

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 April 2012 (12.04.2012)

PCT

(10) International Publication Number
WO 2012/047446 A2

- (51) International Patent Classification:
G06F 15/16 (2006.01) G06F 9/06 (2006.01)
- (21) International Application Number:
PCT/US2011/051019
- (22) International Filing Date:
9 September 2011 (09.09.2011)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
12/892,345 28 September 2010 (28.09.2010) US
- (71) Applicant (for all designated States except US): MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) Inventors: TARTA, Mihail, Gavril; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). KAKI-VAYA, Gopala, Krishna, R.; c/o Microsoft Corporation, LCA - International Patents, 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, 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, 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, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, 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))

Published:

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: PERFORMING COMPUTATIONS IN A DISTRIBUTED INFRASTRUCTURE





Figure 1

(57) Abstract: The present invention extends to methods, systems, and computer program products for performing computations in a distributed infrastructure. Embodiments of the invention include a general purpose distributed computation infrastructure that can be used to perform efficient (in-memory), scalable, failure-resilient, atomic, flow-controlled, long-running state-less and state-full distributed computations. Guarantees provided by a distributed computation infrastructure can build upon existent guarantees of an underlying distributed fabric in order to hide the complexities of fault-tolerance, enable large scale highly available processing, allow for efficient resource utilization, and facilitate generic development of stateful and stateless computations. A distributed computation infrastructure can also provide a substrate on which existent distributed computation models can be enhanced to become failure-resilient.



WO 2012/047446 A2

PERFORMING COMPUTATIONS IN A DISTRIBUTED INFRASTRUCTURE**BACKGROUND**

[0001] Computer systems and related technology affect many aspects of society.

Indeed, the computer system's ability to process information has transformed the way we
5 live and work. Computer systems now commonly perform a host of tasks (e.g., word
processing, scheduling, accounting, etc.) that prior to the advent of the computer system
were performed manually. More recently, computer systems have been coupled to one
another and to other electronic devices to form both wired and wireless computer networks
over which the computer systems and other electronic devices can transfer electronic data.
10 Accordingly, the performance of many computing tasks are distributed across a number of
different computer systems and/or a number of different computing environments.

[0002] In some computing environments, clients route requests for a service to a logical
destination referred to as the service rendezvous point. More specifically, these routed
service requests are processed by federation nodes serving as the rendezvous point owners.

15 A typical service that can be provided at the rendezvous point is storage, such as, for
example, SQL Server Data Services.

[0003] One complexity of this system compared to the traditional client-server model is
that the rendezvous point ownership can change as nodes join and leave a federation, or
faults occur in the system. However, it is desirable to shield clients from this complexity
20 and provide them the illusion that they are getting service from a single, logically
consistent, rendezvous point.

[0004] As such, to provide a single, logically consistent service at a rendezvous point,
state maintained at the rendezvous point must be replicated to ensure that it is not lost
when a node currently serving as the rendezvous point owner faults. This replication
25 enables any node replacing the faulted node as owner to behave in manner consistent with
the actions taken by the faulted node prior to its failure. That is, one can view this as the
replicated state machine problem, which in turn implies data consistency guarantees.

[0005] However, existing distributed system typically lack any solution to the problem
of providing distributed computations as replicated highly available services. Most
30 existing distributed computation models have a different focus. For example, intermediate
results (inputs and outputs) have to use files or specialized stores in order to achieve some
high availability guarantees. Further, these distributed computations models are data flow

oriented, where operators receive inputs, produce outputs and they are mostly functional and stateless. These operators also cannot start unless all of their inputs are available.

[0006] Another limitation of typical distributed computation models is failure to provide strong consistency guarantees for computations that execute in the presence of failures or load balancing, leaving all that work up to the developer/operator of the computation (restart, termination, failover etc.).

BRIEF SUMMARY

[0007] The present invention extends to methods, systems, and computer program products for performing computations in a distributed infrastructure. In some embodiments, local stateful computations are performed. A replica set within a fabric based distributed computing architecture includes a primary compute node and one or more secondary compute nodes. Each of the primary compute node and the one or more secondary compute nodes have the same current state and are configured with the same executable group of operations for performing a local computation. The group of operations including at least one input or output queue related operation and at least one operation that modifies the computing state.

[0008] The primary compute node performs the local computation. To do so, the primary compute node creates a transaction for the local computation. The primary compute node logically performs one or more additional operations in the group of operations. The additional operations include processing the current computing state, including one or more of: logically generating additional computing state and logically modifying existent computing state. The additional operations include logically performing the at least one input or output queue related operation. The additional operations include ending the transaction in accordance with one of a transaction commit operation and a transaction abort operation.

[0009] The primary compute node's performance of the local computation is replicated into system memory to the one or more secondary compute nodes. Replication includes replicating creating a transaction for the local computation, logically performing the one or more additional operations in the group of operations, and ending the transaction. Replication of ending the transaction signals the one or more secondary compute nodes to applying changes in accordance with the transaction commit operation or the transaction abort operation.

[0010] In response to ending the transaction, the primary compute node applying changes in accordance with the transaction commit operation or the transaction abort

operation. When the transaction commit operation ends the transaction, applying physical changes are applied to implement the logical application of each operation in the group of operations. Physical changes include reconciling the current computing state with any additional computing state and any modified existent computing state into a new
5 committed computing state. Physical changes also include altering the state of at least one queue to physically implement the at least one input or output queue related operation. When the transaction abort operation ends the transaction, the at least one input or output queue related operation are logically reversed.

[0011] In other embodiments, local stateless computations are performed. A compute
10 node is configured with an executable group of a plurality of operations for performing a stateless local computation, the group of operations including at least one input or output queue related operation. The compute node performs the local stateless computation.

[0012] The compute node creates a transaction for the local stateless computation. The compute node logically performs one or more operations in the group of operations. A
15 failure of the local computation is reliably detected prior to completing a specified operation after the one or more operations in the group of operations. The local computation is reliably restarted at the specified operation. The compute node logically performs the specified operation and any remaining operations after the specified operation until the plurality of operations is performed. The transaction ends in
20 accordance with one of a transaction commit operation and a transaction abort operation.

[0013] In response to ending the transaction, the compute node applying changes in accordance with the transaction commit operation or the transaction abort operation. When the transaction commit operation ends the transaction, applying physical changes to implement the logical application of each of the plurality of operations in the group of
25 operations. Applying physical changes includes altering the state of at least one queue to physically implement the at least one input or output queue related operation. When the transaction abort operation ends the transaction, logically reversing the at least one input or output queue related operation.

[0014] This summary is provided to introduce a selection of concepts in a simplified
30 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.

[0015] Additional features and advantages of the invention will be set forth in the description which follows, and in part will be obvious from the description, or may be

learned by the practice of the invention. The features and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other features of the present invention will become more fully apparent from the following description and appended claims, or may
5 be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] In order to describe the manner in which the above-recited and other advantages and features of the invention can be obtained, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments
10 thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0017] Figure 1 illustrates an example distributed computation architecture that
15 facilitates performing distributed computation.

[0018] Figure 2A illustrates an additional view of replication for a local stateful computation.

[0019] Figure 2B illustrates a further view of the primary compute instance for the local stateful computation depicted in Figure 2A

20 [0020] Figure 3 illustrates a view of a compute instance for a local stateless computation.

[0021] Figure 4A illustrates an example distributed computation architecture that facilitates performing a distributed computation.

[0022] Figure 4B illustrates the example distributed computation architecture of Figure
25 4A that facilitates acknowledging a received message as part of a stateful distributed computation.

[0023] Figure 5A illustrates an example distributed computation architecture that facilitates acknowledges messages between stateless local computations.

[0024] Figure 5B illustrates an example distributed computation architecture that
30 facilitates acknowledging messages between stateless and stateful local computations.

[0025] Figure 6 illustrates a flow chart of an example method for performing the stateful local computation at a replica set.

[0026] Figure 7 illustrates a flow chart of an example method for performing a stateless local computation.

[0027] Figure 8 illustrates a flow chart of an example method for performing a distributed computation.

[0028] Figure 9 illustrates a flow chart of an example method for acknowledging a received message as part of a stateful distributed computation.

5 [0029] Figure 10 illustrates a flow chart of an example method for acknowledging a received message as part of a stateless distributed computation.

DETAILED DESCRIPTION

[0030] The present invention extends to methods, systems, and computer program products for performing computations in a distributed infrastructure. In some
10 embodiments, local stateful computations are performed. A replica set within a fabric based distributed computing architecture includes a primary compute node and one or more secondary compute nodes. Each of the primary compute node and the one or more secondary compute nodes have the same current state and are configured with the same executable group of operations for performing a local computation. The group of
15 operations including at least one input or output queue related operation and at least one operation that modifies the computing state.

[0031] The primary compute node performs the local computation. To do so, the primary compute node creates a transaction for the local computation. The primary compute node logically performs one or more additional operations in the group of
20 operations. The additional operations include processing the current computing state, including one or more of: logically generating additional computing state and logically modifying existent computing state. The additional operations include logically performing the at least one input or output queue related operation. The additional operations include ending the transaction in accordance with one of a transaction commit
25 operation and a transaction abort operation.

[0032] The primary compute node's performance of the local computation is replicated into system memory to the one or more secondary compute nodes. Replication includes replicating creating a transaction for the local computation, logically performing the one or more additional operations in the group of operations, and ending the transaction.
30 Replication of ending the transaction signals the one or more secondary compute nodes to applying changes in accordance with the transaction commit operation or the transaction abort operation.

[0033] In response to ending the transaction, the primary compute node applying changes in accordance with the transaction commit operation or the transaction abort

operation. When the transaction commit operation ends the transaction, applying physical changes are applied to implement the logical application of each operation in the group of operations. Physical changes include reconciling the current computing state with any additional computing state and any modified existent computing state into a new
5 committed computing state. Physical changes also include altering the state of at least one queue to physically implement the at least one input or output queue related operation. When the transaction abort operation ends the transaction, the at least one input or output queue related operation are logically reversed.

[0034] In other embodiments, local stateless computations are performed. A compute
10 node is configured with an executable group of a plurality of operations for performing a stateless local computation, the group of operations including at least one input or output queue related operation. The compute node performs the local stateless computation.

[0035] The compute node creates a transaction for the local stateless computation. The compute node logically performs one or more operations in the group of operations. A
15 failure of the local computation is reliably detected prior to completing a specified operation after the one or more operations in the group of operations. The local computation is reliably restarted at the specified operation. The compute node logically performs the specified operation and any remaining operations after the specified operation until the plurality of operations is performed. The transaction ends in
20 accordance with one of a transaction commit operation and a transaction abort operation.

[0036] In response to ending the transaction, the compute node applying changes in accordance with the transaction commit operation or the transaction abort operation. When the transaction commit operation ends the transaction, applying physical changes to implement the logical application of each of the plurality of operations in the group of
25 operations. Applying physical changes includes altering the state of at least one queue to physically implement the at least one input or output queue related operation. When the transaction abort operation ends the transaction, logically reversing the at least one input or output queue related operation.

[0037] Embodiments of the present invention may comprise or utilize a special purpose
30 or general-purpose computer including computer hardware, such as, for example, one or more processors and system memory, as discussed in greater detail below. Embodiments within the scope of the present invention also include physical and other computer-readable media for carrying or storing computer-executable instructions and/or data structures. Such computer-readable media can be any available media that can be

accessed by a general purpose or special purpose computer system. Computer-readable media that store computer-executable instructions are computer storage media (devices). Computer-readable media that carry computer-executable instructions are transmission media. Thus, by way of example, and not limitation, embodiments of the invention can
5 comprise at least two distinctly different kinds of computer-readable media: computer storage media (devices) and transmission media.

[0038] Computer storage media (devices) includes RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store desired program code means in the form of
10 computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer.

[0039] A “network” is defined as one or more data links that enable the transport of electronic data between computer systems and/or modules and/or other electronic devices. When information is transferred or provided over a network or another communications
15 connection (either hardwired, wireless, or a combination of hardwired or wireless) to a computer, the computer properly views the connection as a transmission medium.

Transmissions media can include a network and/or data links which can be used to carry or desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer.
20 Combinations of the above should also be included within the scope of computer-readable media.

[0040] Further, upon reaching various computer system components, program code means in the form of computer-executable instructions or data structures can be transferred automatically from transmission media to computer storage media (devices)
25 (or vice versa). For example, computer-executable instructions or data structures received over a network or data link can be buffered in RAM within a network interface module (e.g., a “NIC”), and then eventually transferred to computer system RAM and/or to less volatile computer storage media (devices) at a computer system. Thus, it should be understood that computer storage media (devices) can be included in computer system
30 components that also (or even primarily) utilize transmission media.

[0041] Computer-executable instructions comprise, for example, instructions and data which, when executed at a processor, cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. The computer executable instructions may be, for example, binaries,

intermediate format instructions such as assembly language, or even source code.

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 described features or acts described

5 above. Rather, the described features and acts are disclosed as example forms of implementing the claims.

[0042] Those skilled in the art will appreciate that the invention may be practiced in network computing environments with many types of computer system configurations, including, personal computers, desktop computers, laptop computers, message processors, 10 hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, mobile telephones, PDAs, pagers, routers, switches, and the like. The invention may also be practiced in distributed system environments where local and remote computer systems, which are linked (either by hardwired data links, wireless data links, or by a combination 15 of hardwired and wireless data links) through a network, both perform tasks. In a distributed system environment, program modules may be located in both local and remote memory storage devices.

[0043] Embodiments of the invention include a general purpose distributed computation infrastructure that can be used to perform efficient (in-memory), scalable, failure-resilient, 20 atomic, flow-controlled, long-running state-less and state-full distributed computations. Guarantees provided by a distributed computation infrastructure can build upon existent guarantees of an underlying distributed fabric, such as, for example, an overlay network, in order to hide the complexities of fault-tolerance, enable large scale highly available processing, allow for efficient resource utilization, and facilitate generic development of 25 stateful and stateless computations. Guarantees can include neighborhood and routing consistency, perfect failure detection, and data consistent replication and failover reconfigurations. A distributed computation infrastructure can also provide a substrate on which existent distributed computation models can be enhanced to become failure-resilient.

30 **[0044]** In some embodiments, service state/data/configuration is replicated to provide high availability for distributed servers. Compute nodes (entities involved in replication relationships) form a replica set or partition. Within a replica set or partition, one compute node can be a primary compute node and one or more other compute nodes can be secondary compute nodes. Read operations can be completed by the primary compute

node. Write operations can be seen first by the primary compute node and then forwarded to secondary compute nodes using a quorum commit protocol. Each secondary compute node in a replica set can see and apply changes in the same order as the primary compute node saw and applied them.

5 [0045] A set of failover and reconfiguration algorithms can be employed to provide clients with the illusion of receiving service from a single consistent entity. The algorithms can provide guarantees around maintaining replica sets (partitions) in a functioning state. Maintaining a functioning state can include failover from a primary compute node to a secondary compute node, building a new secondary compute node
10 when too few secondary compute nodes are available, and switching a primary compute node with an existing secondary compute node.

[0046] In order to provide consistency guarantees in a distributed computation infrastructure, failure detectors can be used. A failure detector has responsibility for providing reliable information about which entities within distributed computation
15 architecture (e.g., processes) are down. Different types of failure detectors can be used. In some embodiments, a perfect failure detector, which is strongly accurate (no process that is up is reported to be down) and strongly complete (all processes that are down are reported) is used.

[0047] Embodiments of the invention also include a messaging sub-system. The
20 messaging sub-system can utilize messaging primitives to accomplish sending a message to a node/process or to a replica set within a distributed computation architecture. Utilized messaging paradigms can include point-to-point, multicast, and broadcast, with one or more of best-effort and reliable versions.

[0048] A local computation is defined as being a replica set (partition). Inputs of a local
25 computation can modeled as a set of First In, First Out (“FIFO”) queues (sources). Outputs of a local computation can modeled as a set of FIFO queues (sinks). Each instance of a local computation is a class implementing an interface for exposing its state (if any) and for performing the actual processing. For state-full computations, the state and/or state changes can be made highly available through a fabric (e.g., overlay network)
30 replication layer. Local computations that consume inputs and produce outputs can be made highly available in the same manner.

[0049] As such, a distributed computation can be represented by a Directed Acyclic Graph (“DAG”) of local computations. Outputs from upstream local computations are delivered as input to downstream local computations.

[0050] Figure 1 illustrates an example distributed computation architecture 100 that facilitates performing distributed computation 108. Referring to Figure 1, distributed computation architecture 100 includes replicas 191, 192, 193, and 194. Each of the depicted replicas (which may include one or more computer systems) is connected to one another over (or is part of) a network, such as, for example, a Local Area Network ("LAN"), a Wide Area Network ("WAN"), and even the Internet. Accordingly, each of the depicted replicas as well as any other connected computer systems and their components, can create message related data and exchange message related data (e.g., Internet Protocol ("IP") datagrams and other higher layer protocols that utilize IP datagrams, such as, Transmission Control Protocol ("TCP"), Hypertext Transfer Protocol ("HTTP"), Simple Mail Transfer Protocol ("SMTP"), etc.) over the network.

[0051] Distributed computation 108 is represented by a DAG formed from replicas (local computations) 191, 192, 193, and 194. As depicted, each replica (local computation) includes a plurality of compute instances. For example, replica 191 includes compute instances 101A, 101B, and 101C. Compute instance 101A includes compute state 111, compute processing 121, and sink queue 131. Similarly, replica 192 includes compute instances 102A, 102B, and 102C. Compute instance 102A includes compute state 112, compute processing 122, and sink queue 132.

[0052] Likewise, replica 193 includes compute instances 103A, 103B, and 103C.

Compute instance 103A includes compute state 113, compute processing 123, source queues 133 and 143 and sink queue 153. In turn, replica 194 includes compute instances 104A, 104B, and 104C. Compute instance 104 includes compute state 114, compute processing 124, and source queue 134.

[0053] Generally, compute processing 121 and compute processing 122 perform one or more operations to alter compute state 111 and 112 and output messages 191 and 192 respectively. Compute instance 101A can queue message 191 in sink queue 131. Compute instance 101B can queue message 192 in sink queue 132.

[0054] Changes to compute state 111 and the enqueueing of message 191 in sink queue 131 can be replicated to compute instances 101B and 101C. Similarly, changes to compute state 112 and the enqueueing of message 192 in sink queue 132 can be replicated to compute instances 102B and 102C.

[0055] A queue-to-queue protocol can be used to dequeue messages 191 and 192 from sink queues 131 and 132 respectively and enqueue messages 191 and 192 in source queues 133 and 143 respectively. Dequeueing of message 191 from sink queue 131 can be

replicated to compute instances 101B and 101C. Similarly, dequeuing of message 192 from sink queue 132 can be replicated to compute instances 102B and 102C. Likewise, enqueueing of messages 191 and 192 in source queues 133 and 143 can be replicated to compute instances 103B and 103C.

5 [0056] Compute processing 123 can use messages 191 and 192 to perform one or more operations to alter compute state 113 and output message 193. Compute instance 103A can queue message 193 in sink queue 153. Changes to compute state 113 and the enqueueing of message 193 in sink queue 153 can be replicated to compute instances 103B and 103C.

10 [0057] The queue-to-queue protocol can be used to dequeue message 193 from sink queue 153 and enqueue messages 193 in source queue 134. Dequeueing of message 193 from sink queue 153 can be replicated to compute instances 103B and 103C. Likewise, enqueueing of messages 193 in source queue 134 can be replicated to compute instances 104B and 104C.

15 [0058] Compute processing 124 can use message 193 to perform one or more operations to alter compute state 114. Changes to compute state 113 can be replicated to compute instances 104B and 104C.

[0059] Figure 2A illustrates an additional view of replication for a local stateful computation 212, such as, for example, an edit computation, a sort computation, etc.

20 Generally, local computation 212 receives inputs 216 and generates outputs 217.

[0060] As depicted, primary compute instance 201A includes source queues 206 and 207, local computation 212, and sink queue 208. Local computation 212 further includes computer state 204 and compute processing 203. As compute processing 203 executes, source queues 206 and 207 (e.g., dequeuing messages), compute state 204 (e.g., changing values), and output queue 208 (e.g., enqueueing messages) can be changed. These changes are replicated to corresponding source queues 206 and 207, compute state 204, and output queue 208 at secondary compute instances 201B and 201C. Thus, in the event primary compute instance 201A fails, one of secondary compute instances 201B, 201C, etc can assume the role as the primary compute node. Upon assuming the role as primary
25 computer node, the assuming node can take over compute processing 203.

[0061] Operations that are part of location computation include consumption of input data, production of output data, and modifications to its internal state, is associated with a transaction. The transaction concept provides atomicity guarantees across a set of operations are part of a local computation. A local computation can incrementally create a

set of transactions, sequential or concurrent, in which it records information that needs to be made highly available. Transaction creation can include dequeue operations from the source queues, enqueue operations into the sink queues, and internal state changes. When a transaction completes, all changes created by the operations are committed (or aborted) as an atomic unit. Thus transactions help ensure the processing continues from the same consistent view in the event of a failure. A fail-fast approach can be being taken when a transaction fails to commit (or abort for that matter).

[0062] Source queues provide a dequeue operation and sink queues provide an enqueue operation. Transactions provide the ability to commit or abort the current changes made by a local computation. Each local computation implements a state get/set interface used in case a new replica has to be built for the local computation. Also, each local computation implements an interface that allows for initial processing parameters to be passed in as part of the computation initialization (initial state, configuration options etc.).

[0063] As such, during a local computation, messages are dequeued/consumed from source queues and are processed in an active transaction (e.g., with a given id). As part of the transaction, new state can be created and new output messages produced.

[0064] When a transaction commits, the messages dequeued from the source queues are physically removed. The new output message produced is also enqueued into the sink queue. The new committed state is a committed reconciliation between the old committed state and the uncommitted state change that was produced while processing. For computations is designed in this manner, when failovers occur in the system, the state of the computation and its inputs and outputs are consistent with each other. This results in a simpler experience for the development of new types of highly available local computations.

[0065] Operations performed as part of a local computation can take effect at commit time. That is, when atomically all changes that occurred in processing take place. Logical cursors can be maintained on the source queues and the sink queues may not be touched effectively until commit.

[0066] Transaction operations can be replicated in different modes. In some embodiments, an eager replication mode is used. Using eager replication transaction operations are replicated as they occur and each replica in a replica set buffers operations. In other embodiments, a lazy replication mode is used. Using a lazy replication mode transaction operations on a primary compute instance are seem on second compute instances at a later time. Lazy replication mode can be used when the transaction

operation cost is smaller than the cost of replicating it. Thus, amortizing the cost by batching replication of transaction operations is beneficial. In one specific implementation of lazy mode replication transaction operations are replicated only at commit time.

5 [0067] When a transaction aborts, the transaction can restore elements from the source queues and make them available for subsequent dequeue operations. New state changes can be forgotten and do not necessarily need to be undone. Since it is unlikely sink queues have been touched, outputs can be discarded.

[0068] On a failover, in-flight transactions that exist currently on the new primary (old secondary) can be discarded (as they had no work to do). If a failover does not involve the
10 crash of a primary (e.g., a failover is caused by a load balancing swap role action), transactions in progress on the old primary (new secondary) receive aborting exceptions and processing on the old primary stops.

[0069] Concurrent execution of transactions within a single computation can occur.

[0070] Figure 2B illustrates a further view of the primary compute instance 201A for the
15 local stateful computation 212. As depicted, compute processing 203 can include a number of operations that are executed to implement local stateful computation 212. Included in the operations are operations for dequeuing messages from source queues (e.g., lines 3 and 4 and 6 and 7) and enqueueing messages in sink queues (e.g., lines 11 and 12). Included in the operations are also operations for performing computations based
20 message content (e.g., lines 5 and 8). Include in the operations are also operations for enqueueing a message (e.g., lines 11 and 12). Included in the operation are also operations for implementing state changes (e.g., lines 9 and 10). These operations can be closed within other operations indicating the beginning and end of a transaction (e.g., lines 1 and 2 and 13 and 14).

25 [0071] The performance of operations at primary compute instance 201A can be replicated into system memory 214 to secondary compute nodes 201B, 201C, etc. Replication can occur on a per operation basis or on a batch of operations basis.

[0072] Figure 6 illustrates a flow chart of an example method 600 for performing the stateful local computation at a replica set. Method 600 will be described with respect to
30 the components and data of Figure 2B.

[0073] Method 600 includes an act of a primary compute node performing a local computation (act 601). For example, primary compute instance 201A can perform local computation 212. Performing the local computation includes an act of the creating a

transaction for the local computation (act 602). For example, lines 1 and 2 of compute processing 203 can be executed to create a transaction for location computation 212.

[0074] Performing the local computation includes an act of logically performing one or more additional operations in the group of operations (act 603). For example, lines 3-12
5 of compute processing 203 can be executed to logically perform one or more additional operations of local computation 212.

[0075] Logically performing one or more additional operations includes an act of processing the current computing state, processing the current computing state including one or more of: logically generating additional computing state and logically modifying
10 existent computing state (act 604). For example, lines 5 and 8 can be executed to process compute state 204 along with content from corresponding messages 211 and 212 respectively. Processing compute state 204 can logically generate new state or modify compute state 204. Line 10 can be executed to write out new and/or changed state 214.

[0076] Logically performing one or more additional operations includes an act of
15 logically performing the at least one input or output queue related operation (act 605). For example, lines 3 and 4 or lines 6 and 7 can be executed to perform input queue operations related to source queues 206 and 207 respectively. Lines 11 and 12 can be executed to perform output queue operations related to sink queue 208.

[0077] Performing the local computation includes an act of ending the transaction in
20 accordance with one of a transaction commit operation and a transaction abort operation (act 606). For example, lines 13 and 14 of compute processing 203 can be executed to commit the transaction created at lines 1 and 2. Alternately, lines of instructions for aborting the transaction created at lines 1 and 2 can be executed.

[0078] Method 600 includes an act of replicating the primary compute node's
25 performance of the local computation into system memory to the one or more secondary compute nodes, including replicating creating a transaction for the local computation, logically performing the one or more additional operations in the group of operations, and ending the transaction, replication of ending the transaction signaling the one or more secondary compute nodes to applying changes in accordance with the transaction commit
30 operation or the transaction abort operation (act 607). For example, primary compute instance 201A's performance of local computation 212 can be replicated into system memory 214 to secondary compute instances 201B, 201C, etc. Replication of the performance of local computation 212 can include replicating creating the transaction at lines 1 and 2, performing one or more additional operations at lines 3-12, and committing

the transaction at lines 13 and 14. Alternately, when the transaction is aborted, replication can include replicating aborting the transaction. Replication can occur on a per operation basis (e.g., a line at a time) or on a batch of operations basis (e.g., grouped operations).

5 [0079] Replication of ending (either committing or aborting) the transaction signals the secondary compute instances 201B, 201C, etc, to applying changes in accordance with the transaction commit operation or the transaction abort operation as appropriate.

[0080] Method 600 includes in response to ending the transaction, an act of the primary compute node applying changes in accordance with the transaction commit operation or the transaction abort operation (act 608). For example, primary compute instance 201A
10 can apply changes in accordance with the transaction commit operation (at lines 13 and 14). Alternately, primary compute instance 201A can apply changes in accordance with a transaction abort operation.

[0081] Method 600 includes when the transaction commit operation ends the transaction, an act of applying physical changes to implement the logical application of
15 each operation in the group of operations (act 609). For example, when a transaction commit ends the transaction, primary compute instance 201A can apply physical changes to implement logical application of lines 3-12 of compute processing 203.

[0082] Applying physical changes includes an act of reconciling the current computing state with any additional computing state and any modified existent computing state into a
20 new committed computing state (act 610). For example, primary compute instance 201A can reconcile compute state 204 with changed state 214 (which can include changed and/or new state) into new committed state 204.

[0083] Applying physical changes includes method 600 includes an act of altering the state of at least one queue to physically implement the at least one input or output queue
25 related operation (act 611). For example, primary compute instance 201A can physically dequeue messages 211 and 212 from source queues 206 and 207 respectively and physically enqueue message 213 into sink queue 208.

[0084] Method 600 includes when the transaction abort operation ends the transaction, an act of logically reversing the at least one input or output queue related operation (act
30 612). For example, when a transaction abort ends the transaction, primary compute instance 201A can logically reverse the execution of lines 3-12 of compute processing 203. This can include reversing logical enqueue and dequeue message operations.

[0085] Figure 3 illustrates a view of a compute instance 301 for a local stateless computation 312, such as, for example, a read computation, a filtering computation, etc.

As depicted in Figure 3 includes compute instance 301, failure detector 331, and restart module 332. Compute instance 301 includes source queue 306 local computation 312, and sink queue 307. As compute processing 303 executes, source queue 306 (e.g., dequeuing messages) and output queue 208 (e.g., enqueueing messages) can be changed. Failure
5 detector 331 is configured to detect when and where an operation in compute processing 303 fails. Restart module 332 is configured to restart compute processing 303 from a point of failure.

[0086] Figure 7 illustrates a flow chart of an example method 700 for performing a stateless local computation. Method 700 will be described with respect to the components
10 and data of Figure 3.

[0087] Method 700 includes an act of the compute node performing a local stateless computation (act 701). For example, compute instance 201 can perform local computation 312. Method 700 includes an act of the creating a transaction for the local stateless
15 computation (act 702). For example, lines 0 and 1 of compute processing 303 can be executed to create a transaction for location computation 312. Method 700 includes an act of logically performing one or more operations in the group of operations (act 703). For example, compute instance 301 can logically perform lines 2 and 3 of compute processing 303.

[0088] Method 700 includes an act of reliably detecting a failure of the local
20 computation prior to completing a specified operation after the one or more operations in the group of operations (act 704). For example, failure detector 331 can detect a failure 341 prior to completing line 4 of compute processing 303. Method 700 includes an act of reliably restarting the local computation at the specified operation (act 705). For example, restart module 332 can issue restart 342 to restart local computation 312 at line 4 of
25 compute processing 303.

[0089] Method 700 includes an act of logically performing the specified operation and any remaining operations after the specified operation until the plurality of operations is performed (act 706). For example, compute instance 301 can logically perform line 4 and then lines 5-9 of compute processing 303. Method 700 includes an act of ending the
30 transaction in accordance with one of a transaction commit operation and a transaction abort operation (act 707). For example, lines 8 and 9 of compute processing 3030 can be executed to commit the transaction created at lines 0 and 1. Alternately, lines of instructions for aborting the transaction created at lines 0 and 1 can be executed.

[0090] Method 700 includes in response to ending the transaction, an act of the primary compute node applying changes in accordance with the transaction commit operation or the transaction abort operation (act 708). For example, compute instance 301 can apply changes in accordance with the transaction commit operation (at lines 8 and 9).

5 Alternately, compute instance 301 can apply changes in accordance with a transaction abort operation.

[0091] Method 700 includes when the transaction commit operation ends the transaction, an act of applying physical changes to implement the logical application of each of the plurality of operations in the group of operations (act 709). For example, when
10 a transaction commit ends the transaction, compute instance 301 can apply physical changes to implement logical application of lines 2-7 of compute processing 303.

[0092] Method 700 includes an act of altering the state of at least one queue to physically implement the at least one input or output queue related operation (act 710). For example, compute instance 301 can physically dequeue messages 211 and from source
15 queues 206 and physically enqueue message 312 into sink queue 307.

[0093] Method 700 includes when the transaction abort operation ends the transaction, an act of logically reversing the at least one input or output queue related operation (act 711). For example, when a transaction abort ends the transaction, compute instance 301 can logically reverse the execution of lines 2-8 of compute processing 303. This can
20 include reversing logical enqueue and dequeue message operations.

[0094] Figure 4A illustrates distributed computation architecture 400 that facilitates performing a distributed computation 499. As depicted, distributed computation architecture includes replicas 491 and 492. Replica 491 includes primary compute instance 401A and secondary compute instances 401B and 401C. Likewise, replica 492
25 includes primary compute instance 402A and secondary compute instances 402B and 402C.

[0095] Generally, messages (e.g., message 496) can be transferred from replica 491 to replica 492 in accordance with a queue to queue protocol. The queue to queue protocol can be responsible for implementing a state machine for the exactly once in order transfer
30 of messages from a sink queue to one or more downstream source queues. The queue protocol can limit message transfers to committed messages in the sink queues. Limiting to committed messages can be accomplished by keeping state in each sink queue of a committed mark that gets updated with each commit. The queue to queue protocol can

also implement flow-control policies applied at either or all endpoints in order to prevent over/under utilization in the system.

[0096] The queue to queue protocol can rely on existent fabric (e.g., overlay network) components such as, for example, consistent routing and (e.g., perfect) failure detection.

5 Failure detection can be used to determine a new primary replica of a downstream computation in the case of a failover, and a routing layer can be used to implement reliable message delivery. A receiver compute node's protocol endpoint can implement duplicate detection, message re-ordering, and throttling based on queue quotas. A sender computer node's endpoint can implement a sending window algorithm with retries and exponential
10 back-off.

[0097] Using a windowing mechanism, the queue to queue protocol can achieve flow control in a DAG representing a distributed computation. Temporary performance bottlenecks in a local computation of a downstream partition translates into queue-to-queue protocol exerting back pressure on the sink queue feeding into source queue of the
15 affected local computation. The back pressure in turn affects timely completion of the enqueue operations of the local computations of the upstream partition feeding into the affected sink queue. The reduction in timely completion further in turn translates to the sources queues of the local computation of the upstream partition becoming full. This can potential propagate upstream to the sources of the distributed computation are reached.

20 The windowing mechanism can also automatically restore the DAG representing the distributed computation back to its healthy state after the temporary performance bottleneck in the local computation of the downstream partition is removed.

[0098] Turning back to Figure 4A, primary compute instance 401A includes compute state 411 and compute processing 421 that facilitate performing local computation 441.

25 Local computation 441 can enqueue a message to sink queue 431. Sink queue 431 can include a committed mark 441 indicating a location within sink queue 431. Messages after the location are committed messages and messages before the location are not yet committed (and thus could still be removed on an abort). Primary compute instance 401A is configured to send committed messages in sink queue 431 to primary compute instance
30 402A as part of a distributed computation.

[0099] Primary compute instance 402A includes compute state 412 and compute processing 422 that facilitate performing a local computation 442. Local computation 442 can dequeue a message from source queue 432. Primary compute instance 402A is

configured to receive messages in source queue 432 from primary compute instance 401A as part of the distributed computation.

[0100] Embodiments of the invention include combinations of one or more stateful local computations and/or one or more stateless local computations interconnected with one another in a DAG representing a distributed computation.

[0101] Figure 8 illustrates a flow chart of an example method 800 for performing a distributed computation. Method 800 will be described with respect to the components and data of distributed computation architecture 400 in Figure 4A.

[0102] Method 800 includes an act of a first replica set performing a first local computation representing a portion of the distributed computation (act 801). For example, replica 491 can perform local computation 441 as part of distributed computation 499. Performing a first local computation includes an act of creating a transaction for the first local computation (act 802). For example, operations in compute processing 421 can create a transaction for local computation 441.

[0103] Performing a first local computation an act of logically performing the group of operations for the first local computation, including performing at least one output queue related operation (act 803). For example, primary compute instance 401A can logically perform a group of operations in compute processing 421. The group of operations can include enqueueing message 496 in sink queue 431.

[0104] Performing a first local computation includes an act of committing the transaction (act 804). For example, the transaction created for local computation 441 can be committed.

[0105] In response to committing the transaction, method 800 includes an act of applying physical changes to implement the logical application of each operation in the group of operations, including inserting a message into at least one output queue to physically implement the at least one output queue related operation (act 805). For example, primary compute instance 401A can apply physical changes to implement the logical application of operations in compute processing 421. Physical operations can include reconciling changed and/or new state with compute state 411 and physically enqueueing message 496 in sink queue 431.

[0106] Method 800 includes an act of updating a committed mark within the at least one output queue to indicate that the inserted message is a committed message (act 806). For example, primary compute node 401A can update committed mark 441 to indicate that message 496 is a committed message.

[0107] Method 800 includes an act of transferring the message from the at least one output queue of the first replica set to at least one input queue of at least one second replica set in accordance with a queue-to-queue protocol that detects the inserted message has been committed (act 807). For example, message 496 can be transferred from sink queue 431 to source queue 432 in accordance with queue to queue protocol (having a send end point at replica 491 and a receive endpoint at replica 492) that detects message 496 as committed.

[0108] Method 800 includes an act of the second replica set performing a second local computation representing a next portion of the distributed computation (act 808). For example, replica 492 can perform local computation 442 as part of distributed computation 499. Method 800 includes an act of creating a second transaction for the second local computation (act 809). For example, operations in compute processing 421 can create a transaction for local computation 441.

[0109] Method 800 includes an act of logically performing the group of operations for the second local computation, including performing at least one input queue related operation (act 810). For example, primary compute instance 402A can logically perform a group of operations in compute processing 422. The group of operations can include dequeuing message 496 from source queue 432.

[0110] Method 800 includes an act of ending the second transaction in accordance with one of a transaction commit operation and a transaction abort operation (act 811). For example, the transaction created for local computation 441 can be ended in accordance with a transaction commit operation or a transaction abort operation.

[0111] Method 800 includes when the transaction commit operation ends the transaction, an act of applying physical changes to implement the logical application of each operation in the group of operations, including accessing the message from the at least one input queue to physically implement the at least one input queue related operation (act 812). For example, primary compute instance 402A can apply physical changes to implement the logical application of operations in compute processing 422. Physical operations can include reconciling changed and/or new state with compute state 412 and physically dequeuing message 496 from source queue 431.

[0112] Method 800 includes when the transaction abort operation ends the transaction, an act of logically reversing the at least one input queue related operation (act 813). For example, when a transaction abort ends the transaction from local computation 442,

primary compute instance 402A can logically reverse the execution operations in compute processing 422. This can include reversing the logical dequeuing of message 491.

[0113] Figure 4B illustrates distributed computation architecture 400 that facilitates acknowledging a received message as part of stateful distributed computation 499.

5 [0114] Figure 9 illustrates a flow chart of an example method 900 for acknowledging a received message as part of a stateful distributed computation. Method 900 will be described with respect to the components and data of distributed computation architecture 400 in Figure 4B.

[0115] Method 900 includes an act of receiving a message at an input queue of a first replica set in accordance with a queue-to-queue protocol, the message sent from an output queue at a second replica set in response to the message being committed at the second replica set (act 901). For example, replica 492 receives message 496 at sink queue 432 in accordance with a queue to queue protocol. Message 496 was sent from source queue 431 at replica 491 in response to the message being committed at replica 491.

15 [0116] Method 900 includes an act of replicating the message to the one or more secondary compute nodes included in the second replica set (act 902). For example, replica 492 can replicate message 496 to secondary compute nodes 402B, 402C, etc.

[0117] Method 900 includes an act of receiving acknowledgements from a write quorum of the one or more secondary computer nodes that the message was successfully replicated, acknowledgements from the write quorum indicative of the message being quorum committed within the first replica set (act 903). For example, a write quorum (e.g., a simple majority or other designated threshold) of secondary compute instances 402B, 402C, etc., can send acknowledgments to primary compute instance 402A acknowledging receipt of message 496. For example, secondary compute instance 402B can send acknowledgement 497, secondary compute instance 402C can send acknowledgement 498, etc. Primary compute instance 402A can receive acknowledgments including acknowledgments 497, 498, etc., from the secondary compute nodes.

[0118] In some embodiments, a write quorum is defined as ceiling of $[(N+1)/2]$, N=replica count in the replica set. Thus, if a replica had five secondary compute instances, at least three of the secondary computing instances would have to acknowledge a message to achieve write quorum.

[0119] Method 900 includes an act of the first replica set sending an acknowledgement to the second replica set in response to the message being quorum committed within the

first replica set, the acknowledgement acknowledging that the first replica set successfully received the message and that the second replica set is permitted to dispose of the message (act 904). For example, replica 492 can sent acknowledgement 494 back to replica 491 in response to a sufficient number of secondary compute nodes acknowledging receipt of message 496. Acknowledgment 494 indicates that replica 492 successfully received message 496 and that replica 491 is permitted to dispose of message 496.

[00120] Figure 5A illustrates an example distributed computation architecture that facilitates acknowledges messages between stateless local computations. As depicted, distributed computation architecture 500 includes compute instances 501, 502, and 503 and other compute instances 504 and 506 for performing distributed computation 549. Compute instances 501, 502, and 503 include stateless compute processing 521, 522, and 523 respectively. Each of compute instances 501, 502, and 503 includes a source queue and a sink queue to facilitate the exchange of messages, for example, in accordance with a queue to queue protocol. Compute instance 501 includes source queue 531 and sink queue 541, compute instance 502 includes source queue 532 and sink queue 542, compute instance 503 includes source queue 533 and sink queue 543. Other compute instances 504 can precede (or are upstream from) compute instance 501. Other compute instances 505 can succeed (or are downstream from) compute instance 503.

[0121] Acknowledging a message sent between stateless local computations can differ from acknowledging a message sent between stateful local computations. For example, a stateless local computation can operate on a single compute instance (as opposed to a replica). As such, received messages are not replicated to secondary compute instances. Further, since stateless local computations can fail and restart (without replication), a failure can result in a received message being deleted at a compute instance. However, to restart, the stateless local computation may again need access to the message. Thus, the stateless local computation can re-request the message from an upstream compute instance. That is, if a downstream compute instance sends an acknowledgment upstream to early and then goes down and is restarted, further downstream compute instance can experience input loss.

[0122] Accordingly, in some embodiments, a stateless computation acknowledges input after the outputs that it has created based on those inputs are themselves acknowledged. Acknowledgment in this manner can be recursive, such as, for example, when a plurality of stateless local computations follow one another in a DAG.

[0123] Figure 10 illustrates a flow chart of an example method 1000 for acknowledging a received message as part of a stateless distributed computation. Method 1000 will be described with respect to the components and data of distributed computation architecture 500.

5 [0124] Message 581 can be received in source queue 531 from other compute instances 504. Upon receiving message 581, compute instance 501 can perform compute processing 521. Performing compute processing 521 can cause message 582 to be enqueued in sink queue 541. Upon message 582 being committed, message 582 can be sent to compute instance 502.

10 [0125] Method 1000 includes an act of receiving a message at an input queue of a first compute node in accordance with a queue-to-queue protocol, the message sent from an output queue at a second compute node in response to the message being committed at the second compute node (act 1001). For example, source queue 532 can receive message 582 from sink queue 541 in accordance with a queue to queue protocol and in response to
15 message 582 being committed at compute instance 501.

[0126] Method 1000 includes an act of the creating a transaction for the local stateless computation (act 1002). For example, compute instance 502 can create a transaction for compute processing 522. Method 1000 includes an act of logically performing the group of operations, including generating an additional message based on the message (act
20 1003). For example, compute instance 502 can logically perform a group of operations in compute processing 522. Logically performing the group of operations in compute processing 522 can include generating message 583 based on message 582. Method 1000 includes an act of committing the transaction (act 1004). For example, the transaction created for compute processing 522 can be committed.

25 [0127] In response to committing the transaction, method 1000 includes an act of applying physical changes to implement the logical application of each operation in the group of operations, including inserting the additional message into at least output queue to physically implement the at least one output queue related operation (act 1005). For example, in response to committing the transaction for compute processing 522, compute
30 instance 502 can apply physical operations to implement the logically performed operations. Applying physical operations can include enqueueing message 583 in sink queue 542.

[0128] Method 1000 includes an act of transferring the additional message from the at least one output queue of the first compute node to at least one input queue of a third

compute node in accordance with the queue-to-queue protocol in response to the message being committed at the second compute node (act 1006). For example, message 583 can be transferred from sink queue 542 to source queue 533 in accordance with the queue to queue protocol in response to message 583 being committed at sink queue 542.

5 [0129] Upon receiving message 583, compute instance 503 can perform compute processing 523 and transfer message 584 onto other compute instances 506. Eventually, due to reaching an end compute instance (or a stateful local computation), compute instance 502 can receive acknowledgement 591 for message 584. In turn, compute instance 503 can send acknowledgement 592 to compute instance 502. Acknowledgement
10 592 indicates that compute instance 503 successfully received message 583 and that compute instance 502 is permitted to dispose of message 583.

[0130] Method 1000 includes an act of, subsequent to transferring the additional message to the third compute node, receiving an acknowledgement from the third compute node, the acknowledgement acknowledging that the third compute node successfully
15 received the additional message (act 1007). For example, compute instance 502 can receive acknowledgement 592 subsequent to sending message 583 to compute instance 503.

[0131] Method 1000 includes an act of the first compute node sending a further acknowledgement to second compute node in response to receiving the acknowledgement
20 from the third compute node, the further acknowledgement acknowledging that the first compute node successfully received the message (act 1008). For example, compute instance 502 can send acknowledgment 591 to compute instance 501. Acknowledgment 591 indicates that compute instance 502 successfully received message 582 and that compute instance 501 is permitted to dispose of message 582.

25 [0132] Compute instance 501 can then send acknowledgment 594 upstream to other compute instances 504 as appropriate.

[0133] Figure 5B illustrates an example distributed computation architecture 550 that facilitates acknowledging messages between stateless and stateful local computations. As depicted, distributed computation architecture 550 includes replica 511, compute instance
30 502, and replica 503 and other compute instances 514 and 516 for performing distributed computation 599. Replica 511 includes primary compute instance 511A and secondary compute instances 511B and 511C. Replica 512 includes primary compute instance 512A and secondary compute instances 512B and 512C.

[0134] Primary compute instance 511A includes stateful compute processing 541. Compute instance 512 includes stateless compute processing 542. Primary compute instance 513A includes stateful compute processing 543.

[0135] Each of primary compute instance 511A, compute instance 502, and primary
5 compute instance 513A includes a source queue and a sink queue to facilitate the exchange of messages, for example, in accordance with a queue to queue protocol. Primary Compute instance 511A includes source queue 551 and sink queue 561, Compute instance 502 includes source queue 552 and sink queue 562. Primary compute instance 513A includes source queue 553 and sink queue 563. Other compute instances 514 can
10 precede (or are upstream from) replica 511. Other compute instances 505 can succeed (or are downstream from) replicate 513.

[0136] Message 554 can be received in source queue 551 from other compute instances 554. Upon receiving message 554, primary compute instance 511A, replicates message 554 to secondary compute instance 511B and 511C. Secondary compute instances 511B
15 and 511C (a write quorum) can acknowledge receipt of message 554 by sending acknowledgments 564 and 566 respectively to primacy computer instance 511A. Primary compute instance 511A can send acknowledgement 574 back to other computer instances 514 in response to message 554 being quorum committed. Acknowledgment 575 indicates that replica 511 successfully received message 554 and that other computer
20 instances 514 are permitted to dispose of message 554.

[0137] Compute processing 541 can execute, generating message 556 and enqueue message 556 in sink queue 553. Sink queue 561 includes committed mark 567 indicating a location within sink queue 561. Messages after the location are committed messages and messages before the location are not yet committed (and thus could still be removed on an
25 abort). Primary compute instance 511A is configured to send committed messages in sink queue 553 to compute instance 512 as part of distributed computation 599.

[0138] Source queue 552 can receive message 556 from sink queue 561 in accordance with a queue to queue protocol and in response to message 561 being committed at primary compute instance 511A. Compute processing 542 can be performed to cause
30 message 557 to be enqueued in sink queue 562. Upon message 557 being committed, message 557 can be sent to primary compute instance 513A. Message 557 can be transferred from sink queue 562 to source queue 553 in accordance with the queue to queue protocol in response to message 557 being committed at sink queue 562. Compute

instance 511 is configured to send committed messages in sink queue 562 to primary compute instance 513A as part of distributed computation 599.

[0139] As previously described, compute instance 512 (a stateless local operation) may not immediately acknowledge receipt of message 556.

5 [0140] Message 557 can be received in source queue 553 from compute instance 512. Upon receiving message 557, primary compute instance 513A, replicates message 557 to secondary compute instance 513B and 513C. Secondary compute instances 513B and 513C (a write quorum) can acknowledge receipt of message 557 by sending acknowledgments 568 and 569 respectively to primary compute instance 511A. Primary
10 compute instance 513A can send acknowledgement 572 back to other compute instance 512 in response to message 557 being quorum committed.

[0141] Acknowledgment 572 indicates that replica 513 successfully received message 557 and that compute instance 512 is permitted to dispose of message 557. In response in acknowledgment 572, compute instance 512 can then send acknowledgment 573 upstream
15 to primary compute instance 511A. Acknowledgment 573 indicates that compute instance 512 successfully received message 556 and that primary compute instance 511A is permitted to dispose of message 556.

[0142] Compute processing 543 can also execute, generating message 558 and enqueue message 558 in sink queue 533. Sink queue 563 includes committed mark 571 indicating
20 a location within sink queue 563. Messages after the location are committed messages and messages before the location are not yet committed (and thus could still be removed on an abort). Primary compute instance 513A is configured to send committed messages in sink queue 563 to other computer instances 516 as part of distributed computation 599. Primary compute instance 513A can eventually receive back acknowledgment 576 for
25 message 558 from other compute instances 516.

[0143] As described, embodiments of the invention include a general purpose distributed computation infrastructure that can be used to perform efficient (in-memory), scalable, failure-resilient, atomic, flow-controlled, long-running state-less and state-full distributed computations. Guarantees provided by a distributed computation infrastructure
30 can build upon existent guarantees of an underlying distributed fabric (e.g., neighborhood and routing consistency, perfect failure detection, and data consistent replication and failover reconfigurations) in order to hide the complexities of fault-tolerance, enable large scale highly available processing, allow for efficient resource utilization, and facilitate generic development of stateful and stateless computations. A distributed computation

infrastructure can also provide a substrate on which existent distributed computation models can be enhanced to become failure-resilient.

[0144] The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

CLAIMS

1. At a replica set within a fabric based distributed computing architecture, the replica set including one or more processors and system memory, the replica set including a primary compute node and one or more secondary compute nodes, the primary compute node and the one or more secondary compute nodes having the same current computing state, each of the primary compute node and the one or more secondary compute nodes configured with the same executable group of operations for performing a stateful local computation, the group of operations including at least one input or output queue related operation and at least one operation that modifies the computing state, a method for performing the stateful local computation at the replica set, the method including:
- an act of the primary compute node performing the local stateful computation by:
 - an act of the creating a transaction for the local computation;
 - an act of logically performing one or more additional operations in the group of operations, including:
 - an act of processing the current computing state, processing the current computing state including one or more of: logically generating additional computing state and logically modifying existent computing state;
 - an act of logically performing the at least one input or output queue related operation; and
 - an act of ending the transaction in accordance with one of a transaction commit operation and a transaction abort operation;
 - an act of replicating the primary compute node's performance of the local computation into system memory to the one or more secondary compute nodes, including replicating creating a transaction for the local computation, logically performing the one or more additional operations in the group of operations, and ending the transaction, replication of ending the transaction signaling the one or more secondary compute nodes to applying changes in accordance with the transaction commit operation or the transaction abort operation;
 - in response to ending the transaction, an act of the primary compute node applying changes in accordance with the transaction commit operation or the transaction abort operation, including:

when the transaction commit operation ends the transaction, an act of applying physical changes to implement the logical application of each operation in the group of operations, including:

an act of reconciling the current computing state with any additional
5 computing state and any modified existent computing state into a new committed computing state; and

an act of altering the state of at least one queue to physically implement the at least one input or output queue related operation; and

when the transaction abort operation ends the transaction, an act of

10 logically reversing the at least one input or output queue related operation.

2. The method as recited in claim 1, wherein the act of the primary compute node performing the local computation comprises an act of the primary compute node performing a local computation representing a portion of a distributed computation.

3. The method as recited in claim 1, wherein the act of replicating the primary
15 compute node's performance of the local computation into system memory to the one or more secondary compute nodes comprises an act of replicating the at least one input or output queue related operation.

4. The method as recited in claim 1, further comprising:

an act of primary compute node receiving an indication from a write quorum
20 number of the one or more secondary computer nodes that the message was successfully replicated;

an act of acknowledging receipt of the message to an upstream primary compute instance.

5. The method as recited in claim 1, wherein the act of altering the state of at least
25 one queue to physically implement the at least one input or output queue related operation comprises an act of enqueueing a message in a sink queue.

6. The method as recited in claim 5, further comprising:

an act of marking the messages as committed in the sink queue; and

an act of sending the message to a downstream primary computer node in

30 accordance with a queue to queue protocol.

7. The method as recited in claim 1, wherein the act of altering the state of at least one queue to physically implement the at least one input or output queue related operation comprises an act of altering the state of at least one queue in accordance with a policy that throttles up or down the flow of data between the primary compute node and a compute

node configured to perform a different local computation based on load conditions among the compute nodes.

8. At a compute node within a fabric based distributed computing architecture, the compute node including one or more processors and system memory, the compute node
5 configured with an executable group of a plurality of operations for performing a stateless local computation, the group of operations including at least one input or output queue related operation, a method for performing the stateless local computation, the method including:

an act of the compute node performing the local stateless computation by:

10 an act of the creating a transaction for the local stateless computation;

an act of logically performing one or more operations in the group of operations;

an act of reliably detecting a failure of the local computation prior to completing a specified operation after the one or more operations in the group of
15 operations;

an act of reliably restarting the local computation at the specified operation;

an act of logically performing the specified operation and any remaining operations after the specified operation until the plurality of operations is performed; and

20 an act of ending the transaction in accordance with one of a transaction commit operation and a transaction abort operation;

in response to ending the transaction, an act of the compute node applying changes in accordance with the transaction commit operation or the transaction abort operation, including:

25 when the transaction commit operation ends the transaction, an act of applying physical changes to implement the logical application of each of the plurality of operations in the group of operations, including:

an act of altering the state of at least one queue to physically implement the at least one input or output queue related operation; and

30 when the transaction abort operation ends the transaction, an act of logically reversing the at least one input or output queue related operation.

9. The method as recited in claim 8, wherein the act of the compute node performing the local stateless computation comprises an act of the compute node performing a local stateless computation representing a portion of a distributed computation.

10. The method as recited in claim 8, wherein the act of reliably detecting a failure of the local computation prior to completing a specified operation after the one or more operations in the group of operations comprises an act of a perfect failure detector detecting a failure of the local computation.

5 11. The method as recited in claim 8, wherein the act of altering the state of at least one queue to physically implement the at least one input or output queue related operation comprises:

an act of dequeuing a first message from a source queue;

an act of enqueue a second message in a sink queue;

10 an act of receiving an acknowledgment of the second message from a downstream compute instance; and

an act of sending an acknowledgment of the first message to an upstream compute instance in response to receiving the acknowledgment of the second message from the downstream compute instance.

15 12. The method as recited in claim 8, wherein the act of altering the state of at least one queue to physically implement the at least one input or output queue related operation comprises an act of altering the state of at least one queue in accordance with a policy that throttles up or down the flow of data between the primary compute node and a compute node configured to perform a different local computation based on load conditions among
20 the compute nodes.

13. The method as recited in claim 12, wherein the act of altering the state of at least one queue in accordance with a policy that throttles up or down the flow of data based on load conditions among the compute nodes comprises an act of altering the state of at least one queue based on: one or more of network utilization and available network bandwidth
25 among the compute nodes.

14. The method as recited in claim 12, wherein the act of altering the state of at least one queue in accordance with a policy that throttles up or down the flow of data based on load conditions among the compute nodes comprises an act of altering the state of at least one queue in accordance with a policy that throttles up or down the flow of data between
30 the primary compute node and a compute node configured to perform a different local computation based on load conditions at the primary compute node.

15. A computer program product for use at a replica set within a fabric based distributed computing architecture, the replica set including a primary compute node and one or more secondary compute nodes, the primary compute node and the one or more

secondary compute nodes having the same current computing state, each of the primary compute node and the one or more secondary compute nodes configured with the same executable group of operations for performing a stateful local computation, the group of operations including at least one input or output queue related operation and at least one operation that modifies the computing state, the computer program product for
5 implementing a method for performing the stateful local computation at the replica set, the computer program product comprising one or more computer storage devices have stored thereon computer-executable instructions that, when executed at a processor, cause the primary computer instance to perform the method, including the following:

- 10 perform the local stateful computation by:
- creating a transaction for the local computation;
 - logically performing one or more additional operations in the group of operations, including:
 - 15 processing the current computing state, processing the current computing state including one or more of:
 - logically generating additional computing state and logically modifying existent computing state;
 - logically performing the at least one input or output queue related operation; and
 - 20 ending the transaction in accordance with a transaction commit operation;
 - replicating the primary compute node's performance of each operation in the group of operations on a per operation basis into system memory to the one or more secondary compute nodes, including:
 - replicating creating a transaction for the local computation;
 - 25 replicating logically performing the one or more additional operations in the group of operations; and
 - replicating committing the transaction, wherein replication of committing the transaction signals the one or more secondary compute nodes to applying changes in accordance with the transaction commit operation;
 - 30 in response to committing the transaction, applying physical changes to implement the logical application of each operation in the group of operations, including:
 - reconciling the current computing state with any additional computing state and any modified existent computing state into a new committed computing state;
 - and

altering the state of at least one queue to physically implement the at least one input or output queue related operation.




Figure 1

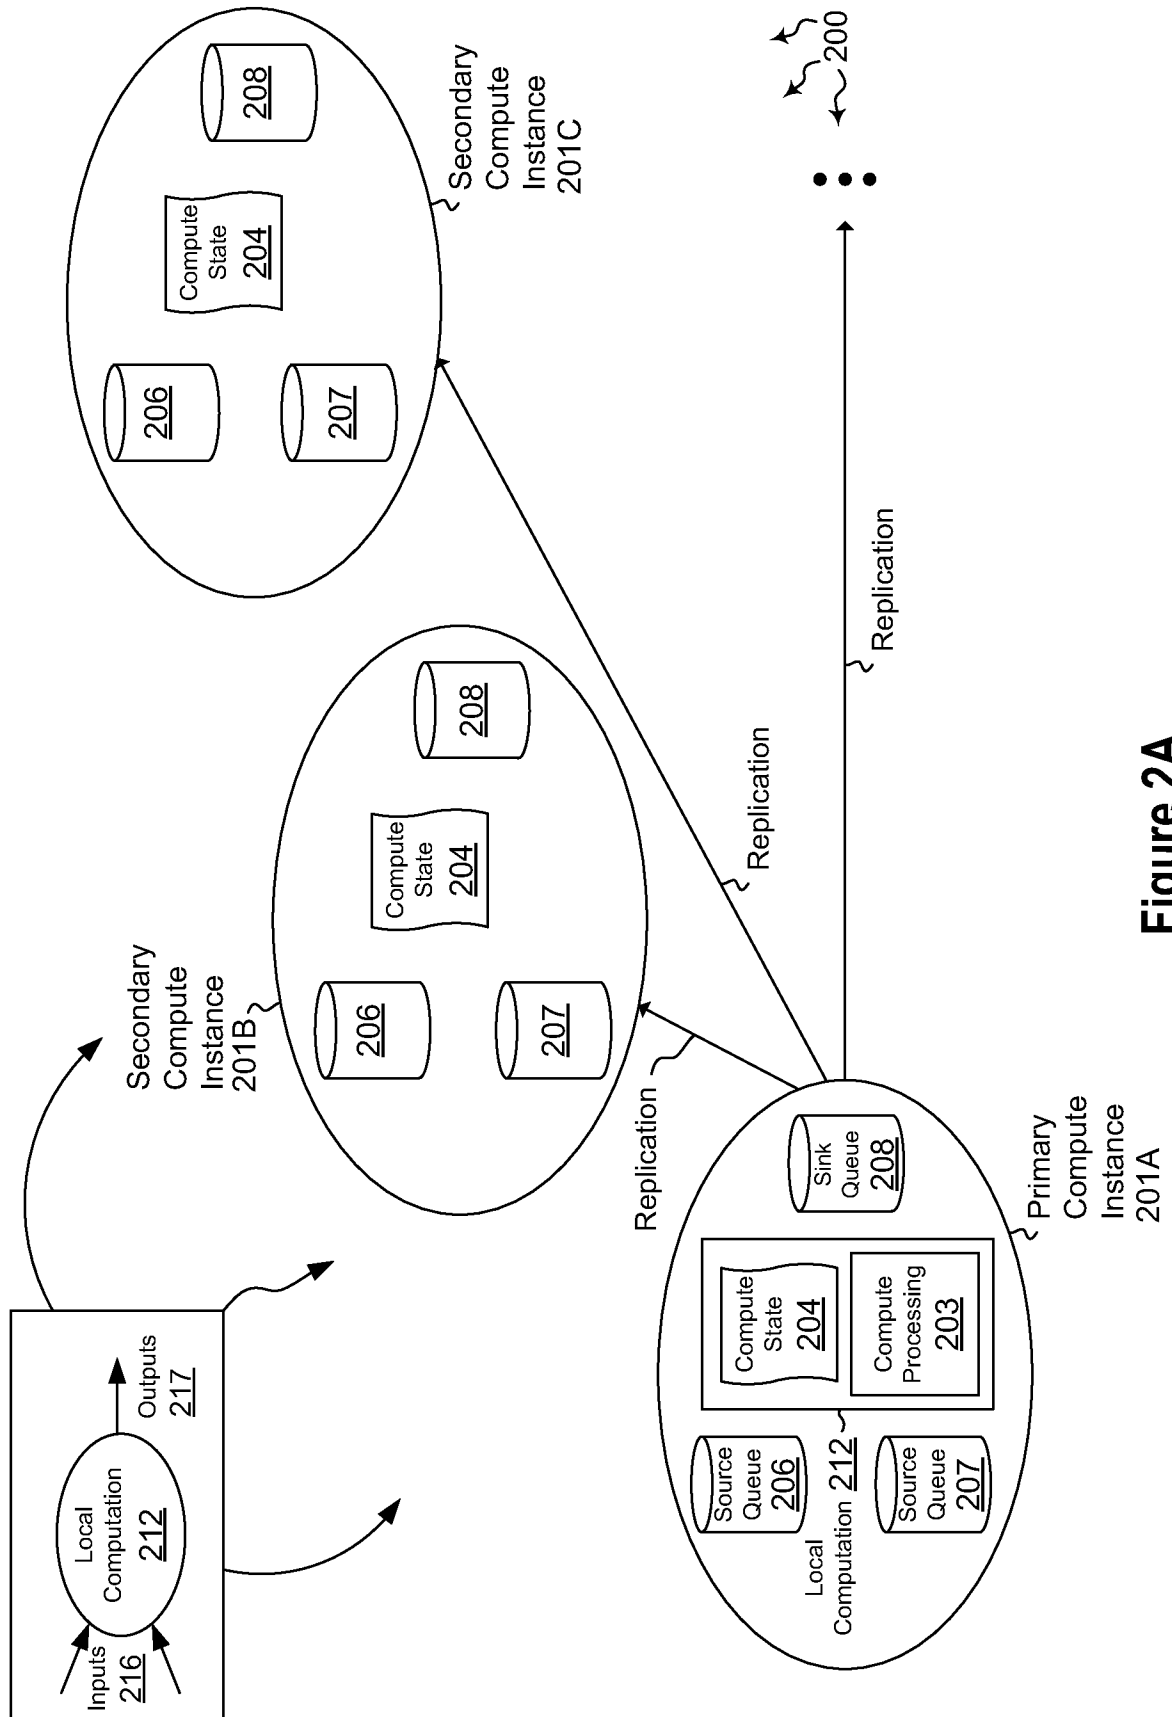


Figure 2A

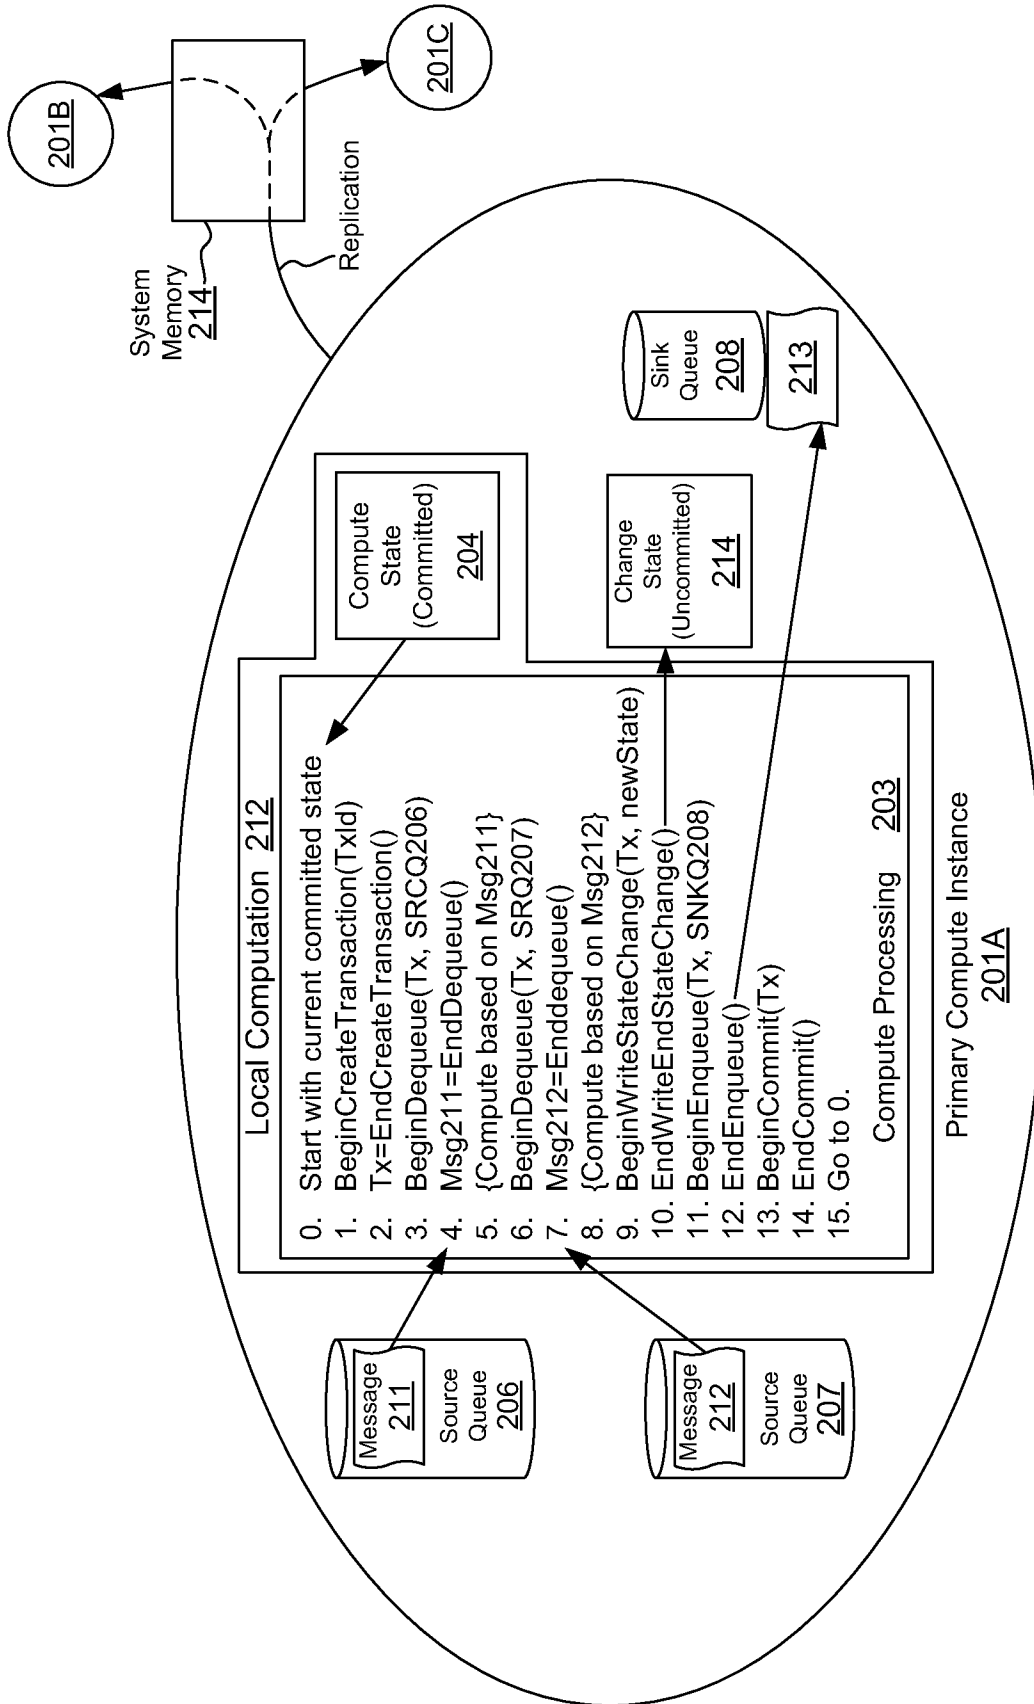


Figure 2B




Figure 3




Figure 4A




Figure 4B




Figure 5A




Figure 5B

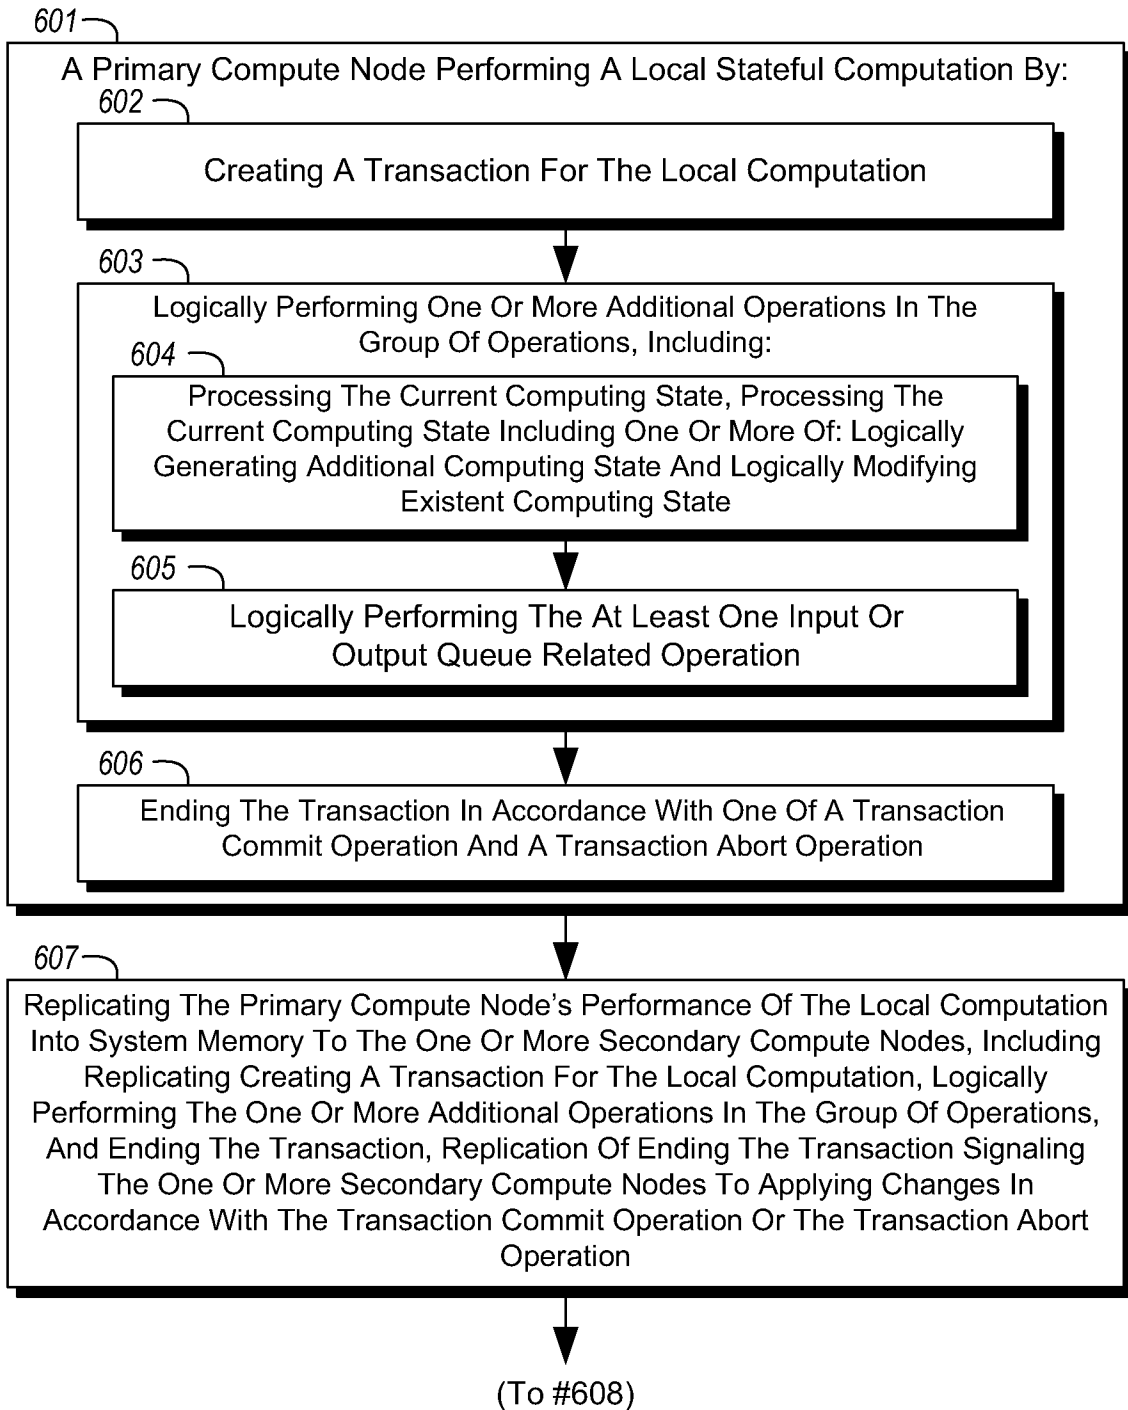

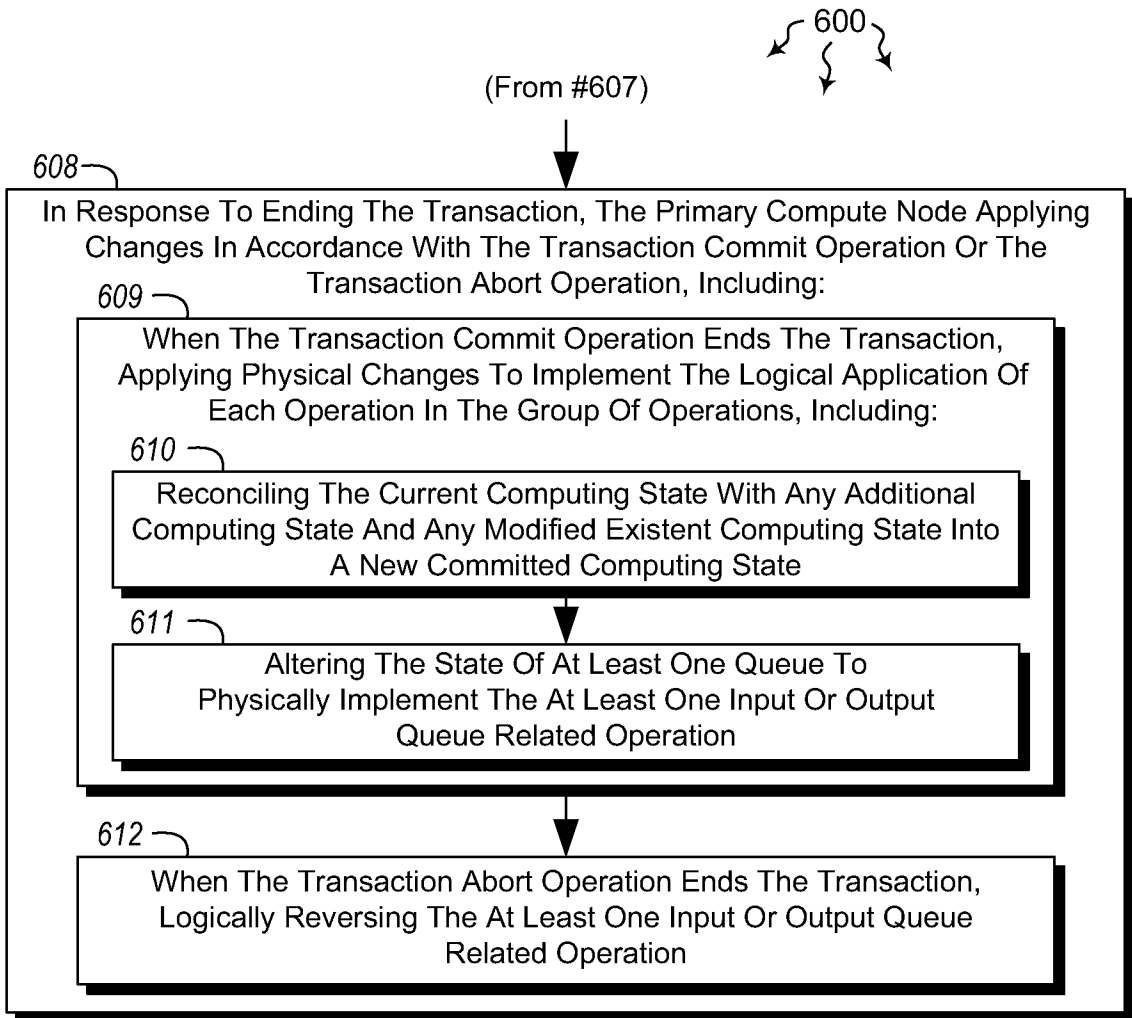


Figure 6



**Figure 6
(Continued)**

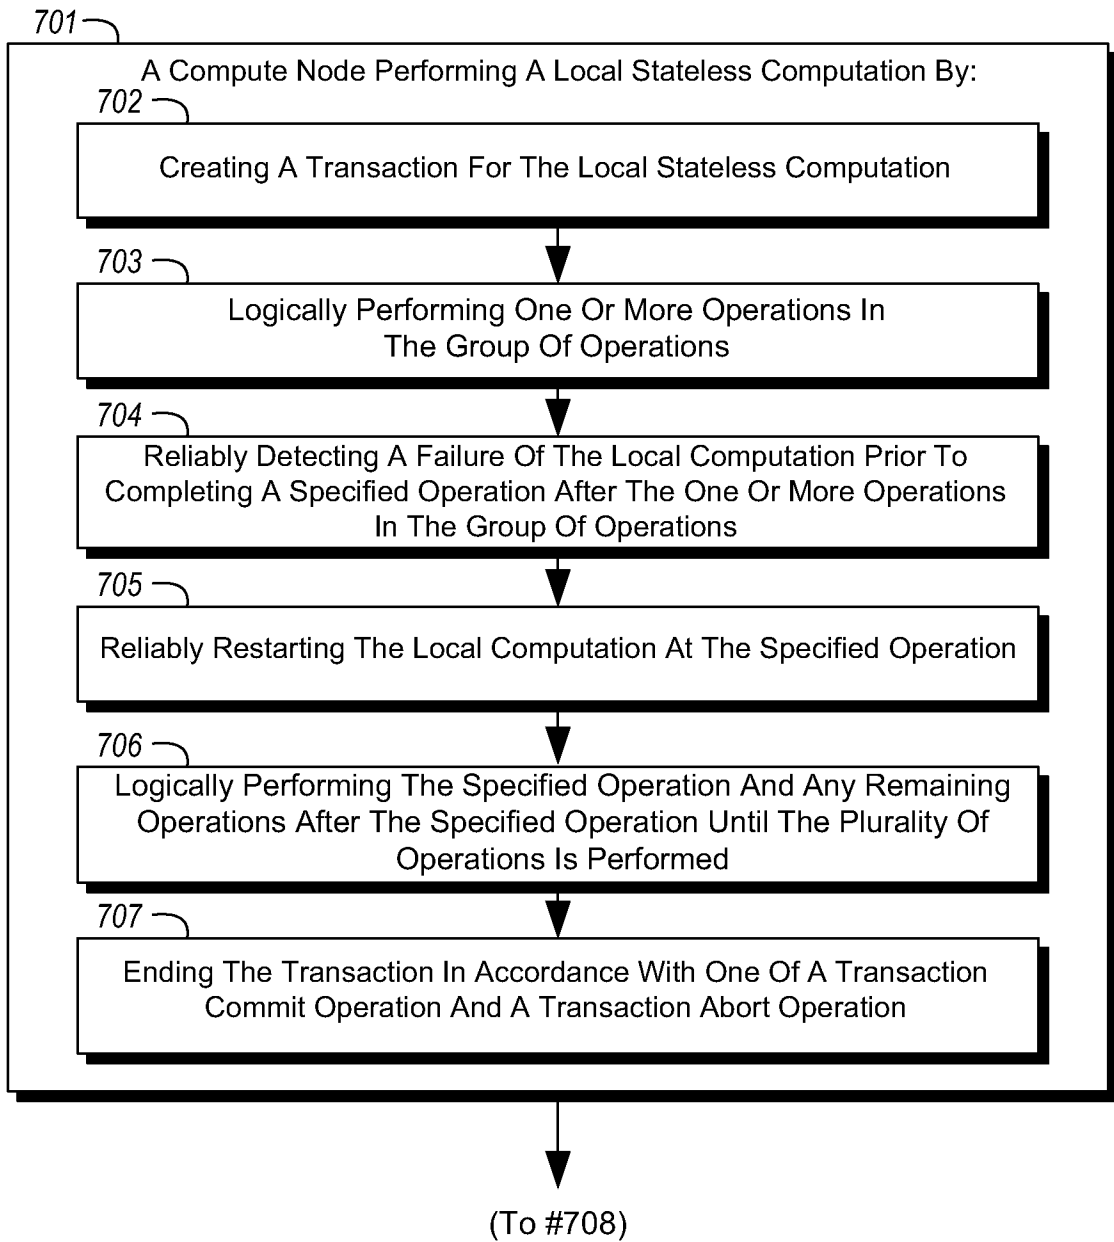



Figure 7

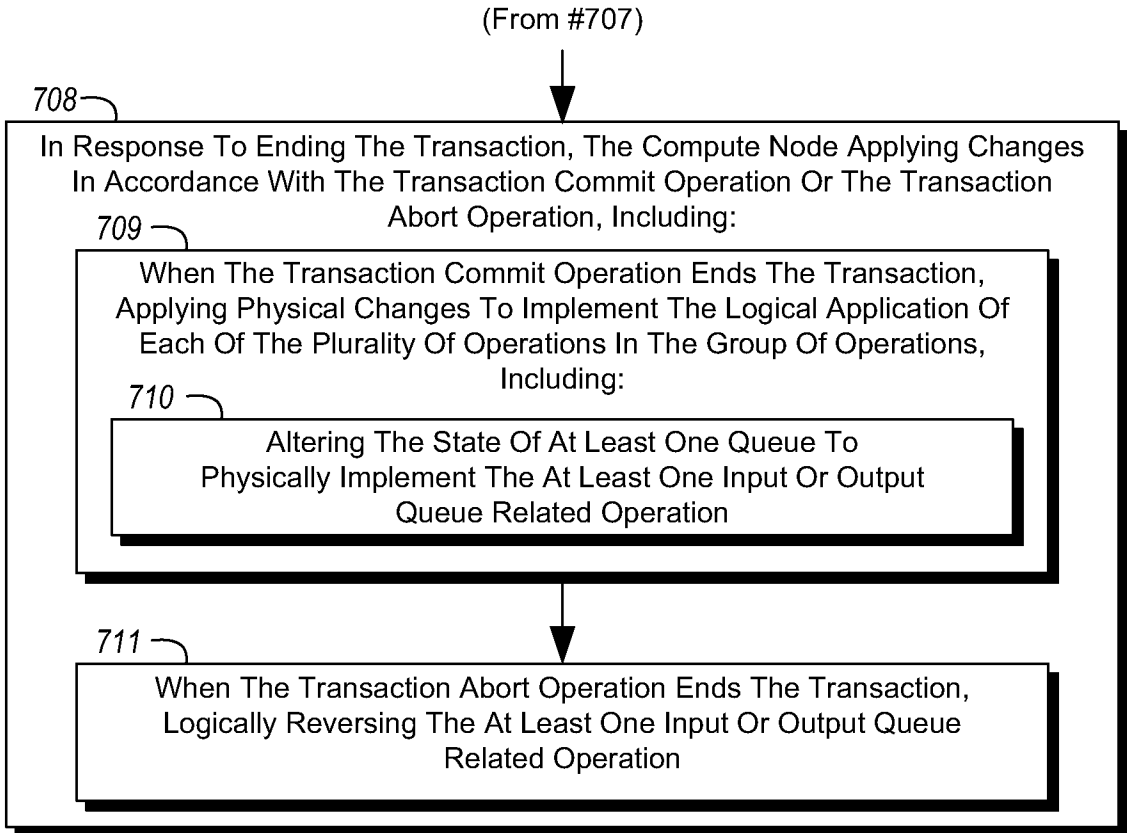



Figure 7
(Continued)

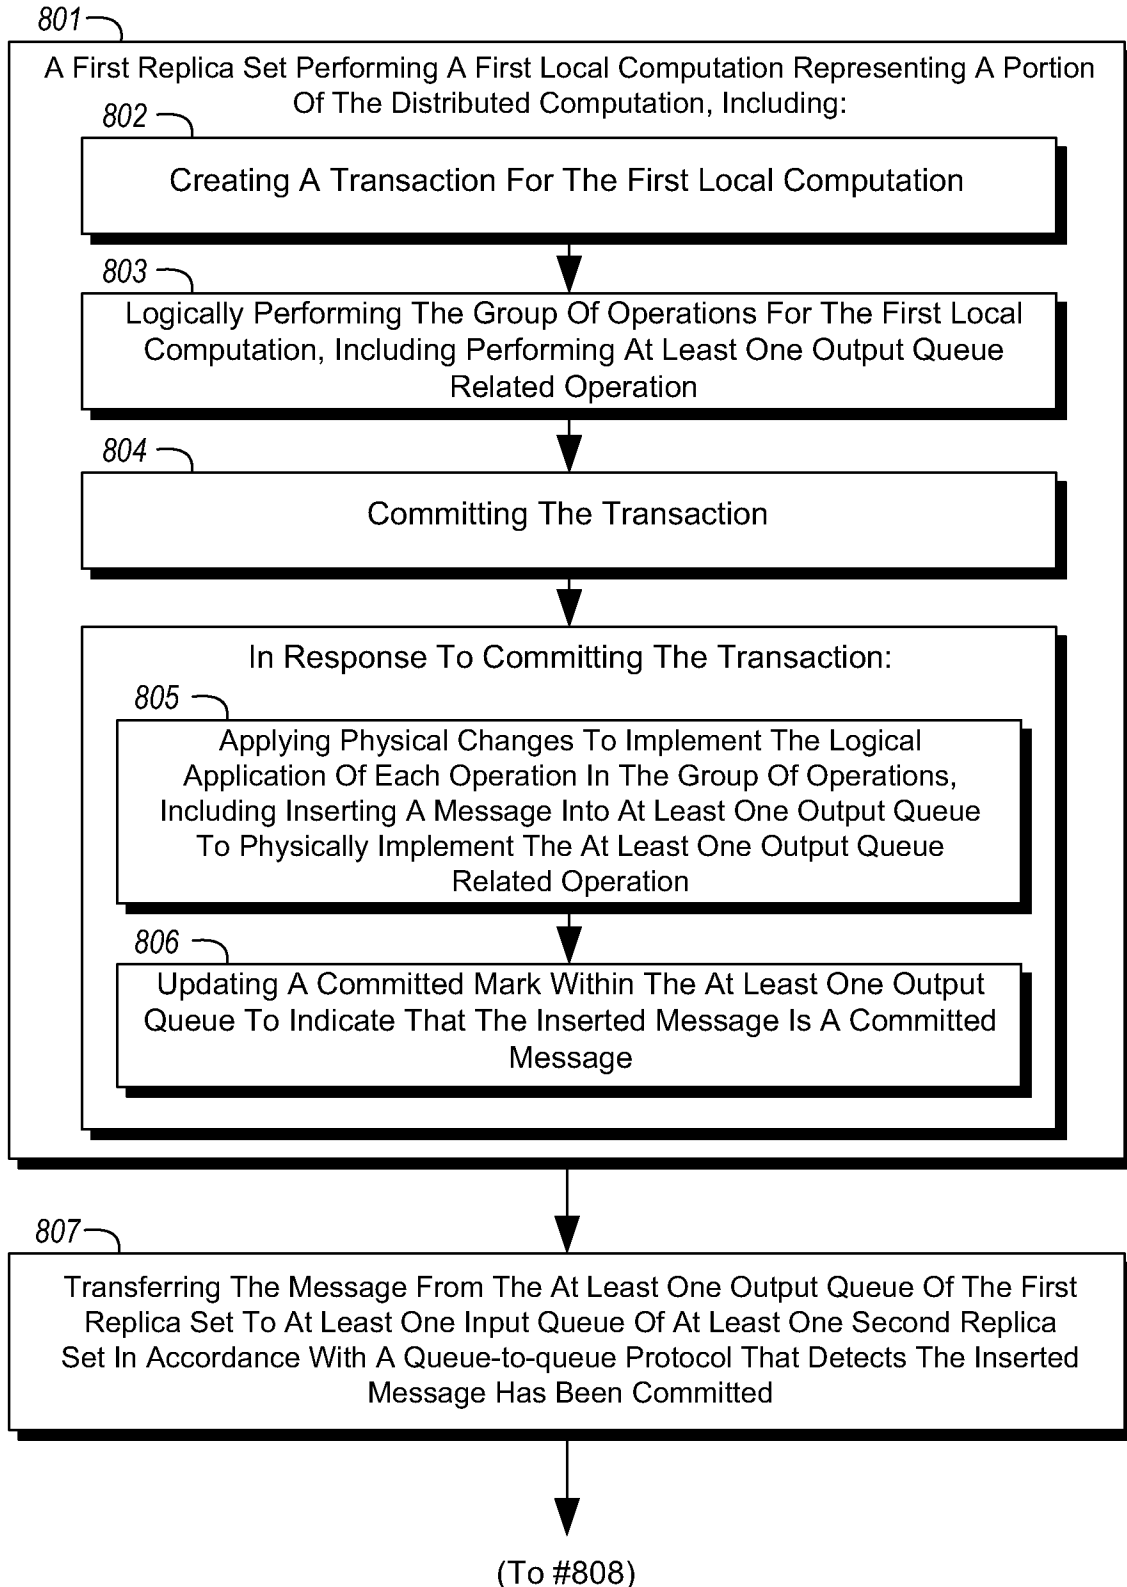



Figure 8

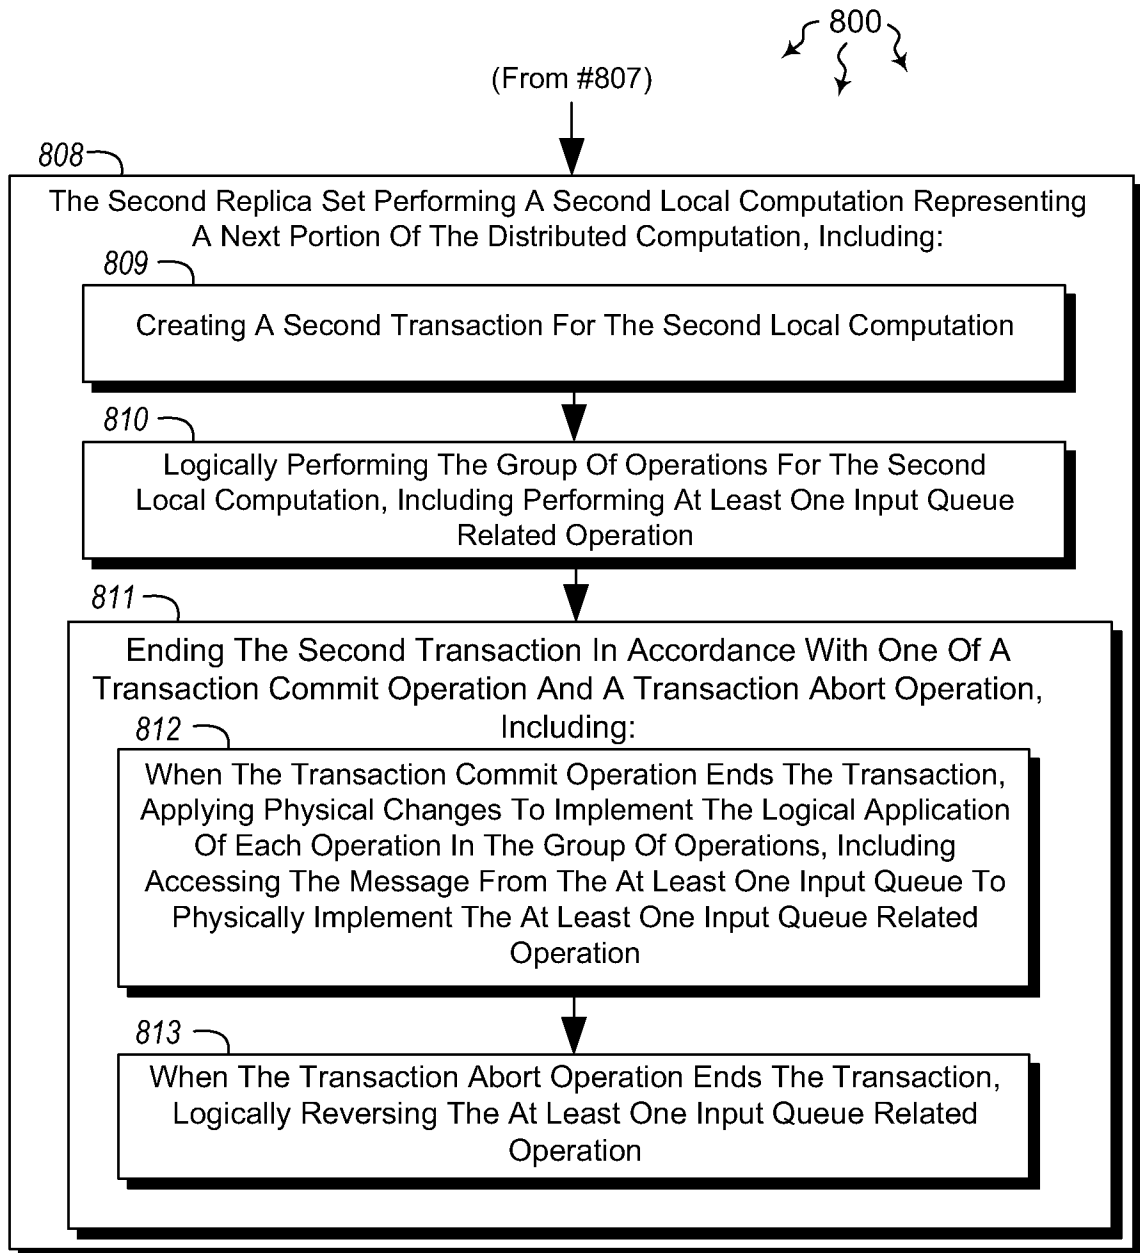
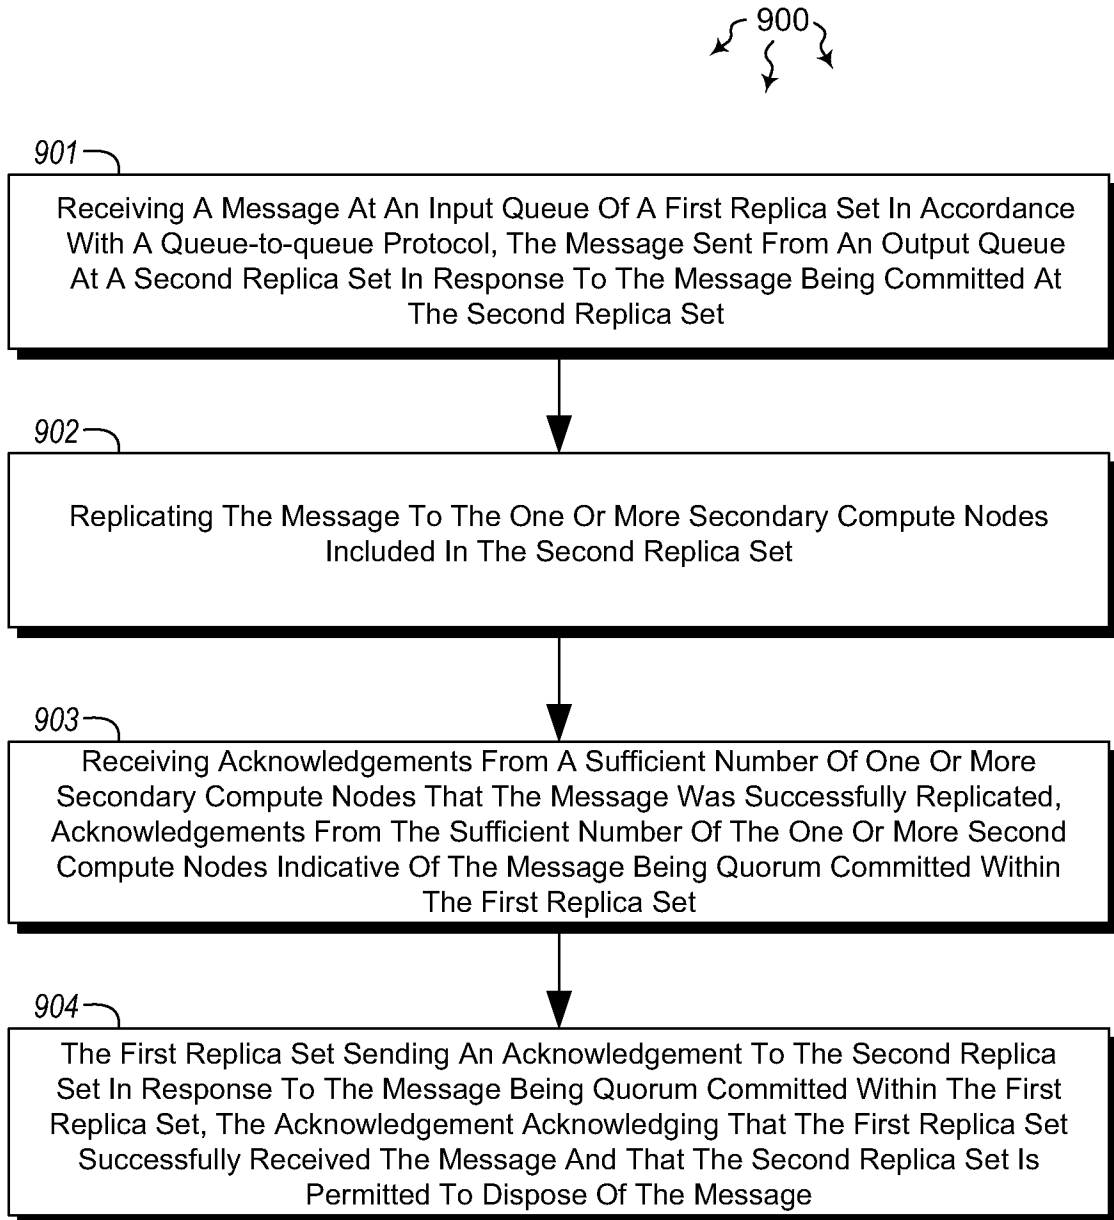


Figure 8
(Continued)

15 / 16

**Figure 9**

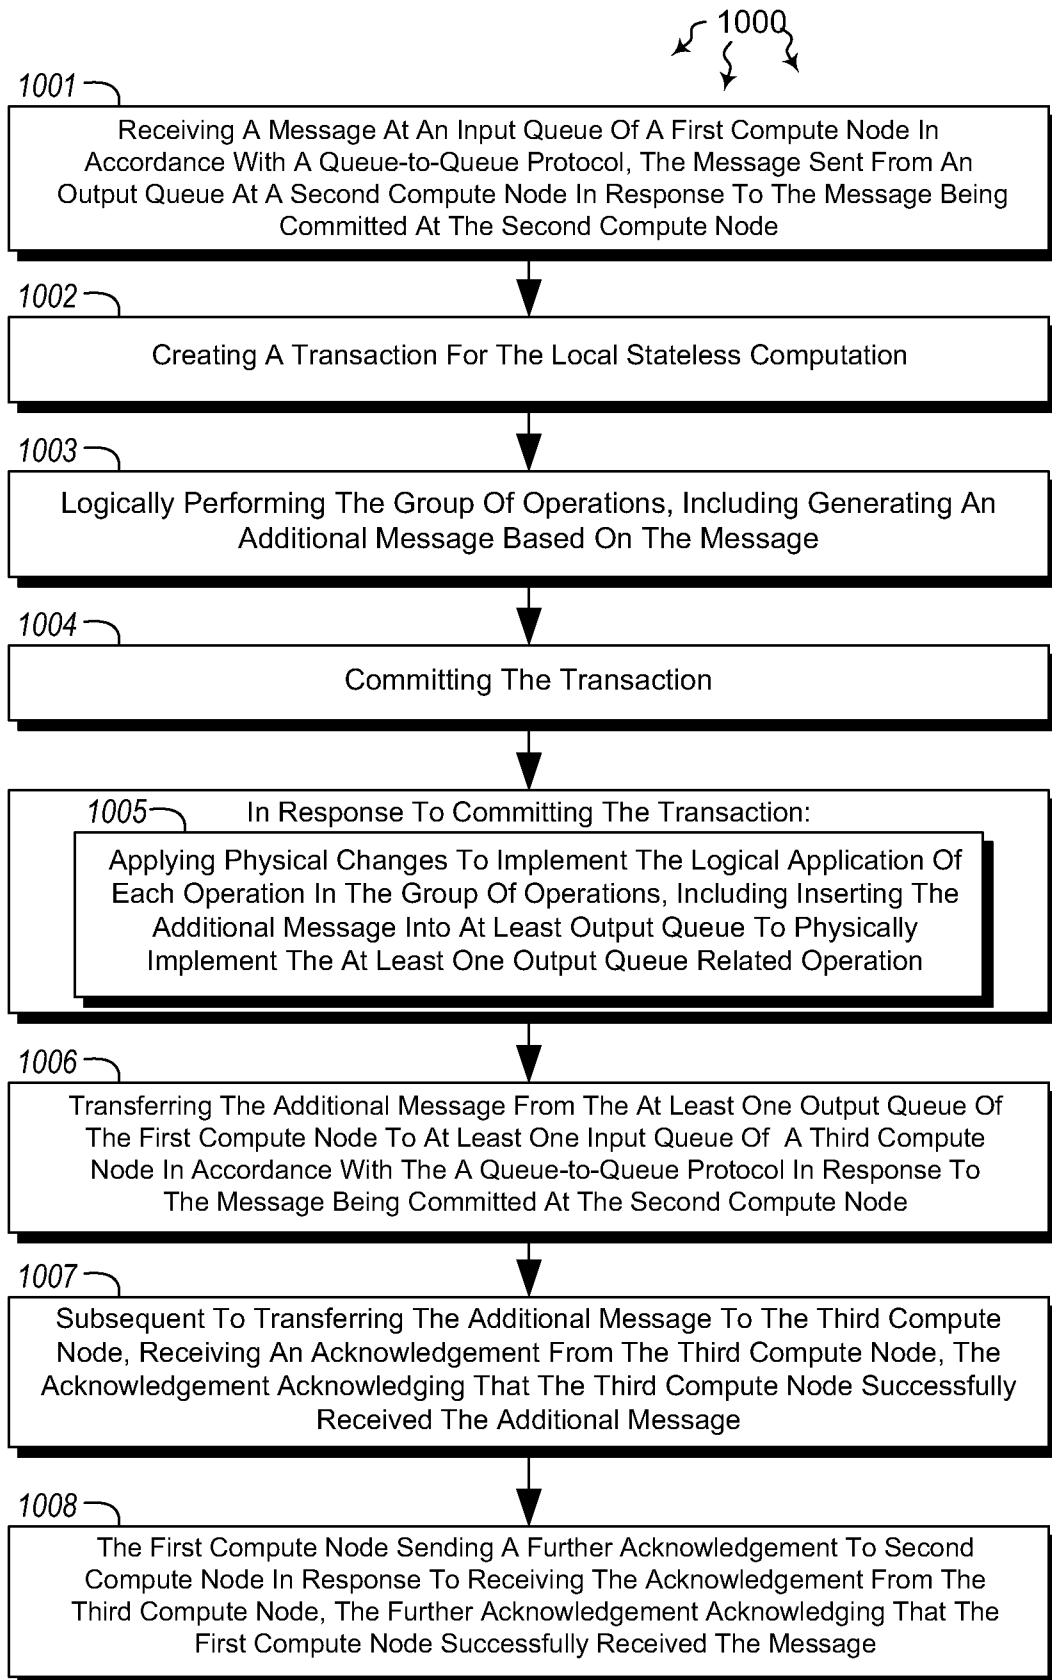


Figure 10