




(86) Date de dépôt PCT/PCT Filing Date: 2015/10/19
 (87) Date publication PCT/PCT Publication Date: 2016/04/28
 (45) Date de délivrance/Issue Date: 2020/01/28
 (85) Entrée phase nationale/National Entry: 2017/04/12
 (86) N° demande PCT/PCT Application No.: US 2015/056159
 (87) N° publication PCT/PCT Publication No.: 2016/064705
 (30) Priorité/Priority: 2014/10/20 (US62/065,941)

(51) Cl.Int./Int.Cl. *G06F 11/07* (2006.01),
G06F 9/52 (2006.01)
 (72) Inventeur/Inventor:
STANFILL, CRAIG W., US
 (73) Propriétaire/Owner:
AB INITIO TECHNOLOGY LLC, US
 (74) Agent: ROBIC

(54) Titre : RECUPERATION ET TOLERANCE AUX PANNES DANS DES CONDITIONS D'INDETERMINISME
COMPUTATIONNEL

(54) Title: RECOVERY AND FAULT-TOLERANCE UNDER COMPUTATIONAL INDETERMINISM



(57) **Abrégé/Abstract:**

A method for promoting fault tolerance and recovery in a computing system including at least one processing node includes promoting availability and recovery of a first processing node, by, at the first processing node, generating first spawn using a spawner that has been assigned a first generation-indicator so that its spawn inherits the first generation indicator, beginning a checkpoint interval to generate nodal recovery information, suspending the spawner from generating spawn, assigning, to the spawner, a second generation-indicator that differs from the first one, resuming the spawner, so that it generates second spawn that inherits the second generation-indicator, controlling an extent to which the second spawn writes to memory, and after committing nodal recovery information acquired during the checkpoint to durable storage, releasing control over the extent to which the second spawn can write to memory.

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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
28 April 2016 (28.04.2016)



(10) International Publication Number
WO 2016/064705 A1

- (51) International Patent Classification:
G06F 11/14 (2006.01) *G06F 9/52* (2006.01)
- (21) International Application Number:
PCT/US2015/056159
- (22) International Filing Date:
19 October 2015 (19.10.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/065,941 20 October 2014 (20.10.2014) US
- (71) Applicant: **AB INITIO TECHNOLOGY LLC** [US/US];
201 Spring Street, Lexington, Massachusetts 02421 (US).
- (72) Inventor: **STANFILL, Craig W.**; 43 Huckleberry Hill
Road, Lincoln, Massachusetts 01773 (US).
- (74) Agents: **MASON, Elliott J., III** et al.; Occhiuti &
Rohlicek LLP, 321 Summer Street, Boston, Massachusetts
02210 (US).
- (81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: RECOVERY AND FAULT-TOLERANCE UNDER COMPUTATIONAL INDETERMINISM




FIG. 1

(57) Abstract: A method for promoting fault tolerance and recovery in a computing system including at least one processing node includes promoting availability and recovery of a first processing node, by, at the first processing node, generating first spawn using a spawner that has been assigned a first generation-indicator so that its spawn inherits the first generation indicator, beginning a checkpoint interval to generate nodal recovery information, suspending the spawner from generating spawn, assigning, to the spawner, a second generation-indicator that differs from the first one, resuming the spawner, so that it generates second spawn that inherits the second generation-indicator, controlling an extent to which the second spawn writes to memory, and after committing nodal recovery information acquired during the checkpoint to durable storage, releasing control over the extent to which the second spawn can write to memory.

WO 2016/064705 A1

1 **RECOVERY AND FAULT-TOLERANCE UNDER**
2 **COMPUTATIONAL INDETERMINISM**

3
4 **BACKGROUND**

5 This description relates to recovery and fault-tolerance in the presence of
6 computational indeterminism.

7 Computational systems occasionally fail for a variety of reasons. When such
8 systems fail, data can be lost. It is desirable to take measures to prevent, or at least
9 minimize, such data loss.

10 Examples of such measures include ACID (Atomic, Consistent, Isolated until
11 committed, Durable when committed) transactions in databases. These known
12 measures are extremely robust. They can be made to meet very high standards of
13 correctness, while also being made fault tolerant.

14 However, all of this robustness comes at a cost. Known methods for guarding
15 against failure have high latency and sometimes cause extended periods during which
16 the apparatus is unavailable. Thus, they are less than optimal for high-volumes of
17 transactions.

18 In addition, some known methods require deterministic computation. In
19 deterministic computation, the order in which tasks are performed is fixed, and the
20 result of a computation remains the same each time it is carried out. It is not clear how
21 these known methods can be adapted to efficiently handle non-deterministic
22 computational environments.

23 Additional complexity arises when a computing apparatus includes multiple
24 processing nodes that cooperate with each other. In such an apparatus, it is possible
25 for one node of the apparatus to fail, and others to keep working. When that failed
26 node recovers, this is no guarantee that it has restored itself to a state that the other
27 nodes expect it to be in.

28 **SUMMARY**

29 In one aspect, the invention features a method for enhancing fault tolerance
30 and recovery in a computing system including at least one processing node, said
31 method including: enhancing availability and recovery of a first processing node,

1 including, at the first processing node, executing a spawner at said node, wherein said
2 spawner, in the course of execution, generates a first spawn, wherein the spawner
3 implements a process, with an associated expected spawner lifetime, configured to
4 generate one or more spawn processes, including the first spawn, that are each
5 associated with a respective expected lifetime shorter than the expected spawner
6 lifetime, and which are configured to be suspended if, following a checkpoint event
7 and suspension of the spawner, the respective one or more spawn processes are still
8 executing after a pre-determined interval measured from occurrence of the checkpoint
9 event and suspension of the spawner, wherein at least one of the one or more spawn
10 processes terminates before the end of the pre-determined interval, wherein executing
11 said spawner includes assigning, to said spawner, a first generation indicator, wherein
12 said first spawn inherits said first generation indicator; beginning a checkpoint
13 interval, at the end of which nodal recovery information, which is usable for recovery
14 of said node, is committed to durable storage, wherein beginning said checkpoint
15 interval includes: suspending said spawner from generating spawn, assigning, to said
16 spawner, a second generation indicator that differs from said first generation
17 indicator, resuming said spawner, thereby enabling said spawner to generate a second
18 spawn, wherein said second spawn inherits said second generation indicator, and
19 controlling an extent to which said second spawn writes to memory; and after
20 committing said nodal recovery information, releasing control over said extent to
21 which said second spawn can write to memory.

22 In some practices, controlling an extent to which the second spawn writes to
23 memory includes preventing the second spawn from consummating a write to the
24 memory. Among these are practices that further include permitting the second spawn
25 to queue the write to memory for eventual consummation thereof after the recovery
26 information has been committed.

27 In other practices, controlling an extent to which the second spawn writes to
28 memory includes determining that the write operation is a commutable operation, and
29 allowing consummation of the commutable operation. Among these practices are
30 those in which determining that the write operation is a commutable operation
31 includes determining that the write operation includes incrementing a variable, and
32 those in which determining that the write operation is a commutable operation

1 includes determining that the write operation includes inserting a record at a specified
2 location.

3 Also among the practices of the invention are those that further include, after
4 suspending the spawner, setting a deadline, thereby providing time for any spawn
5 having the first task generation-indicator to execute to completion, and avoiding
6 overhead associated with having to save states of the spawn having the first
7 generation-indicator. Among these practices are those that include suspending the first
8 spawn if the first spawn is still executing as of the deadline, and those that include
9 enabling the first spawn to avoid suspension as a result of having failed to complete
10 execution by the deadline, for example by changing the first task generation-indicator
11 to the second task generation-indicator in the first spawn if the first spawn is still
12 executing as of the deadline.

13 In those cases in which the first node has a nodal-generation indicator,
14 additional practices of the invention include causing a spawn to become a migrant that
15 migrates to a second node, wherein the second node has a nodal-generation indicator.
16 Among these practices are those in which the nodal-generation counter of the second
17 node indicates that the second node is in a younger generation than the first node, in
18 which case the method further includes youthening the migrant, either by
19 immigration-side youthening of the migrant, or by emigration-side youthening of the
20 migrant.

21 In some cases, the first node is a node in a multi-node system in which each
22 node has a nodal generation-count, and the multi-node system includes at least a
23 second node. In these cases, there are practices of the invention in which, upon
24 recovery following a failure of the second node, the first node rolls back to a state that
25 corresponds to a nodal-generation count of the second node.

26 In other cases, the first node is a node in a multi-node system in which each
27 node has a nodal generation-count, and the multi-node system includes at least a
28 second node. In these cases, some practices of the invention include, upon recovery
29 following a failure of the first node, having the first node roll forward to a state that
30 corresponds to a nodal-generation count of the second node by restoring committed
31 work from a checkpoint and restoring uncommitted work from a journal.

1 In those cases in which the first node is a node in a multi-node system in
2 which each node has a nodal generation-count, practices of the invention include
3 those in which the first node carries out certain acts. These include receiving, from a
4 master node, a message indicating that a checkpoint is to be carried out, in response,
5 youthening a nodal generation count of the first node, suspending spawners from
6 generating spawn, saving spawner recovery information for recovering spawner
7 states, resuming the spawners, determining that no further older-generation
8 immigrants are expected at the first node, and in response to the determination,
9 committing, to the durable storage, the nodal recovery information. Among these
10 practices are those that also include setting a deadline, and, upon lapse of the
11 deadline, suspending all older-generation spawn that are still executing while
12 younger-generation spawn continue to execute.

13 In some cases, the first node is a node in a multi-node system. In such cases,
14 alternative practices of the invention include saving a replica copy of working
15 memory of the first node at the second node, upon failure of the first node,
16 temporarily using the replica copy for processing that would otherwise have been
17 carried out by the first node, and, upon recovery of the first node, communicating, to
18 the first node, information required to update memory in the first node so that
19 subsequent computation can be carried out by the first node.

20 In another aspect, the invention features a computer program product
21 stored in a non-transitory form on a computer-readable medium, for enhancing
22 fault tolerance and recovery in a computing system including at least one
23 processing node, the computer program product including instructions for
24 causing a computing system to: enhance availability and recovery of a first
25 processing node, including, at a first processing node, executing a spawner at
26 said node, wherein said spawner, in the course of execution, generates a first
27 spawn, wherein the spawner implements a process, with an associated
28 expected spawner lifetime, configured to generate one or more spawn
29 processes, including the first spawn, that are each associated with a respective
30 expected lifetime shorter than the expected spawner lifetime, and which are
31 configured to be suspended if, following a checkpoint event and suspension of
32 the spawner, the respective one or more spawn processes are still executing

1 after a pre-determined interval measured from occurrence of the checkpoint
2 event and suspension of the spawner, wherein at least one of the one or more
3 spawn processes terminates before the end of the pre-determined interval,
4 wherein executing said spawner includes assigning, to said spawner, a first
5 generation indicator, wherein said first spawn inherits said first generation
6 indicator; beginning a checkpoint interval, at the end of which nodal recovery
7 information, which is usable for recovery of said node, is committed to
8 durable storage, wherein beginning said checkpoint interval includes:
9 suspending said spawner from generating spawn, assigning, to said spawner, a
10 second generation indicator that differs from said first generation indicator,
11 resuming said spawner, thereby enabling said spawner to generate a second
12 spawn, wherein said second spawn inherits said second generation indicator,
13 and controlling an extent to which said second spawn writes to memory; and
14 after committing said nodal recovery information, releasing control over said
15 extent to which said second spawn can write to memory.

16 In yet another aspect, the invention features a data storage system
17 including durable storage; and one or more processing nodes including least
18 one processor configured to enhance availability and recovery of a first
19 processing node, including, at the first processing node, executing a spawner
20 at said node, wherein said spawner, in the course of execution, generates a first
21 spawn, wherein the spawner implements a process, with an associated
22 expected spawner lifetime, configured to generate one or more spawn
23 processes, including the first spawn, that are each associated with a respective
24 expected lifetime shorter than the expected spawner lifetime, and which are
25 configured to be suspended if, following a checkpoint event and suspension of
26 the spawner, the respective one or more spawn processes are still executing
27 after a pre-determined interval measured from occurrence of the checkpoint
28 event and suspension of the spawner, wherein at least one of the one or more
29 spawn processes terminates before the end of the pre-determined interval,
30 wherein executing said spawner includes assigning, to said spawner, a first generation
31 indicator, wherein said first spawn inherits said first generation indicator; beginning a

1 checkpoint interval, at the end of which nodal recovery information, which is usable
2 for recovery of said node, is committed to durable storage, wherein beginning said
3 checkpoint interval includes suspending said spawner from generating spawn,
4 assigning, to said spawner, a second generation indicator that differs from said first
5 generation indicator, resuming said spawner, thereby enabling said spawner to
6 generate a second spawn, wherein said second spawn inherits said second generation
7 indicator, and controlling an extent to which said second spawn writes to memory;
8 and after committing said nodal recovery information, releasing control over said
9 extent to which said second spawn can write to memory.

10 Aspects can have one or more of the following advantages.

11 The techniques for promoting fault tolerance and recovery described herein
12 enable the computing system to remain highly available. By strategically relaxing
13 certain ACID constraints, the computing system can still provide recoverability, but
14 without the high overhead of more extreme measures. Thus, there are more
15 computing resources available for useful work. Also, by controlling the extent to
16 which spawn (e.g., operating system processes or threads) write to memory, useful
17 work can still be accomplished during a checkpoint interval, while ensuring that
18 integrity of the fault tolerance mechanisms is maintained. Thus, these techniques
19 enhance the internal functioning of the computing system, both in the event of faults,
20 and during normal fault-free operation.

1 **BRIEF DESCRIPTION OF THE DRAWINGS**

2 FIG. 1 shows a single-node computing apparatus for carrying our fault-tolerance and
3 recovery in the presence of computational indeterminism;

4 FIG. 2 shows checkpoint intervals and working intervals encountered during
5 operation of the node of FIG. 1;

6 FIG. 3 shows a method in which spawn are allowed to run to completion during the
7 checkpoint interval of FIG. 2;

8 FIG. 4 shows a method in which a spawning process can continue to generate spawn
9 during the checkpoint interval of FIG. 2;

10 FIG. 5 shows a method in which spawn can continue to work during a checkpoint
11 interval by queuing writes to memory;

12 FIG. 6 shows a multi-node computing apparatus;

13 FIG. 7 shows two generations co-existing in a node from the apparatus of FIG. 7;

14 FIG. 8 shows method steps carried out by a slave node in response to a checkpoint
15 message;

16 FIG. 9 shows method steps involving incrementing a migrant generation count;

17 FIG. 10 shows method steps for recovery after failure;

18 FIG. 11 shows a replica maintained at another node to enable more rapid recovery;

19 FIG. 12 shows a procedure for using the replica shown in FIG. 11 for rapid recovery;
20 and

21 FIG. 13 shows an example of execution of the method shown in FIG. 5 in connection
22 with multiple nodes as illustrated in FIG. 9.

23 **DETAILED DESCRIPTION**

24 FIG. 1 shows an example of a data processing system in which techniques for
25 fault tolerance and recovery in the presence of computational indeterminism can be
26 used. The data processing system includes a single-node computing apparatus 8

1 having a node **10** that includes a working-memory **12**. Processes **14** running on that
2 node **10** use this working-memory **12** to save their respective process states and to
3 store intermediate results of their respective computations. In different embodiments,
4 the processes **14** may be implemented as any of a variety of types of computing
5 resources within an operating system running on the node **10**. For example, the
6 processes **14** may be implemented as operating system ‘processes’ that have their own
7 address space, or as operating system ‘threads’ that have their own execution context
8 (e.g., stage, registers, etc.), or as some other type of ‘task’ that includes a sequence of
9 operations to be performed but does not necessarily have a particular operating system
10 process or thread dedicated to it.

11 Because the working-memory **12** is often volatile, it is prudent to periodically
12 save its state in checkpoint files **18** stored in durable storage **16**. These checkpoint
13 files **18** can be used to recover state information in case of an interruption in operation
14 of the node **10**.

15 Among the processes **14** running on the node are spawners **20**. A “spawner” is
16 a type of process that has, among its properties, the ability to generate one or more
17 other processes. The processes that are generated by a spawner are referred to herein,
18 both in the singular and in the plural, as “spawn.” The act of generating such spawn is
19 referred to by appropriate variants of the verb “to spawn.” FIG. 1 shows a spawner **20**
20 that has generated spawn **22A**, **22B**. A spawner **20** is generally a long-lived process,
21 whereas the spawn **22A**, **22B**, although numerous, tend to be much shorter-lived. In
22 some instances, a spawner is a process that lives longer than spawn generated by that
23 spawner. The spawn **22A**, **22B** are also independent of and asynchronous relative to
24 each other so that the extent to which a spawn **22A**, **22B** has completed its
25 computation is unrelated to when the spawner **20** generated that spawn **22A**, **22B** in
26 the first place. As a result, the order in which spawn **22A**, **22B** carry out computations
27 is indeterminate. Since the order in which computations are carried out can often
28 affect the results, this makes the computation as a whole indeterminate.

29 During the course of its operation, the computing apparatus **8** communicates
30 with the outside world. For example, the computing apparatus **8** may receive one or
31 more streams of incoming messages **24** and produce one or more streams of outgoing
32 messages **28**. As will be described in greater detail below, these messages **24**, **28** are

1 temporarily stored by the apparatus **8** within the durable storage **16**. These messages
2 **24, 28** may be temporarily stored for escrow within respective areas that are
3 physically and/or logically separate. Incoming messages **24** can be stored in an
4 incoming-message escrow area **26** in durable storage **16**, and outgoing messages **28**
5 can be stored in an outgoing-message escrow area **30** in durable storage **16**.

6 Referring to FIG. 2, the node's operation is marked by checkpoint intervals **32**
7 and working intervals **34**. During a working interval **34**, the node performs useful
8 work as processes advance towards completion. During the checkpoint interval **32**,
9 the node **10** suspends all processes, serializes them, and saves the result to durable
10 storage **16**. The node **10** then saves anything else that is in working-memory into the
11 durable storage **16**. At this point, the checkpoint is said to have been "committed" and
12 the processes **14** are said to have been "checkpointed."

13 It is preferable that the checkpoint interval **32** be much shorter than the
14 working interval **34**. The apparatus and methods described herein are intended to
15 increase the amount of processing that can be performed in the working intervals **34**
16 by reducing the length of the checkpoint interval **32**.

17 Once a checkpoint is committed, the node **10** allows processes **14** to resume
18 and releases outgoing messages **28** that are in the outgoing-message escrow area **30**.

19 The storing of outgoing messages **28** in an outgoing-message escrow area **30**,
20 rather than sending them immediately, is useful to guard against inconsistency that
21 may result from a failure of the node **10**. For example, it would be quite possible for a
22 process **14** to send an outgoing message **28** announcing the result of some
23 computation. If the node **10** were to fail after this message **28** has already been sent
24 but before the computational result is committed to durable storage **16**, the node **10**
25 would re-start and re-execute the uncommitted computation. Upon completion,
26 another message **28** would be sent announcing the result of this second computation.
27 If the two results are different, which is not unlikely in the case of non-deterministic
28 computations, one of the messages will be invalid.

29 As a concrete example, consider the case in which a process **14** awards a prize
30 to a customer based on the result of a random number generator. Without the
31 outgoing-message escrow area **30**, the process **14** would send a message **28** to a first

1 customer announcing that a prize was forthcoming. The node **10** would then crash and
2 re-start. Since the state of the process **14** was never saved, there is no record of that
3 customer having been awarded any prize, or of the process **14** having successfully
4 completed execution. The node **10** may then re-execute the process **14**, which would
5 then generate a different random number, thus causing a second message **28** to be sent
6 to another customer announcing that a prize was forthcoming. This would require
7 either awarding two prizes where only one was intended, or managing at least one
8 disappointed customer.

9 To recover after a failure, the node **10** retrieves, from durable storage **16**, the
10 state of all processes **14** and the state of working-memory **12**. It then retransmits any
11 outgoing messages **28** that are stored in the outgoing-message escrow area **30**,
12 retrieves incoming messages from incoming-message escrow area **16** for processing,
13 and then resumes normal operation on live data.

14 The retransmission of all messages **28** in the outgoing-message escrow area **30**
15 can result in recipients receiving duplicate messages. In one embodiment, the
16 recipient is configured to ignore repeated messages **28**. In another embodiment, upon
17 recovery, the recipient and the recovered node **10** communicate to identify messages
18 **28** that have been received. This permits the node **10** to avoid sending duplicate
19 messages **28**.

20 The procedure described above introduces results in a lengthy checkpoint
21 interval **32**, which may, in turn, result in either low throughput or an increased latency
22 between the receipt of an incoming message and the production of a corresponding
23 outgoing message. A variety of methods can be used to reduce this latency, and/or
24 increase the throughput.

25 A first optimization method features the maintenance of a journal **36** in the
26 background. Every time an item in the working-memory **12** is altered a corresponding
27 entry will be written to the journal **36** such that the journal entry may be used to
28 reproduce the alteration at recovery time. The journal **36** may be asynchronously
29 transferred to durable storage **16**. During the checkpoint interval **32** the node **10** will
30 ensure that all journal entries have indeed been made durable. Recovery can then be

1 achieved by using an older checkpoint file **18** containing a snapshot of working-
2 memory **12** and applying the changes as set forth in the journal **36**.

3 This optimization reduces the length of the checkpoint interval **32** but at the
4 cost of increasing the time to recover. In particular, the longer it has been since the
5 last full image of working-memory **12** was obtained, the greater the number of entries
6 there will be in the journal **36**. This will increase recovery time.

7 A second optimization method relies on the fact that the node **10** has only a
8 few spawning processes **20**, each of which generates (or “spawns”) multitudes of
9 short-lived processes, referred to herein in both the singular and plural as the “spawn
10 **22A, 22B**.”

11 The lifetime of spawn **22A, 22B** is random, but has an expected value that is
12 much shorter than the expected value of the lifetime of a spawner **20**. As such, it
13 makes little sense to spend time serializing spawn **22A, 22B** during a checkpoint. In
14 fact, in some cases, the time required to serialize spawn **22A, 22B** is an appreciable
15 fraction the spawn’s expected lifetime. It is therefore often advantageous to instead
16 suspend the spawner **20**, thus preventing generating of new spawn and to then allow
17 the existing spawn **22A, 22B** to terminate naturally.

18 To save time, the checkpoint-generating method, as shown in FIG. 3, includes
19 suspending a spawner **20** (step **38**) and serializing it (step **40**). However, existing
20 spawn **22A** continue executing (step **42**). Then, after a straggler’s deadline, which is
21 selected to define a sufficiently long idle-down interval be long enough to allow most
22 spawn **22A** to finish executing but not so long as to appreciably impact latency (step
23 **46**), spawn that are still executing, referred to as “straggling spawn,” are suspended
24 (step **48**), and serialized (step **50**).

25 The foregoing method thus reduces the length of the checkpoint interval **32** by
26 reducing the number of processes **14** that need to be serialized. It does so by allowing
27 processes **22** that are expected to terminate quickly to do so, thus eliminating the need
28 to serialize those processes **22** as part of creating a valid checkpoint.

29 A third optimization procedure arises from the recognition that the evil to be
30 avoided is actually a change to the working-memory **12** during the checkpoint interval

1 **32**. Therefore, if a process **14** does not actually have to write to working-memory **12**,
2 but instead only has to read working-memory **12**, it makes no sense to suspend it.

3 To implement this third optimization procedure, the node **10** relies on a
4 generation count associated with each process. Since the spawner **20** is a process, it
5 has a generation count. Since the spawn **22A**, **22B** of a spawner is also a process, it
6 too has a generation count. The generation counts of a spawn **22A** and a spawner **20**
7 that spawned that spawn **22A** are related. In particular, the generation count of a
8 spawn **22A**, **22B** is equal to the generation count of the spawner **20** that spawned it.
9 The act of causing the generate count of a spawn **22A**, **22B** to be related to, or
10 derivable from, the generation count of the spawner that spawned it is described by
11 the appropriate form of the verb “to inherit.” When a spawner **20** with a particular
12 generation count generates spawn **22A**, **22B**, the spawn **22A**, **22B** is said to have
13 inherited the spawner’s generation count.

14 In operation, prior to the onset of a checkpoint interval **32**, a spawner **20** will
15 have generated older-generation spawn **22A**. At the beginning of the checkpoint
16 interval **32**, the spawner **20** is “youthened.”

17 The verb “to youthen” and its variants and cognates describes a particular
18 computational operation that can be carried out on an integer. As used herein, the
19 particular integer upon which the youthening operation operates is the generation
20 count. describes an operation that can be carried out on a generation count.

21 In the particular example described herein, the act of youthening a spawner **20**
22 means the act of incrementing its generation count. After having been youthened, the
23 spawner **20** then continues to generate spawn during the checkpoint interval **32**, only
24 now, it generates younger-generation spawn **22B**. The result of this is that two kinds
25 of spawn **22** coexist within the node **10**: older-generation spawn **22A**, which the
26 spawner **20** generated before having been youthened, and a younger-generation spawn
27 **22B**, which the spawned **20** generated after having been youthened.

28 Referring to FIG. 4, at the beginning of a checkpoint interval, the spawner **20**
29 is suspended (step **52**), such that it does not generate any new spawn while being
30 suspended, and its process state saved (step **54**). The spawner **20** then has its
31 generation count incremented (step **56**), after which the spawner **20** resumes (i.e., is

1 un-suspended). After resuming, the spawner **20** is once again able to generate spawn
2 **22**, though this time, all its spawn **22B** will be in the younger-generation spawn.

3 Any younger-generation spawn **22B** that attempt to write to the working-
4 memory **12**, are blocked until the checkpoint interval **32** is completed. Thus, younger-
5 generation spawn **22** cannot run to completion. They can only run until it is time to
6 actually write to working-memory **12** for the first time. Nevertheless, younger-
7 generation spawn **22** can at least run partway to completion. This allows some
8 processing to occur even during a checkpoint interval **32**.

9 In general, during the checkpoint interval **32**, all processes **14** in memory **12**
10 will be serialized. However, in the optimization method of FIG. 4, it is desirable to
11 serialize only older-generation spawn **22A**.

12 The generation count enables the node **10** to identify which spawn is younger-
13 generation spawn **22B** and to therefore avoid saving their state.

14 Unfortunately, having to wait for older-generation spawn **22A** to complete
15 processing tended to increase latency because younger-generation spawn **22B** could
16 not proceed full bore until either all older-generation spawn **22A** were done or until
17 the straggler deadline triggered suspension of stragglers from the older-generation
18 spawn **22A**.

19 In a variant of the second optimization method, instead of blocking younger-
20 generation spawn **22B** that attempt to modify working-memory **12**, and thus losing
21 the opportunity to continue doing useful work, the node **10** can tag each data item in
22 working-memory **12** with a generation number. If a younger-generation spawn **22B**
23 modifies a memory location, rather than blocking until after the checkpoint, the node
24 **10** will youthen the memory location by updating its generation number. Then, if an
25 old-generation spawn **22A** attempts to read or write such a youthened memory
26 location the older-generation spawn **22A** will spontaneously youthen itself by
27 suspending itself, writing its state to the checkpoint, updating its generation number,
28 and resuming execution in a youthened state. The node **10** also tags the entries written
29 to the journal with the generation number so that it can distinguish journal entries
30 corresponding to the older generation of spawn from journal entries corresponding to
31 the younger generation of spawn.

1 A fourth optimization method relies on the idea of allowing younger-
2 generation spawn **22B** to continue processing even past the first attempted write to
3 working-memory **12**. This method relies on the fact that sometimes, the order in
4 which computational results are written into working-memory **12** does not matter. If
5 this is the case, writes to working-memory **12** can simply be queued until later. This
6 method allows a younger-generation spawn **22B** to keep working during a checkpoint
7 interval **32** even after the first time attempts to write to working-memory **12**.

8 In general, whenever one carries out a sequence of operations, a question that
9 arises is whether or not the order of operations in the sequence makes a difference in
10 the result of the sequence. An operation within this sequence is said to be
11 “commutable” if the location of that operation within the sequence does not affect the
12 result. Otherwise, the operation is “non-commutable.” Examples of commutable
13 operations are instructions to increment or decrement a value, instructions to insert a
14 record into a list at some defined location, and in general, any operation that does not
15 require reading a value to be carried out. The fourth optimization method exploits
16 these commutable operations.

17 Referring now to FIG. 5, in this fourth optimization method, a node **10**
18 receives a write request from a younger-generation spawn **22B** (step **60**) at a time
19 when normally the younger-generation spawn **22B** would not be permitted to write to
20 memory **12**. However, in this method, the node **10** distinguishes between commutable
21 operations and non-commutable operations (step **62**). If the proposed write is
22 commutable, the node **10** queues it (step **64**). The younger-generation spawn **22B** then
23 continues execution (step **66**). This allows the younger-generation spawn **22B** to
24 continue processing past the first time it tries to write to working-memory **12**. As a
25 result, younger-generation spawn **22B** continue to execute during the checkpoint
26 interval **32** for as long as any write operations carried out by that younger-generation
27 spawn **22B** are commutable. On the other hand, if the proposed write is a non-
28 commutable write, then the node **10** suspends execution of the younger-generation
29 spawn **22B** (step **68**).

30 In addition to non-commutable writes, there may be other conditions in which
31 a spawn **22B** may be allowed to write under conditions when it would normally not be
32 able to do so. One other example arises when a younger-generation spawn **22B**, after

1 having inspected memory **12**, recognizes that no further memory access by an older-
2 generation spawn **22A** is possible.

3 A fifth optimization method is one that reduces the latency that arises because
4 the outgoing-message escrow area **30** does not release outgoing messages **28** until a
5 checkpoint interval **32** is complete and all computations associated with generating
6 the outgoing messages **28** have been committed to durable storage **16**. The idea of
7 waiting until the end of a checkpoint interval **32** before releasing messages **28** from
8 outgoing-message escrow area **30** is useful where the consequences of sending the
9 wrong message are severe. However, there are times when the consequence of
10 sending an incorrect message is minimal, but the consequence of sending a delayed
11 message is severe.

12 As an example, consider the case in which the outgoing message **28** is a
13 coupon for goods in a particular retail store. Suppose the apparatus has detected that a
14 user is in the vicinity of that particular retail store at a particular instant. Obviously, it
15 would be desirable to transmit the message **28** immediately, before the user has had a
16 chance to leave the retail store. If this message **28** were to languish on the outgoing-
17 message escrow area **30** waiting to be sent, the opportunity for the coupon to be useful
18 would be lost. On the other hand, if that coupon were the result of a computation that
19 was subsequently lost because of a failure in the node **10**, it is unlikely anybody
20 would complain. After all, the store would have made a sale it might not otherwise
21 have made, and the user would have obtained a good at some discount.

22 This fifth optimization method, in which outgoing messages **28** are released
23 without waiting for the underlying data to be committed to durable storage **16**,
24 presupposes that time is of the essence in delivering an outgoing message **28**, and that
25 the cost of an incorrect or inconsistent outgoing message **28** is minimal in comparison
26 with adverse consequences of its late delivery. In the fifth optimization method,
27 outgoing messages **28** are released from the outgoing-message escrow area **30** prior to
28 completion of the checkpoint interval **32** or bypass the outgoing-message escrow area
29 **30** completely.

30 FIG. 6 shows a multi-node apparatus **70** in which multiple nodes **72**, **74**, **76**,
31 **78** of the type described in connection with FIGS. 1-6 are in communication with

1 each other and cooperate with each other in carrying out data processing. In such a
2 case, a task may send a message from a first node **72** to a second node **74**.

3 In some cases the message may have the effect of migrating a task from the
4 first node **72** to the second node **74**. A task that migrates from one node to another is
5 referred to as a “migrant” task. Depending on point-of-view, a migrant task is either
6 an “immigrant” task or an “emigrant” task. From the point of view of the first node
7 **72**, the migrant task is an “emigrant” task because the task is leaving the first node.
8 Conversely, from the point of view of the second node **74**, the migrant task is an
9 “immigrant” task because it is arriving at the second node **74**.

10 In other cases the message might be a remote procedure call or a remote data
11 access request such that the requesting task cannot proceed until it receives a message
12 in return. In other cases the task may simply asynchronously transmit information
13 from the first node **72** to the second node **74** using the message. A computing system
14 such as that described in U.S. Patent Application Serial No. 14/842,956, entitled
15 “EXECUTING GRAPH-BASED PROGRAM SPECIFICATIONS,” filed on
16 September 2, 2015, incorporated herein by reference, for example, can be configured
17 using the techniques for promoting fault tolerance and recovery described herein.

18 In such a case, application of the foregoing methods would be less than
19 optimal in part because a message **80** from the first node **72** to the second node **76**
20 cannot be transmitted until it is released from escrow at the completion of the next
21 checkpoint interval **32**. This introduces considerable latency. While this latency could
22 potentially be reduced by exempting messages transmitted from node to node within
23 the multi-node apparatus **70** from being escrowed, such an exemption is not sufficient
24 due to indeterminism.

25 For example, this and other difficulties arise when multiple nodes **72**, **74**, **76**,
26 **78** are present because many computations are non-deterministic. Examples of such
27 non-deterministic computations are those in which results depend on the order in
28 which reads and writes occur, those that rely on real time clocks, and those that rely
29 on the outcome of a random number generator, an example of which has already been
30 described above in connection with the desirability of an outgoing-message escrow
31 area **30**.

1 If a first node **72** communicates with a second node **74**, then loses contact
2 (e.g., due to failure) before the next checkpoint interval **32**, the apparatus **70** may end
3 up with inconsistencies owing to this non-determinism as follows. After the failure
4 the apparatus **70** will recover the first node **72** from the most recent checkpoint and
5 restart the computation. The computation may be restarted from a point in the
6 computation before transmission of a message from the first node **72** to the second
7 node **74**. Owing to the non-deterministic nature of the computations, the first node **72**
8 may well send a completely different message to the second node **74** after recovery
9 from the checkpoint. But, that second node **74** may have already received the original
10 message, potentially placing the two nodes **72** and **74** in an inconsistent state. For
11 example, node **72** is in a state in which it has sent node **74** the “new” version of the
12 message, but node **74** is in a state in which it has already acted on the “old” version of
13 the message. Furthermore, node **74** may have sent a message to yet another node **76**,
14 based on that original message received from node **72**, so node **72** and node **76** are
15 also may also be in an inconsistent state. Thus, inconsistency may spread through all
16 nodes in the apparatus **70** like a virus.

17 One way to avoid the foregoing difficulty is to ensure that all the nodes **72**, **74**,
18 **76**, **78** synchronize their checkpoints, for example, using a ‘barrier sync’ operation, as
19 follows. A ‘checkpoint leader’ transmits a message to all nodes commanding them to
20 begin a checkpoint interval. Then, after each checkpoint is complete, each node
21 responds to the checkpoint leader affirming that the checkpoint is complete. When the
22 checkpoint leader has received affirmations from all nodes, it will then command all
23 nodes to commit the checkpoint and then resume processing.

24 This approach forms the basis of a solution to the multi-node checkpoint
25 problem, but does not completely solve it for two reasons. First, in a multi-node
26 apparatus it is possible that some nodes survive a failure, in which case the surviving
27 nodes must be rolled back from their current state to the checkpoint state (rather than
28 being rolled forward to the checkpoint state). Second, when the checkpoint is
29 performed there may be messages in transit, which might allow non-determinism to
30 leak from the old processing interval, across the checkpoint, and into the new
31 processing interval.

1 In a single-node apparatus, if the node **10** fails, it only has to roll forward to
2 recover uncommitted work. But in a multi-node apparatus **70**, other nodes **72, 74, 76**,
3 which did not fail, may have to roll backward when a node **78** fails. This mechanism,
4 in which a distributed apparatus **70** recovers by having some nodes **78** roll forward
5 and other nodes **72, 74, 76** roll backward means that in effect, all nodes **72, 74, 76, 78**
6 can be made to restart at the same checkpoint. The resulting apparatus **70** thus
7 achieves the effect of simultaneous checkpoints across all nodes. It does not, however,
8 do so by trying to actually synchronize operation across all nodes, which as noted
9 above is difficult. Instead, it does so by manipulating the states of the nodes **72, 74,**
10 **76, 78** to reap the benefits of synchronized checkpoints without actually having to
11 provide such checkpoints.

12 To implement the foregoing recovery method, the nodes **72, 74, 76, 78** execute
13 a distributed checkpoint method as described in detail below. Referring to FIG. 7,
14 when implementing the distributed checkpoint method, every process and every
15 message acquires a generation count **82**. In addition, a running count **84** is maintained
16 of the tasks associated with each checkpoint. Each node also maintains a spawner-
17 registry **86** of its spawners **20**. Additionally, each node **72** maintains a nodal
18 generation count **88**.

19 The nodal generation count **88** enables a node **76** to enforce a generation gap
20 in which work carried out by younger-generation spawn **22A** and work carried out by
21 older-generation spawn **22B** do not interfere with each other. As a result of the
22 generation gap, the older generation and the younger generation can more or less
23 ignore each other. In effect, the node **76** becomes two virtual machines, one seen by
24 the older-generation spawn **22A** and another seen by the younger-generation spawn
25 **22B**. These two virtual machines coexist on the same physical platform but are
26 otherwise orthogonal to each other.

27 In addition, each node **76** also implements a bidirectional journal **90** that
28 enables that node **76** to roll forward or backward to a particular state as needed. The
29 bidirectional journal **90** includes changes to working storage **92**, a listing of
30 checkpointed task states **94**, and checkpointed messages **96**. These elements provide a
31 way to roll forward in time. In addition, the bidirectional journal **90** features an undo
32 log **98** in memory **12** to enable the node **76** to roll backward in time. In general,

1 rolling forward in time is how a failed node recovers. Rolling backward in time is
2 what a node does when another node in the apparatus **70** has failed.

3 In operation, as shown in FIG. 6, a master node **72** transmits a checkpoint
4 message **100** to all other nodes **74, 76, 78**, i.e. “slave nodes,” indicating that a
5 checkpoint is due. However, there is no requirement that this checkpoint occur at the
6 same time in all nodes **72, 74, 76, 78**.

7 FIG. 8 shows a flowchart for an exemplary fault-tolerance and recovery
8 procedure in the presence of computational indeterminism. In response to receiving a
9 checkpoint message (step **102**), a slave node **76** will not immediately begin a
10 checkpoint interval. As noted above, this is impractical. Instead, the slave node **76**
11 increments its nodal generation count **88** (step **104**) and create a journal entry
12 indicating the incrementing of its nodal generation count **88** (step **106**).

13 The slave node **76** then suspends all of its spawners **20** (step **108**), and writes
14 their states to the bidirectional journal **90** (step **110**). Then, for each of its spawners
15 **20**, the slave node **76** increments that spawner’s generation count **82** (step **112**). With
16 its generation count **82** having been incremented, the spawner **20** is allowed to resume
17 operation (step **114**). However, since the spawner’s generation count **82** will have
18 been incremented, any resulting spawn **22B** will be in the younger generation.

19 At this point, two generations will co-exist in the slave node **76**. The older-
20 generation spawn **22A**, namely those having a generation count that is one less than
21 the node’s generation count, can continue to process to completion, writing to
22 memory **12** as necessary. The younger-generation spawn **22B**, namely those whose
23 generation counts match the nodal generation count **88**, may process until it is time to
24 write to memory **12**. At that point, younger generation spawn **22B** are blocked.

25 It should be noted that in the description thus far, there are only two
26 generations of spawn involved: older-generation spawn **22A**, whose generation count
27 **82** is one less than the nodal generation count **88**, and a younger-generation spawn
28 **22B**, whose generation count **82** matches the nodal generation count **88**. However, in
29 principle there is no reason that more than two generations cannot coexist on the same
30 platform.

1 Referring to FIG. 9, in a multi-node apparatus **70**, it is possible for a task **79** to
2 emigrate from a sending node **78** and immigrate into a receiving node **76**. As
3 described in connection with FIG. 6, such a task **79** is referred to as a “migrant task”
4 or a “migrant.”

5 In the following discussion, it is necessary to refer to values associated with
6 particular objects. To avoid ambiguity with reference numerals in the figures, and in a
7 manner consistent with standard mathematical notation, the parentheses will be used
8 to mean “of.” Thus, since “**88**” has been assigned to “nodal generation count” and
9 “**76**” is a node, the nodal generation count **88** of node **76** will be written as **88(76)**.

10 A difficulty can arise when the migrant’s generation count **82(79)** is not the
11 same as the nodal generation count **88(76)** of the receiving node **76**. These difficulties
12 can be avoided by implementing message escrow areas between nodes. But this
13 would reintroduce the latency that the distributed checkpoint method was intended to
14 avoid in the first place.

15 According to the distributed checkpoint method, there are three possibilities:
16 the sending node’s nodal count **88(78)** is the same as the receiving node’s nodal count
17 **88(76)**; the sending node **78** has a lower nodal generation count **88(78)** than the
18 receiving node’s nodal count **88(76)**; and the sending node **78** has a higher nodal
19 generation count **88(78)** than the receiving node’s **88(76)**.

20 In the first possibility, a migrant will have the same generation count **82(79)** as
21 the nodal count **88(78)** of the sending node **78**. Therefore, the sending node **78**, the
22 receiving node **76**, and the migrant **79** all have the same generation count. In that case,
23 nothing special has to be done.

24 The second possibility can arise when the receiving node **76** increments its
25 generation count **88(76)** while the migrant **79** is in transit. This means that, upon
26 immigrating into the receiving node **76**, the migrant **79** presents itself as a member of
27 what has now become the older generation of the receiving node **76**. In that case, the
28 receiving node **76** will youthen the migrant **79** by incrementing the migrant’s
29 generation count **82(79)**. As a result, the migrant task **79** will be able to continue
30 processing, but, like the rest of the younger generation spawn **22B**, it will be blocked
31 from writing to memory **12**. The youthening of the migrant **79** is then journalized at

1 the receiving node **76**. Since the act of youthening takes place at the receiving node
2 **76**, it is referred to as “immigrant-side youthening.”

3 The third possibility can arise when the sending node **78** will increment its
4 generation count **88(78)** before the migrant **79** has emigrated. In that case, the sending
5 node **78** youthens the migrant **79** by incrementing the migrant’s generation count
6 **82(79)** before it is sent, and journalizes the youthening event at the sending node.
7 Since the act of youthening takes place at the sending node **78**, it is referred to as
8 “emigrant-side youthening.”

9 In either case, a node **76** that has received a checkpoint message from a master
10 node will set a deadline to allow the older-generation spawn **22A** to finish execution,
11 thereby insuring near-extinction of the older generation, and avoiding the need to
12 record their states (step **116**). Nevertheless, there may be spawn **22A** of the older
13 generation that are slow to terminate. It is impractical for a node **76** to wait for an
14 extended period for once the deadline is reached, any older generation spawn **22** that
15 is still running will be suspended, serialized, journaled, and youthened, after which it
16 is allowed to resume execution subject to the constraint that it not write to the
17 working-memory **12** until after the working-memory **12** has been committed to
18 durable storage **16**.

19 The slave node **76** will not begin the actual checkpoint until it knows that no
20 more older-generation immigrants are expected to arrive. In order to implement this,
21 whenever a node **72** recognizes that all older-generation emigrants have successfully
22 emigrated, it broadcasts a flush message to all other nodes **74, 76, 78**. Once the slave
23 node **76** has received flush messages from all nodes **72, 74, 78**, it knows that the flow
24 of older-generation immigrants has been quenched (step **118**). Younger-generation
25 immigrants may still arrive at a slave node **76**, just as younger-generation emigrants
26 may still leave from the slave node **76**. However, these younger-generation emigrants
27 are not pertinent to the checkpoint process.

28 At this point, the slave node **76** is now ready to commit its working-memory
29 **12** to durable storage **16** (step **120**). This is carried out in the same manner described
30 above for the single-node case.

1 The procedure for restarting after failure of a node, shown in FIG. 10, depends
2 on whether the node involved is one that failed or not. After receiving an instruction
3 to restart (step **122**), the node determines if it is the node that failed, or if another node
4 in the apparatus **70** failed (step **124**). If the node is one that failed, it retrieves the log
5 and rolls forward from its last valid checkpoint (step **126**). If the node is not one that
6 failed (i, it rolls back to its last checkpoint (step **128**).

7 An example of a “roll back” operation involves the following steps: (1)
8 terminate all tasks currently running (including both spawners and spawn); (2) use the
9 bidirectional journal entries to undo any changes to memory.

10 After any failed nodes have been rolled forward and any surviving nodes have
11 been rolled back, the apparatus **70** may also perform other operations as part of
12 restarting the tasks. For example, the apparatus **70** may perform the following
13 operations: (1) flush the communications network to ensure that all messages
14 predating the failure have been discarded, (2) restart all tasks that were part of the
15 checkpoint by retrieving their saved state from the journal and restarting them, (3)
16 retransmit any messages that were not sent prior to the checkpoint, and (4) process
17 any messages were received but not yet processed as of the checkpoint.

18 The task of rolling forward from the last valid checkpoint on a failed node is
19 one that is potentially time-consuming. Referring to FIG. 11, in some practices, it is
20 useful to maintain a replica **128** of memory **130** from a first node **132** on a second
21 node **134**. Preferably, the second node **134** does not have the same failure mode as the
22 first node **132**. In normal operation, the replica **128** is synchronized with the memory
23 **130** at the first node **132** at each checkpoint. The replica **128** also has an associated
24 undo log **136** to enable it to roll backward to its state at the most recent checkpoint.

25 Referring now to FIG. 12, upon failure of the first node **132**, the replica **128** at
26 the second node **124** is designated a master (step **138**). All processes on the second
27 node **124** are killed (step **140**), after which the second node **134** is restarted (step **142**).
28 The former replica **128**, which now serves as a master copy, is rolled back to the last
29 checkpoint with the aid of the undo log **136** (step **144**). Operation of the multi-node
30 apparatus **70** can then resume with the wait for recovery being on the order of the roll-
31 back time. This is typically much shorter than the roll-forward time. Meanwhile, the

1 recovered first node **132** can proceed to roll-forward to the correct state without
2 slowing down the overall recovery of the multi-node apparatus **70**. Once the first node
3 **132** is ready, it takes over as master again, and the former replica **128** becomes a
4 replica again.

5 Although FIG. 11 shows only one second node **134**, it is understood that there
6 can be more than one second node, each of which has a replica **128** and an undo log
7 **136**. In that case, upon failure of the first node **132**, one of the second nodes must be
8 elected to serve as proprietor of the new master copy of the first node's memory.

9 In some cases, there may be many idempotent operations. In such cases,
10 instead of rolling forward it is not unreasonable to simply repeat computations that
11 would carry out idempotent operations since those computations would not cause any
12 harm.

13 The end result of recovery is that all points are at a state consistent with the
14 transition from one generation to the next. As a result, no work from older-generation
15 processes is lost, but all work done by younger generation processes is lost. This
16 ensures a state that is consistent across all nodes. In this context, a state is "consistent"
17 if it could have been arrived at in the absence of any fault. In contrast, a state is
18 "inconsistent" if it can only be explained by the occurrence of one or more faults.

19 FIG. 13 illustrates the states of several spawned processes in both the sending
20 node **78** and receiving node **76** in the multi-node apparatus **70** referred to in
21 connection with FIGS. 6 and 9. In FIG. 13, time increases downward along the
22 vertical axis. The time axis shows a first interval **146**, a second interval **148** following
23 the first interval **146**, and a third interval **150** following the second interval **148**.

24 FIG. 13 shows several spawned processes **22A-H**, each of which has an
25 associated generation count. Spawn having a generation count of N will be referred to
26 herein as "first-generation spawn." Spawn having a generation count of $N+1$ will be
27 referred to herein as "second-generation spawn." The adjectives "first-generation" and
28 "second-generation" will also be used to refer to other entities that are tagged with a
29 generation count, including nodes, migrant tasks, and spawned process.

1 During the first interval **146**, the sending node **78** is a first-generation node.
2 During the second and third interval **150**, the sending node **78** is a second-generation
3 node. It should be noted that this progression of nodes is cyclic so that the third
4 interval **150** will be followed by an interval that plays the same role, for the second
5 generation, that the second interval **148** played for the first generation. This same
6 progression occurs on the receiving node **76**, though not necessarily in synchrony
7 with the progression occurring at the sending node **78**. For convenience, the same
8 reference numbers are used to designate intervals in both the sending and receiving
9 nodes **78**, **76**. However, this is not meant to imply that they are synchronized.

10 During the first interval **146**, a spawning process **20** spawns various first-
11 generation spawned processes **22A-22E**. Throughout this first interval **146**, any first-
12 generation spawned process **22A-22E** is free to write to a sending-node memory **12A**.

13 During the second interval **148**, the sending node **78** becomes a second-
14 generation node. As such, the spawning process **20** now spawns only second-
15 generation spawned processes. During this second interval **148**, any first-generation
16 spawned processes **22A-22E** remain free to write to the sending-node memory **12A**.
17 Second-generation spawned processes **22F-22G** are free to execute, but are forbidden
18 from writing to the sending-node memory **12A**. The purpose of this second interval
19 **148** is therefore to allow any residual first-generation spawn **22C**, **22D**, **22E** some
20 time to finish executing before a checkpoint interval **32** occurs.

21 During the third interval **150**, the spawning process **20** spawns another second-
22 generation spawned process **22H**. During this third interval **150**, no first-generation
23 spawn remain, and all second-generation spawn **22F-22H** are free to write to the
24 sending-node memory **12A**.

25 At the sending node **78**, a first first-generation spawned process **22A**, a second
26 first-generation spawned process **22B**, a third first-generation spawned process **22C**, a
27 fourth first-generation spawned process **22D**, and a fifth first-generation spawned
28 process **22E** all begin execution during the first interval **146**. However, of these, only
29 the first first-generation spawned process **22A** and the second first-generation
30 spawned process **22B** manage to finish execution during the first interval **146**. The
31 third first-generation spawned process **22C** manages to finish during the second

1 interval **148**. The fourth first-generation spawned process **22D** takes so long it cannot
2 finish until the third interval **150** has already begun. The fifth first-generation
3 spawned process **22E** never actually finishes at the sending node **78**. Instead, it
4 migrates to the receiving node **76** part way through the second interval **148**. It does so
5 while the receiving node **76** is still in its own second interval **148**.

6 During execution, the first first-generation spawned process **22A** writes to the
7 sending-node memory **12A** during the first interval **146** and the third first-generation
8 spawned process **22C** writes to the sending-node memory **12A** during the second
9 interval **148**. The second first-generation spawned process **22B** does not write to the
10 sending-node memory **12A** at all during execution. The fifth first-generation spawned
11 process **22E** eventually writes to the sending-node memory **12A**, but only at the
12 receiving node **76**.

13 During the second interval **148**, a first second-generation spawned process
14 **22F** and a second second-generation spawned process **22G** both begin execution.
15 Sometime during the second interval **148**, the first second-generation spawned
16 process **22F** reaches a point at which it must write to the sending-node memory **12A**.
17 However, since it is still the second interval **148**, it is forbidden from writing to the
18 sending-node memory **12A**. Therefore, it becomes suspended, as indicated by the
19 dashed lines. Once the third interval **150** begins, the first second-generation spawned
20 process **22F** writes to the sending-node memory **12A** and completes execution.

21 Meanwhile, the second second-generation spawned process **22G** has started
22 late enough during the second interval **148** so that by the time it actually has to write
23 to the sending-node memory **12A**, the third interval **150** has already begun.
24 Accordingly, the second second-generation spawned process **22G** executes without
25 interruption.

26 A third second-generation spawned process **22H** begins during the third
27 interval **150**. This is essentially a mirror image of the first first-generation spawned
28 process **22A**.

29 In the course of execution, the first first-generation spawned process **22A**
30 causes a first task **152** to migrate to the receiving node **76**. The first task **152** inherits
31 the generation number of the first first-generation spawned process **22A**. As such, it

1 begins its existence as a first-generation task. This first task **152** arrives at the
2 receiving node **76** while the receiving node **76** is still operating in the first interval
3 **146**. The receiving node **76** is thus acting as a first-generation node. Accordingly, the
4 first task **152** is free to execute and to write to a receiving-node memory **12B** provided
5 it does so before a third interval **150** begins on the receiving node **76**.

6 Also in the course of execution, the second first-generation spawned process
7 **22B** causes a second task **154** to migrate to the receiving node **76**. The second task
8 **154** inherits the generation number of the first first-generation spawned process **22A**.
9 As such, it starts its existence as a first-generation task. However, this second task **154**
10 arrives at the receiving node **76** while the receiving node **76** is already operating in its
11 second interval **148**. Accordingly, the second task **154** is changed into a second-
12 generation task from a first-generation task. This includes an accompanying step of
13 journalizing the second task **154** in a receiving node journal file **156**.

14 A similar event occurs in connection with the fifth first-generation spawned
15 process **22E** at the sending node **78**. This fifth first-generation spawned process **22E**
16 migrates to the receiving node **76** midway through execution. However, by the time it
17 arrives at the receiving node **76**, the receiving node **76** has already begun its own
18 second interval **148**. As such, the second node is has become a second-generation
19 node. Therefore, the fifth first-generation spawned process **22E** is changed into a
20 second-generation spawned process. This change is accompanied by journalizing the
21 fifth first-generation spawned process **22E** in a sending-node journal file **158**. The
22 fifth first-generation spawned process **22E** then continues execution on the receiving
23 node **76**, though as a second-generation spawned process.

24 Meanwhile, back at the sending node **78**, the fourth first-generation spawned
25 process **22D** has not yet finished execution by the end of the second interval **148**. At
26 this point, the fourth first-generation spawned process **22D** is both journalized at the
27 sending-node journal file **158** and has its generation count incremented so that it now
28 becomes a second-generation spawned process. The fourth first-generation spawned
29 process **22D** then continues to execute during the third interval **150**.

30 It should be noted that the fourth first-generation spawned process **22D**
31 sustained the same two steps that were sustained by the fifth first-generation spawned

1 process **22E** during its migration to the receiving node **76**, namely a journalizing step,
2 and a generation change. Thus, it is not unreasonable to say that fourth first-
3 generation spawned process **22D** in some sense also migrated. The main difference is
4 that the fifth first-generation spawned process **22E** underwent an inter-node migration
5 whereas the fourth first-generation spawned process **22D** underwent an intra-node
6 migration.

7 The checkpoint and recovery method described herein is thus based on the
8 recognition that the desirability of simultaneously executing checkpoints across
9 multiple nodes does not stem from temporal synchronicity but rather from a side
10 effect of temporal synchronicity. The method thus reproduces the side effect of
11 temporal synchronicity of checkpoints across multiple nodes without actually having
12 to achieve it.

13 The fault-tolerance and recovery approach described above can be
14 implemented, for example, using a programmable computing system executing
15 suitable software instructions or it can be implemented in suitable hardware such as a
16 field-programmable gate array (FPGA) or in some hybrid form. For example, in a
17 programmed approach the software may include procedures in one or more computer
18 programs that execute on one or more programmed or programmable computing
19 system (which may be of various architectures such as distributed, client/server, or
20 grid) each including at least one processor, at least one data storage system (including
21 volatile and/or non-volatile memory and/or storage elements), at least one user
22 interface (for receiving input using at least one input device or port, and for providing
23 output using at least one output device or port). The software may include one or
24 more modules of a larger program, for example, that provides services related to the
25 design, configuration, and execution of dataflow graphs. The modules of the program
26 (e.g., elements of a dataflow graph) can be implemented as data structures or other
27 organized data conforming to a data model stored in a data repository.

28 The software may be stored in non-transitory form, such as being embodied in
29 a volatile or non-volatile storage medium, or any other non-transitory medium, using
30 a physical property of the medium (e.g., surface pits and lands, magnetic domains, or
31 electrical charge) for a period of time (e.g., the time between refresh periods of a
32 dynamic memory device such as a dynamic RAM). In preparation for loading the

1 instructions, the software may be provided on a tangible, non-transitory medium, such
2 as a CD-ROM or other computer-readable medium (e.g., readable by a general or
3 special purpose computing system or device), or may be delivered (e.g., encoded in a
4 propagated signal) over a communication medium of a network to a tangible, non-
5 transitory medium of a computing system where it is executed. Some or all of the
6 processing may be performed on a special purpose computer, or using special-purpose
7 hardware, such as coprocessors or field-programmable gate arrays (FPGAs) or
8 dedicated, application-specific integrated circuits (ASICs). The processing may be
9 implemented in a distributed manner in which different parts of the computation
10 specified by the software are performed by different computing elements. Each such
11 computer program is preferably stored on or downloaded to a computer-readable
12 storage medium (e.g., solid state memory or media, or magnetic or optical media) of a
13 storage device accessible by a general or special purpose programmable computer, for
14 configuring and operating the computer when the storage device medium is read by
15 the computer to perform the processing described herein. The inventive system may
16 also be considered to be implemented as a tangible, non-transitory medium,
17 configured with a computer program, where the medium so configured causes a
18 computer to operate in a specific and predefined manner to perform one or more of
19 the processing steps described herein.

20 A number of embodiments of the invention have been described. Nevertheless,
21 it is to be understood that the foregoing description is intended to illustrate and not to
22 limit the scope of the invention, which is defined by the scope of the following
23 claims. Accordingly, other embodiments are also within the scope of the following
24 claims. For example, various modifications may be made without departing from the
25 scope of the invention. Additionally, some of the steps described above may be order
26 independent, and thus can be performed in an order different from that described.

27 Having described the invention, what we claim as new, and secured by letters
28 patent is:

1 **CLAIMS**

2

3 1. A method for enhancing fault tolerance and recovery in a computing
4 system including at least one processing node, said method including:
5 enhancing availability and recovery of a first processing node,
6 including, at the first processing node,

7 executing a spawner at said node, wherein said spawner, in the course of
8 execution, generates a first spawn, wherein the spawner implements a
9 process, with an associated expected spawner lifetime, configured to
10 generate one or more spawn processes, including the first spawn, that are
11 each associated with a respective expected lifetime shorter than the
12 expected spawner lifetime, and which are configured to be suspended if,
13 following a checkpoint event and suspension of the spawner, the
14 respective one or more spawn processes are still executing after a pre-
15 determined interval measured from occurrence of the checkpoint event
16 and suspension of the spawner, wherein at least one of the one or more
17 spawn processes terminates before the end of the pre-determined interval,

18 wherein executing said spawner includes assigning, to said spawner, a
19 first generation indicator,

20 wherein said first spawn inherits said first generation indicator;

21 beginning a checkpoint interval, at the end of which nodal recovery
22 information, which is usable for recovery of said node, is committed to
23 durable storage, wherein beginning said checkpoint interval includes

24 suspending said spawner from generating spawn,

25 assigning, to said spawner, a second generation indicator that differs
26 from said first generation indicator,

27 resuming said spawner, thereby enabling said spawner to generate a
28 second spawn, wherein said second spawn inherits said second
29 generation indicator, and

- 1 controlling an extent to which said second spawn writes to memory;
2 and
- 3 after committing said nodal recovery information, releasing control over said
4 extent to which said second spawn can write to memory.
- 5 **2.** The method of claim 1, wherein controlling an extent to which said
6 second spawn writes to memory includes preventing said second
7 spawn from consummating a write to said memory.
- 8 **3.** The method of claim 2, further including permitting said second spawn
9 to queue said write to memory for eventual consummation thereof after
10 said recovery information has been committed.
- 11 **4.** The method of claim 1, wherein controlling an extent to which said
12 second spawn writes to memory includes determining that said write
13 operation is a commutable operation whose sequence with respect to
14 other operations does not affect a result of said commutable operation,
15 and allowing consummation of said commutable operation.
- 16 **5.** The method of claim 4, wherein determining that said write operation
17 is a commutable operation includes determining that said write
18 operation includes incrementing a variable.
- 19 **6.** The method of claim 4, wherein determining that said write operation
20 is a commutable operation includes determining that said write
21 operation includes inserting a record at a specified location.
- 22 **7.** The method of claim 1, further including, after suspending said
23 spawner, setting a deadline, thereby providing time for any spawn
24 having said first task generation-indicator to execute to completion,
25 and avoiding overhead associated with having to save states of said
26 spawn having said first generation-indicator.
- 27 **8.** The method of claim 7, further including suspending said first spawn if
28 said first spawn is still executing as of said deadline.

- 1 **9.** The method of claim 7, further including enabling said first spawn to
2 avoid suspension as a result of having failed to complete execution by
3 said deadline.
- 4 **10.** The method of claim 9, wherein enabling said first spawn to avoid
5 suspension as a result of having failed to complete execution by said
6 deadline includes changing said first task generation-indicator to said
7 second task generation-indicator in said first spawn if said first spawn
8 is still executing as of said deadline.
- 9 **11.** The method of claim 1, wherein said first node has a nodal-generation
10 indicator, said method further including causing a spawn to become a
11 migrant that migrates to a second node, wherein said second node has a
12 nodal-generation indicator.
- 13 **12.** The method of claim 11, wherein said nodal-generation counter of said
14 second node indicates that said second node is in a younger generation
15 than said first node, wherein said method further includes youthening
16 said migrant to cause a generation indicator value of the migrant to be
17 changed to indicate a younger generation value associated with the
18 migrant.
- 19 **13.** The method of claim 12, wherein youthening said migrant includes
20 immigration-side youthening of said migrant.
- 21 **14.** The method of claim 12, wherein youthening said migrant includes
22 emigration-side youthening of said migrant.
- 23 **15.** The method of claim 1, wherein said first node is a node in a multi-
24 node system in which each node has a nodal generation-count, wherein
25 said multi-node system includes at least a second node, wherein, upon
26 recovery following a failure of said second node, said first node rolls
27 back to a state that corresponds to a nodal-generation count of said
28 second node.

- 1 **16.** The method of claim 1, wherein said first node is a node in a multi-
2 node system in which each node has a nodal generation-count, wherein
3 said multi-node system includes at least a second node, wherein, upon
4 recovery following a failure of said first node, said first node rolls
5 forward to a state that corresponds to a nodal-generation count of said
6 second node by restoring committed work from a checkpoint and
7 restoring uncommitted work from a journal.
- 8 **17.** The method of claim 1, wherein said first node is a node in a multi-
9 node system in which each node has a nodal generation-count, said
10 method including, at said first node,
11 receiving, from a master node, a message indicating that a checkpoint is to be
12 carried out,
13 in response, youthening a nodal generation count of said first node to cause
14 the nodal-generation count of said first node be changed to indicate a
15 younger generation value associated with the first node,
16 suspending spawners from generating spawn,
17 saving spawner recovery information for recovering spawner states,
18 resuming said spawners,
19 determining that no further older-generation immigrants are expected at said
20 first node, and
21 in response to said determination, committing, to said durable storage, said
22 nodal recovery information.
- 23 **18.** The method of claim 17, further including setting a deadline, and, upon
24 lapse of said deadline, suspending all older-generation spawn that are
25 still executing while younger-generation spawn continue to execute.
- 26 **19.** The method of claim 1, wherein said first node is a node in a multi-
27 node system, said method including saving a replica copy of working
28 memory of said first node at said second node, upon failure of said first

1 node, temporarily using said replica copy for processing that would
2 otherwise have been carried out by said first node, and, upon recovery
3 of said first node, communicating, to said first node, information
4 required to update memory in said first node so that subsequent
5 computation can be carried out by said first node.

6 **20.** A computer-readable storage medium having recorded thereon
7 instructions for enhancing fault tolerance and recovery in a computing
8 system including at least one processing node, the instructions being
9 executable by a computing system to:

10 enhance availability and recovery of a first processing node, including,
11 at a first processing node,

12 executing a spawner at said node, wherein said spawner, in the course of
13 execution, generates a first spawn, wherein the spawner implements a
14 process, with an associated expected spawner lifetime, configured to
15 generate one or more spawn processes, including the first spawn, that are
16 each associated with a respective expected lifetime shorter than the
17 expected spawner lifetime, and which are configured to be suspended if,
18 following a checkpoint event and suspension of the spawner, the
19 respective one or more spawn processes are still executing after a pre-
20 determined interval measured from occurrence of the checkpoint event
21 and suspension of the spawner, wherein at least one of the one or more
22 spawn processes terminates before the end of the pre-determined interval,

23 wherein executing said spawner includes assigning, to said spawner, a
24 first generation indicator,

25 wherein said first spawn inherits said first generation indicator;

26 beginning a checkpoint interval, at the end of which nodal recovery
27 information, which is usable for recovery of said node, is committed to
28 durable storage, wherein beginning said checkpoint interval includes

29 suspending said spawner from generating spawn,

- 1 assigning, to said spawner, a second generation indicator that differs
2 from said first generation indicator,
- 3 resuming said spawner, thereby enabling said spawner to generate a
4 second spawn, wherein said second spawn inherits said second
5 generation indicator, and
- 6 controlling an extent to which said second spawn writes to memory;
7 and
- 8 after committing said nodal recovery information, releasing control over said
9 extent to which said second spawn can write to memory.
- 10 **21.** The computer-readable storage medium of claim **20**, wherein
11 controlling an extent to which said second spawn writes to memory
12 includes preventing said second spawn from consummating a write to
13 said memory.
- 14 **22.** The computer-readable storage medium of claim **21**, further including
15 instructions for causing the computing system to permit said second
16 spawn to queue said write to memory for eventual consummation
17 thereof after said recovery information has been committed.
- 18 **23.** The computer-readable storage medium of claim **20**, wherein
19 controlling an extent to which said second spawn writes to memory
20 includes determining that said write operation is a commutable
21 operation whose sequence with respect to other operations does not
22 affect a result of said commutable operation, and allowing
23 consummation of said commutable operation.
- 24 **24.** The computer-readable storage medium of claim **23**, wherein
25 determining that said write operation is a commutable operation
26 includes determining that said write operation includes incrementing a
27 variable.
- 28 **25.** The computer-readable storage medium of claim **23**, wherein
29 determining that said write operation is a commutable operation

- 1 includes determining that said write operation includes inserting a
2 record at a specified location.
- 3 **26.** The computer-readable storage medium of claim **20**, further including
4 instructions for causing the computing system to, after suspending said
5 spawner, set a deadline, thereby providing time for any spawn having
6 said first task generation-indicator to execute to completion, and
7 avoiding overhead associated with having to save states of said spawn
8 having said first generation-indicator.
- 9 **27.** The computer-readable storage medium of claim **26**, further including
10 instructions for causing the computing system to suspend said first
11 spawn if said first spawn is still executing as of said deadline.
- 12 **28.** The computer-readable storage medium of claim **26**, further including
13 instructions for causing the computing system to enable said first
14 spawn to avoid suspension as a result of having failed to complete
15 execution by said deadline.
- 16 **29.** The computer-readable storage medium of claim **28**, wherein enabling
17 said first spawn to avoid suspension as a result of having failed to
18 complete execution by said deadline includes changing said first task
19 generation-indicator to said second task generation-indicator in said
20 first spawn if said first spawn is still executing as of said deadline.
- 21 **30.** The computer-readable storage medium of claim **20**, wherein the first
22 node has a nodal-generation indicator, said computer-readable storage
23 medium further including instructions for causing the computing
24 system to cause a spawn to become a migrant that migrates to a second
25 node, wherein said second node has a nodal-generation indicator.
- 26 **31.** The computer-readable storage medium of claim **30**, wherein said
27 nodal-generation counter of said second node indicates that said
28 second node is in a younger generation than said first node, wherein
29 said computer-readable storage medium further includes instructions
30 for causing the computing system to youthen said migrant to cause a

- 1 generation indicator value of the migrant to be changed to indicate a
2 younger generation value associated with the migrant.
- 3 **32.** The computer-readable storage medium of claim **31**, wherein
4 youthening said migrant includes immigration-side youthening of said
5 migrant.
- 6 **33.** The computer-readable storage medium of claim **31**, wherein
7 youthening said migrant includes emigration-side youthening of said
8 migrant.
- 9 **34.** The computer-readable storage medium of claim **20**, wherein said first
10 node is a node in a multi-node system in which each node has a nodal
11 generation-count, wherein said multi-node system includes at least a
12 second node, wherein, upon recovery following a failure of said second
13 node, said first node rolls back to a state that corresponds to a nodal-
14 generation count of said second node.
- 15 **35.** The computer-readable storage medium of claim **20**, wherein said first
16 node is a node in a multi-node system in which each node has a nodal
17 generation-count, wherein said multi-node system includes at least a
18 second node, wherein, upon recovery following a failure of said first
19 node, said first node rolls forward to a state that corresponds to a
20 nodal-generation count of said second node by restoring committed
21 work from a checkpoint and restoring uncommitted work from a
22 journal.
- 23 **36.** The computer-readable storage medium of claim **20**, wherein said first
24 node is a node in a multi-node system in which each node has a nodal
25 generation-count, said computer-readable storage medium including
26 instructions for causing the computing system to, at said first node,
27 receive, from a master node, a message indicating that a checkpoint is to be
28 carried out,

- 1 in response, you then a nodal generation count of said first node to cause the
2 nodal-generation count of said first node be changed to indicate a younger
3 generation value associated with the first node,
4 suspend spawners from generating spawn,
5 save spawner recovery information for recovering spawner states,
6 resume said spawners,
7 determine that no further older-generation immigrants are expected at said
8 first node, and
9 in response to said determination, commit, to said durable storage, said nodal
10 recovery information.
- 11 **37.** The computer-readable storage medium of claim **36**, further including
12 instructions for causing the computing system to set a deadline, and,
13 upon lapse of said deadline, suspend all older-generation spawn that
14 are still executing while younger-generation spawn continue to
15 execute.
- 16 **38.** The computer-readable storage medium of claim **20**, wherein said first
17 node is a node in a multi-node system, said computer-readable storage
18 medium including instructions for causing the computing system to
19 save a replica copy of working memory of said first node at said
20 second node, upon failure of said first node, temporarily using said
21 replica copy for processing that would otherwise have been carried out
22 by said first node, and, upon recovery of said first node, communicate,
23 to said first node, information required to update memory in said first
24 node so that subsequent computation can be carried out by said first
25 node.
- 26 **39.** A computing system including:
27 a data storage system including durable storage; and
28 one or more processing nodes including least one processor configured
29 to enhance availability and recovery of a first processing node,
30 including, at the first processing node,
31 executing a spawner at said node, wherein said spawner, in the course of
32 execution, generates a first spawn, wherein the spawner implements a

1 process, with an associated expected spawner lifetime, configured to
2 generate one or more spawn processes, including the first spawn, that are
3 each associated with a respective expected lifetime shorter than the
4 expected spawner lifetime, and which are configured to be suspended if,
5 following a checkpoint event and suspension of the spawner, the
6 respective one or more spawn processes are still executing after a pre-
7 determined interval measured from occurrence of the checkpoint event
8 and suspension of the spawner, wherein at least one of the one or more
9 spawn processes terminates before the end of the pre-determined interval,
10 wherein executing said spawner includes assigning, to said spawner, a
11 first generation indicator,
12 wherein said first spawn inherits said first generation indicator;
13 beginning a checkpoint interval, at the end of which nodal recovery
14 information, which is usable for recovery of said node, is committed to
15 durable storage, wherein beginning said checkpoint interval includes
16 suspending said spawner from generating spawn,
17 assigning, to said spawner, a second generation indicator that differs
18 from said first generation indicator,
19 resuming said spawner, thereby enabling said spawner to generate a
20 second spawn, wherein said second spawn inherits said second
21 generation indicator, and
22 controlling an extent to which said second spawn writes to memory;
23 and
24 after committing said nodal recovery information, releasing control over said
25 extent to which said second spawn can write to memory.

26 **40.** The computing system of claim **39**, wherein controlling an extent to
27 which said second spawn writes to memory includes preventing said
28 second spawn from consummating a write to said memory.

- 1 **41.** The computing system of claim **40**, said processor further configured
2 to permit said second spawn to queue said write to memory for
3 eventual consummation thereof after said recovery information has
4 been committed.
- 5 **42.** The computing system of claim **39**, wherein controlling an extent to
6 which said second spawn writes to memory includes determining that
7 said write operation is a commutable operation sequence with respect
8 to other operations does not affect a result of said commutable
9 operation, and allowing consummation of said commutable operation.
- 10 **43.** The computing system of claim **42**, wherein determining that said
11 write operation is a commutable operation includes determining that
12 said write operation includes incrementing a variable.
- 13 **44.** The computing system of claim **42**, wherein determining that said
14 write operation is a commutable operation includes determining that
15 said write operation includes inserting a record at a specified location.
- 16 **45.** The computing system of claim **39**, said processor further configured
17 to, after suspending said spawner, set a deadline, thereby providing
18 time for any spawn having said first task generation-indicator to
19 execute to completion, and avoiding overhead associated with having
20 to save states of said spawn having said first generation-indicator.
- 21 **46.** The computing system of claim **45**, said processor further configured
22 to suspend said first spawn if said first spawn is still executing as of
23 said deadline.
- 24 **47.** The computing system of claim **45**, said processor further configured
25 to enable said first spawn to avoid suspension as a result of having
26 failed to complete execution by said deadline.
- 27 **48.** The computing system of claim **47**, wherein enabling said first spawn
28 to avoid suspension as a result of having failed to complete execution
29 by said deadline includes changing said first task generation-indicator

- 1 to said second task generation-indicator in said first spawn if said first
2 spawn is still executing as of said deadline.
- 3 **49.** The computing system of claim **39**, wherein said first node has a
4 nodal-generation indicator, said processor further configured to cause a
5 spawn to become a migrant that migrates to a second node, wherein
6 said second node has a nodal-generation indicator.
- 7 **50.** The computing system of claim **49**, wherein said nodal-generation
8 counter of said second node indicates that said second node is in a
9 younger generation than said first node, said processor further
10 configured to youthen said migrant to cause a generation indicator
11 value of the migrant to be changed to indicate a younger generation
12 value associated with the migrant.
- 13 **51.** The computing system of claim **50**, wherein youthening said migrant
14 includes immigration-side youthening of said migrant.
- 15 **52.** The computing system of claim **50**, wherein youthening said migrant
16 includes emigration-side youthening of said migrant.
- 17 **53.** The computing system of claim **39**, wherein said first node is a node in
18 a multi-node system in which each node has a nodal generation-count,
19 wherein said multi-node system includes at least a second node,
20 wherein, upon recovery following a failure of said second node, said
21 first node rolls back to a state that corresponds to a nodal-generation
22 count of said second node.
- 23 **54.** The computing system of claim **39**, wherein said first node is a node in
24 a multi-node system in which each node has a nodal generation-count,
25 wherein said multi-node system includes at least a second node,
26 wherein, upon recovery following a failure of said first node, said first
27 node rolls forward to a state that corresponds to a nodal-generation
28 count of said second node by restoring committed work from a
29 checkpoint and restoring uncommitted work from a journal.

- 1 **55.** The computing system of claim **39**, wherein said first node is a node in
2 a multi-node system in which each node has a nodal generation-count,
3 said the processor further configured to, at said first node,
4 receive, from a master node, a message indicating that a checkpoint is to be
5 carried out,
6 in response, youthen a nodal generation count of said first node to cause the
7 nodal-generation count of said first node be changed to indicate a younger
8 generation value associated with the first node,
9 suspend spawners from generating spawn,
10 save spawner recovery information for recovering spawner states,
11 resume said spawners,
12 determine that no further older-generation immigrants are expected at said
13 first node, and
14 in response to said determination, commit, to said durable storage, said nodal
15 recovery information.
- 16 **56.** The computing system of claim **55**, said processor further configured
17 to set a deadline, and, upon lapse of said deadline, suspend all older-
18 generation spawn that are still executing while younger-generation
19 spawn continue to execute.
- 20 **57.** The computing system of claim **39**, wherein said first node is a node in
21 a multi-node system, said processor further configured to save a replica
22 copy of working memory of said first node at said second node, upon
23 failure of said first node, temporarily using said replica copy for
24 processing that would otherwise have been carried out by said first
25 node, and, upon recovery of said first node, communicate, to said first
26 node, information required to update memory in said first node so that
27 subsequent computation can be carried out by said first node.

28




FIG. 1




FIG. 2

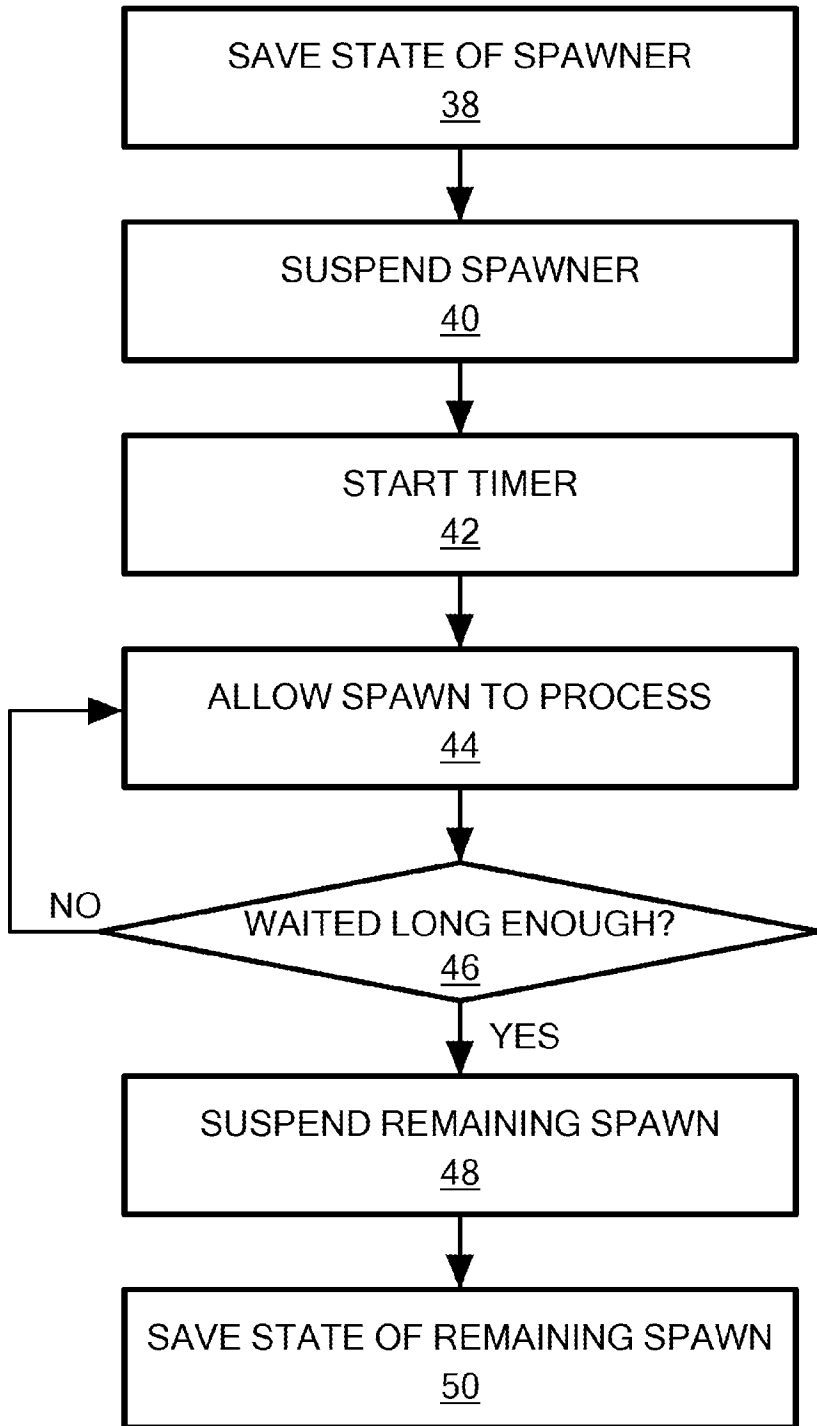


FIG. 3

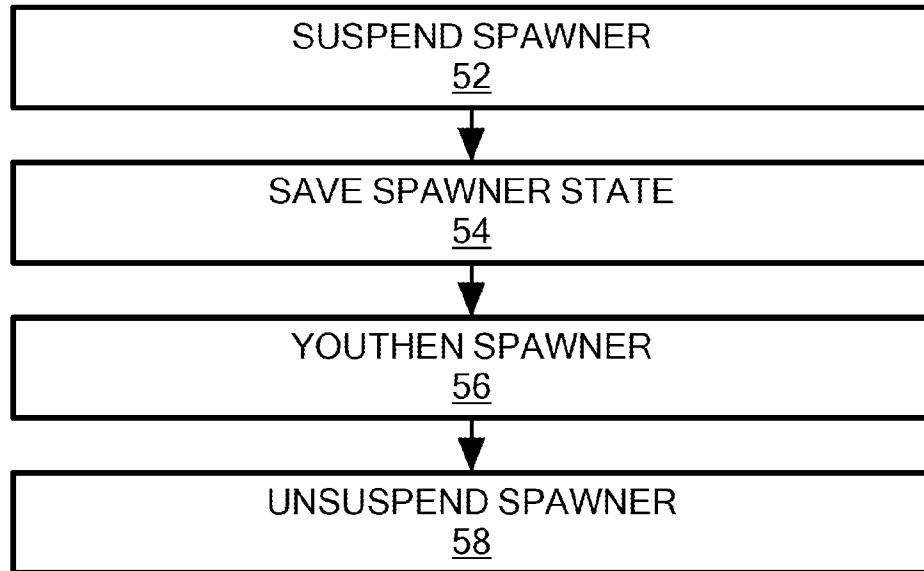


FIG. 4

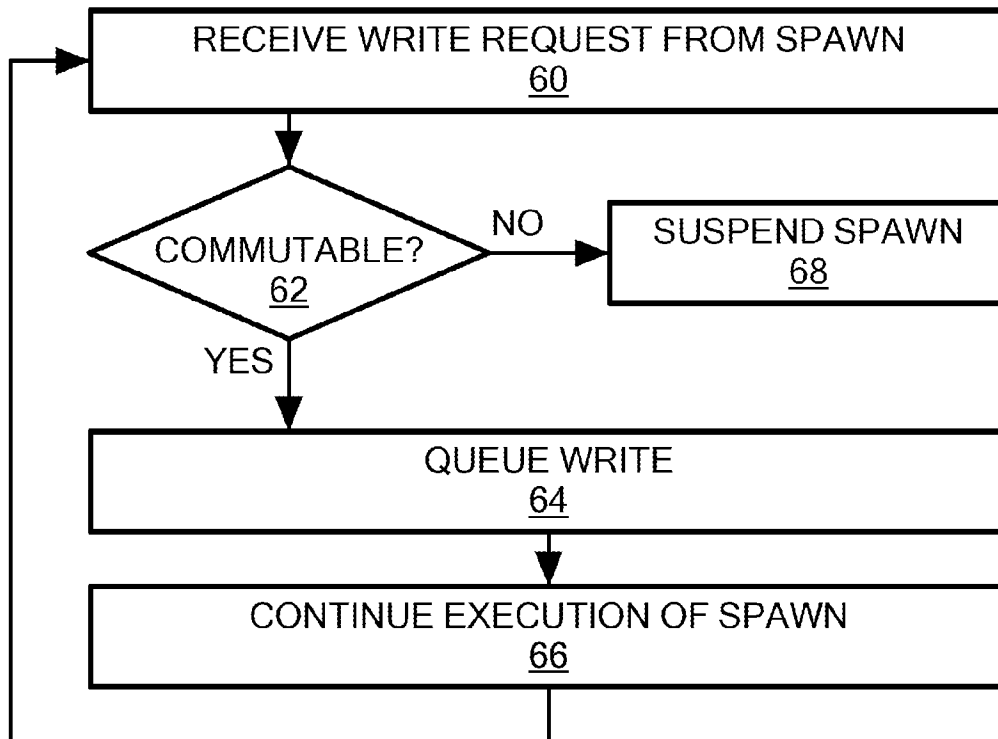


FIG. 5




FIG. 6


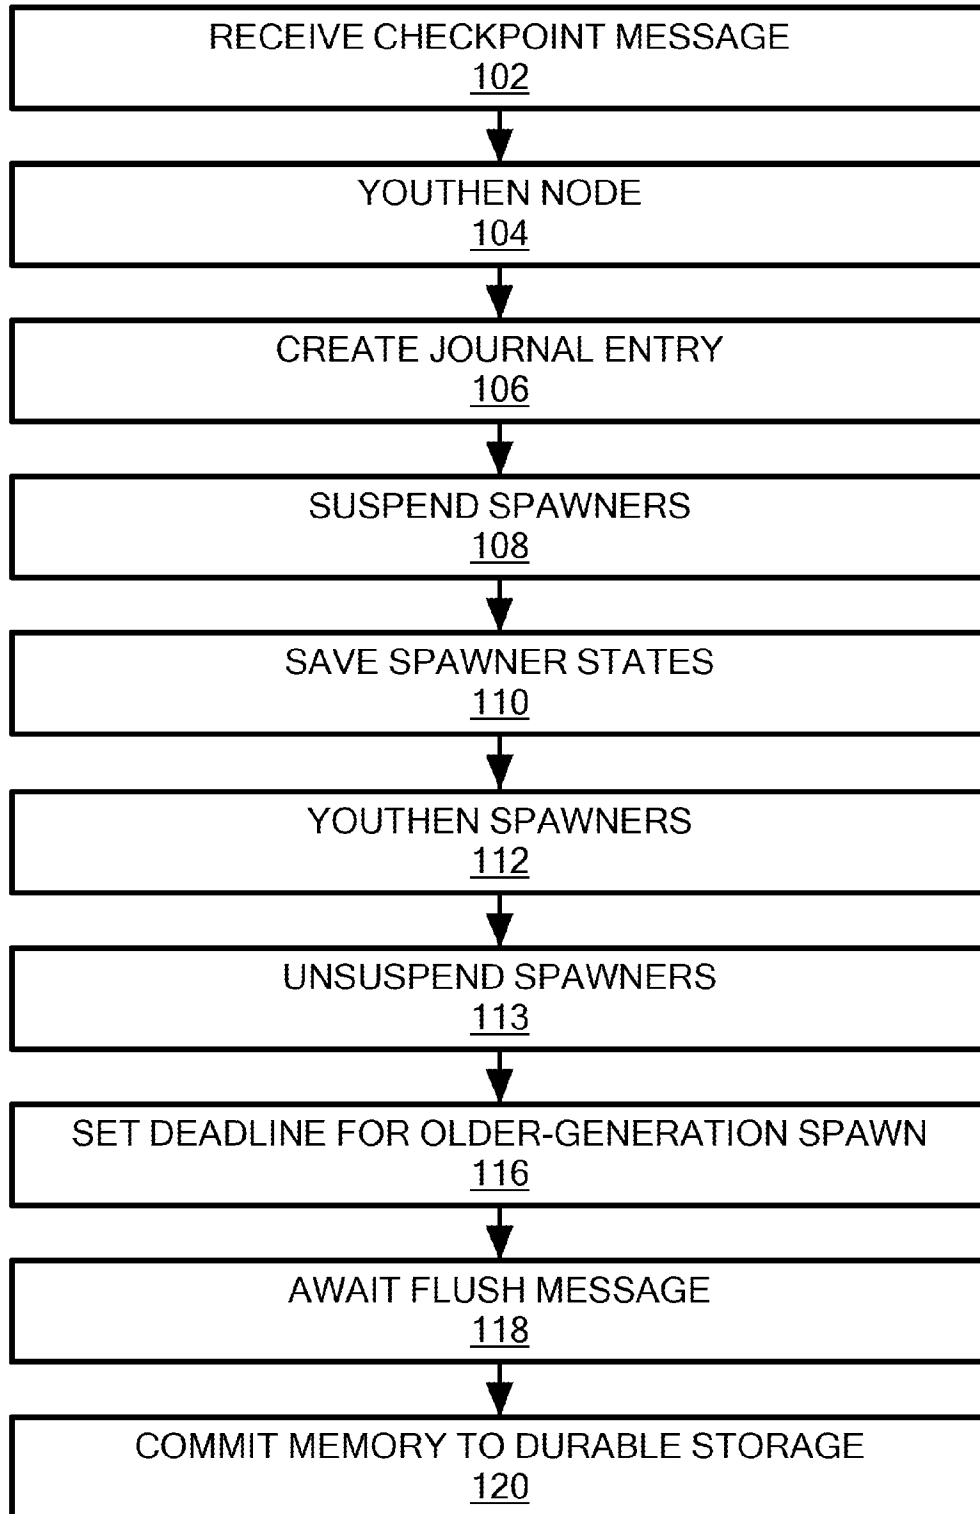



FIG. 7

**FIG. 8**



```
IF 82(79)<88(79)
  {82(79)++;
  JOURNALIZE(79)}
END IF
```

```
IF 82(79)<88(76)
  {82(79)++;
  JOURNALIZE(79)}
END IF
```

FIG. 9

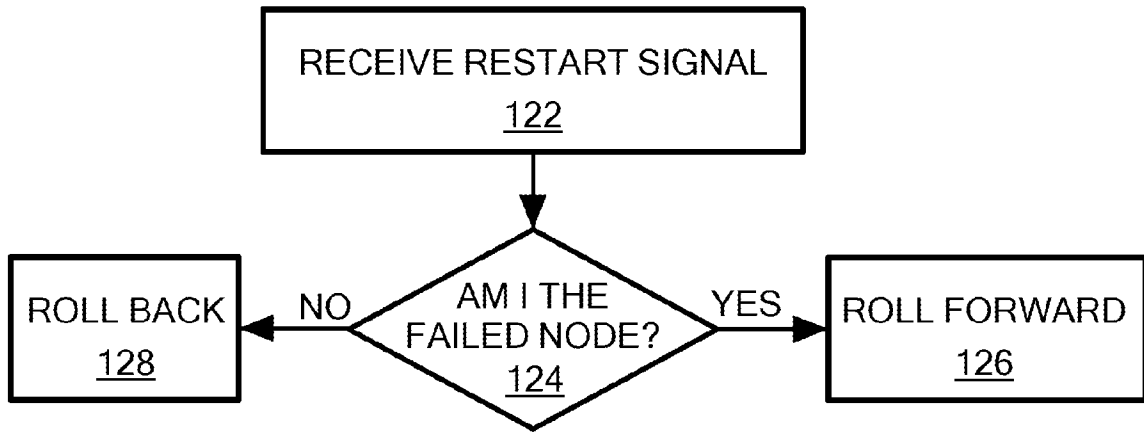


FIG. 10


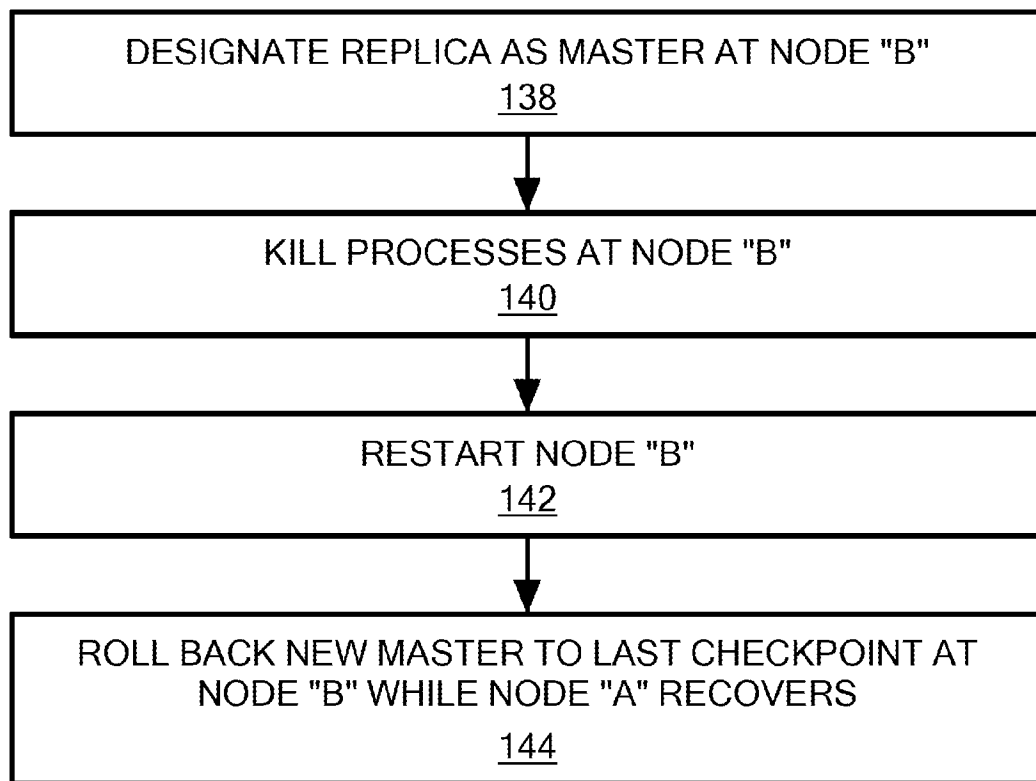


FIG. 11

***FIG. 12***

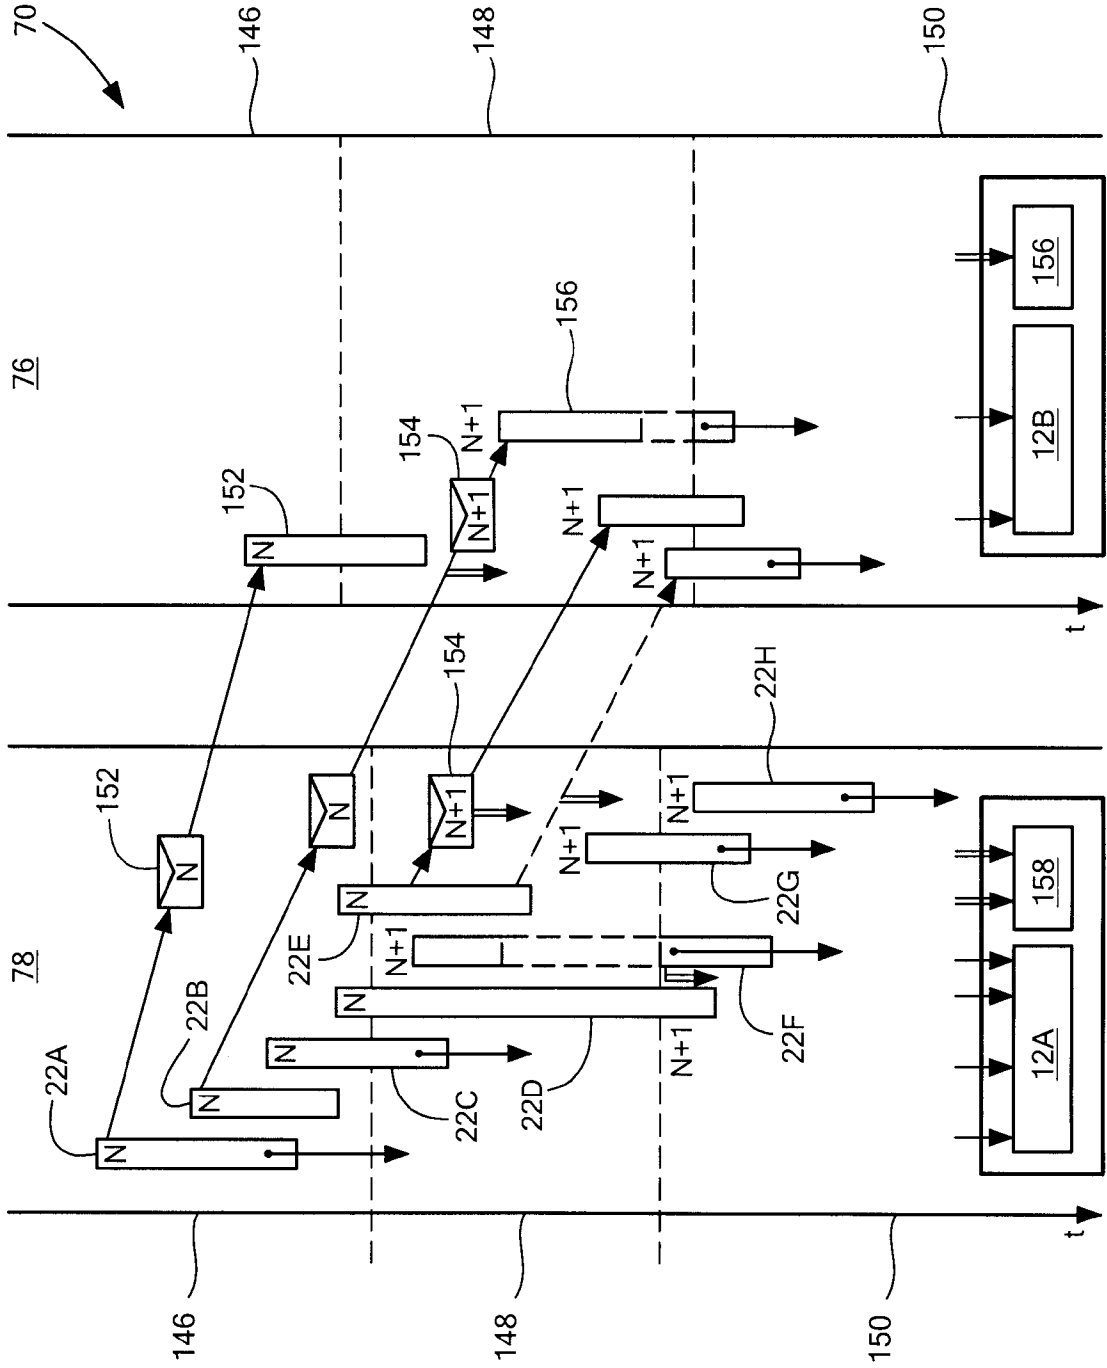


FIG. 13

8

