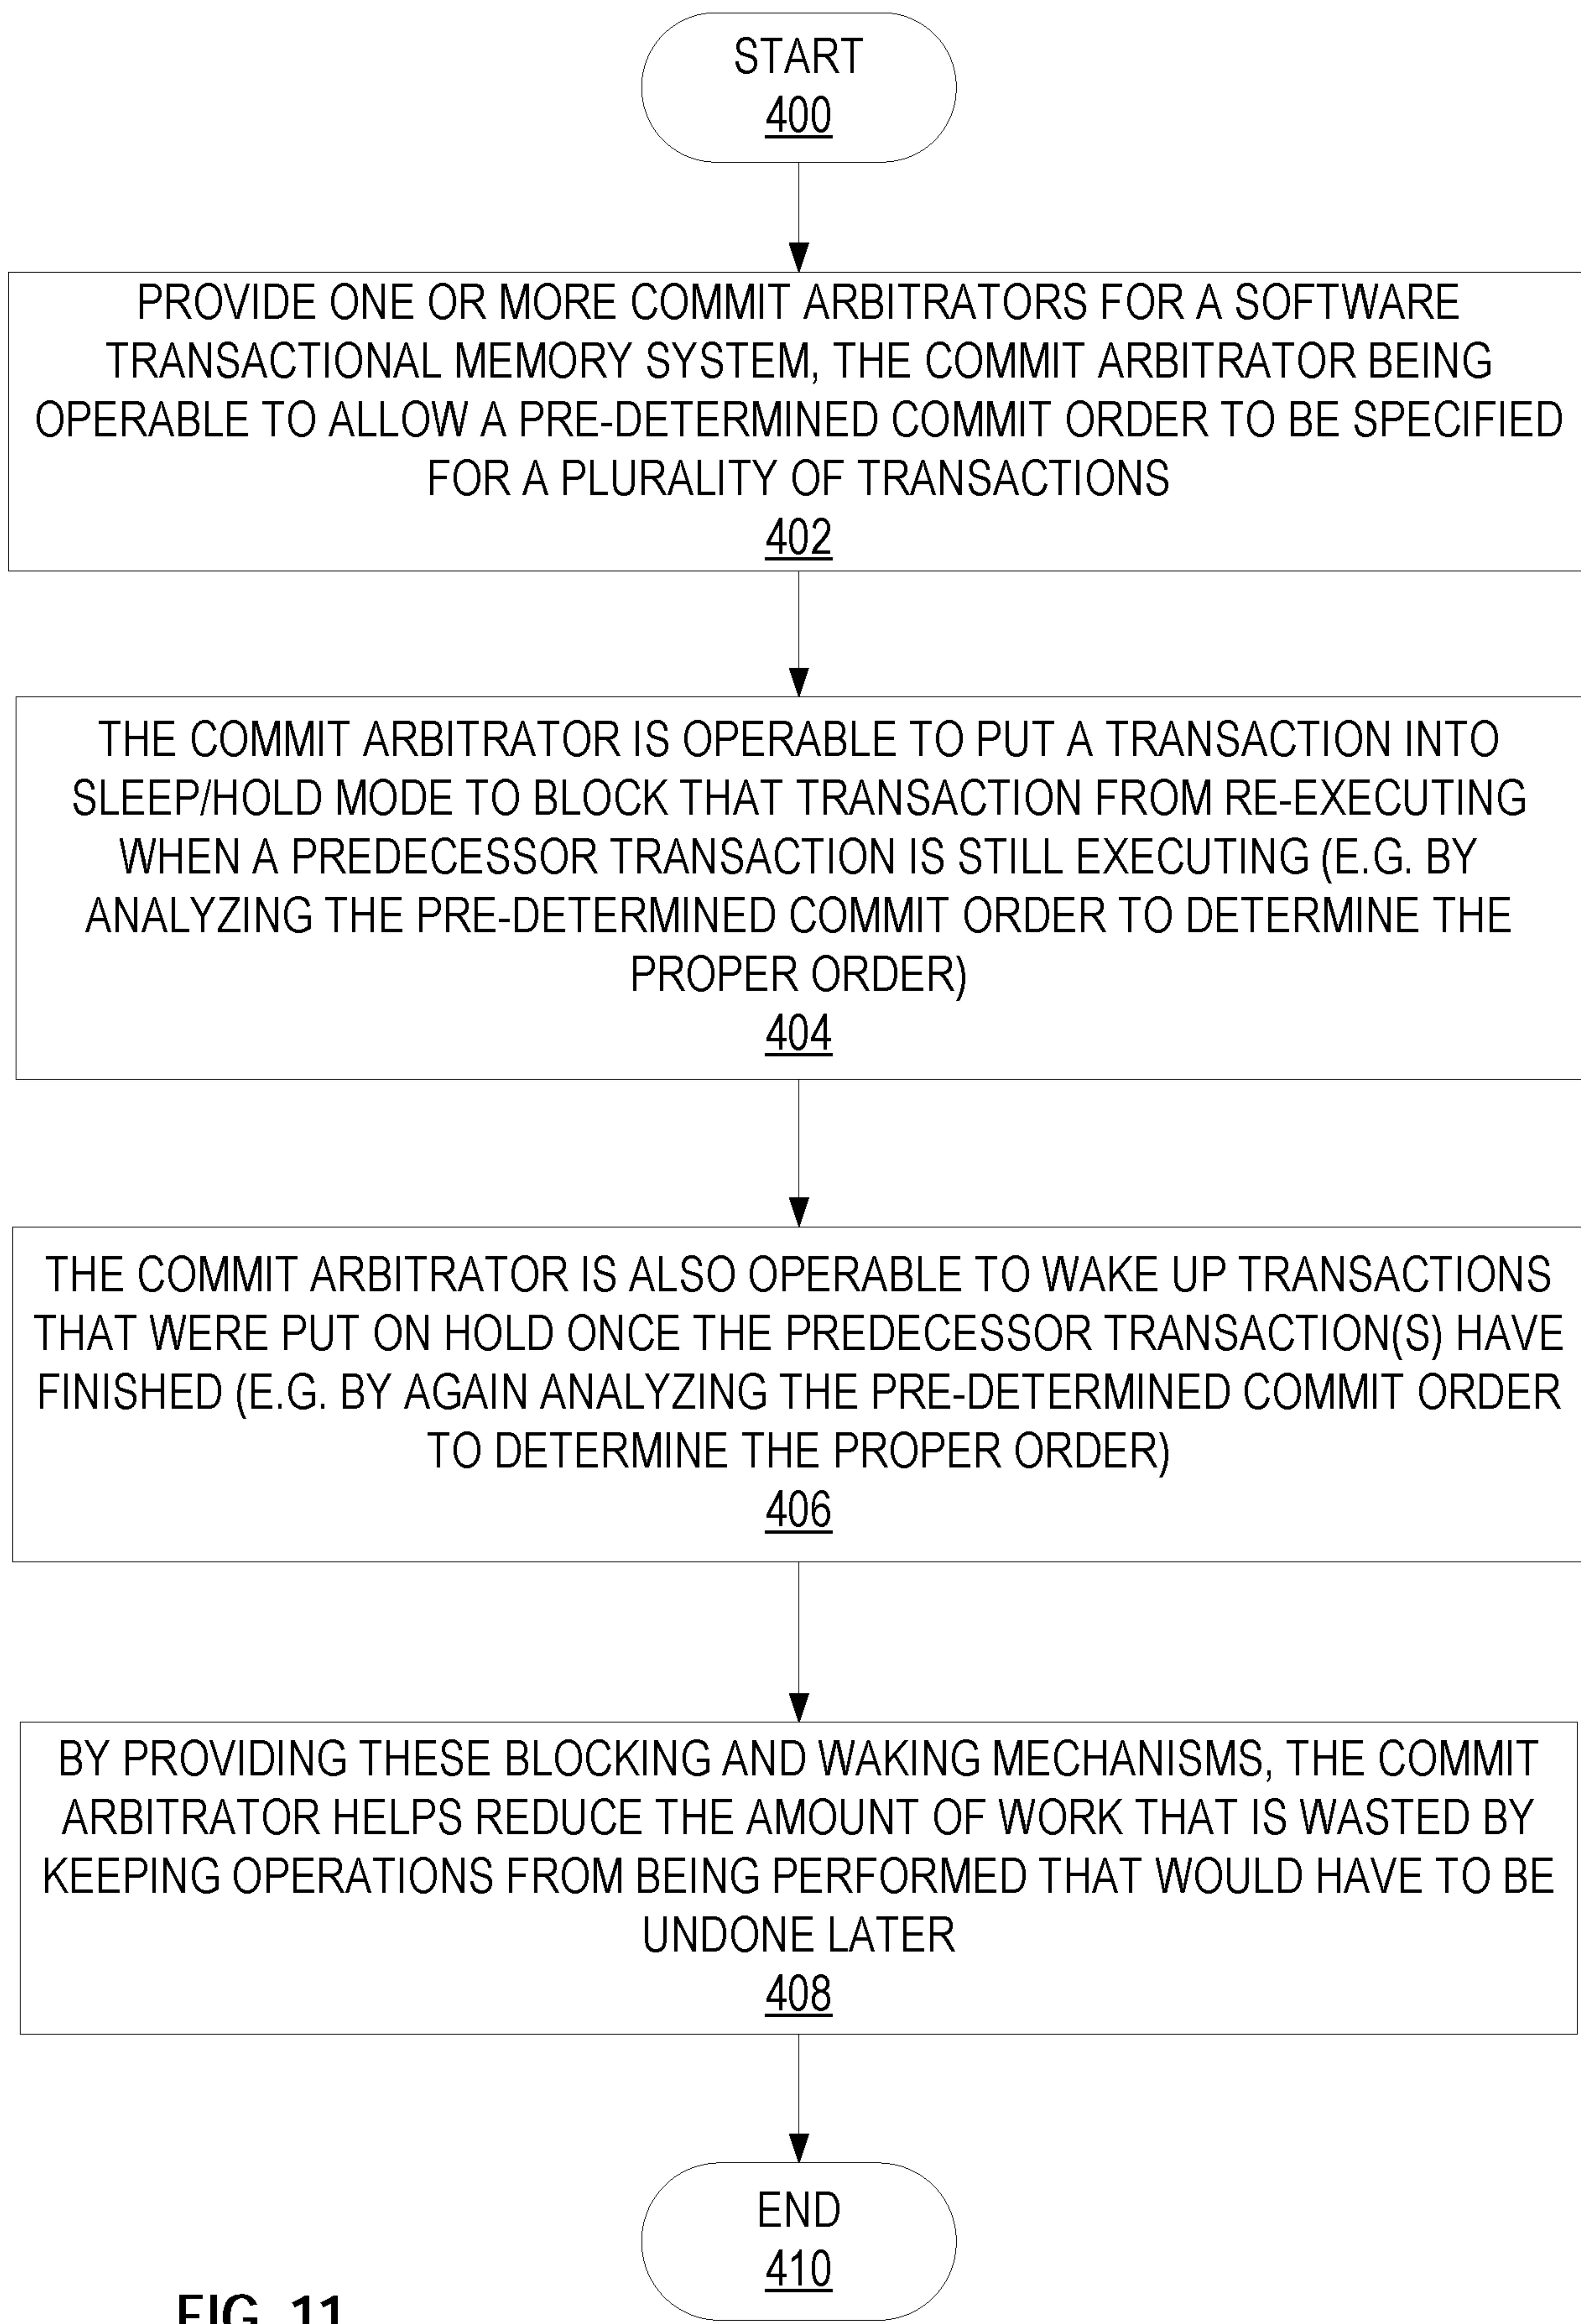


FIG. 11

11 / 11

START
430

PROVIDE A CONTENTION MANAGEMENT PROCESS THAT IS INVOKED WHEN A CONFLICT OCCURS BETWEEN A FIRST TRANSACTION AND A SECOND TRANSACTION

432

USE A PRE-DETERMINED COMMIT ORDER IN THE CONTENTION MANAGEMENT PROCESS TO AID IN DETERMINING WHETHER THE FIRST TRANSACTION OR THE SECOND TRANSACTION SHOULD WIN THE CONFLICT AND BE ALLOWED TO PROCEED

434

ANALYZE AN ENTIRE ANCESTOR CHAIN OF A PRE-DETERMINED COMMIT ORDER TO HELP DETERMINE PROPER CONFLICT MANAGEMENT (E.G. SUPPOSE THERE ARE FOUR TRANSACTIONS, TWO PARENTS AND TWO CHILDREN, WHERE B IS NESTED WITHIN A AND D IS NESTED WITHIN C, AND GIVEN AN ORDERING RELATIONSHIP BETWEEN A AND C WHERE A SHOULD COMMIT BEFORE C; THEN IF B AND D CONFLICT, THE CONTENTION MANAGEMENT PROCESS SHOULD FAVOR B BECAUSE FAVORING D IS USELESS GIVEN THAT A MUST COMMIT BEFORE C)

436

END
438

FIG. 12

SOFTWARE TRANSACTIONAL MEMORY APPLICATION

200

PROGRAM LOGIC

204

LOGIC FOR PROVIDING A SOFTWARE TRANSACTIONAL MEMORY (STM) SYSTEM 206

LOGIC FOR PROVIDING A COMMIT ARBITRATOR THAT ALLOWS A PRE-DETERMINED COMMIT ORDER TO BE SPECIFIED (STATICALLY OR DYNAMICALLY) FOR A PLURALITY OF TRANSACTIONS IN THE STM SYSTEM 208

LOGIC FOR ALLOWING THE COMMIT ARBITRATOR TO USE THE PRE-DETERMINED COMMIT ORDER AT RUNTIME TO AID IN DETERMINING AN ORDER IN WHICH TO COMMIT THE PLURALITY OF TRANSACTIONS IN THE SOFTWARE TRANSACTIONAL MEMORY SYSTEM 210

LOGIC FOR PROVIDING A CONTENTION MANAGEMENT PROCESS THAT IS INVOKED WHEN A CONFLICT OCCURS BETWEEN A FIRST TRANSACTION AND A SECOND TRANSACTION 212

LOGIC FOR USING THE PRE-DETERMINED COMMIT ORDER IN THE CONTENTION MANAGEMENT PROCESS TO AID IN DETERMINING WHETHER THE FIRST TRANSACTION OR THE SECOND TRANSACTION SHOULD WIN THE CONFLICT AND BE ALLOWED TO PROCEED (E.G. DEPENDING ON WHICH ONE HAS THE LOWER COMMIT ORDER NUMBER OF TWO TRANSACTIONS IN THE SAME TRANSACTION GROUP) 214

LOGIC FOR ALLOWING THE COMMIT ARBITRATOR TO BE OPERABLE TO USE THE PRE-DETERMINED ORDERING TO TRACK ONE OR MORE ORDERING VALUES (E.G. IN TOTAL ORDERING – A NEXT-TO-COMMIT FIELD THAT REPRESENTS A NEXT TRANSACTION OF THE PLURALITY OF TRANSACTIONS THAT SHOULD BE ALLOWED TO COMMIT) AND FOR COMPARING THE ONE OR MORE ORDERING VALUES TO A PARTICULAR COMMIT ORDER NUMBER OF A GIVEN TRANSACTION TO SEE IF THE COMMIT OF THE TRANSACTION IS PROPER GIVEN THE ORDERING THAT SHOULD BE ENFORCED 216

OTHER LOGIC FOR OPERATING THE APPLICATION 220