



- (51) International Patent Classification:
G06F 9/54 (2006.01) G06Q 10/06 (2012.01)
- (21) International Application Number:
PCT/US2014/044642
- (22) International Filing Date:
27 June 2014 (27.06.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
61/841,334 29 June 2013 (29.06.2013) US
14/145,127 31 December 2013 (31.12.2013) US
- (63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 14/145,127 (CON)
Filed on 31 December 2013 (31.12.2013)
- (71) Applicant: GOOGLE INC. [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

- (72) Inventors: FLEURY, Eduardo Madeira; Google Inc., 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). BANADAKI, Seyed Vahab Mirrokni; Google Inc., 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). HAJAJ, Nissan; Google Inc., 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). DING, Jerry Yi; Google Inc., 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). LATTANZI, Silvio; Google Inc., 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).
- (74) Agents: GROVER, Melanie et al.; Brake Hughes Beller-mann LLP, PO Box 52050, Minneapolis, Minnesota 55402 (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,

[Continued on next page]

(54) Title: ASYNCHRONOUS MESSAGE PASSING FOR LARGE GRAPH CLUSTERING

(57) Abstract: Systems and methods for sending asynchronous messages include receiving, using at least one processor, at a node in a distributed graph, a message with a first value and determining, at the node, that the first value replaces a current value for the node. In response to determining that the first value replaces the current value, the method also includes setting a status of the node to active and sending messages including the first value to neighboring nodes. The method may also include receiving the messages to the neighboring nodes at a priority queue. The priority queue propagates messages in an intelligently asynchronous manner, and the priority queue propagates the messages to the neighboring nodes, the status of the node is set to inactive. The first value may be a cluster identifier or a shortest path identifier.

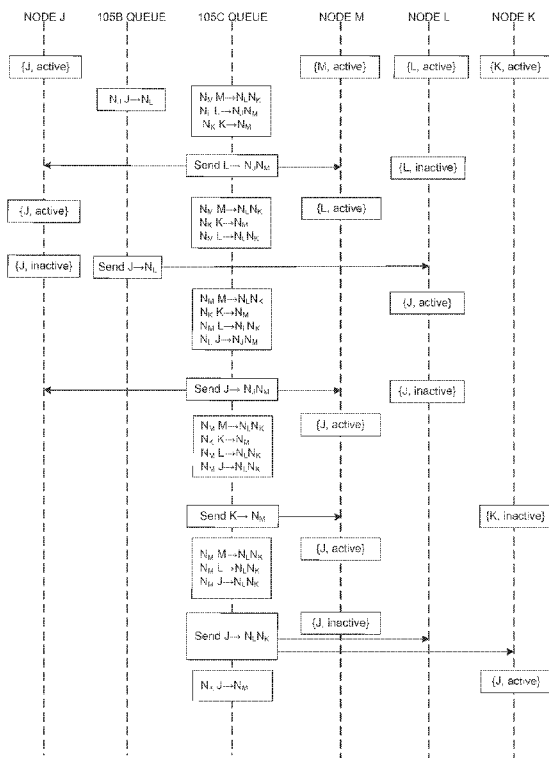


FIG. 4

WO 2014/210501 A1

KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, 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.

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).

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ,

Published:
— *with international search report (Art. 21(3))*

ASYNCHRONOUS MESSAGE PASSING FOR LARGE GRAPH CLUSTERING

RELATED APPLICATION

[0001] This application claims priority to, and is a continuation of, U.S. Nonprovisional Patent Application Serial No. 14/143,894, entitled “COMPUTING CONNECTED COMPONENTS IN LARGE GRAPHS,” filed on December 30, 2013, which claims priority to U.S. Provisional Patent Application Serial No. 61/841,334, entitled “ASYNCHRONOUS MESSAGE PASSING FOR LARGE GRAPH CLUSTERING” filed on June 29, 2013. This application also claims priority to U.S. Provisional Patent Application Serial No. 61/841,334, entitled “ASYNCHRONOUS MESSAGE PASSING FOR LARGE GRAPH CLUSTERING” filed on June 29, 2013. The subject matter of these earlier filed applications are hereby incorporated by reference.

BACKGROUND

[0002] Graphs are a basic modeling tool to model social, communication, and information networks. A graph $G(V, E)$ consists of a set of nodes V , and a set of edges $E \in V^2$ where each edge connects two nodes in the graph. In many applications, analysis is performed on large graphs that do not fit on one machine. Consequently, the graph is stored in several machines and mined in a distributed manner, for example by applying distributed programming tools like Map-Reduce or Hadoop. A basic analysis tool for graphs is to compute connected components of the graph. A connected component of a graph $G(V, E)$ is a maximal set of nodes that can be reached from each other via sequences of edges of the graph. Computing connected components of graph G results in a partitioning of the nodes V into one of several clusters, where each cluster is a connected component. For example, FIG. 2 illustrates a graph G with three connected components. Connected component 205 includes nodes A, B, C, and D, connected component 210 includes nodes F, G, I, and H, and connected component 215 includes nodes J, K, L, and M. The connected components may also be referred to as a cluster of nodes.

[0003] Computing connected components in graphs is a basic tool for computing coherent clusters of nodes and also to perform hierarchical clustering. But computing clusters of nodes distributed across multiple machines can be time and cost prohibitive as the running time of the hashing functions are dependent on the size of the graph, the number of messages sent between machines during the rounds of Map-Reduce, and the number of rounds of Map-Reduce performed. It is a challenge is to compute connected components for a large graph in a small number of rounds of Map-Reduce.

SUMMARY

[0004] Implementations provide an asynchronous message passing framework for computing connected components in a large distributed graph. The asynchronous message passing frameworks utilizes a priority queue to pass messages between nodes. Each node is able to perform its share of the computations and send messages regardless of the state of its peers, unlike synchronous message passing frameworks such as Map-Reduce and Hadoop. Using the priority queue allows the system to send blocks of requests between distributed computing devices for more efficient processing. The framework also uses a smaller memory footprint because each node tracks only a minimum identifier for a potential connected component, rather than a set of possible identifiers. This also makes message sizes smaller. Because messaging is asynchronous, each node may independently checkpoint its state and can restart after failure at the last checkpointed state. Recovering from failure is also more efficient because the number of messages that need to be replayed to recover a single machine is significantly less than synchronous message passing frameworks.

[0005] One aspect of the disclosure can be embodied in a system that includes distributed computing devices represented by leaf servers and memory storing a graph of nodes and edges, the graph being distributed across the leaf servers. A leaf server can include memory storing a cluster identifier for each node assigned to the leaf server, at least one processor, and memory storing instructions that, when executed by the at least one processor, cause the leaf server to send asynchronous messages between neighboring nodes, the messages comprising an asynchronous message from a first node to a second node including the cluster

identifier for the first node, wherein sending the asynchronous message is triggered by an update to the cluster identifier. The graph may include more than one billion nodes.

[0006] The system can include one or more of the following features, for example, the memory may further store instructions that, when executed by the at least one processor, cause the leaf server to propagate the message to the second node and compare the cluster identifier from the message with the cluster identifier for the second node to determine whether to update the cluster identifier for the second node. When it is determined that the cluster identifier of the second node is to be updated, the instructions may cause the leaf server to update the cluster identifier for the second node with the cluster identifier from the message and generate messages to neighboring nodes of the second node, the messages including the updated cluster identifier. In such implementations, the leaf server may set a status of the second node to active as part of the updating and set the status of the second node to inactive in response to propagation of the messages to neighboring nodes. In some such implementations the instructions may further cause the leaf server to store the updated cluster identifier in persistent memory.

[0007] As another example, the leaf server may also include a priority queue engine that propagates messages between neighboring nodes. In such an implementation, the priority queue engine may be capable of bundling together messages directed to nodes on another leaf server before propagation of the messages and/or of intelligently skipping redundant messages rather than propagating the messages. As another example, the cluster identifier may represent the smallest identifier seen by the first node and the cluster identifier for the second node is to be updated when the cluster identifier from the message is smaller than the cluster identifier for the second node. In another example, at least one leaf server includes a plurality of processors and the at least one leaf server uses the plurality of processors to concurrently send multiple messages departing from nodes and to receive multiple messages arriving at nodes.

[0008] Another aspect of the disclosure can be embodied in a method that includes propagating, using at least one processor, messages sent between nodes in a distributed graph in an asynchronous manner, the messages including respective cluster identifiers. In response to a first node of the distributed graph receiving one

of the messages, the method may also include comparing a cluster identifier from the received message with a cluster identifier for the first node to determine whether to update the cluster identifier for the first node and, when it is determined that the cluster identifier of the first node is to be updated, updating the cluster identifier for the first node with the cluster identifier from the message and generating messages to neighboring nodes of the first node, the messages including the updated cluster identifier.

[0009] The method can include one or more of the following features. For example, a priority queue engine may control the propagating. The priority queue engine may propagate the messages in an arbitrary manner rather than a first-in-first-out or last-in-last-out manner. In some implementations, nodes in the distributed graph are assigned to one of a plurality of leaf servers with each leaf server having a respective priority queue engine, the priority queue engine being one of the respective priority queue engines, and wherein the priority queue engine bundles messages directed to nodes assigned to a remote leaf server of the plurality of leaf servers prior to propagating the messages. Propagating the messages at a particular leaf server can continue despite a failure of another leaf server.

[00010] As another example, the method may further include storing the updated cluster identifier in persistent memory and, in response to the first node determining that it experienced a failure, obtaining the cluster identifier from the persistent memory and generating messages to the neighboring nodes of the first node, the messages requesting a cluster identifier from the respective neighboring nodes. In some implementations, the cluster identifier for the first node represents the smallest identifier seen by the first node and the cluster identifier for the first node is to be updated when the cluster identifier from the message is smaller than the cluster identifier for the first node.

[00011] As another example, the method may also include setting a status of the first node to active as part of the updating, and setting the status of the first node to inactive in response to propagating the messages to the neighboring nodes. In some implementations, the propagating is performed without regard to respective states of the nodes in the distributed graph.

[00012] In another aspect, a method includes receiving, using at least one processor, at a node in a distributed graph, a message with a first value and

determining, at the node, that the first value replaces a current value for the node. Responsive to the determining, the method also includes setting a status of the node to active and sending messages that include the first value to neighboring nodes. The method may also include receiving, using the at least one processor, the messages to the neighboring nodes at a priority queue, wherein the priority queue propagates messages in an intelligently asynchronous manner, and wherein when the message is propagated to the neighboring nodes, the status of the node is set to inactive.

[00013] The method may include one or more of the following features, for example, the node may be a first node and the method may further include receiving, at the first node, a request from a second node for messages sent after a first time and determining, at the first node, whether the messages to neighboring nodes were sent after the first time. When the messages were sent after the first time, the method includes re-sending the message to the second node and when the messages were not sent after the first time, the method includes ignoring the request. In some implementations the second node sends the request after recovering a current value for the second node from a checkpoint. As another example, the first value and the current value may be cluster identifiers or the first value and the current value may be shortest path identifiers.

[00014] Another aspect of the disclosure can be embodied on a computer-readable medium having recorded and embodied thereon instructions that, when executed by a processor of a computer system, cause the computer system to perform any of the methods disclosed herein.

[00015] One or more of the implementations of the subject matter described herein can be implemented so as to realize one or more of the following advantages. As one example, the system may reduce the time and cost of performing operations on a large distributed graph, such as determining connected components in a graph, computing the shortest path between two nodes, search, etc. Implementations not only reduce the number of messages sent, but also reduce the size of the messages. Also, because message passing is asynchronous, the processing time is reduced as nodes need not wait for other nodes to finish processing. The asynchronous nature also permits efficient recovery when a node or leaf stops responding temporarily, making the system robust under frequent failures. To confirm the fault-tolerance property of the system, an experiment was conducted in which a subset of machines

were made to artificially fail periodically over time and the increase in total running time and number of messages exchanged was recorded. Surprisingly, the running times increased only 10% to 20% with a 20% machine failure rate. Even with a 70% failure rate, the running time was multiplied by a factor of 2, demonstrating the high level of fault tolerance achieved with implementations. Connected component generation is a building block for graph clustering, which can be used in various fields, such as network analysis, vision and image processing, machine learning, knowledge discovery, search, etc.

[00016] The details of one or more implementations are set forth in the accompanying drawings and the description below. Other features will be apparent from the description and drawings, and from the claims.

BRIEF DESCRIPTION OF DRAWINGS

[00017] FIG. 1 illustrates an example system in accordance with the disclosed subject matter.

[00018] FIG. 2 illustrates an example graph distributed across three leaves and including three clusters.

[00019] FIG. 3 illustrates a flow diagram of an example process for computing connected components of a large distributed graph.

[00020] FIG. 4 illustrates an example of asynchronous message passing, in accordance with an implementation.

[00021] FIG. 5 shows an example of a computer device that can be used to implement the described techniques.

[00022] FIG. 6 shows an example of a distributed computer device that can be used to implement the described techniques.

[00023] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[00024] FIG. 1 is a block diagram of a distributed graph system 100 in accordance with an example implementation. The system 100 may be used to calculate connected components of a large distributed graph using the techniques described herein. The graph system 100 may include root 120 and graph cluster 160.

Root 120 and graph cluster 160 may be computing devices that take the form of a number of different devices, for example a standard server, a group of such servers, or a rack server system. In some implementations, the root 120 and the graph cluster 160 may be distributed systems implemented in a series of computing devices, such as a group of servers. In some implementations, the servers may be organized into a tree structure, with at least a root server 120 and leaf servers 150A to 150n. In some implementations (not shown), the tree may include intermediate servers, so that there are one or more layers between the root 120 and the leaf servers 150A to 150n. The root 120 and graph cluster 160 may be examples of computer device 600, as depicted in FIG. 6.

[00025] The graph system 100 illustrated in FIG. 1 operates over a large graph with, for example, billions of nodes. The root 120 may include one or more servers that operate with the graph cluster 160 to perform operations on the data graph represented by nodes and edges 154. The nodes and edges 154 may be stored in one or more data structures that represent each node of the graph and an adjacency list for each node. The data structure may also represent other information, for example the attributes of the edges that link a node to its neighbors in the adjacency list. The root 120 may include one or more servers that receive commands or requests from a requester, such as client 170. The root 120 may initiate and monitor calculations performed on the graph and may manage the results of the calculations. In some implementations, the root 120 may facilitate searches and queries on the graph. The root 120 may also store instructions in memory 144 that, when executed by the hardware processor 142, cause the system 100 to perform operations. For example, the root 120 and/or graph cluster 160 may include logic to process messages sent from one node to its neighbors.

[00026] System 100 may also include a graph cluster 160. Graph cluster 160 may be a collection of distributed computing devices each with its own hardware processor and memory. The number of computing devices that comprise graph cluster 160 can vary. The graph cluster 160 may be divided into one or more leaf servers, such as leaf 150A, leaf 150B, leaf 150n, with n representing any positive integer. A leaf server may be a logical division of nodes in the graph, with each graph node being assigned to a leaf server. Thus, a leaf server may correspond to one computing device, or a leaf server may be a logical computing device and may

share a physical computing device with other leaves. In some implementations, a node's assigned leaf may change as the graph is updated, making leaf assignments flexible. The root 120 may determine which nodes are assigned to each leaf as the nodes are added to the graph or updated.

[00027] The root 120 may route processing requests to the leaf servers and act as the primary means of coordination between the leaves at processing time. In some implementations, leaves may send messages directly to each other rather than going through the root 120. The messages may be messages from a node in the graph to its neighbors. In addition to the leaf servers 150, the graph cluster 160 may include one or more layers of intermediate servers between the root node 120 and the leaf servers 150, but are not shown in FIG. 1 for the sake of brevity. Thus, for example, an intermediate server may be associated with, for example, 20 leaf servers. The intermediate server may be connected directly to the root 120, or there may be one or more additional layers between the intermediate server and the root 120. Thus, although FIG. 1 shows communications directly between the root 120 and leaf servers 150, it is to be understood that intermediate devices may be used to direct communications and aggregate results using known methods, such as remote procedure calls. The root, intermediate, and leaf servers that make up the tree may, collectively, be referred to as the graph.

[00028] Each of the leaf servers 150 that make up graph cluster 160 can include node states 152 and nodes and edges 154. A state for a node may be used in calculating connected components for the graph and may include a status and a cluster identifier. The cluster identifier identifies the cluster, or potential connected component, that a node belongs to. During the calculation of connected components, the cluster identifier in the state of a node V may change, but it generally represents the smallest identifier that the node V has seen. This is similar to a Hash-Only-Min hashing algorithm. Alternatively, in some implementations the identifier may represent an identifier for the node with the greatest identifier received. At the conclusion of the computation, the state for a node will contain the identifier of the connected component to which the node belongs. The state of a node may not be altered by messages sent between nodes, only by the node itself or the priority queue engine 156. In addition to the identifier, each node state may include a status. The status may indicate whether the node is active or inactive. An active status indicates

that the identifier in the state has changed and the node needs to send a message with the new identifier to its neighbors. After the message has been propagated the state status may be changed to inactive. When all nodes in the graph are inactive, and all messages have been sent by the priority queue engine 156, the system has computed the connected components for the graph. At this point the state identifier for each node contains the identifier of the connected component to which the node belongs.

[00029] Each leaf server 150 may also include a priority queue engine 156. Messages sent between nodes may be handled by the priority queue engine 156. The priority queue engine 156 may track which messages have been sent and which still need to be sent, but does not need to send messages in the order at which they arrive. In some implementations, the priority queue engine 156 may arbitrarily choose which message to propagate next. In some implementations, the priority queue engine 156 may bundle messages destined for other leaf servers together, and send the bundled messages together. This makes messaging more efficient. Furthermore, if a leaf is not responding, the priority queue engine 156 can keep track of messages not successfully acknowledged so that these messages can be resent when the leaf begins responding again. In some implementations, when the priority queue engine 156 propagates the messages from a node it may set the status of the node to inactive.

[00030] Graph system 100 may be in communication with clients 170 over network 180. The network 180 may be for example, the Internet or the network 180 can be a wired or wireless local area network (LAN), wide area network (WAN), etc., implemented using, for example, gateway devices, bridges, switches, and/or so forth. Via the network 180, the system 100 may communicate with and transmit data to/from clients 170 or other computing devices.

[00031] FIG. 3 illustrates a flow diagram of an example process 300 for computing connected components of a large distributed graph. The process 300 shown in FIG. 3 may be performed by a graph system distributed across multiple computing devices, such as graph system 100 of FIG. 1. The process 300 may speed the calculation of connected components in a distributed graph by using an asynchronous message passing protocol to calculate the cluster each node belongs to. A root, such as root 120, may initiate process 300 to determine which cluster each node in the graph belongs to. Process 300 is described for one node, but it is

understood that the process 300 occurs for all nodes simultaneously. Furthermore, while FIG. 3 illustrates a cluster identifier being passed in the message, the message may pass other values, such as a shortest path identifier.

[00032] The process may begin with the system initializing the state of the node (305). The initial state for a node will be an active status and an identifier set equal to the identifier of the node. Because each node is active, it will send a message to each neighbor, the message including its state identifier (310). This message may be sent to the priority queue and distributed by the priority queue engine. Any known or later developed method of sending messages between nodes may be used. Once the message is propagated, the status of the node may change from active to inactive (315). In some implementations this can be done by the priority queue engine.

[00033] The node may then enter a waiting stage, represented by the dashed line in FIG. 3. The node is waiting for a message with a cluster identifier. If no message ever arrives, the node will remain in the waiting stage. When it receives such a message (320) it may compare the value, e.g., the cluster identifier, from the message with the value, e.g., the cluster identifier, in its state (325). If the message cluster identifier is greater than or equal to the identifier in the state (325, No), the node does nothing and continues waiting for another message. While waiting the node has an inactive status. If the message cluster identifier is less than the identifier in the state (325, Yes), the node may change its status to active and change the cluster identifier in its state to the message cluster identifier (330). Because the node's status is active, it sends a message to its neighbors, via the priority queue (335). The message contains the updated cluster identifier received in the message. When the message is propagated by the priority queue engine, the status may change back to inactive (340). The node then returns to a waiting stage, waiting for another message. The system may also checkpoint the states of the nodes at various points while nodes are performing process 300. Checkpointing the state may cause the system to store the states for the nodes in persistent memory so that if the leaf fails or needs to restart for any reason, the node can determine the most recent state, and begin from that point. Different implementations may use different criteria to decide when to write checkpoints, as this can be a relatively expensive procedure. For example, each leaf server may checkpoint the states of the nodes assigned to the leaf

after a predetermined number of messages have been sent or propagated. As another example, the leaf may checkpoint the states after a predetermined amount of time. Other checkpointing criteria may also be used.

[00034] If a node or a leaf server fails (e.g., does not respond), the connected component computation process need not restart. Because the messaging is asynchronous and nodes can independently checkpoint their state, each node affected by the failure may pick up where it left off. For example, after a fault a node may initialize its state to the last checkpointed state for the node (345). This checkpointed state may be stored in persistent memory, as described above. As part of the initialization process, the node may use information in the checkpoint to determine whether to set its status to active or inactive. The node may then request from its neighbors any messages sent since the time of the checkpoint (350). The node may then set its status to inactive (355) and wait for messages, as indicated by the dashed line. The neighbors that receive the recovery request may ignore the request if they have no information to send. For example the neighbor node may determine whether it sent a message since the time of the checkpoint. If it has not, the request may be ignored. If the neighbor node did send a message, and the neighbor node is inactive, the neighbor node may send its cluster identifier to the node that sent the recovery request. If the neighbor nodes status is active, the priority queue engine already includes a message for the node that sent the recovery message, so another message does not need to be sent. Once the node that sent the recovery request begins receiving messages, it may process the messages as indicated above, with regard to steps 320 through 340.

[00035] The priority queue engine may be propagating messages arbitrarily as it receives them. When the priority queue has no other messages to send, it may check the state of the nodes. If all nodes are inactive, the leaf may report that its nodes are inactive. Of course, if the priority queue receives a message from another leaf, it may report that its nodes are once again active (e.g., no longer inactive). A root node may track the status of each leaf in the distributed system. When each leaf is inactive, the clusters are stable and the system has computed the connected components for the graph. Process 300 has thus ended for all nodes in the graph and the value of the state cluster identifier for each node indicates which connected component in the graph the node belongs to. The system may store the cluster

identifier from the node's state as an attribute of the node. The cluster identifier or other value sent via message can be used in various ways for additional graph processing such as suggesting new connections in a social networking graph, load distribution for traffic in a computer network graph, vision or image processing, machine learning, responding to search queries, knowledge discovery, etc. Additionally, although the process 300 is illustrated as sending a lowest identifier, it is understood that the asynchronous messaging techniques may be used with a highest identifier with appropriate modifications.

[00036] FIG. 4 illustrates an example of asynchronous messaging using cluster 215 of FIG. 2. While cluster 215 is used an example for brevity, it is understood that the computation of connected components would occur for all nodes of the graph simultaneously, and not just for the nodes of cluster 215. It is also understood that the representation of the states below uses alphanumeric values, e.g., {value, status}, for ease of explanation and the actual data structure implementation may use other representations that are more memory conscious. Upon initialization of the state, the nodes in cluster 215 may begin with the following states (S_v):

$$S_J = \{J, \text{active}\}$$

$$S_K = \{K, \text{active}\}$$

$$S_L = \{L, \text{active}\}$$

$$S_M = \{M, \text{active}\}$$

[00037] The identifier in each state may be propagated to its neighbors. Thus, Node J sends J to Node L, Node L sends L to Node J and Node M, Node M sends M to Node L and Node K, and Node K sends K to Node M. The messages can be sent simultaneously and independently from each other. Node J is on a different leaf than Nodes K, L, and M. Thus, the priority queue of Leaf 150B may delay sending Node J's message to Node L so it can bundle the message with other messages directed to Leaf 150C. For example, Node B may also be sending messages to Nodes D and C. These messages may be bundled with Node J's message to Node L and delivered together to the priority queue engine of Leaf 150C, thus making more efficient use of network resources.

[00038] The priority queue engine on Leaf 150C may choose to propagate Node L's messages to Nodes M and J first, leaving the messages of M to Node L, M to Node K, and K to Node M in the queue. The priority queue engine may also set

the status of L to inactive. Node M receives L's message and determines that L is less than M. Thus, Node M changes its cluster identifier to L and marks itself as active. Node M then sends a message to Node L and Node K with the identifier of L. Node J also receives Node L's message, but J is less than L, so Node J takes no action.

[00039] In the meantime, the priority queue of Leaf 150B may propagate Node J's message to Node L to the priority queue of Leaf 150C, causing Node L to receive the identifier J. Because J is less than L, Node L changes its status to active, and sends a message to Node J and Node M with J as an identifier. Thus, the priority queue on Leaf 150C may have messages of M to Node L, M to Node K, K to Node M, L to Node L, L to Node K, J to Node J and J to Node M in the queue. The queue engine may propagate the messages J to Node M and J to Node J next, setting L's status to inactive. Because J is lower than L Node M's current identifier, Node M changes its status to active, changes its identifier to J and sends a message to Node L and Node K with identifier J. The priority queue engine may propagate the message K to Node M next, setting K's status to inactive. Node M does nothing because K is not less than J, M's current status. Although not discussed above in the example, the nodes may also checkpoint their state at appropriate times.

[00040] As demonstrated, the order of message propagation can be random without affecting the outcome of the computation. Because it is safe to change the order of message propagation, the priority queue engine may intelligently choose which messages to send first, so that the total number of messages sent by the system is reduced. As one example, if the system uses a smallest identifier, the priority queue engine may propagate messages that include smaller identifiers and skip messages with higher identifiers to the same nodes because the latter would be redundant. Accordingly, using the example above the queue includes three messages to Nodes K and L, one with J, one with M, and one with L. Because J is lower than both M and L, the priority queue engine may propagate the message with J but skip the messages with M and L. When the priority queue engine skips the messages sending M and L it may still change the status of Node M, which sent the message to Nodes L and K, to inactive.

[00041] Furthermore, the asynchronous message passing reduces contention and increases parallelism. The tradeoff for the ability to perform asynchronous

message passing is that messages sent between two nodes must be idempotent, e.g., not alter the state of the source nor target nodes if the framework decides to resent them one or more times, and the messages must supersede all previous messages sent between the source and target nodes. In other words, the system is allowed to suppress any but the last message sent between two nodes without changing the behavior of the implemented algorithm.

[00042] The processes described above may be used with additional parameters. For example, in considering what constitutes a neighbor node, the system may consider only edges in the graph that meet certain criteria. This will, of course, result in a higher number of clusters, or connected components, for the graph. Other such parameters may be used in conjunction with this disclosure.

[00043] The asynchronous message passing framework described above may also be used with other techniques of graph analysis beyond computing connected components. For example, the asynchronous message passing framework can be used to calculate a shortest path between nodes using the priority queue engine and checkpointing. In such implementations the information included in the message sent between nodes and the operations performed at the nodes may differ from those discussed above, but the role of the priority queue engine and the fault tolerant techniques are similar.

[00044] FIG. 5 shows an example of a generic computer device 500, which may be system 100 or client 170 of FIG. 1, which may be used with the techniques described here. Computing device 500 is intended to represent various example forms of computing devices, such as laptops, desktops, workstations, personal digital assistants, cellular telephones, smart phones, tablets, servers, and other computing devices, including wearable devices. The components shown here, their connections and relationships, and their functions, are meant to be exemplary only, and are not meant to limit implementations of the inventions described and/or claimed in this document.

[00045] Computing device 500 includes a processor 502, memory 504, a storage device 506, and expansion ports 510 connected via an interface 508. In some implementations, computing device 500 may include transceiver 546, communication interface 544, and a GPS (Global Positioning System) receiver module 548, among other components, connected via interface 508. Device 500 may

communicate wirelessly through communication interface 544, which may include digital signal processing circuitry where necessary. Each of the components 502, 504, 506, 508, 510, 540, 544, 546, and 548 may be mounted on a common motherboard or in other manners as appropriate.

[00046] The processor 502 can process instructions for execution within the computing device 500, including instructions stored in the memory 504 or on the storage device 506 to display graphical information for a GUI on an external input/output device, such as display 516. Display 516 may be a monitor or a flat touchscreen display. In some implementations, multiple processors and/or multiple buses may be used, as appropriate, along with multiple memories and types of memory. Also, multiple computing devices 500 may be connected, with each device providing portions of the necessary operations (e.g., as a server bank, a group of blade servers, or a multi-processor system).

[00047] The memory 504 stores information within the computing device 500. In one implementation, the memory 504 is a volatile memory unit or units. In another implementation, the memory 504 is a non-volatile memory unit or units. The memory 504 may also be another form of computer-readable medium, such as a magnetic or optical disk. In some implementations, the memory 504 may include expansion memory provided through an expansion interface.

[00048] The storage device 506 is capable of providing mass storage for the computing device 500. In one implementation, the storage device 506 may be or contain a computer-readable medium, such as a floppy disk device, a hard disk device, an optical disk device, or a tape device, a flash memory or other similar solid state memory device, or an array of devices, including devices in a storage area network or other configurations. A computer program product can be tangibly embodied in such a computer-readable medium. The computer program product may also contain instructions that, when executed, perform one or more methods, such as those described above. The computer- or machine-readable medium is a storage device such as the memory 504, the storage device 506, or memory on processor 502.

[00049] The interface 508 may be a high speed controller that manages bandwidth-intensive operations for the computing device 500 or a low speed controller that manages lower bandwidth-intensive operations, or a combination of

such controllers. An external interface 540 may be provided so as to enable near area communication of device 500 with other devices. In some implementations, controller 508 may be coupled to storage device 506 and expansion port 514. The expansion port, which may include various communication ports (e.g., USB, Bluetooth, Ethernet, wireless Ethernet) may be coupled to one or more input/output devices, such as a keyboard, a pointing device, a scanner, or a networking device such as a switch or router, e.g., through a network adapter.

[00050] The computing device 500 may be implemented in a number of different forms, as shown in the figure. For example, it may be implemented as a standard server 530, or multiple times in a group of such servers. It may also be implemented as part of a rack server system. In addition, it may be implemented in a personal computer such as a laptop computer 522, or smart phone 536. An entire system may be made up of multiple computing devices 500 communicating with each other. Other configurations are possible.

[00051] FIG. 6 shows an example of a generic computer device 600, which may be system 100 of FIG. 1, which may be used with the techniques described here. Computing device 600 is intended to represent various example forms of large-scale data processing devices, such as servers, blade servers, datacenters, mainframes, and other large-scale computing devices. Computing device 600 may be a distributed system having multiple processors, possibly including network attached storage nodes, that are interconnected by one or more communication networks. The components shown here, their connections and relationships, and their functions, are meant to be exemplary only, and are not meant to limit implementations of the inventions described and/or claimed in this document.

[00052] Distributed computing system 600 may include any number of computing devices 680. Computing devices 680 may include a server or rack servers, mainframes, etc. communicating over a local or wide-area network, dedicated optical links, modems, bridges, routers, switches, wired or wireless networks, etc.

[00053] In some implementations, each computing device may include multiple racks. For example, computing device 680a includes multiple racks 658a – 658n. Each rack may include one or more processors, such as processors 652a-652n and 662a-662n. The processors may include data processors, network attached

storage devices, and other computer controlled devices. In some implementations, one processor may operate as a master processor and control the scheduling and data distribution tasks. Processors may be interconnected through one or more rack switches 658, and one or more racks may be connected through switch 678. Switch 678 may handle communications between multiple connected computing devices 600.

[00054] Each rack may include memory, such as memory 654 and memory 664, and storage, such as 656 and 666. Storage 656 and 666 may provide mass storage and may include volatile or non-volatile storage, such as network-attached disks, floppy disks, hard disks, optical disks, tapes, flash memory or other similar solid state memory devices, or an array of devices, including devices in a storage area network or other configurations. Storage 656 or 666 may be shared between multiple processors, multiple racks, or multiple computing devices and may include a computer-readable medium storing instructions executable by one or more of the processors. Memory 654 and 664 may include, e.g., volatile memory unit or units, a non-volatile memory unit or units, and/or other forms of computer-readable media, such as a magnetic or optical disks, flash memory, cache, Random Access Memory (RAM), Read Only Memory (ROM), and combinations thereof. Memory, such as memory 654 may also be shared between processors 652a-652n. Data structures, such as an index, may be stored, for example, across storage 656 and memory 654. Computing device 600 may include other components not shown, such as controllers, buses, input/output devices, communications modules, etc.

[00055] An entire system, such as system 100, may be made up of multiple computing devices 600 communicating with each other. For example, device 680a may communicate with devices 680b, 680c, and 680d, and these may collectively be known as system 100. As another example, system 100 of FIG. 1 may include one or more computing devices 600 as root 120 and one or more of leaves 150, a separate computing device 600 as root 120, and one or more computing devices 600 as graph cluster 160. Furthermore, some of the computing devices may be located geographically close to each other, and others may be located geographically distant. The layout of system 600 is an example only and the system may take on other layouts or configurations.

[00056] Various implementations can include implementation in one or more computer programs that are executable and/or interpretable on a programmable system including at least one programmable processor, which may be special or general purpose, coupled to receive data and instructions from, and to transmit data and instructions to, a storage system, at least one input device, and at least one output device.

[00057] These computer programs (also known as programs, software, software applications or code) include machine instructions for a programmable processor, and can be implemented in a high-level procedural and/or object-oriented programming language, and/or in assembly/machine language. As used herein, the terms “machine-readable medium” “computer-readable medium” refers to any non-transitory computer program product, apparatus and/or device (e.g., magnetic discs, optical disks, memory (including Read Access Memory), Programmable Logic Devices (PLDs)) used to provide machine instructions and/or data to a programmable processor but not to transitory signals.

[00058] The systems and techniques described here can be implemented in a computing system that includes a back end component (e.g., as a data server), or that includes a middleware component (e.g., an application server), or that includes a front end component (e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the systems and techniques described here), or any combination of such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication (e.g., a communication network). Examples of communication networks include a local area network (“LAN”), a wide area network (“WAN”), and the Internet.

[00059] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[00060] A number of implementations have been described. Nevertheless, various modifications may be made without departing from the spirit and scope of the invention. In addition, the logic flows depicted in the figures do not require the

particular order shown, or sequential order, to achieve desirable results. In addition, other steps may be provided, or steps may be eliminated, from the described flows, and other components may be added to, or removed from, the described systems. Accordingly, other implementations are within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A system comprising:
distributed computing devices represented by leaf servers; and
memory storing a graph of nodes and edges, the graph being distributed
across the leaf servers,
wherein a leaf server includes:
memory storing a cluster identifier for each node assigned to the leaf
server,
at least one processor, and
memory storing instructions that, when executed by the at least one
processor, cause the leaf server to send asynchronous messages
between neighboring nodes, the messages including the cluster
identifier for the first node, wherein sending the asynchronous message
is triggered by an update to the cluster identifier.
2. The system of claim 1 wherein the memory further stores instructions that,
when executed by the at least one processor, cause the leaf server to:
propagate the message to the second node;
compare the cluster identifier from the message with the cluster identifier for
the second node to determine whether to update the cluster identifier for
the second node; and
when it is determined that the cluster identifier of the second node is to be
updated:
update the cluster identifier for the second node with the cluster identifier
from the message, and
generate messages to neighboring nodes of the second node, the messages
including the updated cluster identifier.
3. The system of claim 2 wherein memory further stores instructions that, when
executed by the at least one processor, cause the leaf server to:
set a status of the second node to active as part of the updating.

4. The system of claim 3 wherein the memory further stores instructions that, when executed by the at least one processor, cause the leaf server to: set the status of the second node to inactive in response to propagation of the messages to neighboring nodes.
5. The system of any of claims claim 2 through 4, wherein the memory further stores instructions that, when executed by the at least one processor, cause the leaf server to: store the updated cluster identifier in persistent memory.
6. The system of any of claims 2 through 5, wherein the leaf server includes a priority queue engine that propagates messages between neighboring nodes.
7. The system of claim 6 wherein the priority queue engine is capable of bundling together messages directed to nodes on another leaf server before the propagation of the messages.
8. The system of claim 6, wherein the priority queue engine is capable of intelligently skipping redundant messages rather than propagating the redundant messages.
9. The system of any of claims 2 through 8, wherein the cluster identifier represents the smallest identifier seen by the first node and the cluster identifier for the second node is to be updated when the cluster identifier from the message is smaller than the cluster identifier for the second node.
10. The system of any of claims 1 through 9 wherein the graph includes more than one billion nodes.
11. The system of any of claims 1 through 10, wherein at least one leaf server includes a plurality of processors and the at least one leaf server uses the plurality of processors to concurrently send multiple messages departing from nodes and to receive multiple messages arriving at nodes.

12. A computer-implemented method comprising:
propagating, using at least one processor, messages sent between nodes in a distributed graph in an asynchronous manner, the messages including respective cluster identifiers; and
in response to a first node of the distributed graph receiving one of the messages:
comparing, using the at least one processor, a cluster identifier from the received message with a cluster identifier for the first node to determine whether to update the cluster identifier for the first node, and
when it is determined that the cluster identifier of the first node is to be updated:
updating the cluster identifier for the first node with the cluster identifier from the message, and
generating messages to neighboring nodes of the first node, the messages including the updated cluster identifier.
13. The method of claim 12, wherein a priority queue engine controls the propagating.
14. The method of claim 13, wherein the priority queue engine propagates the messages in an arbitrary manner rather than a first-in-first-out or last-in-last-out manner.
15. The method of any of claims 13 through 14, wherein nodes in the distributed graph are assigned to one of a plurality of leaf servers with each leaf server having a respective priority queue engine, the priority queue engine being one of the respective priority queue engines, and wherein the priority queue engine bundles messages directed to nodes assigned to a remote leaf server of the plurality of leaf servers prior to propagating the messages.
16. The method of claim 15, wherein propagating the messages at a particular leaf server continues despite a failure of another leaf server.

17. The method of any of claims 12 through 16, further comprising storing the updated cluster identifier in persistent memory.
18. The method of claim 17, further comprising, in response to the first node determining that it experienced a failure:
obtaining the cluster identifier from the persistent memory; and
generating messages to the neighboring nodes of the first node, the messages requesting a cluster identifier from the respective neighboring nodes.
19. The method of claim 18 wherein the cluster identifier for the first node represents the smallest identifier seen by the first node and the cluster identifier for the first node is to be updated when the cluster identifier from the message is smaller than the cluster identifier for the first node.
20. The method of any of claims 12 through 19, further comprising:
setting a status of the first node to active as part of the updating.
21. The method of claim 20, further comprising:
setting the status of the first node to inactive in response to propagation of the messages to the neighboring nodes.
22. The method of any of claims 12 through 21, wherein the propagating is performed without regard to respective states of the nodes in the distributed graph.
23. A computer-implemented method comprising:
receiving, using at least one processor, at a node in a distributed graph, a message with a first value;
determining, at the node, that the first value replaces a current value for the node;
responsive to the determining, setting a status of the node to active and sending messages that include the first value to neighboring nodes;
receiving, using the at least one processor, the messages to the neighboring nodes at a priority queue, wherein the priority queue propagates messages

in an intelligently asynchronous manner, and wherein when the message is propagated to the neighboring nodes, the status of the node is set to inactive.

24. The method of claim 23, wherein the node is a first node and the method further comprises:
receiving, at the first node, a request from a second node for messages sent after a first time;
determining, at the first node, whether the messages to neighboring nodes were sent after the first time;
when the messages were sent after the first time, re-sending the message to the second node; and
when the messages were not sent after the first time, ignoring the request.
25. The method of claim 24, wherein the second node sends the request after recovering a current value for the second node from a checkpoint.
26. The method of claim 23, wherein the first value and the current value are cluster identifiers.
27. The method of claim 23, wherein the first value and the current value are shortest path identifiers.

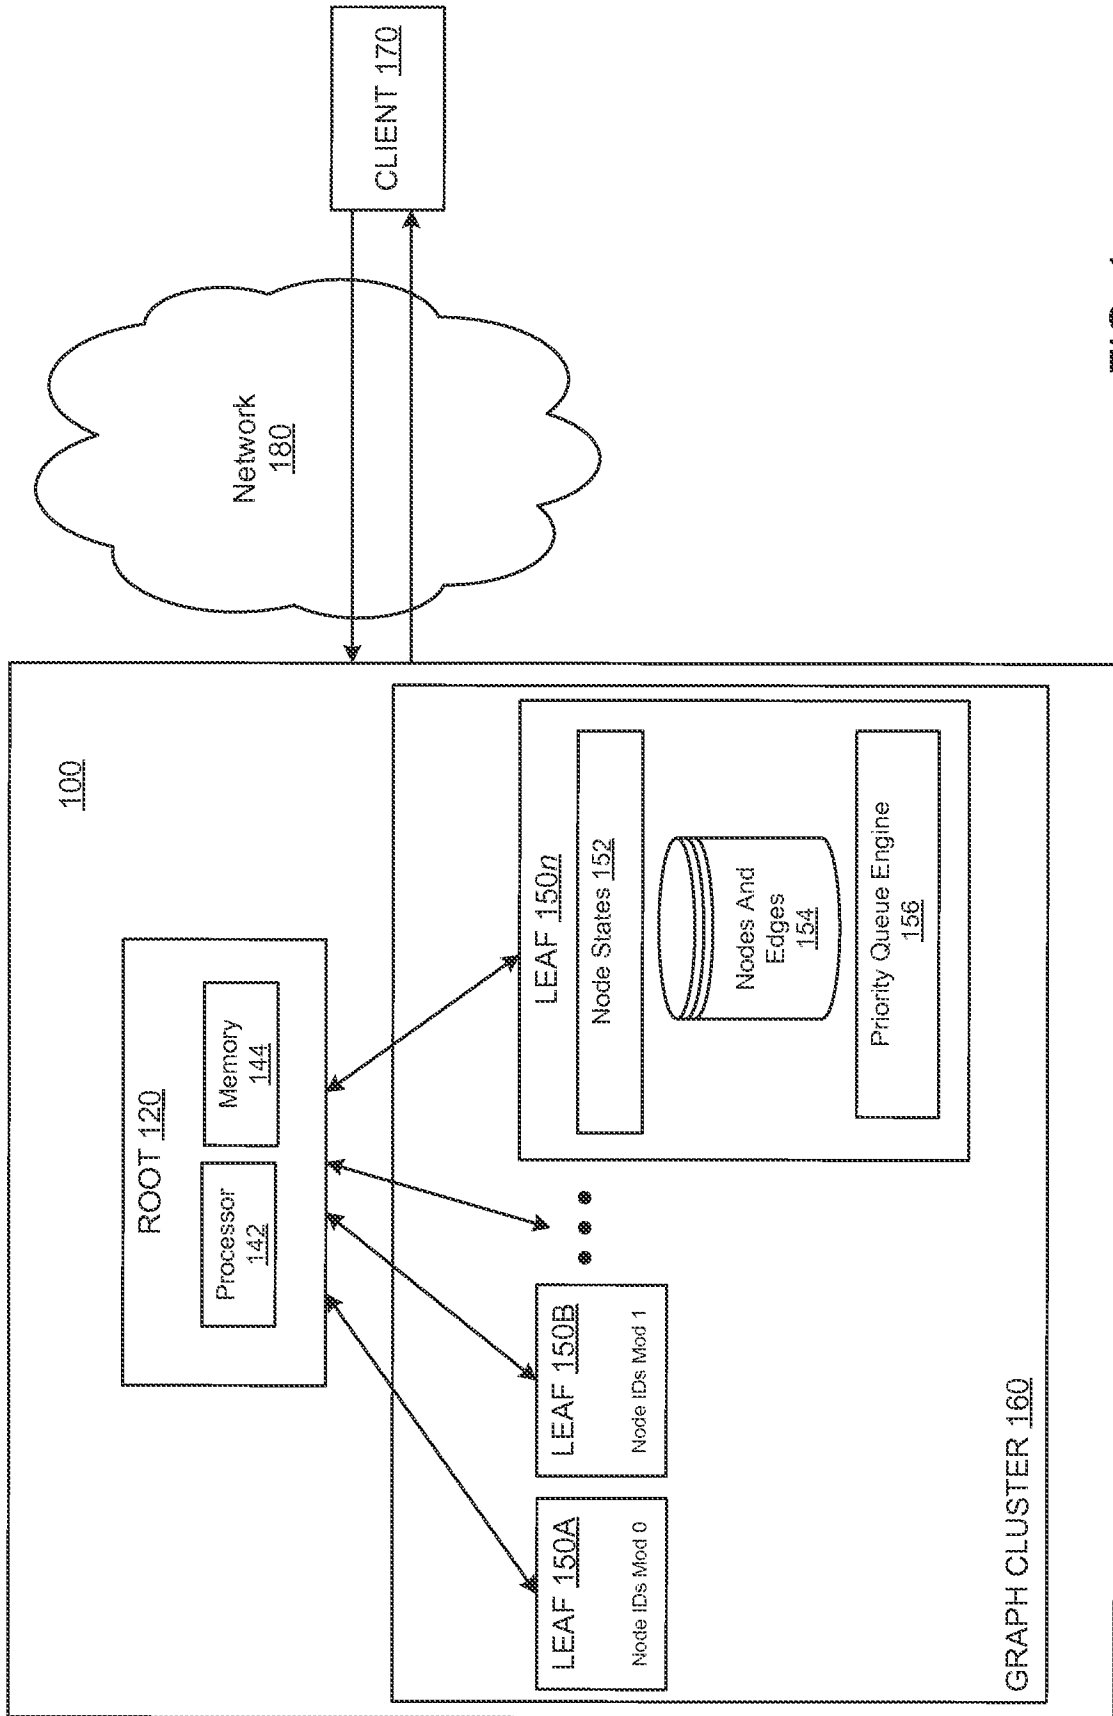
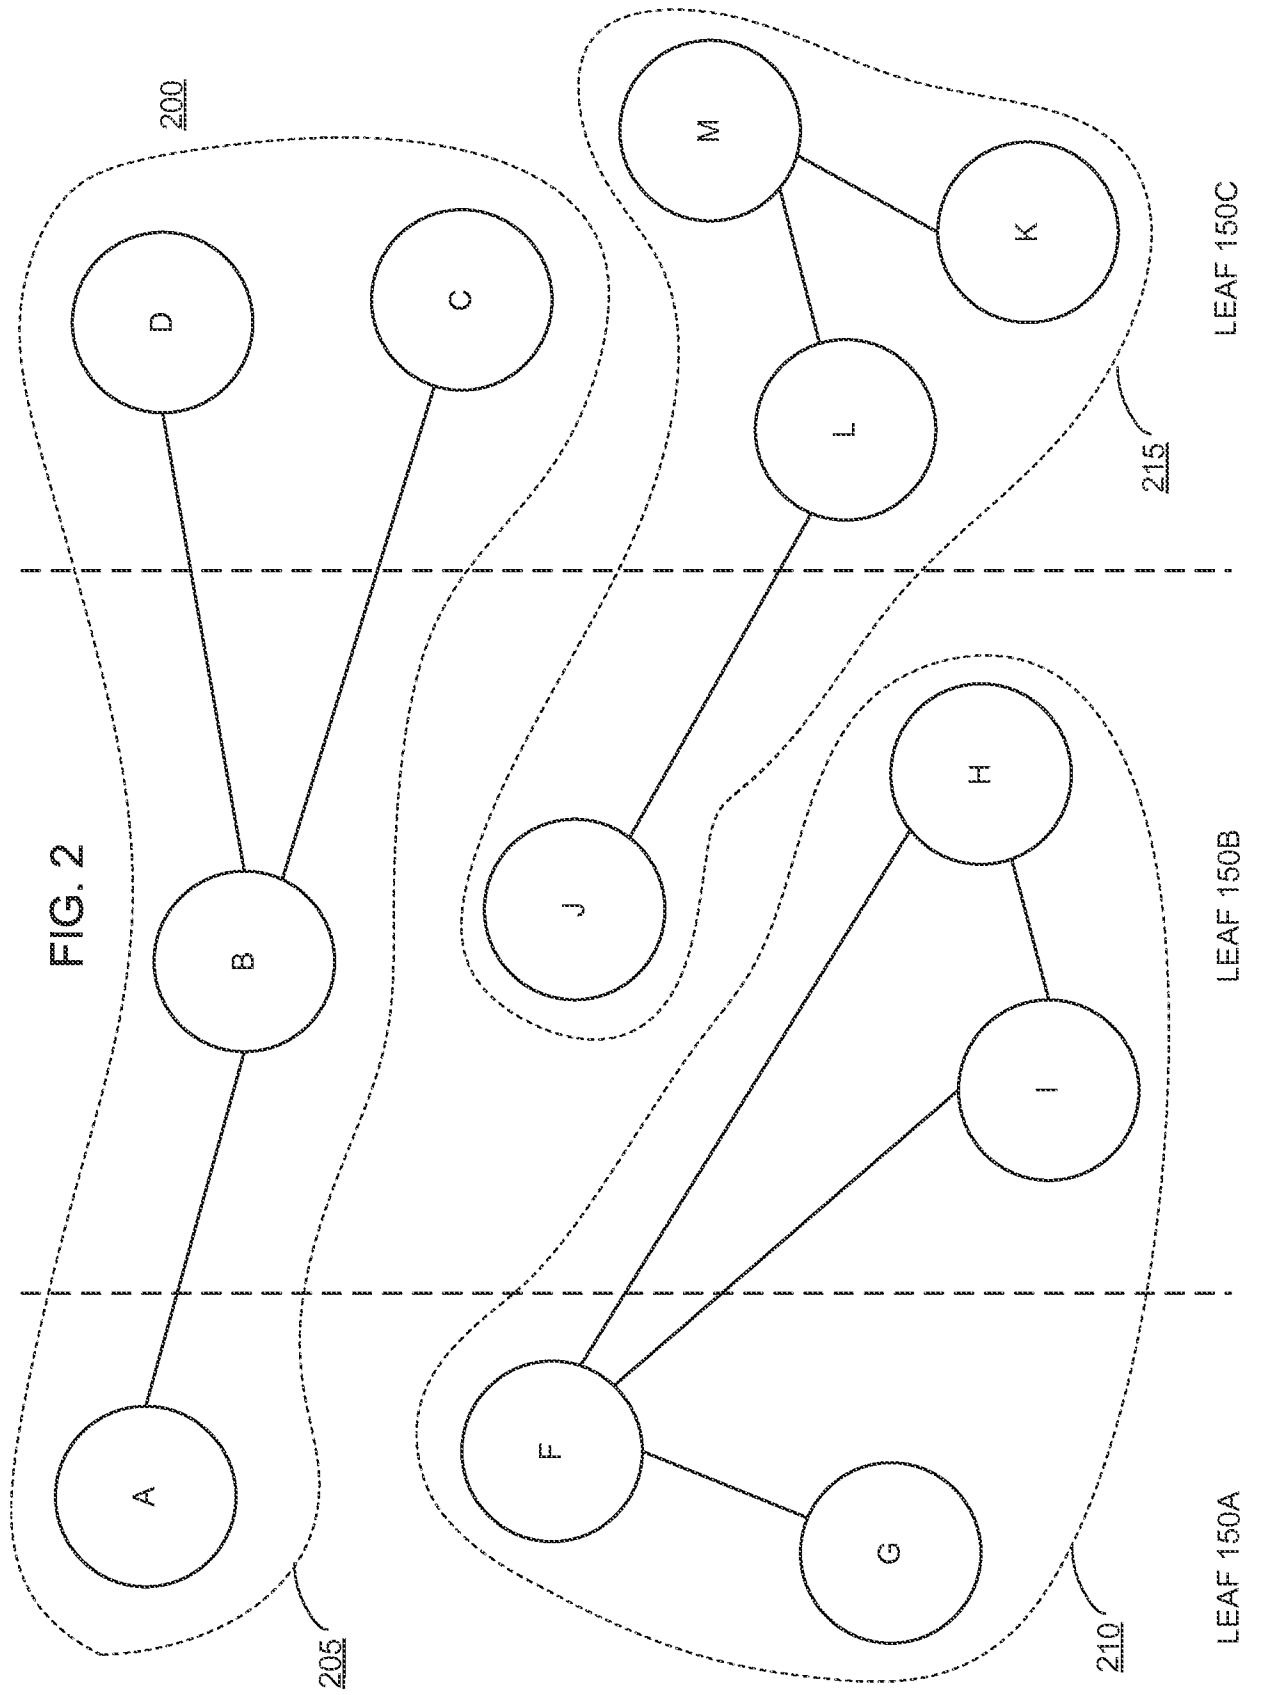


FIG. 1



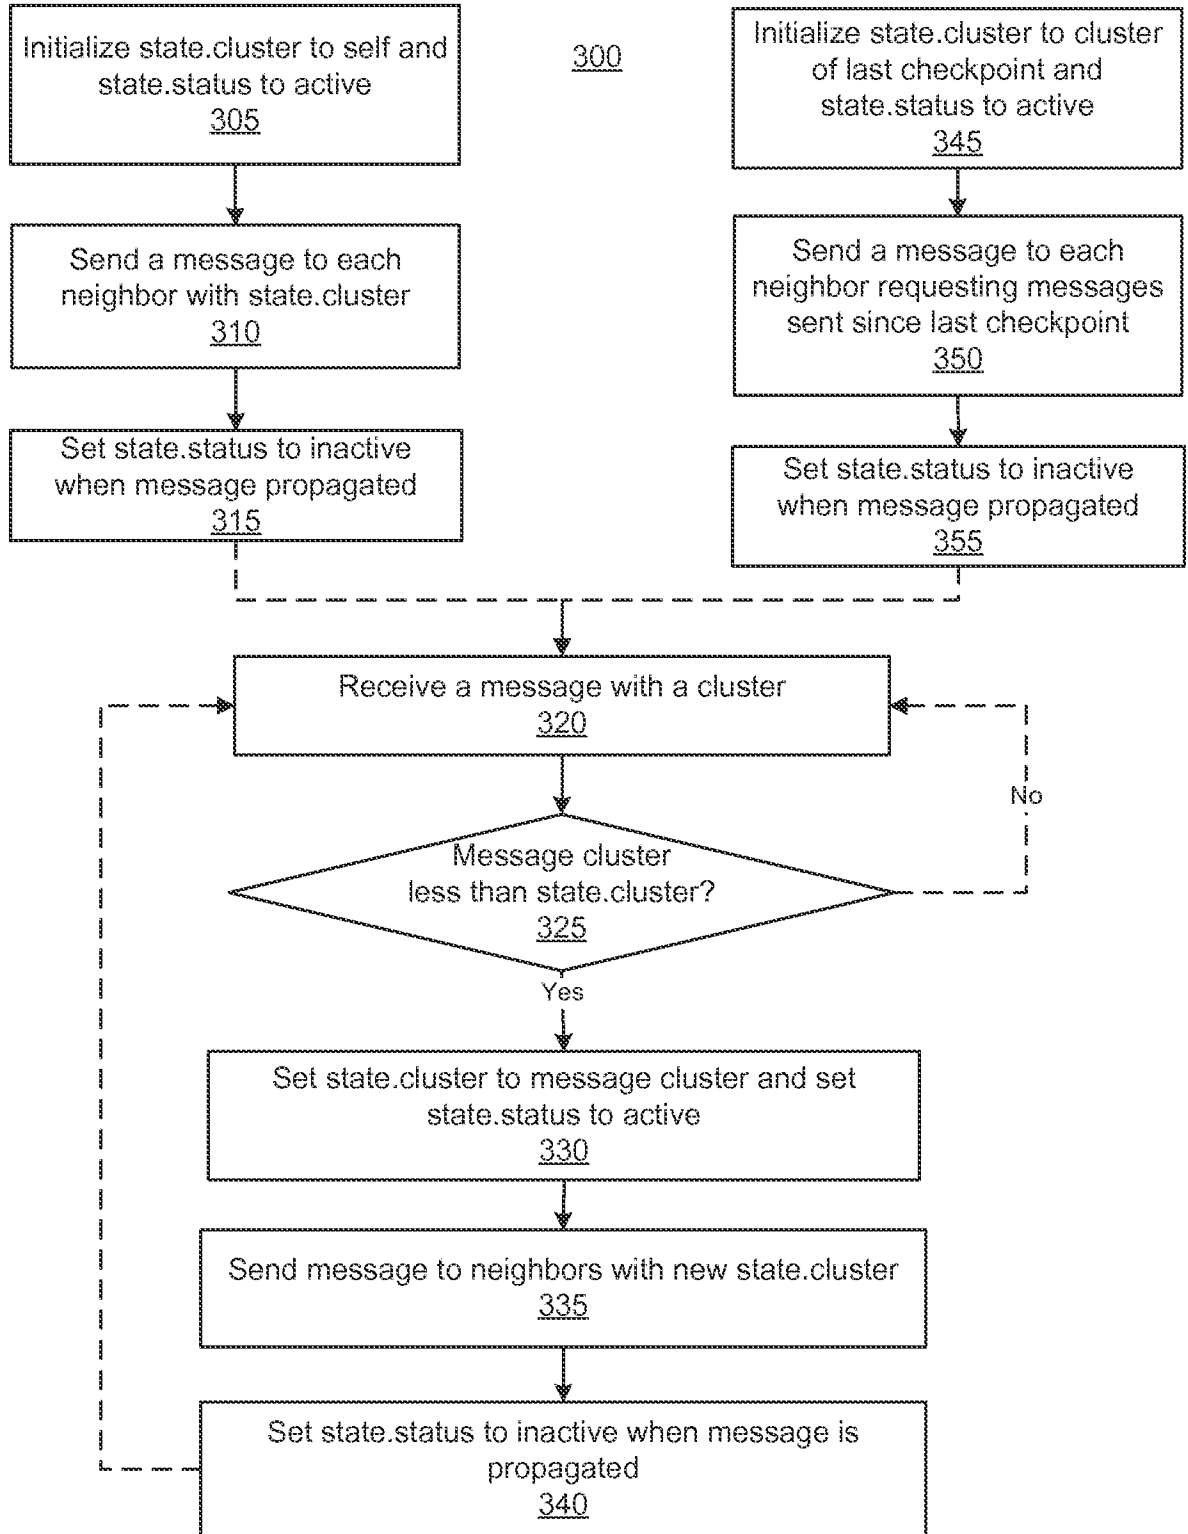


FIG. 3

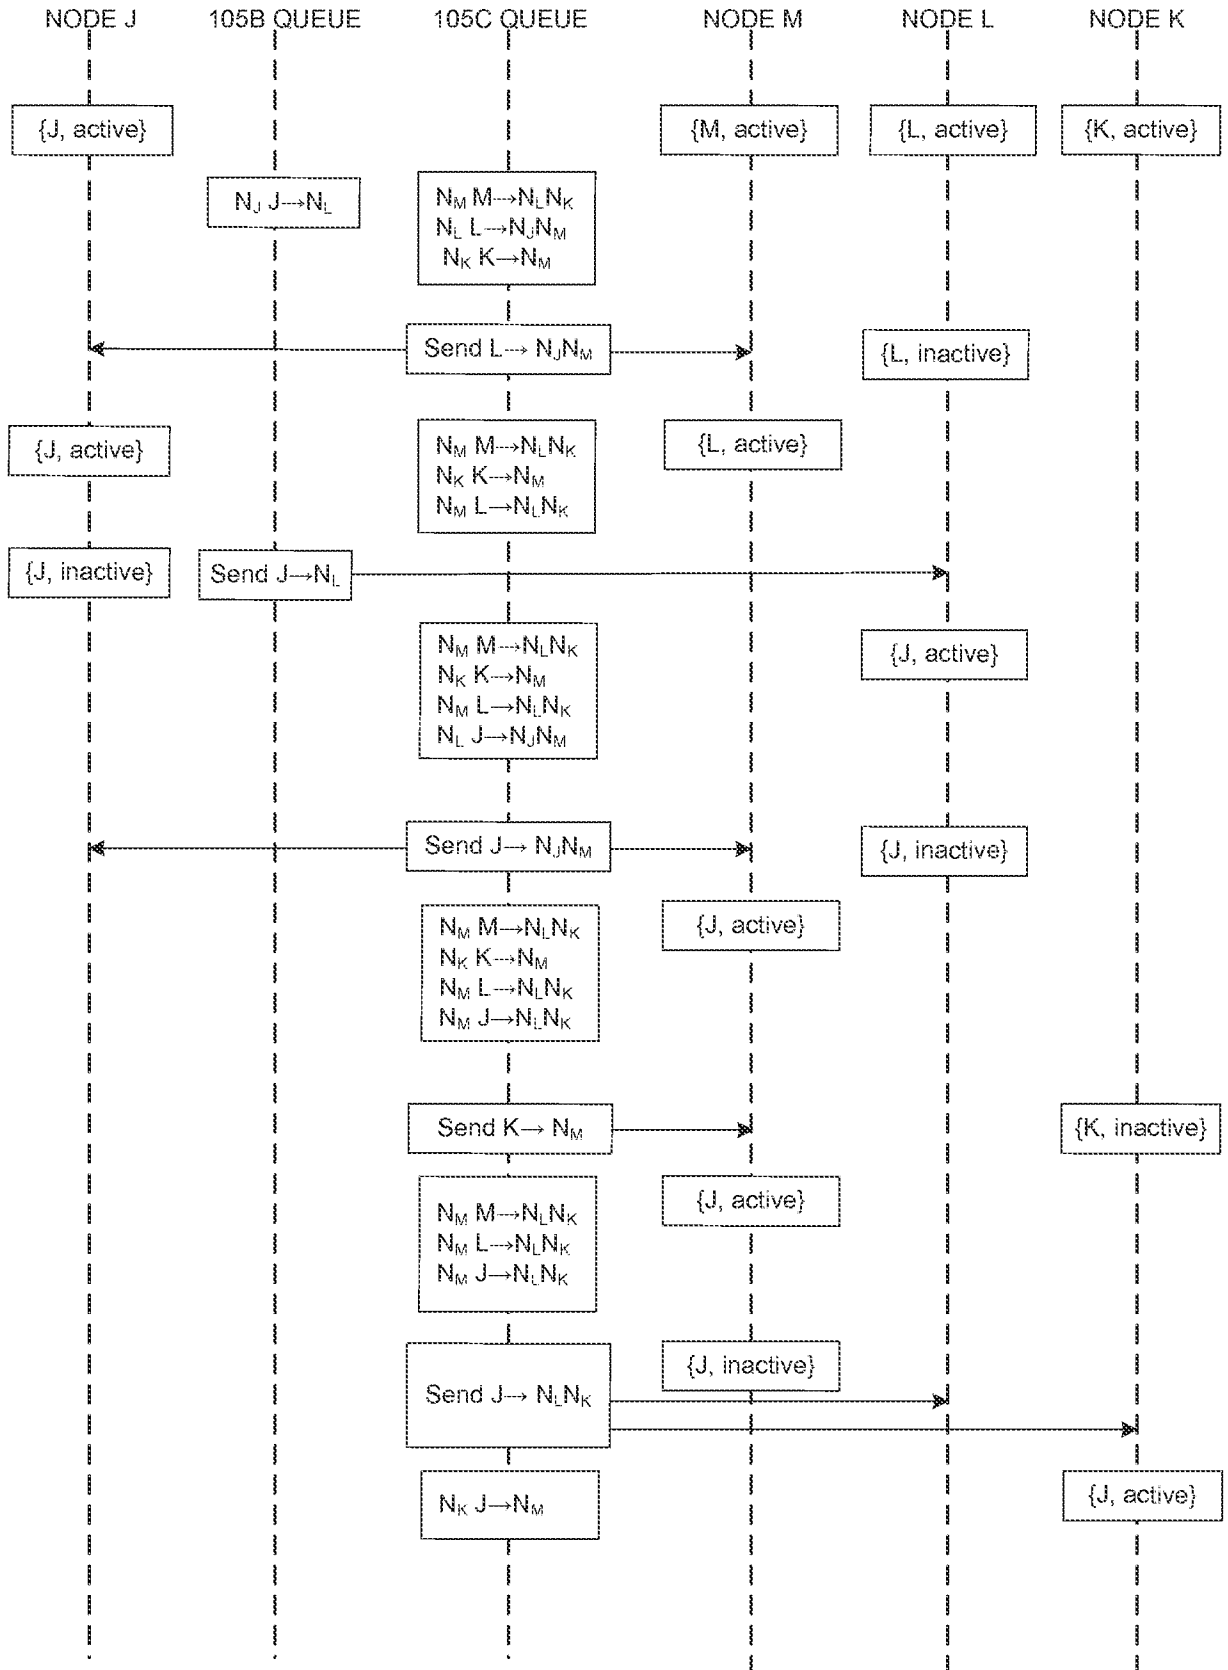


FIG. 4

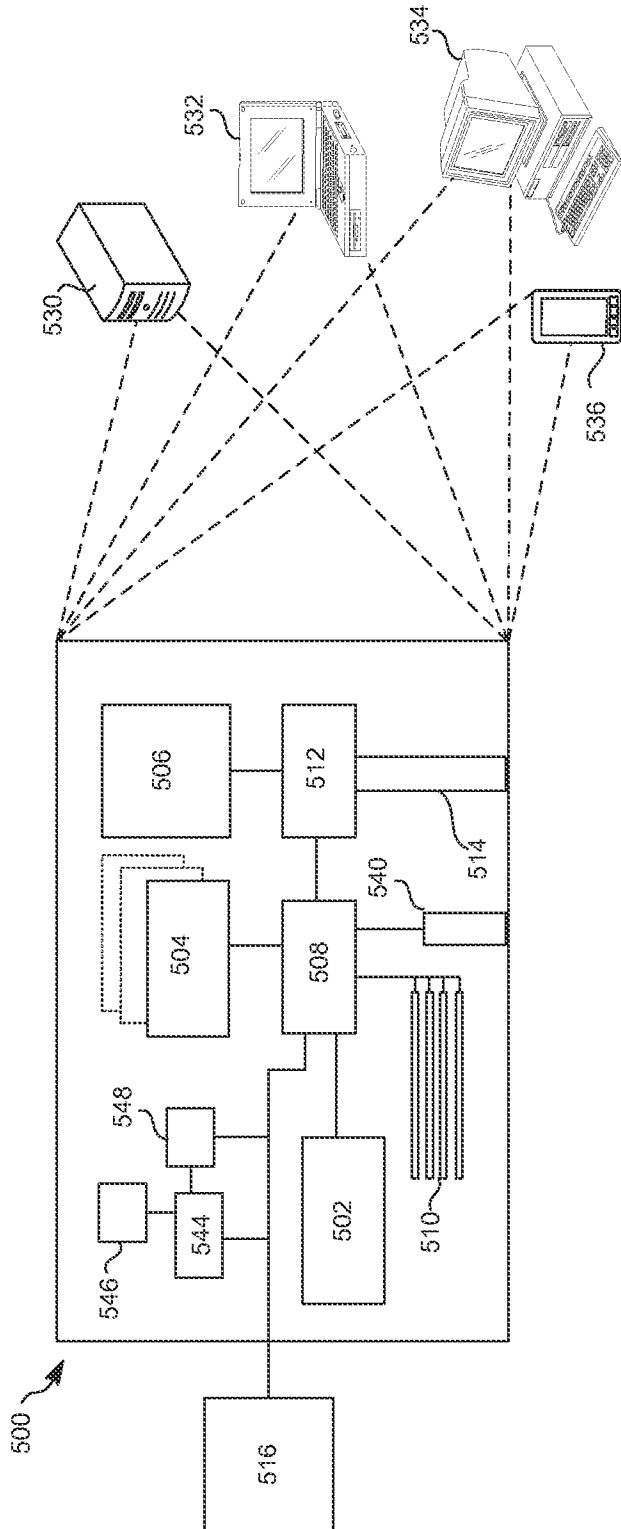


FIG. 5

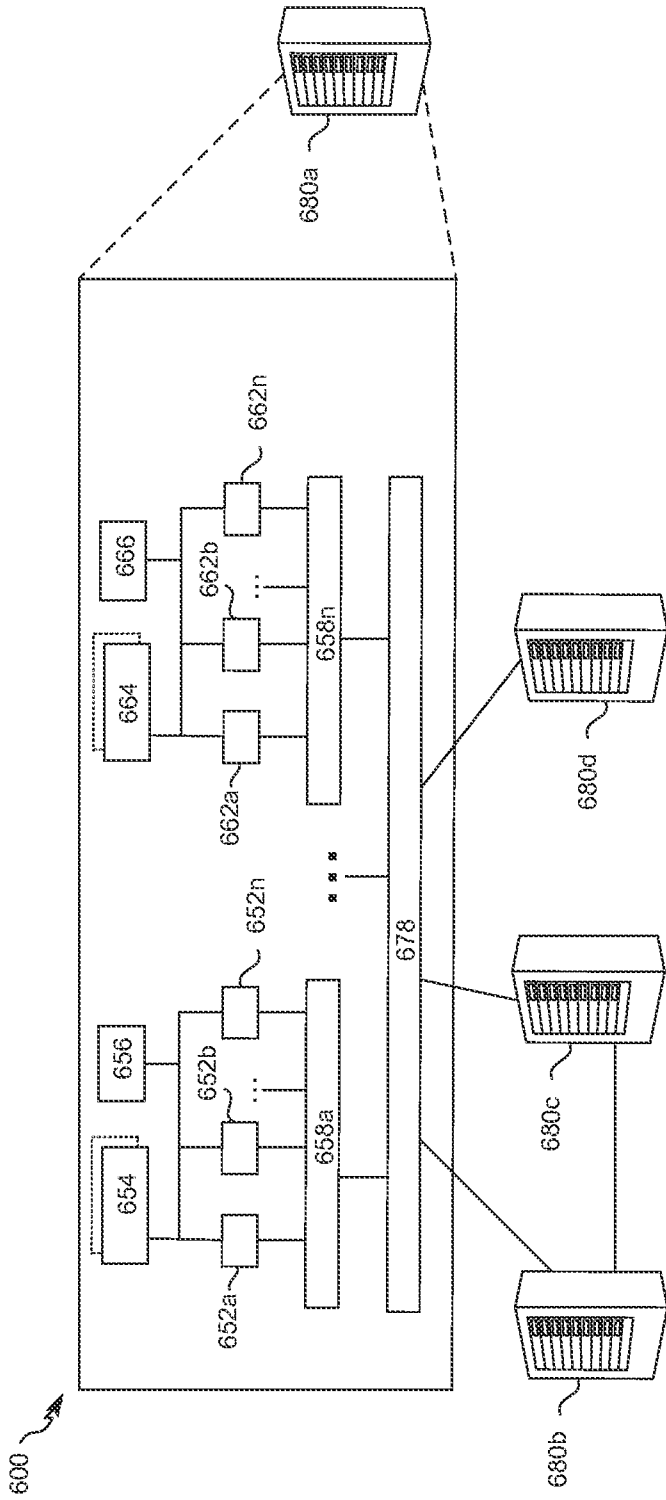


FIG. 6

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/044642

A. CLASSIFICATION OF SUBJECT MATTER INV. G06F9/54 G06Q10/06 ADD.		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols) G06F G06Q		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) EPO-Internal		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A	US 2010/083194 A1 (BAGHERJEIRAN ABRAHAM [US] ET AL) 1 April 2010 (2010-04-01) abstract paragraph [0005] - paragraph [0006] paragraph [0023] - paragraph [0024] paragraph [0030] - paragraph [0032] paragraph [0034] - paragraph [0036] -----	1, 12, 23 2-11, 13-22, 24-27
X A	US 2013/024479 A1 (GONG NAN [SE] ET AL) 24 January 2013 (2013-01-24) abstract paragraph [0008] paragraph [0026] - paragraph [0027] paragraph [0061] - paragraph [0064] ----- -/--	1, 12, 23 2-11, 13-22, 24-27
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
* Special categories of cited documents :		
"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family	
Date of the actual completion of the international search <p style="text-align: center; font-size: 1.2em;">30 September 2014</p>	Date of mailing of the international search report <p style="text-align: center; font-size: 1.2em;">08/10/2014</p>	
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer <p style="text-align: center; font-size: 1.2em;">Archontopoulos, E</p>	

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2014/044642

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2011/066649 A1 (BERLYANT MIKHAIL [US] ET AL) 17 March 2011 (2011-03-17)	1,12,23
A	abstract paragraph [0035] - paragraph [0037] paragraph [0044] - paragraph [0052] -----	2-11, 13-22, 24-27
A	NUUTILA ESKO ET AL: "On finding the strongly connected components in a directed graph", INFORMATION PROCESSING LETTERS, AMSTERDAM, NL, vol. 49, no. 1, 14 January 1994 (1994-01-14), pages 9-14, XP002712565, ISSN: 0020-0190, DOI: 10.1016/0020-0190(94)90047-7 [retrieved on 2013-09-09] the whole document -----	1-27
A	THOMAS SEIDL ET AL: "CC-MR Finding Connected Components in Huge Graphs with MapReduce", 24 September 2012 (2012-09-24), MACHINE LEARNING AND KNOWLEDGE DISCOVERY IN DATABASES, SPRINGER BERLIN HEIDELBERG, BERLIN, HEIDELBERG, PAGE(S) 458 - 473, XP047016935, ISBN: 978-3-642-33459-7 the whole document -----	1-27
A	DIP SANKAR BANERJEE ET AL: "Hybrid algorithms for list ranking and graph connected components", HIGH PERFORMANCE COMPUTING (HIPC), 2011 18TH INTERNATIONAL CONFERENCE ON, IEEE, 18 December 2011 (2011-12-18), pages 1-10, XP032116109, DOI: 10.1109/HIPC.2011.6152655 ISBN: 978-1-4577-1951-6 the whole document -----	1-27
A	HILLEL GAZIT ED - RICHARD COLE: "An optimal randomized parallel algorithm for finding connected components in a graph", FOUNDATIONS OF COMPUTER SCIENCE, 1985., 27TH ANNUAL SYMPOSIUM ON, IEEE, PISCATAWAY, NJ, USA, 27 October 1986 (1986-10-27), pages 492-501, XP031288052, ISBN: 978-0-8186-0740-0 the whole document ----- ----- -/--	1-27

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/044642

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	VĀ CLAV KOUBEK ET AL: "Parallel algorithms for connected components in a graph", 9 September 1985 (1985-09-09), FUNDAMENTALS OF COMPUTATION THEORY, SPRINGER BERLIN HEIDELBERG, BERLIN, HEIDELBERG, PAGE(S) 208 - 217, XP019192859, ISBN: 978-3-540-15689-5 the whole document	1-27
A	----- ABRAMOWSKI S ET AL: "Searching connected components in very large grid graphs", 17 June 1986 (1986-06-17), GRAPH-THEORETIC CONCEPTS IN COMPUTER SCIENCE, SPRINGER BERLIN HEIDELBERG, BERLIN, HEIDELBERG, PAGE(S) 118 - 130, XP019187497, ISBN: 978-3-540-17218-5 the whole document -----	1-27

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2014/044642

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010083194	A1	01-04-2010	NONE
US 2013024479	A1	24-01-2013	NONE
US 2011066649	A1	17-03-2011	US 2011066649 A1 17-03-2011 US 2011066894 A1 17-03-2011