

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
8 January 2009 (08.01.2009)

PCT

(10) International Publication Number
WO 2009/006023 A2(51) International Patent Classification:
G06F 12/00 (2006.01)

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.

(21) International Application Number:
PCT/US2008/067343

(22) International Filing Date: 18 June 2008 (18.06.2008)

(25) Filing Language: English

(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, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(30) Priority Data:
11/824,379 29 June 2007 (29.06.2007) US

(71) Applicant (for all designated States except US): MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventor: TAILLEFER, Martin; 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, AO, 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

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))

Published:

- without international search report and to be republished upon receipt of that report

(54) Title: MEMORY TRANSACTION GROUPING

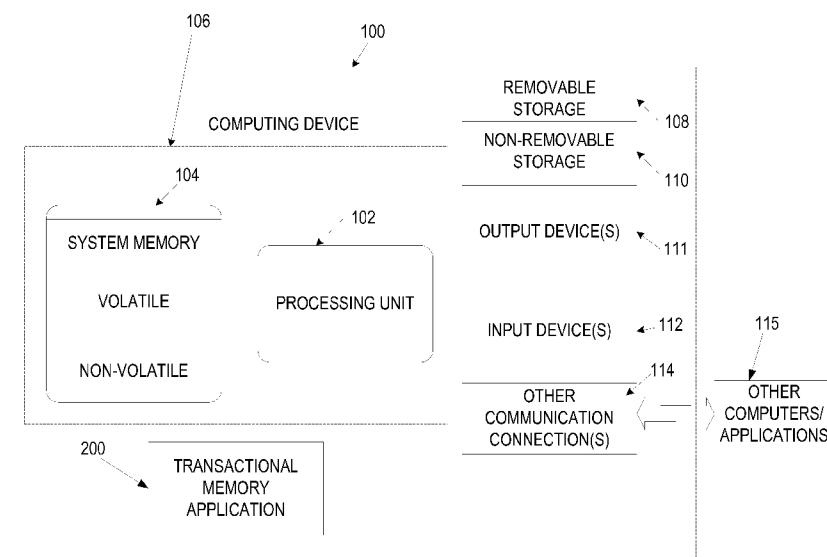


FIG. 1

(57) **Abstract:** Various technologies and techniques are disclosed for providing a transaction grouping feature for use in programs operating under a transactional memory system. The transaction grouping feature is operable to allow transaction groups to be created that contain related transactions. The transaction groups are used to enhance performance and/or operation of the programs. For example, different locking and versioning mechanisms can be used with different transaction groups. When running transactions, a hardware transactional memory execution mechanism can be used for one transaction group while a software transactional memory execution mechanism used for another transaction group.

WO 2009/006023 A2

MEMORY TRANSACTION GROUPING

BACKGROUND

[001] Steadily over time computers have become much more powerful, with more processing power and memory to handle advanced operations. This trend has

5 recently shifted away from a focus on ever-increasing single-processor clock rates and towards an increase in the number of processors available in a single computer. Software developers want to take advantage of improvements in computer processing power, enabling their software programs to be executed faster as new hardware is adopted. With the new hardware trends, however, this requires a
10 different approach: developers must arrange for one or more tasks of a particular software program to be executed “concurrently” (sometimes called “in parallel”), so that the same logical operation can utilize many processors at one time, and deliver better performance as more processors are added to the computers on which such software runs.

15 [002] Transactional memory is designed to ease development of concurrent programs by providing atomicity and isolation to regions of program code. Transactional memory (TM) 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
20 that executes a series of reads and writes to shared memory. TM is used as an alternative to traditional locking mechanisms. TM allows concurrent programs to be written more simply. A transaction specifies a sequence of code that is supposed to execute as if it were executing in isolation, whereas in reality it executes in a normal multithreaded environment with many concurrent activities.

25 This illusion of isolation may be achieved by fine-grained locking of objects or memory ranges, and by executing in a mode that allows the effects of the transaction to be rolled back if the transaction is discovered to be in conflict with some other transaction. We say that a data access is “transacted” if the access is protected by these locking and rollback mechanisms.

30 [003] Different locking and versioning mechanisms are possible, including several software-based and hardware-based approaches. Different mechanisms have

features and qualities making each suitable or preferable in different situations. Combining different mechanisms within a single process generally is not possible, leading to the selection of generic mechanisms which typically compromise on performance in order to achieve general applicability.

5

SUMMARY

[004] Various technologies and techniques are disclosed for providing a transaction grouping feature for use in programs operating under a transactional memory system. The transaction grouping feature is operable to allow transaction groups to be created that contain related transactions. The transaction groups are used to enhance operation of the programs. Transaction groups are defined such that the transactions in each group are known to operate on disjoint data, which enables incompatible locking and versioning mechanisms within each such group, in turn allowing fine-tuning of the specific mechanisms for each particular group.

[005] 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

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

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

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

[009] Figure 4 is a process flow diagram for one implementation of the system of Figure 1 illustrating the stages involved in allowing a programmer to group transactions.

[010] Figure 5 is a process flow diagram for one implementation of the system of Figure 1 illustrating the stages involved in providing a language compiler that automatically groups transactions based on specific heuristics.

[011] Figure 6 is a process flow diagram for one implementation of the system of Figure 1 illustrating the stages involved in providing a runtime environment that automatically groups transactions.

[012] Figure 7 is a process flow diagram for one implementation of the system of Figure 1 that illustrates the stages involved in providing specialized contention management for different transaction groups.

[013] Figure 8 is a process flow diagram for one implementation of the system of Figure 1 that illustrates the stages involved in providing specialized locking and versioning mechanisms for different transaction groups.

[014] Figure 9 is a process flow diagram for one implementation of the system of Figure 1 that illustrates the stages involved in naming a grouping of related transactions to enhance debugging or other processes.

[015] Figure 10 is a diagrammatic view of multiple transaction groups.

DETAILED DESCRIPTION

[016] The technologies and techniques herein may be described in the general context as a transactional memory system, but the technologies and techniques also serve 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.

[017] In one implementation, a transaction grouping feature is provided for use in programs operating under a transactional memory system. The transaction grouping feature allows transactions to be placed into groups. If it can be determined that a set of transactions access data (e.g. read/write data) which is demonstrably disjoint from the data accessed by any other transactions, then this set can be considered as a “transaction group”.

[018] By the above definition, transactions that are part of a group are known to operate on read/write data which is disjoint from the read/write data accessed by

other transactions within other groups. As a result, it becomes possible to implement distinct locking and versioning mechanisms for each such group, allowing each transaction group to leverage specially-selected locking and versioning algorithms most appropriate for the data accessed by the transactions in the group.

[019] In addition to the particular data accessed by the transactions in a group, many other factors can influence the particular selection of locking and versioning algorithms used within a group. For example, the duration of the transactions or the nature of the code within the transactions are two other such factors. In one implementation, locking and versioning mechanisms that are normally incompatible can be used concurrently within a process, leading to potentially increased performance.

[020] Determining when transactions can be grouped can be accomplished through a plurality of means. One implementation may leverage programmer-supplied annotations to mark the groups, as described in Figure 4. Another implementation may use compiler heuristics to automatically infer groups and group membership, as described in Figure 5. Still another implementation may use a runtime environment to dynamically and automatically infer groups and group membership, as described in Figure 6. It should be appreciated that the specific mechanisms involved in creating groups and assigning group membership are many and can be combined in various ways.

[021] 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 (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.

[022] 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, 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 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.

[023] 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 transactional memory application 200. Transactional memory application 200 will be described in further detail in Figure 2.

[024] Turning now to Figure 2 with continued reference to Figure 1, a transactional memory application 200 operating on computing device 100 is illustrated. Transactional memory application 200 is one of the application programs that reside on computing device 100. However, it will be understood that 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 30 of transactional memory application 200 can be part of system memory 104, on

other computers and/or applications 115, or other such variations as would occur to one in the computer software art.

[025] Transactional memory application 200 includes program logic 204, which is responsible for carrying out some or all of the techniques described herein.

5 Program logic 204 includes logic for providing a transaction grouping feature that allows related transactions in a particular program to be grouped together 206 (as described below with respect to Figures 3-6); logic for providing specialized contention management using transaction groups 210 (as described below with respect to Figure 7); logic for providing different locking and versioning mechanisms for different transaction groups 212 (as described below with respect to Figure 8); logic for naming a transaction group to enhance debugging or other processes 214 (as described below with respect to Figure 9); and other logic for operating the transactional memory application 220.

10 [026] Turning now to Figures 3-10 with continued reference to Figures 1-2, the stages for implementing one or more implementations of transactional memory application 200 are described in further detail. In some implementations, the processes of Figures 3-10 are at least partially implemented in the operating logic of computing device 100. Figure 3 is a high level process flow diagram for transactional memory application 200. The process begins at start point 240 with 15 providing a transactional memory system using software, hardware, and/or combinations thereof (stage 242). The system provides a transaction grouping feature that allows related transactions in a particular program to be grouped together either manually (as described in Figure 4), and/or programmatically (as described in Figures 5 and 6) (stage 244). The system uses the transaction groups 20 to improve program performance or otherwise enhance program operation, as described in Figures 7-10 (stage 246). The process ends at end point 248.

25 [027] Figure 4 illustrates one implementation of the stages involved in allowing a programmer to group transactions. The process begins at start point 270 with receiving input from a programmer to access the source code of a particular 30 program that executes under a transactional memory system (stage 272). The programmer adds declarations or otherwise assigns groups to transactions (stage

274). The system uses the specified groups to enhance program operation, as described in further detail in Figures 7-9 (stage 276). The process ends at end point 278.

[028] Figure 5 illustrates one implementation of the stages involved in providing a language compiler that automatically groups transactions. The process begins at start point 290 with providing a language compiler for compiling programs that execute under a transactional memory system (stage 292). At compile time of a particular program, logic is used to determine if any transactions may be grouped together (stage 294). As a few non-limiting examples, transactions may be grouped together by a compiler by leveraging the particular semantics of the programming language or through global analysis that demonstrates that groups are possible. The system creates the identified groups in the program (stage 296) and then uses the specified groups to enhance the program operation (stage 298). The process ends at end point 300.

[029] Figure 6 illustrates one implementation of the stages involved in providing a runtime environment that groups transactions. The process begins at start point 310 with providing a runtime environment for running programs under a transactional memory system (stage 312). When running a particular program, logic is used to identify transactions that should be grouped together (stage 314).

As a few non-limiting examples, transactions may be grouped together by a runtime environment by identifying sets of transactions which must by construction operate on disjoint read/write data. The runtime may operate on hints supplied by the compiler to hone its analysis process. The runtime groups the identified transactions together (stage 316) and uses the grouped transactions to improve operation of the particular program (stage 318). The process ends at end point 320.

[030] Turning now to Figures 7-9, examples will be used to describe how program operation can be enhanced using grouped transactions. Figure 7 illustrates one implementation of the stages involved in providing specialized contention management using grouped transactions. The process begins at start point 340 with allowing one or more specialized contention management policies to be defined by

a programmer (stage 342). For each policy, the system allows transaction execution scheduling settings, transaction abort handling settings, and/or other settings to be specified (stage 344). The previously defined policy can then be assigned to each transaction group in a process (stage 346). The system uses the 5 policies to implement specialized contention management for the grouped transactions (stage 348).

[031] Contention management is the mechanism used by a runtime system to select appropriate behavior whenever a conflict is detected between multiple concurrently executing transactions. Contention management decides which of the 10 conflicting transactions, if any, to preserve and which to abort. It further decides how to reschedule execution of the individual transactions such that they can run to completion. Specialized contention management as used herein refers to the ability to apply contention management heuristics which are distinct from the default heuristics of a given runtime.

[032] In one implementation, by applying different policies to different groups, enhanced performance of the program can be achieved. For example, one type of specialized contention management policy can be assigned to a particular transaction group that will give the best performance for the types of operations those transactions contain. Another specialized contention management policy can 20 be assigned to another transaction group that will get the best performance for the types of operations that the other transaction group contains. The process ends at end point 350.

[033] Figure 8 illustrates one implementation of the stages involved in providing different locking and versioning mechanisms for different groups of transactions.

The process begins at start point 370. The system uses logic to determine what 25 type of locking and versioning to use for each transaction group (stage 372) and then adjusts the locking and versioning used on each transaction group as necessary (stage 374). The process ends at end point 376.

[034] As an example, one transactional memory locking and versioning 30 mechanism can be used with one particular transaction group while another

potentially incompatible transactional memory locking and versioning mechanism is used with another transaction group.

[035] Let's look at a non-limiting example to further illustrate how different transactional memory mechanisms can be combined together when there are multiple transaction groups. One transaction group could use a buffered update scheme for versioning, while another group could use an in-place update scheme with undo logging. By grouping transactions, the data can be isolated in ways that enable the combinations of such incompatible transactional memory locking and versioning mechanisms to be used, which may enable improved overall performance. As another example, fast but limited hardware-based transactional memory mechanisms can be used in some transaction groups, while incompatible and slower software transactional memory mechanisms can be used in other transaction groups when the hardware limitations are not acceptable for the transactions of the group.

[036] While the previous hypothetical example just mentioned using this technique with two transaction groups, the concept can be used with more than two groups and with various combinations of transactional memory locking and versioning mechanisms, including both hardware and software approaches.

[037] Figure 9 illustrates one implementation of the stages involved in naming a grouping of related transactions to enhance debugging or other processes. The process begins at start point 450 with allowing related transaction to be grouped together by a user and/or programmatically (stage 452). A naming feature is provided that allows each transaction group to be given a name by a user and/or programmatically (stage 454). The system then displays or uses the group names to enhance debugging, profiling, or other processes (stage 456). For example, the group names can be displayed in a debugger or profiler to allow a user to more easily identify the particular transaction group. The process ends at end point 458.

[038] Figure 10 is a diagrammatic view of multiple transaction groups. In the example shown, a first transaction group 500 contains four transactions, one of which is nested within the other. The second transaction group 502 just contains a single transaction. Numerous other transactional grouping scenarios are also

possible that contain a different number of transaction groups and/or a different number of transactions within each group.

[039] 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, 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.

5 [040] For example, a person of ordinary skill in the computer software art will recognize that the examples discussed herein could be organized differently on one or more computers to include fewer or additional options or features than as 10 portrayed in the examples.

What is claimed is:

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

5 providing a transaction grouping feature for use with a transactional memory system, the transaction grouping feature being operable to allow transaction groups to be created where all transactions within a particular group are known to only access data which is disjoint from that accessed by any transaction outside of the group (206).

2. The computer-readable medium of claim 1, further having computer-10 executable instructions for causing a computer to perform steps comprising:

providing different locking mechanisms for at least some of the transaction groups (212).

3. The computer-readable medium of claim 2, where some locking mechanisms are implemented through software methods while others are 15 implemented using hardware methods (372).

4. The computer-readable medium of claim 2, further having computer- executable instructions for causing a computer to perform steps comprising:

using locking mechanisms which are not functionally compatible with one another (372).

20 5. The computer-readable medium of claim 1, further having computer- executable instructions for causing a computer to perform steps comprising:

providing different versioning mechanisms for at least some of the transaction groups (212).

25 6. The computer-readable medium of claim 5, where some versioning mechanisms are implemented using hardware methods while others are implemented using software methods (372).

7. The computer-readable medium of claim 5, further having computer- executable instructions for causing a computer to perform steps comprising:

30 using versioning mechanisms which are not functionally compatible with one another (372).

8. The computer-readable medium of claim 1, further having computer-executable instructions for causing a computer to perform steps comprising:
providing specialized contention management using the transaction groups (210).

5 9. The computer-readable medium of claim 1, further having computer-executable instructions for causing a computer to perform steps comprising:
provide a naming feature that is operable to allow transaction groups to be named (454).

10. 10. The computer-readable medium of claim 9, wherein the naming feature is used in a debugger (456).

11. 11. The computer-readable medium of claim 9, wherein the naming feature is used in a profiler (456).

12. 12. The computer-readable medium of claim 1, wherein at least some of the transaction groups are assigned manually (274).

15 13. 13. The computer-readable medium of claim 12, wherein the transaction grouping feature is operable to allow a programmer to manually specify the transaction groups within source code (274).

14. 14. The computer-readable medium of claim 1, wherein at least some of the transaction groups are assigned automatically (294).

20 15. 15. The computer-readable medium of claim 14, wherein the transaction grouping feature is operable to allow the transaction groups to be automatically identified by a compiler (294).

25 16. 16. The computer-readable medium of claim 14, wherein the transaction grouping feature is operable to allow the transaction groups to be automatically identified by a runtime environment (294).

17. 17. The computer-readable medium of claim 1, wherein the transaction grouping feature is operable to improve performance of the programs (246).

18. 18. A method for using different locking and versioning mechanisms with different transaction groups comprising the steps of:

30 creating a plurality of transaction groups (500);

using a hardware transactional memory execution mechanism for executing transactions within a first one of the transaction groups (500); and

using a software transactional memory execution mechanism for executing transactions within a second one of the transaction groups (502).

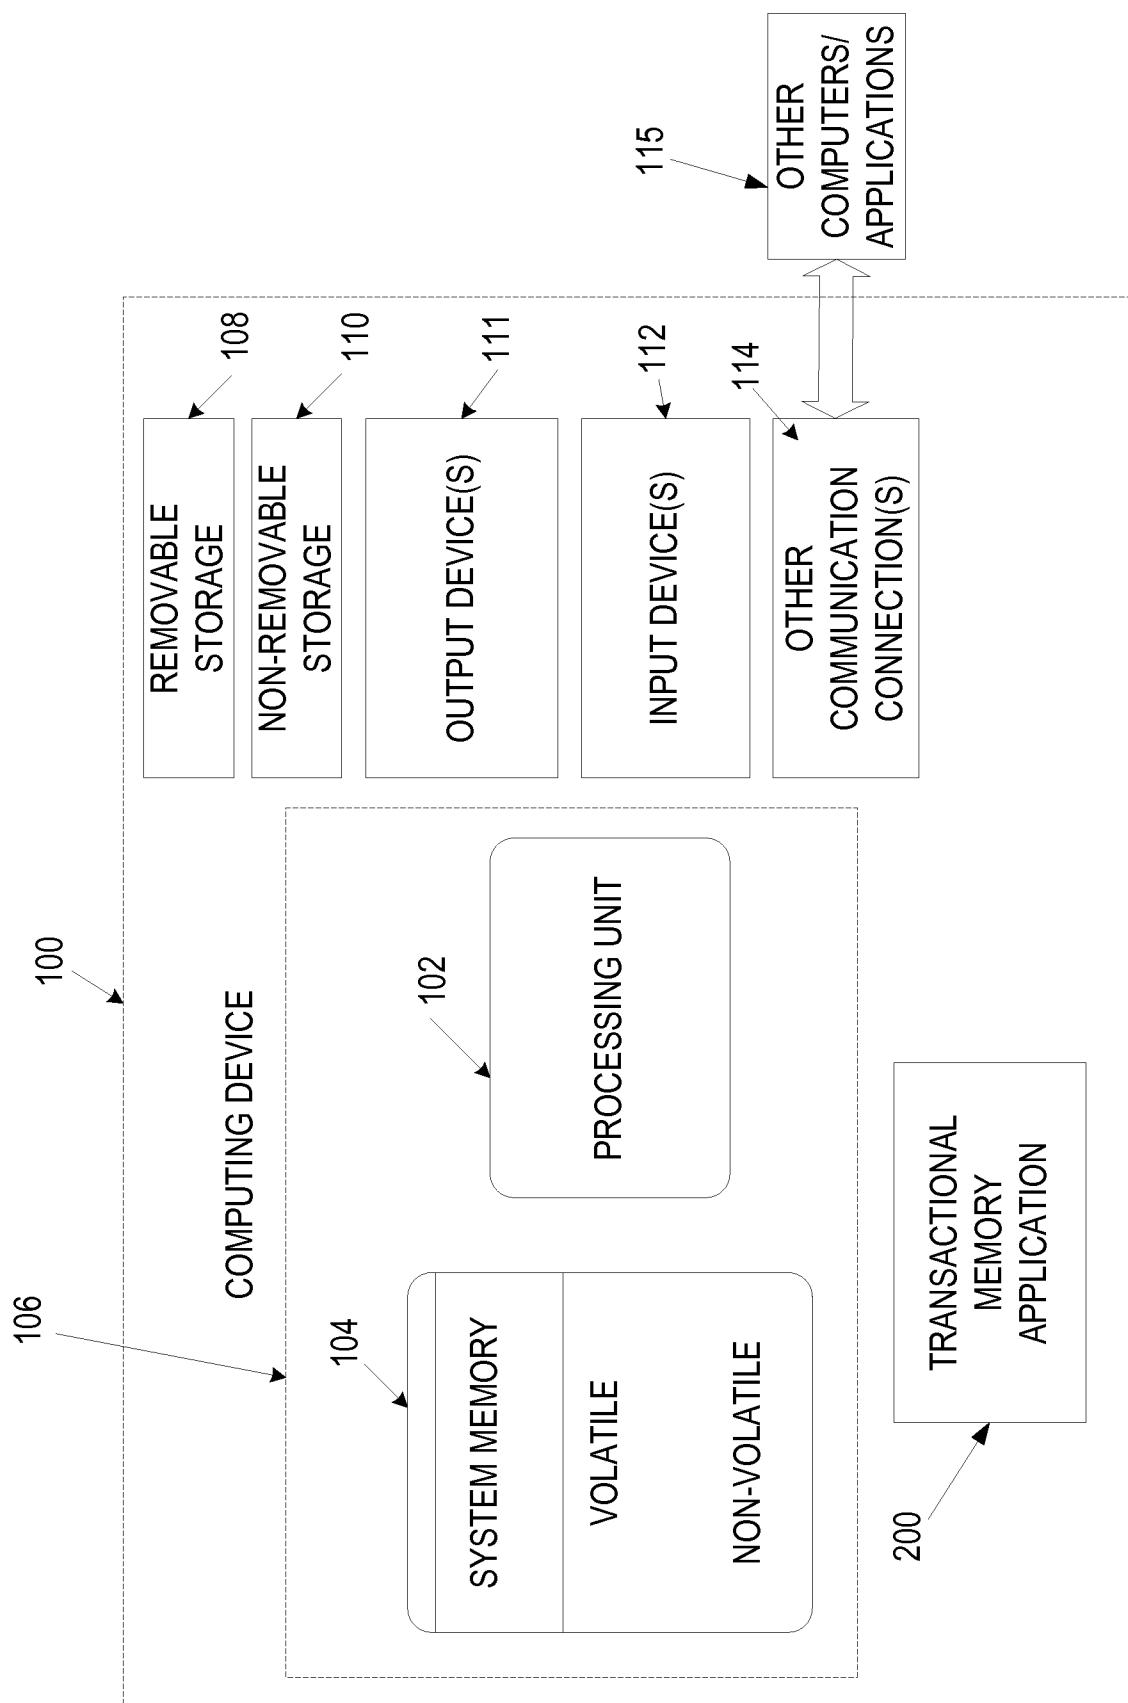
5 19. A computer-readable medium having computer-executable instructions for causing a computer to perform the steps recited in claim 18 (200).

20. A method for supporting transaction grouping comprising the steps of:

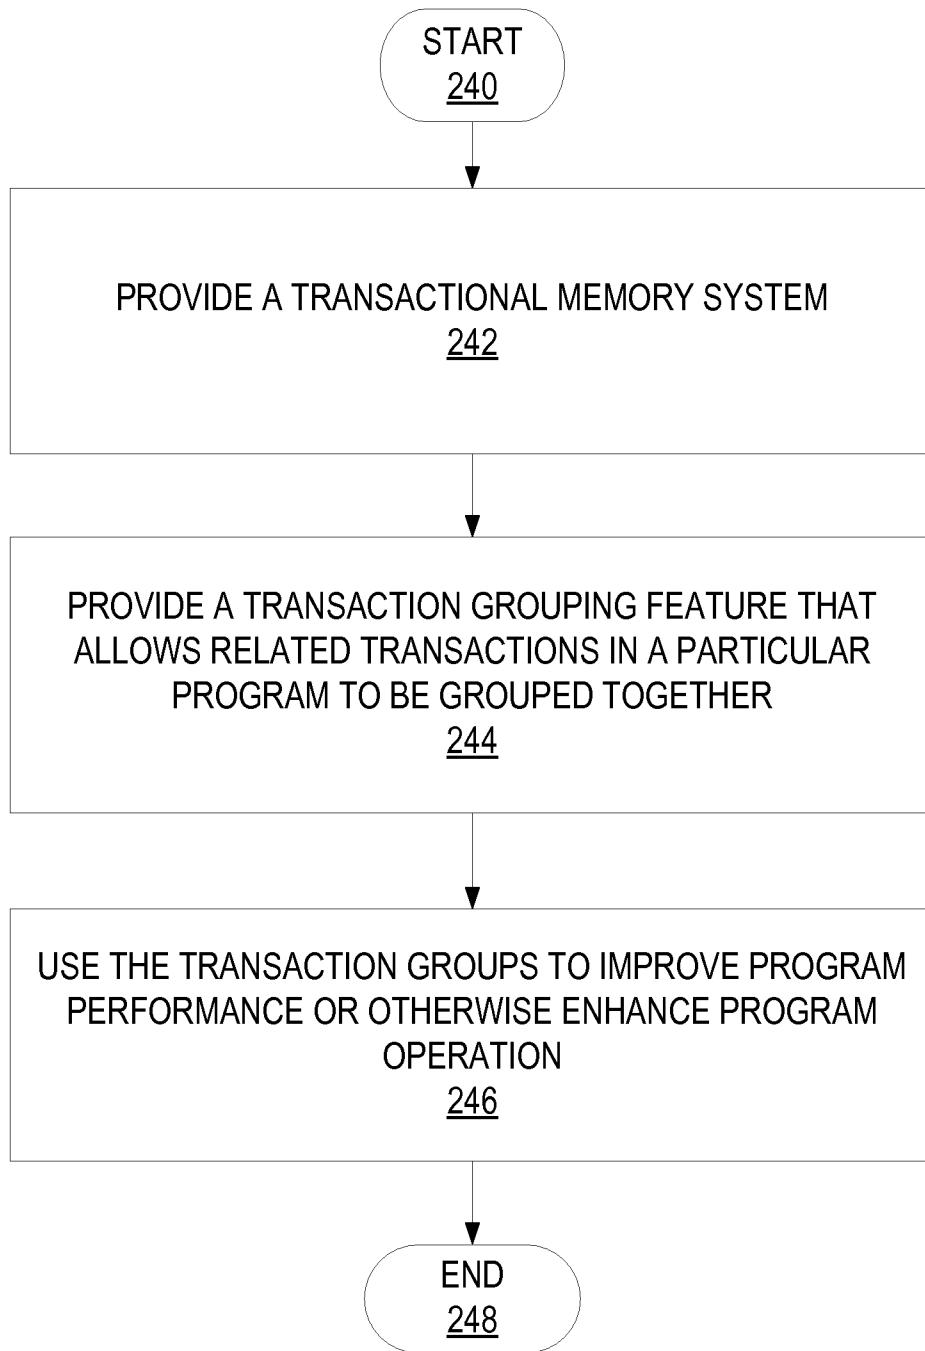
providing a transaction grouping feature that allows transactions which access data disjoint from all other transactions in programs operating under a

10 transactional memory system to be grouped together into transaction groups (244); and

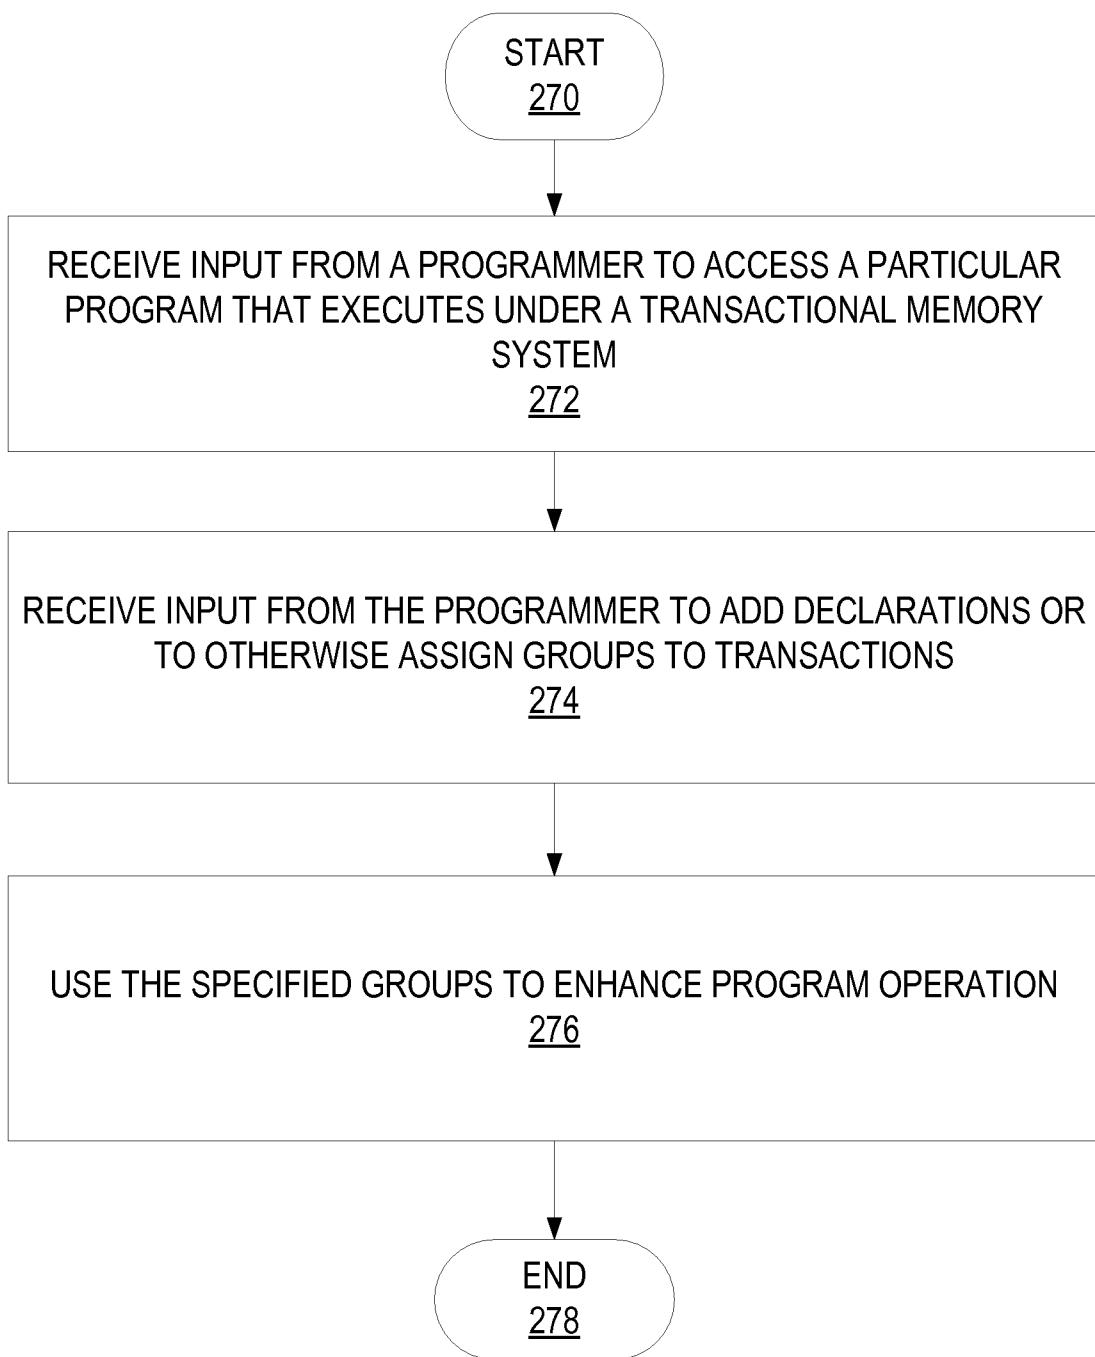
using the transaction groups to make performance of the programs faster than without using the transaction groups (246).

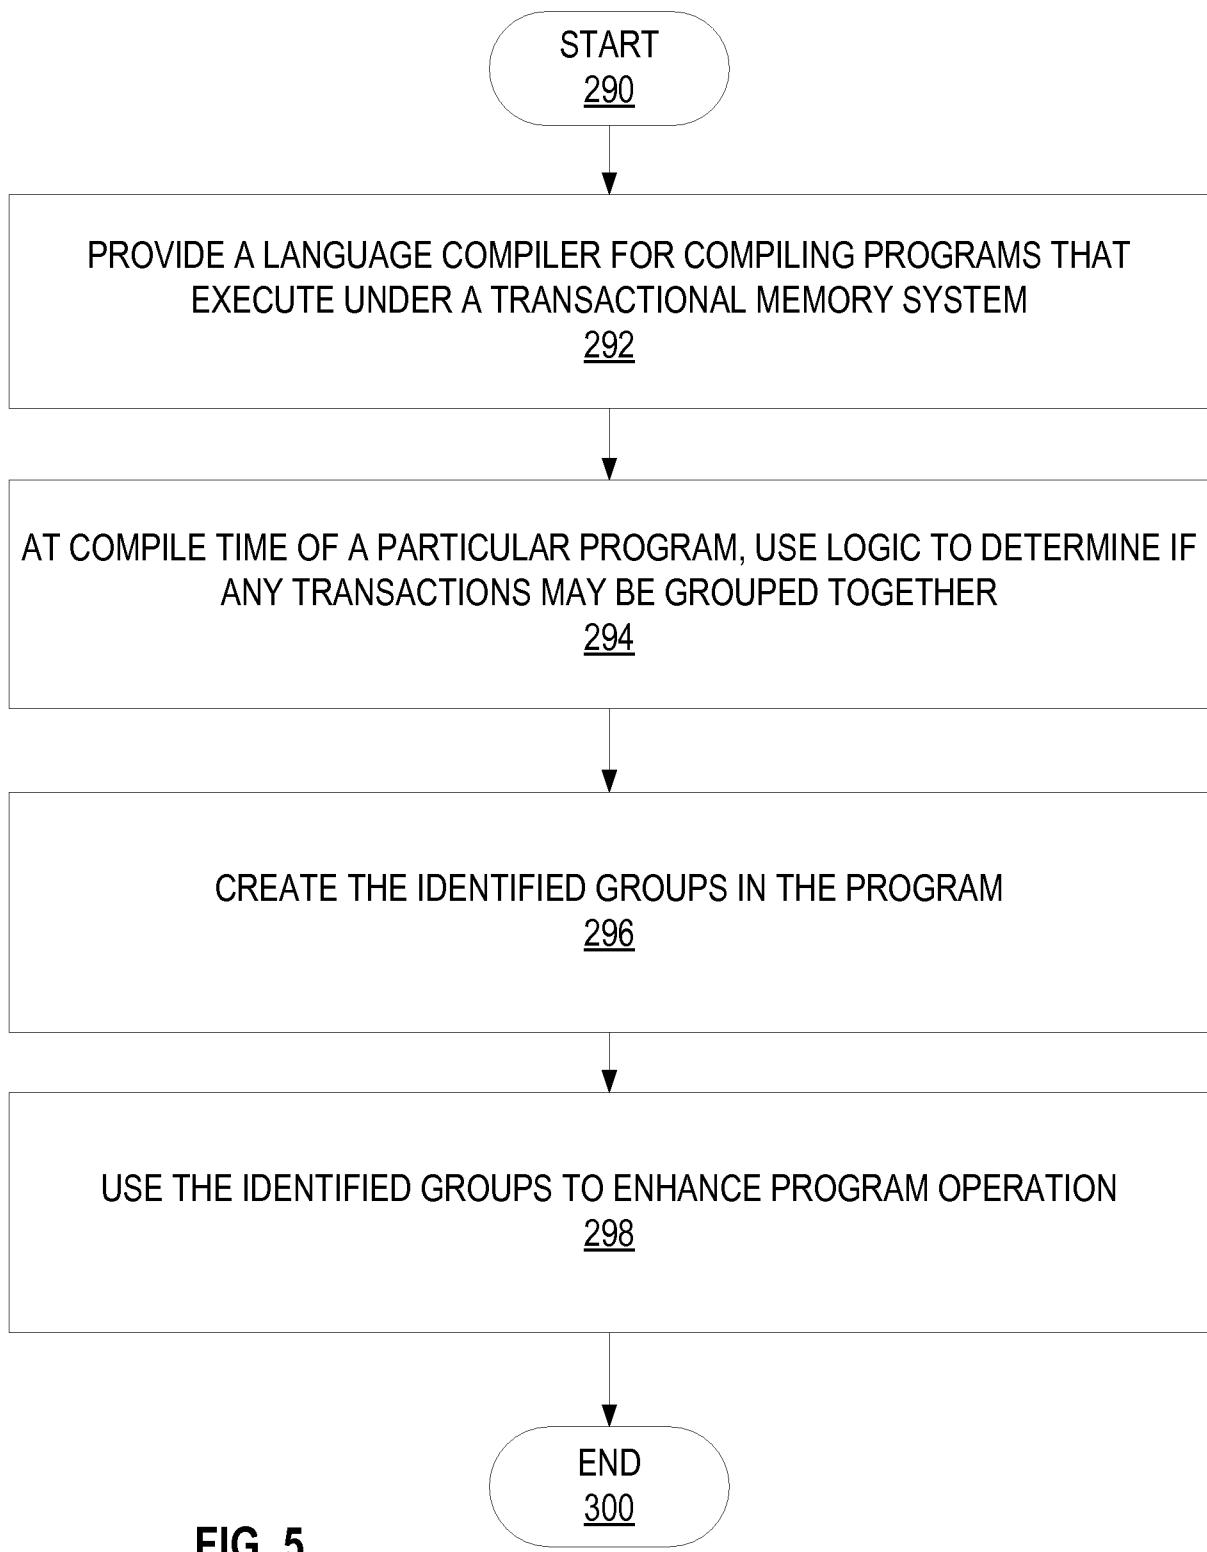
**FIG. 1**

| | | | | | | |
|--|-----------------------------|--|--|---|--|---|
| TRANSACTIONAL MEMORY APPLICATION <u>200</u> | PROGRAM LOGIC <u>204</u> | LOGIC FOR PROVIDING A TRANSACTION GROUPING FEATURE THAT ALLOWS RELATED TRANSACTIONS IN A PARTICULAR PROGRAM TO BE GROUPED TOGETHER <u>206</u> | LOGIC FOR PROVIDING SPECIALIZED CONTENTION MANAGEMENT FOR DIFFERENT TRANSACTION GROUPS <u>210</u> | LOGIC FOR PROVIDING SPECIALIZED LOCKING AND VERSIONING MECHANISMS FOR DIFFERENT TRANSACTION GROUPS <u>212</u> | LOGIC FOR NAMING A TRANSACTION GROUP TO ENHANCE DEBUGGING OR OTHER PROCESSES <u>214</u> | OTHER LOGIC FOR OPERATING THE APPLICATION <u>220</u> |
|--|-----------------------------|--|--|---|--|---|

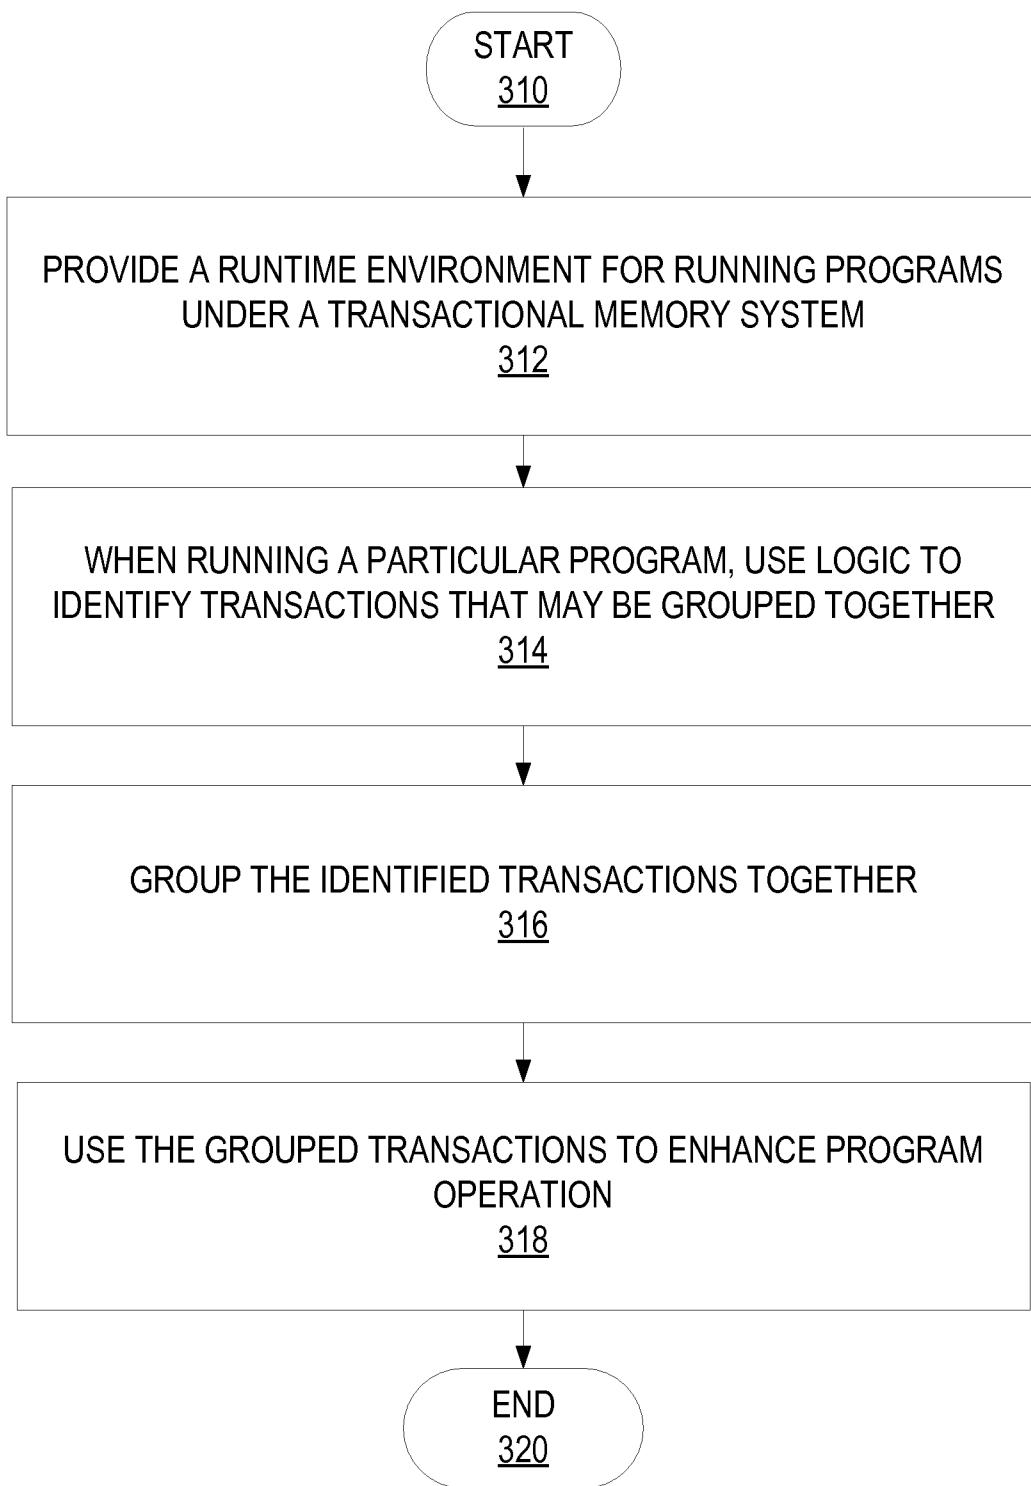
**FIG. 3**

4 / 10

**FIG. 4**

**FIG. 5**

6 / 10

**FIG. 6**

7 / 10

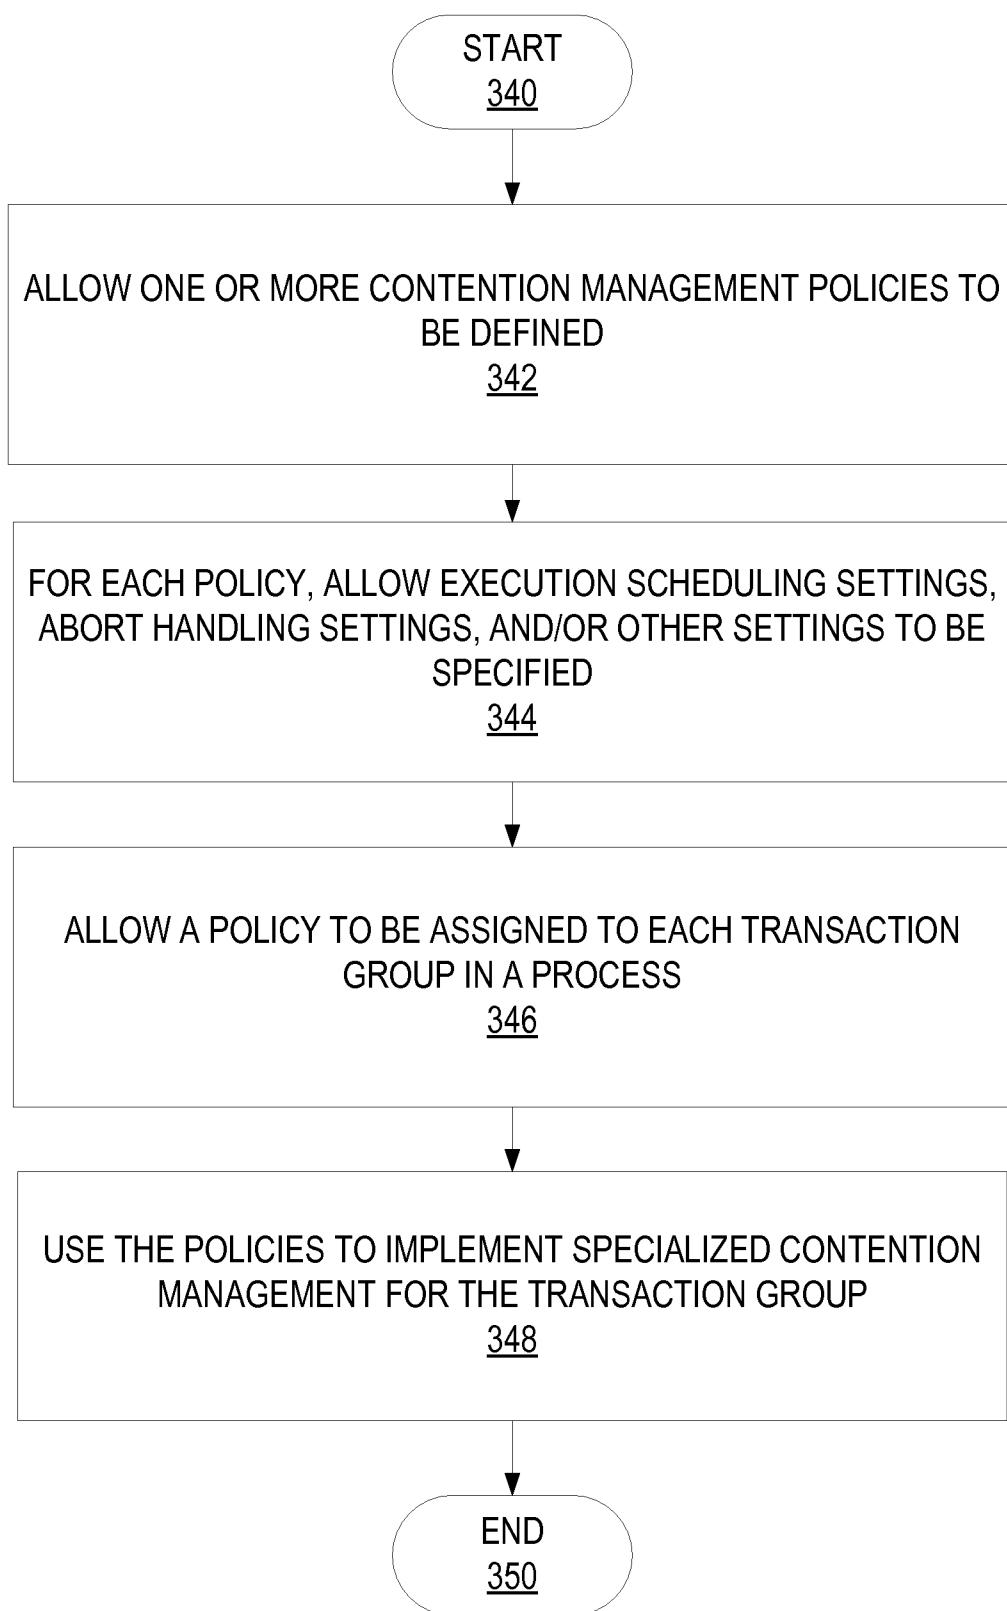


FIG. 7

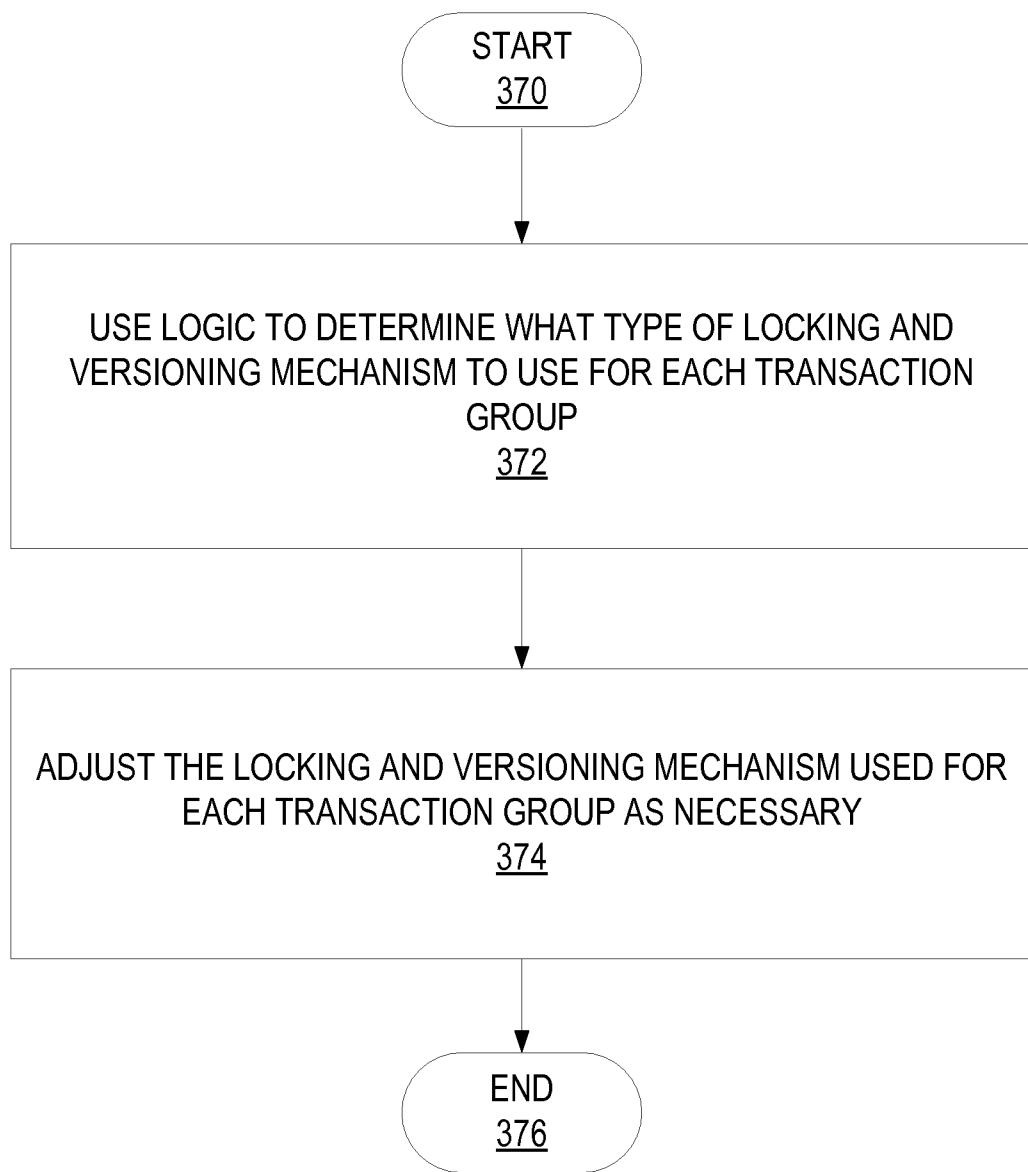
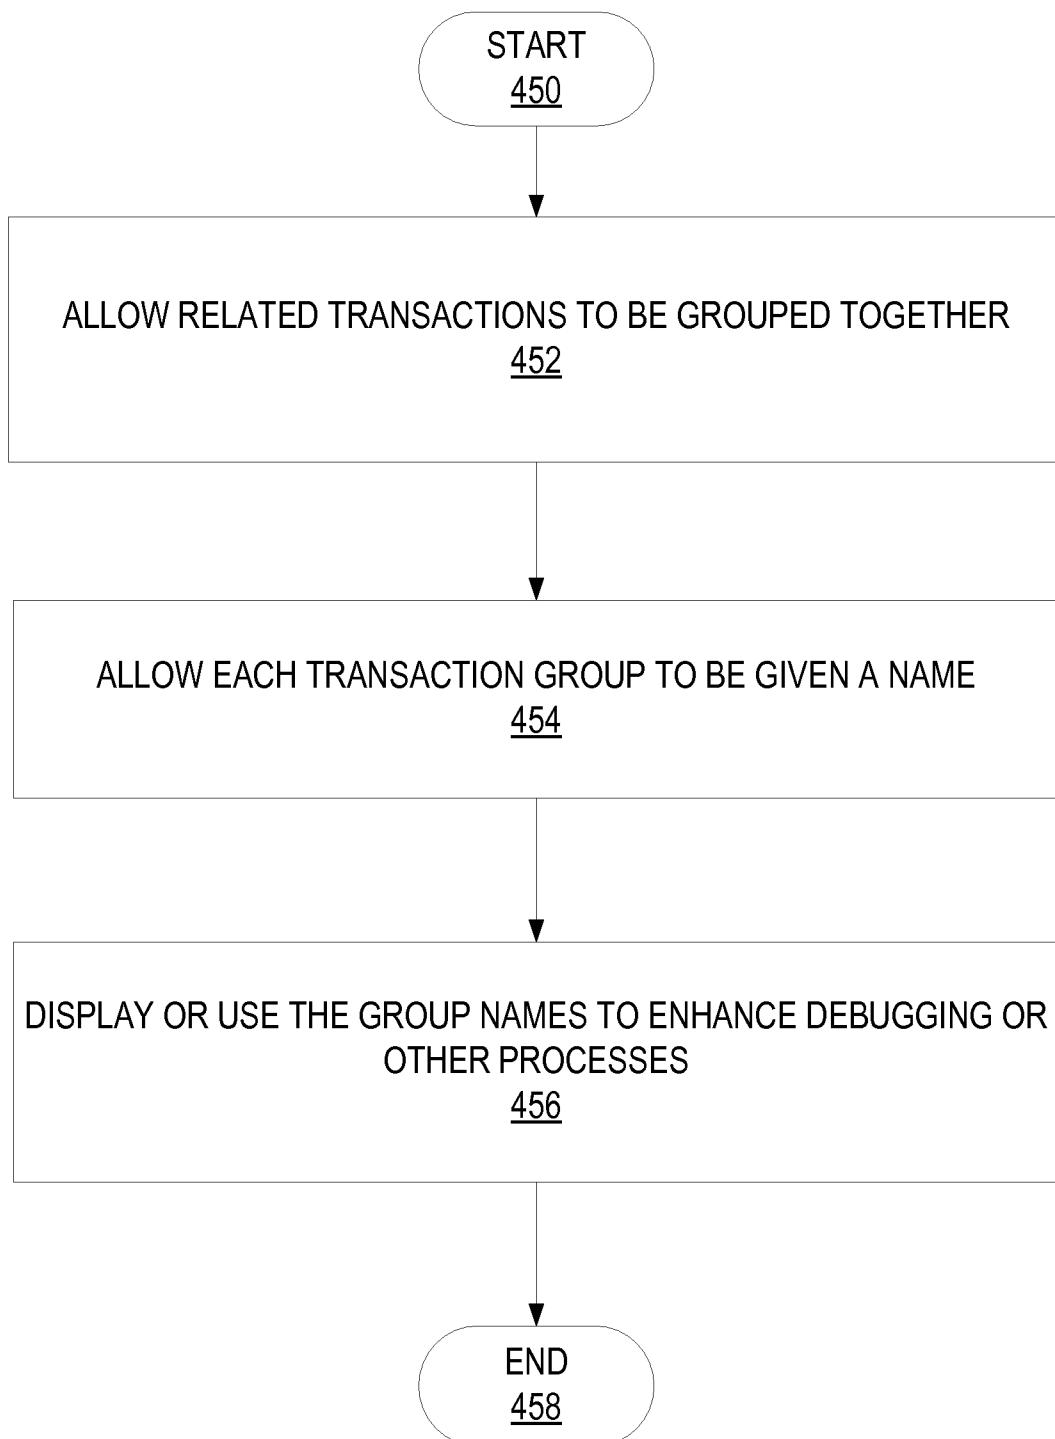


FIG. 8

**FIG. 9**

