



- (51) **International Patent Classification:**
H04L 12/26 (2006.01) *H04L 12/56* (2006.01)
- (21) **International Application Number:**
PCT/CA2010/001531
- (22) **International Filing Date:**
30 September 2010 (30.09.2010)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/247,343 30 September 2009 (30.09.2009) US
- (72) **Inventors; and**
- (71) **Applicants :** **CHRAPKO, Evan V** [CA/CA]; #128, 14 - 9977- 178 Street, Edmonton, Alberta T5T 6J6 (CA).
CHAN, Leo M. [CA/CA]; #30, 4755 Terwillegar Common, Edmonton, Alberta T6R 3V6 (CA).
- (74) **Agent: SMART & BIGGAR;** Box 2999, Station D, 900-55 Metcalfe Street, Ottawa, Ontario K1P 5Y6 (CA).
- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM,

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

- (84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) **Title:** SYSTEMS AND METHODS FOR SOCIAL GRAPH DATA ANALYTICS TO DETERMINE CONNECTIVITY WITHIN A COMMUNITY

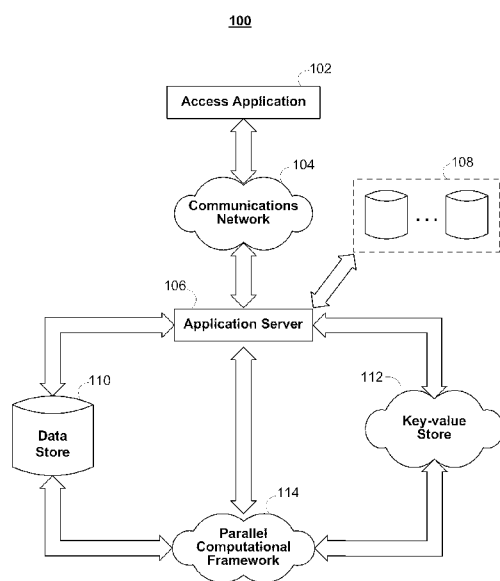


FIG. 1

(57) **Abstract:** Systems and methods for social graph data analytics to determine the connectivity between nodes within a community are provided. A user may assign user connectivity values to other members of the community, or connectivity values may be automatically harvested or assigned from third parties or based on the frequency of interactions between members of the community. Connectivity values may represent such factors as alignment, reputation, status, and/or influence within a social graph of a network community, or the degree of trust. The paths connecting a first node to a second node may be retrieved, and social graph data analytics may be performed on the retrieved paths. For example, a network connectivity value may be determined from all or a subset of all of the retrieved paths. A parallel computational framework may operate in connection with a key-value store to perform some or all of the computations related to the connectivity determinations. Network connectivity values and/or other social graph data may be outputted to third-party processes and services for use in initiating automatic transactions or making automated network-based or real-world decisions.

107617-0002-WO1

SYSTEMS AND METHODS FOR SOCIAL GRAPH DATA ANALYTICS
TO DETERMINE CONNECTIVITY WITHIN A COMMUNITY

Background of the Invention

This invention relates generally to networks of individuals and/or entities and network communities and, more particularly, to systems and methods for
5 determining trust scores or connectivity within or between individuals and/or entities or networks of individuals and/or entities.

The connectivity, or relationships, of an individual or entity within a network community may be used to infer attributes of that individual or entity. For example, an individual or entity's connectivity within a network community may
10 be used to determine the identity of the individual or entity (e.g., used to make decisions about identity claims and authentication), the trustworthiness or reputation of the individual or entity, or the membership, status, and/or influence of that individual or entity in a particular community or subset of a particular community.

15 An individual or entity's connectivity within a network community, however, is difficult to quantify. For example, network communities may include hundreds, thousands, millions, billions or more members. Each member may possess varying degrees of connectivity information about itself and possibly about other members of the community. Some of this information may be highly
20 credible or objective, while other information may be less credible and subjective. In addition, connectivity information from community members may come in various forms and on various scales, making it difficult to meaningfully compare one member's "trustworthiness" or "competence" and connectivity information with another member's "trustworthiness" or "competence" and connectivity
25 information. Also, many individuals may belong to multiple communities, further complicating the determination of a quantifiable representation of trust and connectivity within a network community. Even if a quantifiable representation of an individual's connectivity is determined, it is often difficult to use this

representation in a meaningful way to make real-world decisions about the individual (e.g., whether or not to trust the individual).

Further, it may be useful for these real-world decisions to be made prospectively (i.e., in advance of an anticipated event). Such prospective analysis may be difficult as an individual or entity's connectivity within a network community may change rapidly as the connections between the individual or entity and others in the network community may change quantitatively or qualitatively. This analysis becomes increasingly complex as if applied across multiple communities.

Summary of the Invention

In view of the foregoing, systems and methods are provided for determining the connectivity between nodes within a network community and inferring attributes, such as trustworthiness or competence, from the connectivity.

Connectivity may be determined, at least in part, using various graph traversal and normalization techniques described in more detail below.

In an embodiment, a path counting approach may be used where processing circuitry is configured to count the number of paths between a first node n_1 and a second node n_2 within a network community. A connectivity rating $R_{n_1 n_2}$ may then be assigned to the nodes. The assigned connectivity rating may be proportional to the number of subpaths, or relationships, connecting the two nodes, among other possible measures. Using the number of subpaths as a measure, a path with one or more intermediate nodes between the first node n_1 and the second node n_2 may be scaled by an appropriate number (e.g., the number of intermediate nodes) and this scaled number may be used to calculate the connectivity rating.

In some embodiments, weighted links are used in addition or as an alternative to the subpath counting approach. Processing circuitry may be configured to assign a relative user weight to each path connecting a first node n_1 and a second node n_2 within a network community. A user connectivity value may be assigned to each link. For example, a user or entity associated with node n_1 may assign user connectivity values for all outgoing paths from node n_1 . In some embodiments, the connectivity values assigned by the user or entity may be

indicative of that user or entity's trust in the user or entity associated with node n_2 .

The link values assigned by a particular user or entity may then be compared to each other to determine a relative user weight for each link.

The relative user weight for each link may be determined by first
 5 computing the average of all the user connectivity values assigned by that user (i.e., the out-link values). If t_i is the user connectivity value assigned to link i , then the relative user weight, w_i , assigned to that link may be given in accordance with:

$$w_i = 1 + (t_i - \bar{t}_i)^2 \quad (1)$$

To determine the overall weight of a path, in some embodiments, the
 10 weights of all the links along the path may be multiplied together. The overall path weight may then be given in accordance with:

$$w_{path} = \prod(w_i) \quad (2)$$

The connectivity value for the path may then be defined as the minimum user
 connectivity value of all the links in the path multiplied by the overall path weight
 15 in accordance with:

$$t_{path} = w_{path} \times t_{min} \quad (3)$$

To determine path connectivity values, in some embodiments, a parallel
 computational framework or distributed computational framework (or both) may
 be used. For example, in one embodiment, a number of core processors implement
 20 an Apache Hadoop or Google MapReduce cluster. This cluster may perform some or all of the distributed computations in connection with determining new path link values and path weights.

The processing circuitry may identify a changed node within a network
 community. For example, a new outgoing link may be added, a link may be
 25 removed, or a user connectivity value may have been changed. In response to identifying a changed node, in some embodiments, the processing circuitry may re-compute link, path, and weight values associated with some or all nodes in the implicated network community or communities.

In some embodiments, only values associated with affected nodes in the
 30 network community are recomputed after a changed node is identified. If there exists at least one changed node in the network community, the changed node or nodes may first undergo a prepare process. The prepare process may include a

"map" phase and "reduce" phase. In the map phase of the prepare process, the prepare process may be divided into smaller sub-processes which are then distributed to a core in the parallel computational framework cluster. For example, each node or link change (e.g., tail to out-link change and head to in-link change) may be mapped to a different core for parallel computation. In the reduce phase of the prepare process, each out-link's weight may be determined in accordance with equation (1). Each of the out-link weights may then be normalized by the sum of the out-link weights (or any other suitable value). The node table may then be updated for each changed node, its in-links, and its out-links.

10 After the changed nodes have been prepared, the paths originating from each changed node may be calculated. Once again, a "map" and "reduce" phase of this process may be defined. During this process, in some embodiments, a depth-first search may be performed of the node digraph or node tree. All affected ancestor nodes may then be identified and their paths recalculated.

15 In some embodiments, to improve performance, paths may be grouped by the last node in the path. For example, all paths ending with node n_1 may be grouped together, all paths ending with node n_2 may be grouped together, and so on. These path groups may then be stored separately (e.g., in different columns of a single database table). In some embodiments, the path groups may be stored in columns of a key-value store implementing an HBase cluster (or any other compressed, high performance database system, such as BigTable).

20 In some embodiments, one or more threshold functions may be defined. The threshold function or functions may be used to determine the maximum number of links in a path that will be analyzed in a connectivity determination or connectivity computation. Threshold factors may also be defined for minimum link weights, path weights, or both. Weights falling below a user-defined or system-defined threshold may be ignored in a connectivity determination or connectivity computation, while only weights of sufficient magnitude may be considered.

30 In some embodiments, a user connectivity value may represent the degree of trust between a first node and a second node. In one embodiment, node n_1 may assign a user connectivity value of l_1 to a link between it and node n_2 . Node n_2

may also assign a user connectivity value of l_2 to a reverse link between it and node n_1 . The values of l_1 and l_2 may be at least partially subjective indications of the trustworthiness of the individual or entity associated with the node connected by the link. For example, one or more of the individual or entity's reputation, status, and/or influence within the network community (or some other community), the individual or entity's alignment with the trusting party (e.g., political, social, or religious alignment), past dealings with the individual or entity, and the individual or entity's character and integrity (or any other relevant considerations) may be used to determine a partially subjective user connectivity value indicative of trust.

5 A user (or other individual authorized by the node) may then assign this value to an outgoing link connecting the node to the individual or entity. Objective measures (e.g., data from third-party ratings agencies or credit bureaus) may also be used, in some embodiments, to form composite user connectivity values indicative of trust. The subjective, objective, or both types of measures may be automatically

10 harvested or manually inputted for analysis.

In some embodiments, a decision-making algorithm may access the connectivity values in order to make automatic decisions (e.g., automatic network-based decisions, such as authentication or identity requests) on behalf of a user. Connectivity values may additionally or alternatively be outputted to external

20 systems and processes located at third-parties. The external systems and processes may be configured to automatically initiate a transaction (or take some particular course of action) based, at least in part, on received connectivity values. For example, electronic or online advertising may be targeted to subgroups of members of a network community based, at least in part, on network connectivity values.

25 In some embodiments, a decision-making algorithm may access the connectivity values to make decisions prospectively (e.g., before an anticipated event like a request for credit). Such decisions may be made at the request of a user, or as part of an automated process (e.g., a credit bureau's periodic automated analysis of a database of customer information). This prospective analysis may

30 allow for the initiation of a transaction (or taking of some particular action) in a fluid and/or dynamic manner.

Brief Description of the Drawings

The above and other features of the present invention, its nature and various advantages will be more apparent upon consideration of the following detailed description, taken in conjunction with the accompanying drawings, and in which:

5 FIG. 1 is an illustrative block diagram of a network architecture used to support connectivity within a network community in accordance with one embodiment of the invention;

 FIG. 2 is another illustrative block diagram of a network architecture used to support connectivity within a network community in accordance with one
10 embodiment of the invention;

 FIGS. 3A and 3B show illustrative data tables for supporting connectivity determinations within a network community in accordance with one embodiment of the invention;

 FIGS. 4A-4D show illustrative processes for supporting connectivity
15 determinations within a network community in accordance with one embodiment of the invention; and

 FIG. 5 shows an illustrative process for querying all paths to a target node and computing a network connectivity value in accordance with one embodiment of the invention.

20

Detailed Description

Systems and methods for determining the connectivity between nodes in a network community are provided. As defined herein, a "node" may include any user terminal, network device, computer, mobile device, access point, or any other
25 electronic device. In some embodiments, a node may also represent an individual human being, entity (e.g., a legal entity, such as a public or private company, corporation, limited liability company (LLC), partnership, sole proprietorship, or charitable organization), concept (e.g., a social networking group), animal, or inanimate object (e.g., a car, aircraft, or tool). As also defined herein, a "network
30 community" may include a collection of nodes and may represent any group of devices, individuals, or entities.

For example, all or some subset of the users of a social networking website or social networking service (or any other type of website or service, such as an online gaming community) may make up a single network community. Each user may be represented by a node in the network community. As another example, all
5 the subscribers to a particular newsgroup or distribution list may make up a single network community, where each individual subscriber may be represented by a node in the network community. Any particular node may belong in zero, one, or more than one network community, or a node may be banned from all, or a subset of, the community. To facilitate network community additions, deletions, and link
10 changes, in some embodiments a network community may be represented by a directed graph, or digraph, weighted digraph, tree, or any other suitable data structure.

FIG. 1 shows illustrative network architecture 100 used to support the connectivity determinations within a network community. A user may utilize
15 access application 102 to access application server 106 over communications network 104. For example, access application 102 may include a standard web browser, application server 106 may include a web server, and communication network 106 may include the Internet. Access application 102 may also include proprietary applications specifically developed for one or more platforms or
20 devices. For example, access application 102 may include one or more instances of an Apple iOS, Android, or WebOS application or any suitable application for use in accessing application 106 over communications network 104. Multiple users may access application service 106 via one or more instances of access application 102. For example, a plurality of mobile devices may each have an
25 instance of access application 102 running locally on the devices. One or more users may use an instance of access application 102 to interact with application server 106.

Communication network 104 may include any wired or wireless network, such as the Internet, WiMax, wide area cellular, or local area wireless network.
30 Communication network 104 may also include personal area networks, such as Bluetooth and infrared networks. Communications on communications network

104 may be encrypted or otherwise secured using any suitable security or encryption protocol.

Application server 106, which may include any network server or virtual server, such as a file or web server, may access data sources 108 locally or over
5 any suitable network connection. Application server 106 may also include processing circuitry (e.g., one or more microprocessors), memory (e.g., RAM, ROM, and hybrid types of memory), storage devices (e.g., hard drives, optical drives, and tape drives). The processing circuitry included in application server 106 may execute a server process for supporting the network connectivity
10 determinations of the present invention, while access application 102 executes a corresponding client process. The processing circuitry included in application server 106 may also perform any of the calculations and computations described herein in connection with determining network connectivity. In some
embodiments, a computer-readable medium with computer program logic recorded
15 thereon is included within application server 106. The computer program logic may determine the connectivity between two or more nodes in a network community and it may or may not output such connectivity to a display screen or data

For example, application server 106 may access data sources 108 over the
20 Internet, a secured private LAN, or any other communications network. Data sources 108 may include one or more third-party data sources, such as data from third-party social networking services and third-party ratings bureaus. For example, data sources 108 may include user and relationship data (e.g., "friend" or "follower" data) from one or more of Facebook, MySpace, openSocial, Friendster,
25 Bebo, hi5, Orkut, PerfSpot, Yahoo! 360, LinkedIn, Twitter, Google Buzz, Really Simple Syndication readers or any other social networking website or information service. Data sources 108 may also include data stores and databases local to application server 106 containing relationship information about users accessing application server 106 via access application 102 (e.g., databases of addresses,
30 legal records, transportation passenger lists, gambling patterns, political and/or charity donations, political affiliations, vehicle license plate or identification

numbers, universal product codes, news articles, business listings, and hospital or university affiliations).

Application server 106 may be in communication with one or more of data store 110, key-value store 112, and parallel computational framework 114. Data store 110, which may include any relational database management system (RDBMS), file server, or storage system, may store information relating to one or more network communities. For example, one or more of data tables 300 (FIG. 3A) may be stored on data store 110. Data store 110 may store identity information about users and entities in the network community, an identification of the nodes in the network community, user link and path weights, user configuration settings, system configuration settings, and/or any other suitable information. There may be one instance of data store 110 per network community, or data store 110 may store information relating to a plural number of network communities. For example, data store 110 may include one database per network community, or one database may store information about all available network communities (e.g., information about one network community per database table).

Parallel computational framework 114, which may include any parallel or distributed computational framework or cluster, may be configured to divide computational jobs into smaller jobs to be performed simultaneously, in a distributed fashion, or both. For example, parallel computational framework 114 may support data-intensive distributed applications by implementing a map/reduce computational paradigm where the applications may be divided into a plurality of small fragments of work, each of which may be executed or re-executed on any core processor in a cluster of cores. A suitable example of parallel computational framework 114 includes an Apache Hadoop cluster.

Parallel computational framework 114 may interface with key-value store 112, which also may take the form of a cluster of cores. Key-value store 112 may hold sets of key-value pairs for use with the map/reduce computational paradigm implemented by parallel computational framework 114. For example, parallel computational framework 114 may express a large distributed computation as a sequence of distributed operations on data sets of key-value pairs. User-defined map/reduce jobs may be executed across a plurality of nodes in the cluster. The

processing and computations described herein may be performed, at least in part, by any type of processor or combination of processors. For example, various types of quantum processors (e.g., solid-state quantum processors and light-based quantum processors), artificial neural networks, and the like may be used to
5 perform massively parallel computing and processing.

In some embodiments, parallel computational framework 114 may support two distinct phases, a "map" phase and a "reduce" phase. The input to the computation may include a data set of key-value pairs stored at key-value store 112. In the map phase, parallel computational framework 114 may split, or divide,
10 the input data set into a large number of fragments and assign each fragment to a map task. Parallel computational framework 114 may also distribute the map tasks across the cluster of nodes on which it operates. Each map task may consume key-value pairs from its assigned fragment and produce a set of intermediate key-value pairs. For each input key-value pair, the map task may invoke a user defined map
15 function that transmutes the input into a different key-value pair. Following the map phase, parallel computational framework 114 may sort the intermediate data set by key and produce a collection of tuples so that all the values associated with a particular key appear together. Parallel computational framework 114 may also partition the collection of tuples into a number of fragments equal to the number of
20 reduce tasks.

In the reduce phase, each reduce task may consume the fragment of tuples assigned to it. For each such tuple, the reduce task may invoke a user-defined reduce function that transmutes the tuple into an output key-value pair. Parallel computational framework 114 may then distribute the many reduce tasks across the
25 cluster of nodes and provide the appropriate fragment of intermediate data to each reduce task.

Tasks in each phase may be executed in a fault-tolerant manner, so that if one or more nodes fail during a computation the tasks assigned to such failed nodes may be redistributed across the remaining nodes. This behavior may allow
30 for load balancing and for failed tasks to be re-executed with low runtime overhead.

Key-value store 112 may implement any distributed file system capable of storing large files reliably. For example key-value store 112 may implement Hadoop's own distributed file system (DFS) or a more scalable column-oriented distributed database, such as HBase. Such file systems or databases may include
5 BigTable-like capabilities, such as support for an arbitrary number of table columns.

Although FIG. 1, in order to not over-complicate the drawing, only shows a single instance of access application 102, communications network 104, application server 106, data source 108, data store 110, key-value store 112, and
10 parallel computational framework 114, in practice network architecture 100 may include multiple instances of one or more of the foregoing components. In addition, key-value store 112 and parallel computational framework 114 may also be removed, in some embodiments. As shown in network architecture 200 of FIG. 2, the parallel or distributed computations carried out by key-value store 112
15 and/or parallel computational framework 114 may be additionally or alternatively performed by a cluster of mobile devices 202 instead of stationary cores. In some embodiments, cluster of mobile devices 202, key-value store 112, and parallel computational framework 114 are all present in the network architecture. Certain application processes and computations may be performed by cluster of mobile
20 devices 202 and certain other application processes and computations may be performed by key-value store 112 and parallel computational framework 114. In addition, in some embodiments, communication network 104 itself may perform some or all of the application processes and computations. For example, specially-configured routers or satellites may include processing circuitry adapted to carry
25 out some or all of the application processes and computations described herein.

Cluster of mobile devices 202 may include one or more mobile devices, such as PDAs, cellular telephones, mobile computers, or any other mobile computing device. Cluster of mobile devices 202 may also include any appliance (e.g., audio/video systems, microwaves, refrigerators, food processors) containing
30 a microprocessor (e.g., with spare processing time), storage, or both. Application server 106 may instruct devices within cluster of mobile devices 202 to perform computation, storage, or both in a similar fashion as would have been distributed to

multiple fixed cores by parallel computational framework 114 and the map/reduce computational paradigm. Each device in cluster of mobile devices 202 may perform a discrete computational job, storage job, or both. Application server 106 may combine the results of each distributed job and return a final result of the
5 computation.

FIG. 3A shows illustrative data tables 300 used to support the connectivity determinations of the present invention. One or more of tables 300 may be stored in, for example, a relational database in data store 110 (FIG. 1). Table 302 may store an identification of all the nodes registered in the network community. A
10 unique identifier may be assigned to each node and stored in table 302. In addition, a string name may be associated with each node and stored in table 302. As described above, in some embodiments, nodes may represent individuals or entities, in which case the string name may include the individual or person's first and/or last name, nickname, handle, or entity name.

15 Table 304 may store user connectivity values. In some embodiments, user connectivity values may be assigned automatically by the system (e.g., by application server 106 (FIG. 1)). For example, application server 106 (FIG. 1) may monitor all electronic interaction (e.g., electronic communication, electronic transactions, or both) between members of a network community. In some
20 embodiments, a default user connectivity value (e.g., the link value 1) may be assigned initially to all links in the network community. After electronic interaction is identified between two or more nodes in the network community, user connectivity values may be adjusted upwards or downwards depending on the type of interaction between the nodes and the result of the interaction. For
25 example, each simple email exchange between two nodes may automatically increase or decrease the user connectivity values connecting those two nodes by a fixed amount. More complicated interactions (e.g., product or service sales or inquires) between two nodes may increase or decrease the user connectivity values connecting those two nodes by some larger fixed amount. In some embodiments,
30 user connectivity values between two nodes may always be increased unless a user or node indicates that the interaction was unfavorable, not successfully completed, or otherwise adverse. For example, a transaction may not have been timely

executed or an email exchange may have been particularly displeasing. Adverse interactions may automatically decrease user connectivity values while all other interactions may increase user connectivity values (or have no effect). In addition, user connectivity values may be automatically harvested using outside sources.

5 For example, third-party data sources (such as ratings agencies and credit bureaus) may be automatically queried for connectivity information. This connectivity information may include completely objective information, completely subjective information, composite information that is partially objective and partially subjective, any other suitable connectivity information, or any combination of the
10 foregoing.

In some embodiments, user connectivity values may be manually assigned by members of the network community. These values may represent, for example, the degree or level of trust between two users or nodes or one node's assessment of another node's competence in some endeavor. As described above, user
15 connectivity values may include a subjective component and an objective component in some embodiments. The subjective component may include a trustworthiness "score" indicative of how trustworthy a first user or node finds a second user, node, community, or subcommunity. This score or value may be entirely subjective and based on interactions between the two users, nodes, or
20 communities. A composite user connectivity value including subjective and objective components may also be used. For example, third-party information may be consulted to form an objective component based on, for example, the number of consumer complaints, credit score, socio-economic factors (e.g., age, income, political or religious affiliations, and criminal history), or number of citations/hits
25 in the media or in search engine searches. Third-party information may be accessed using communications network 104 (FIG. 1). For example, a third-party credit bureau's database may be polled or a personal biography and background information, including criminal history information, may be accessed from a third-party database or data source (e.g., as part of data sources 108 (FIG. 1) or a
30 separate data source) or input directly by a node, user, or system administrator.

Table 304 may store an identification of a link head, link tail, and user connectivity value for the link. Links may or may not be bidirectional. For

example, a user connectivity value from node n_1 to node n_2 may be different (and completely separate) than a link from node n_2 to node n_1 . Especially in the trust context described above, each user can assign his or her own user connectivity value to a link (i.e., two users need not trust each other an equal amount in some embodiments).

Table 306 may store an audit log of table 304. Table 306 may be analyzed to determine which nodes or links have changed in the network community. In some embodiments, a database trigger is used to automatically insert an audit record into table 306 whenever a change of the data in table 304 is detected. For example, a new link may be created, a link may be removed, or a user connectivity value may be changed. This audit log may allow for decisions related to connectivity values to be made prospectively (i.e., before an anticipated event). Such decisions may be made at the request of a user, or as part of an automated process, such as the processes described below with respect to FIG. 5. This prospective analysis may allow for the initiation of a transaction (or taking of some particular action) in a fluid and/or dynamic manner. After such a change is detected, the trigger may automatically create a new row in table 306. Table 306 may store an identification of the changed node, and identification of the changed link head, changed link tail, and the user connectivity value to be assigned to the changed link. Table 306 may also store a timestamp indicative of the time of the change and an operation code. In some embodiments, operation codes may include "insert," "update," or "delete" operations, corresponding to whether a link was inserted, a user connectivity value was changed, or a link was deleted, respectively. Other operation codes may be used in other embodiments.

FIG. 3B shows illustrative data structure 310 used to support the connectivity determinations of the present invention. In some embodiments, data structure 310 may be stored using key-value store 112 (FIG. 1), while tables 300 are stored in data store 110 (FIG. 1). As described above, key-value store 112 (FIG. 1) may implement an HBase storage system and include BigTable support. Like a traditional relational database management system, the data shown in FIG. 3B may be stored in tables. However, the BigTable support may allow for an

arbitrary number of columns in each table, whereas traditional relational database management systems may require a fixed number of columns.

Data structure 310 may include node table 312. In the example shown in FIG. 3B, node table 312 includes several columns. Node table 312 may include
5 row identifier column 314, which may store 64-bit, 128-bit, 256-bit, 512-bit, or 1024-bit integers and may be used to uniquely identify each row (e.g., each node) in node table 312. Column 316 may include a list of all the incoming links for the current node. Column 318 may include a list of all the outgoing links for the current node. Column 320 may include a list of node identifiers to which the
10 current node is connected. A first node may be connected to a second node if outgoing links may be followed to reach the second node. For example, for A -> B, A is connected to B, but B may not be connected to A. As described in more detail below, column 320 may be used during the portion of process 400 (FIG. 4A) shown in FIG. 4B. Node table 312 may also include one or more "bucket"
15 columns 322. These columns may store a list of paths that connect the current node to a target node. As described above, grouping paths by the last node in the path (e.g., the target node) may facilitate connectivity computations. As shown in FIG. 3B, in some embodiments, to facilitate scanning, bucket column names may include the target node identifier appended to the end of the "bucket:" column

20 FIGS. 4A-4D show illustrative processes for determining the connectivity of nodes within a network community. FIG. 4A shows process 400 for updating a connectivity graph (or any other suitable data structure) associated with a network community. As described above, in some embodiments, each network community is associated with its own connectivity graph, digraph, tree, or other suitable data
25 structure. In other embodiments, a plurality of network communities may share one or more connectivity graphs (or other data structure).

In some embodiments, the processes described with respect to FIG. 4A-4D may be executed to make decisions prospectively (i.e., before an anticipated event). Such decisions may be made at the request of a user, or as part of an
30 automated process, such as the processes described below with respect to FIG. 5. This prospective analysis may allow for the initiation of a transaction (or taking of some particular action) in a fluid and/or dynamic manner.

At step 402, a determination is made whether at least one node has changed in the network community. As described above, an audit record may be inserted into table 306 (FIG. 3) after a node has changed. By analyzing table 306 (FIG. 3), a determination may be made (e.g., by application server 106 of FIG. 1) that a new link has been added, an existing link has been removed, or a user connectivity value has changed. If, at step 404, it is determined that a node has changed, then process 400 continues to step 410 (shown in FIG. 4B) to prepare the changed nodes, step 412 (shown in FIG. 4C) to calculate paths originating from the changed nodes, step 414 (shown in FIG. 4D) to remove paths that go through a changed node, and step 416 (shown in FIG. 4E) to calculate paths that go through a changed node. It should be noted that more than one step or task shown in FIGS. 4B, 4C, 4D, and 4E may be performed in parallel using, for example, a cluster of cores. For example, multiple steps or tasks shown in FIG. 4B may be executed in parallel or in a distributed fashion, then multiple steps or tasks shown in FIG. 4C may be executed in parallel or in a distributed fashion, then multiple steps or tasks shown in FIG. 4D may be executed in parallel or in a distributed fashion, and then multiple steps or tasks shown in FIG. 4E may be executed in parallel or in a distributed fashion. In this way, overall latency associated with process 400 may be reduced.

If a node change is not detected at step 404, then process 400 enters a sleep mode at step 406. For example, in some embodiments, an application thread or process may continuously check to determine if at least one node or link has changed in the network community. In other embodiments, the application thread or process may periodically check for changed links and nodes every n seconds, where n is any positive number. After the paths are calculated that go through a changed node at step 416 or after a period of sleep at step 406, process 400 may determine whether or not to loop at step 408. For example, if all changed nodes have been updated, then process 400 may stop at step 418. If, however, there are more changed nodes or links to process, then process 400 may loop at step 408 and return to step 404.

In practice, one or more steps shown in process 400 may be combined with other steps, performed in any suitable order, performed in parallel (e.g., simultaneously or substantially simultaneously), or removed.

FIGS. 4B-4E each include processes with a "map" phase and "reduce" phase. As described above, these phases may form part of a map/reduce computational paradigm carried out by parallel computational framework 114 (FIG. 1), key-value store 112 (FIG. 1), or both. As shown in FIG. 4B, in order to prepare any changed nodes, map phase 420 may include determining if there are any more link changes at step 422, retrieving the next link change at step 440, mapping the tail to out-link change at step 442, and mapping the head to in-link change at step 444.

If there are no more link changes at step 422, then, in reduce phase 424, a determination may be made at step 426 that there are more nodes and link changes to process. If so, then the next node and its link changes may be retrieved at step 428. The most recent link changes may be preserved at step 430 while any intermediate link changes are replaced by more recent changes. For example, the timestamp stored in table 306 (FIG. 3) may be used to determine the time of every link or node change. At step 432, the average out-link user connectivity value may be calculated. For example, if node n_l has eight out-links with assigned user connectivity values, these eight user connectivity values may be averaged at step 432. At step 434, each out-link's weight may be calculated in accordance with equation (1) above. All the out-link weights may then be summed and used to normalize each out-link weight at step 436. For example, each out-link weight may be divided by the sum of all out-link weights. This may yield a weight between 0 and 1 for each out-link. At step 438, the existing buckets for the changed node, in-links, and out-links may be saved. For example, the buckets may be saved in key-value store 112 (FIG. 1) or data store 110 (FIG. 1). If there are no more nodes and link changes to process at step 426, the process may stop at step 446.

As shown in FIG. 4C, in order to calculate paths originating from changed nodes, map phase 448 may include determining if there are any more changed nodes at step 450, retrieving the next changed node at step 466, marking existing

buckets for deletion by mapping changed nodes to the NULL path at step 468, recursively generating paths by following out-links at step 470, and if the path is a qualified path, mapping the tail to the path. Qualified paths may include paths that satisfy one or more predefined threshold functions. For example, a threshold
5 function may specify a minimum path weight. Paths with path weights greater than the minimum path weight may be designated as qualified paths.

If there are no more changed nodes at step 450, then, in reduce phase 452, a determination may be made at step 454 that there are more nodes and paths to process. If so, then the next node and its paths may be retrieved at step 456. At
10 step 458, buckets may be created by grouping paths by their head. If a bucket contains only the NULL path at step 460, then the corresponding cell in the node table may be deleted at step 462. If the bucket contains more than the NULL path, then at step 464 the bucket is saved to the corresponding cell in the node table. If there are no more nodes and paths to process at step 456, the process may stop at
15 step 474.

As shown in FIG. 4D, in order to remove paths that go through a changed node, map phase 476 may include determining if there are any more changed nodes at step 478 and retrieving the next changed node at step 488. At step 490, the "bucket:" column in the node table (e.g., column 322 of node table 312 (both of
20 FIG. 3B)) corresponding to the changed node may be scanned. For example, as described above, the target node identifier may be appended to the end of the "bucket:" column name. Each bucket may include a list of paths that connect the current node to the target node (e.g., the changed node). At step 492, for each matching node found by the scan and the changed node's old buckets, the matching
25 node may be matched to a (changed node, old bucket) deletion pair.

If there are no more changed nodes at step 478, then, in reduce phase 480, a determination may be made at step 484 that there are more node and deletion pairs to process. If so, then the next node and its deletion pairs may be retrieved at step 484. At step 486, for each deletion pair, any paths that go through the changed
30 node in the old bucket may be deleted. If there are no more nodes and deletion pairs to process at step 482, the process may stop at step 494.

As shown in FIG. 4E, in order to calculate paths that go through a changed node, map phase 496 may include determining if there are any more changed nodes at step 498 and retrieving the next changed node at step 508. At step 510, the "bucket:" column in the node table (e.g., column 322 of node table 312 (both of
5 FIG. 3B)) corresponding to the changed node may be scanned. At step 512, for each matching node found in the scan and the changed node's paths, all paths in the scanned bucket may be joined with all paths of the changed bucket. At step 514, each matching node may be mapped to each qualified joined

If there are no more changed nodes at step 498, then, in reduce phase 500, a
10 determination may be made at step 502 that there are more node and paths to process. If so, then the next node and its paths may be retrieved at step 504. Each path may then be added to the appropriate node bucket at step 506. If there are no more nodes and paths to process at step 502, the process may stop at step 516.

FIG. 5 shows illustrative process 520 for supporting a user query for all
15 paths from a first node to a target node. For example, a first node (representing, for example, a first individual or entity) may wish to know how connected the first node is to some second node (representing, for example, a second individual or entity) in the network community. In the context of trust described above (and where the user connectivity values represent, for example, at least partially
20 subjective user trust values), this query may return an indication of how much the first node may trust the second node. In general, the more paths connecting the two nodes may yield a greater (or lesser if, for example, adverse ratings are used) network connectivity value (or network trust amount).

At step 522, the node table cell where the row identifier equals the first
25 node identifier and the column equals the target node identifier appended to the "bucket:" column name prefix is accessed. All paths may be read from this cell at step 524. The path weights assigned to the paths read at step 524 may then be summed at step 526. At step 528, the path weights may be normalized by dividing each path weight by the computed sum of the path weights. A network
30 connectivity value may then be computed at step 530. For example, each path's user connectivity value may be multiplied by its normalized path weight. The

network connectivity value may then be computed in some embodiments in accordance with:

$$t_{network} = \sum t_{path} \times w_{path} \quad (4)$$

where t_{path} is the user connectivity value for a path (given in accordance with equation (3)) and w_{path} is the normalized weight for that path. The network connectivity value may then be held or outputted (e.g., displayed on a display device, output by processing circuitry of application server 106, and/or stored on data store 110 (FIG. 1)). In addition, a decision-making algorithm may access the network connectivity value in order to make automatic decisions (e.g., automatic network-based decisions, such as authentication or identity requests) on behalf of the user. Network connectivity values may additionally or alternatively be outputted to external systems and processes located at third-parties. The external systems and processes may be configured to automatically initiate a transaction (or take some particular course of action) based, at least in part, on the received network connectivity values. Process 520 may stop at step 532.

In practice, one or more steps shown in process 520 may be combined with other steps, performed in any suitable order, performed in parallel (e.g., simultaneously or substantially simultaneously), or removed. In addition, as described above, various threshold functions may be used in order to reduce computational complexity. For example, a threshold function defining the maximum number of links to traverse may be defined. Paths containing more than the threshold specified by the threshold function may not be considered in the network connectivity determination. In addition, various threshold functions relating to link and path weights may be defined. Links or paths below the threshold weight specified by the threshold function may not be considered in the network connectivity determination.

Although process 520 describes a single user query for all paths from a first node to a target node, in actual implementations groups of nodes may initiate a single query for all the paths from each node in the group to a particular target node. For example, multiple members of a network community may all initiate a group query to a target node. Process 520 may return an individual network connectivity value for each querying node in the group or a single composite

network connectivity value taking into account all the nodes in the querying group. For example, the individual network connectivity values may be averaged to form a composite value or some weighted average may be used. The weights assigned to each individual network connectivity value may be based on, for example,

5 seniority in the community (e.g., how long each node has been a member in the community), rank, or social stature. In addition, in some embodiments, a user may initiate a request for network connectivity values for multiple target nodes in a single query. For example, node n_i may wish to determine network connectivity values between it and multiple other nodes. For example, the multiple other nodes

10 may represent several candidates for initiating a particular transaction with node n_i . By querying for all the network connectivity values in a single query, the computations may be distributed in a parallel fashion to multiple cores so that some or all of the results are computed substantially simultaneously.

In addition, queries may be initiated in a number of ways. For example, a

15 user (represented by a source node) may identify another user (represented by a target node) in order to automatically initiate process 520. A user may identify the target node in any suitable way, for example, by selecting the target node from a visual display, graph, or tree, by inputting or selecting a username, handle, network address, email address, telephone number, geographic coordinates, or unique

20 identifier associated with the target node, or by speaking a predetermined command (e.g., "query node 1" or "query node group 1, 5, 9" where 1, 5, and 9 represent unique node identifiers). After an identification of the target node or nodes is received, process 520 may be automatically executed. The results of the process (e.g., the individual or composite network connectivity values) may then

25 be automatically sent to one or more third-party services or processes as described above.

In an embodiment, a user may utilize access application 102 to generate a user query that is sent to access application server 106 over communications network 104 (*see also*, FIG. 1) and automatically initiate process 520. For

30 example, a user may access an Apple iOS, Android, or WebOS application or any suitable application for use in accessing application 106 over communications network 104. The application may display a searchable list of relationship data

related to that user (e.g., "friend" or "follower" data) from one or more of Facebook, MySpace, openSocial, Friendster, Bebo, hi5, Orkut, PerfSpot, Yahoo! 360, LinkedIn, Twitter, Google Buzz, Really Simple Syndication readers or any other social networking website or information service. In some embodiments, a user may search for relationship data that is not readily listed – i.e., search Facebook, Twitter, or any suitable database of information for target nodes that are not displayed in the searchable list of relationship data. A user may select a target node as described above (e.g., select an item from a list of usernames representing a "friend" or "follower") to request a measure of how connected the user is to the target node. Using the processes described with respect to FIGs. 3 and 4A-D, this query may return an indication of how much the user may trust the target node. The returned indication may be displayed to the user using any suitable indicator. In some embodiments, indicator may be a percentage that indicates how trustworthy the target node is to the user.

In some embodiments, a user may utilize access application 102 to provide manual assignments of at least partially subjective indications of how trustworthy the target node is. For example, the user may specify that he or she trusts a selected target node (e.g., a selected "friend" or "follower") to a particular degree. The particular degree may be in the form of a percentage that represents the user's perception of how trustworthy the target node is. The user may provide this indication before, after, or during process 520 described above. The indication provided by the user (e.g., the at least partially subjective indications of trustworthiness) may then be automatically sent to one or more third-party services or processes as described above. In some embodiments, the indications provided by the user may cause a node and/or link to change in a network community. This change may cause a determination to be made that at least one node and/or link has changed in the network community, which in turn triggers various processes as described with respect to FIGs. 3 and 4A-4D.

In some embodiments, a path counting approach may be used in addition to or in place of the weighted link approach described above. Processing circuitry (e.g., of application server 106) may be configured to count the number of paths between a first node n_1 and a second node n_2 within a network community. A

connectivity rating $R_{n_1n_2}$ may then be assigned to the nodes. The assigned connectivity rating may be proportional to the number of paths, or relationships, connecting the two nodes. A path with one or more intermediate nodes between the first node n_1 and the second node n_2 may be scaled by an appropriate number
5 (e.g., the number of intermediate nodes) and this scaled number may be used to calculate the connectivity rating.

Each equation presented above should be construed as a class of equations of a similar kind, with the actual equation presented being one representative example of the class. For example, the equations presented above include all
10 mathematically equivalent versions of those equations, reductions, simplifications, normalizations, and other equations of the same degree.

The above described embodiments of the invention are presented for purposes of illustration and not of limitation. The following numbered paragraphs give additional embodiments of the present invention.

What is Claimed is:

1. A method for determining the network connectivity between a first node and a second node connected to the first node by at least one path, the method comprising:
 - identifying paths to the second node from the first node within a
5 network community;
 - using processing circuitry to:
 - determine a normalized path weight for each identified path;
 - determine a user connectivity value for each identified path;
 - for each identified path, sum the product of the user
10 connectivity value and the normalized path weight to produce a network connectivity indication; and
 - output the network connectivity indication.
2. The method of claim 1 further comprising accessing a link threshold value, wherein identifying paths to the second node from the first node within a network community comprises identifying only those paths containing fewer links than the accessed link threshold value.
3. The method of claim 1 further comprising accessing a path weight threshold value, wherein identifying paths to the second node from the first node within a network community comprises identifying only those paths with a normalized path weight above the accessed path weight threshold value.
4. The method of claim 1 wherein the control circuitry determines a normalized path weight for each identified path by dividing the product of the link weights of each link in an identified path by the sum of the path weights of all identified paths.
5. The method of claim 1 wherein the control circuitry determines a user connectivity value for each identified path by multiplying the path weight for an identified path and the minimum connectivity value assigned to a link in the identified path.

6. The method of claim 1 wherein the user connectivity value represents at least one of a subjective user trust value or a competency assessment.

7. The method of claim 6 further comprising using the control circuitry to

access third-party ratings data; and

5 in part, on the third-party ratings information.

8. The method of claim 1 wherein identifying paths to the second node from the first node within a network community comprises accessing data from a social networking service.

9. The method of claim 1 wherein identifying paths to the second node from the first node within a network community comprises retrieving a pre-stored identification of the paths to the second node from the first node from a table in a database.

10. The method of claim 1 further comprising automatically making at least one network-based decision based, at least in part, on the network connectivity indication.

11. A system for determining the network connectivity between a first node and a second node connected to the first node by at least one path, the system comprising processing circuitry configured to:

5 identify paths to the second node from the first node within a network community;

determine a normalized path weight for each identified path;

determine a user connectivity value for each identified path;

10 for each identified path, sum the product of the user connectivity value and the normalized path weight to produce a network connectivity indication; and

output the network connectivity indication.

12. The system of claim 11 wherein the processing circuitry is further configured to access a link threshold value, wherein the control circuitry identifies paths to the second node from the first node within a network community by identifying only those paths containing fewer links than the accessed link
5 threshold value.

13. The system of claim 11 wherein the processing circuitry is further configured to access a path weight threshold value, wherein the control circuitry identifies paths to the second node from the first node within a network community by identifying only those paths with a normalized path weight above
5 the accessed path weight threshold value.

14. The system of claim 11 wherein the control circuitry determines a normalized path weight for each identified path by dividing the product of the link weights of each link in an identified path by the sum of the path weights of all identified paths.

15. The system of claim 11 wherein the control circuitry determines a user connectivity value for each identified path by multiplying the path weight for an identified path and the minimum connectivity value assigned to a link in the identified path.

16. The system of claim 11 wherein the user connectivity value represents at least one of a subjective user trust value or a competency assessment.

17. The system of claim 16 wherein the control circuitry is further configured to
access third-party ratings data; and
determine the network connectivity indication based, at least
5 in part, on the third-party ratings information.

18. The system of claim 11 wherein the control circuitry identifies paths to the second node from the first node within a network community by accessing data from a social networking service.

19. The system of claim 11 wherein the control circuitry identifies paths to the second node from the first node within a network community by retrieving a pre-stored identification of the paths to the second node from the first node from a table in a database.

5 20. The system of claim 11 wherein the control circuitry is further configured to automatically making at least one network-based decision based, at least in part, on the network connectivity indication.

21. A system for determining the network connectivity between a first node and a second node connected to the first node by at least one path, the system comprising:

means for identifying paths to the second node from the first node
5 within a network community;
means for determining a normalized path weight for each identified path;
means for determining a user connectivity value for each identified path;
10 means for summing, for each identified path, the product of the user connectivity value and the normalized path weight to produce a network connectivity indication; and
means for outputting the network connectivity indication.

22. The system of claim 21 further comprising means for accessing a link threshold value, wherein the means for identifying paths to the second node from the first node within a network community comprise means for identifying only those paths containing fewer links than the accessed link threshold
5 value.

23. The system of claim 21 further comprising means for accessing a path weight threshold value, wherein the means for identifying paths to the second node from the first node within a network community comprises means for identifying only those paths with a normalized path weight above the accessed
5 path weight threshold value.

24. The system of claim 21 wherein the means for determining a normalized path weight for each identified path comprises means for dividing the product of the link weights of each link in an identified path by the sum of the path weights of all identified paths.

25. The system of claim 21 wherein the means for determining a user connectivity value for each identified path comprises means for multiplying the path weight for an identified path and the minimum connectivity value assigned to a link in the identified path.

26. The system of claim 21 wherein the user connectivity value represents at least one of a subjective user trust value or a competency assessment.

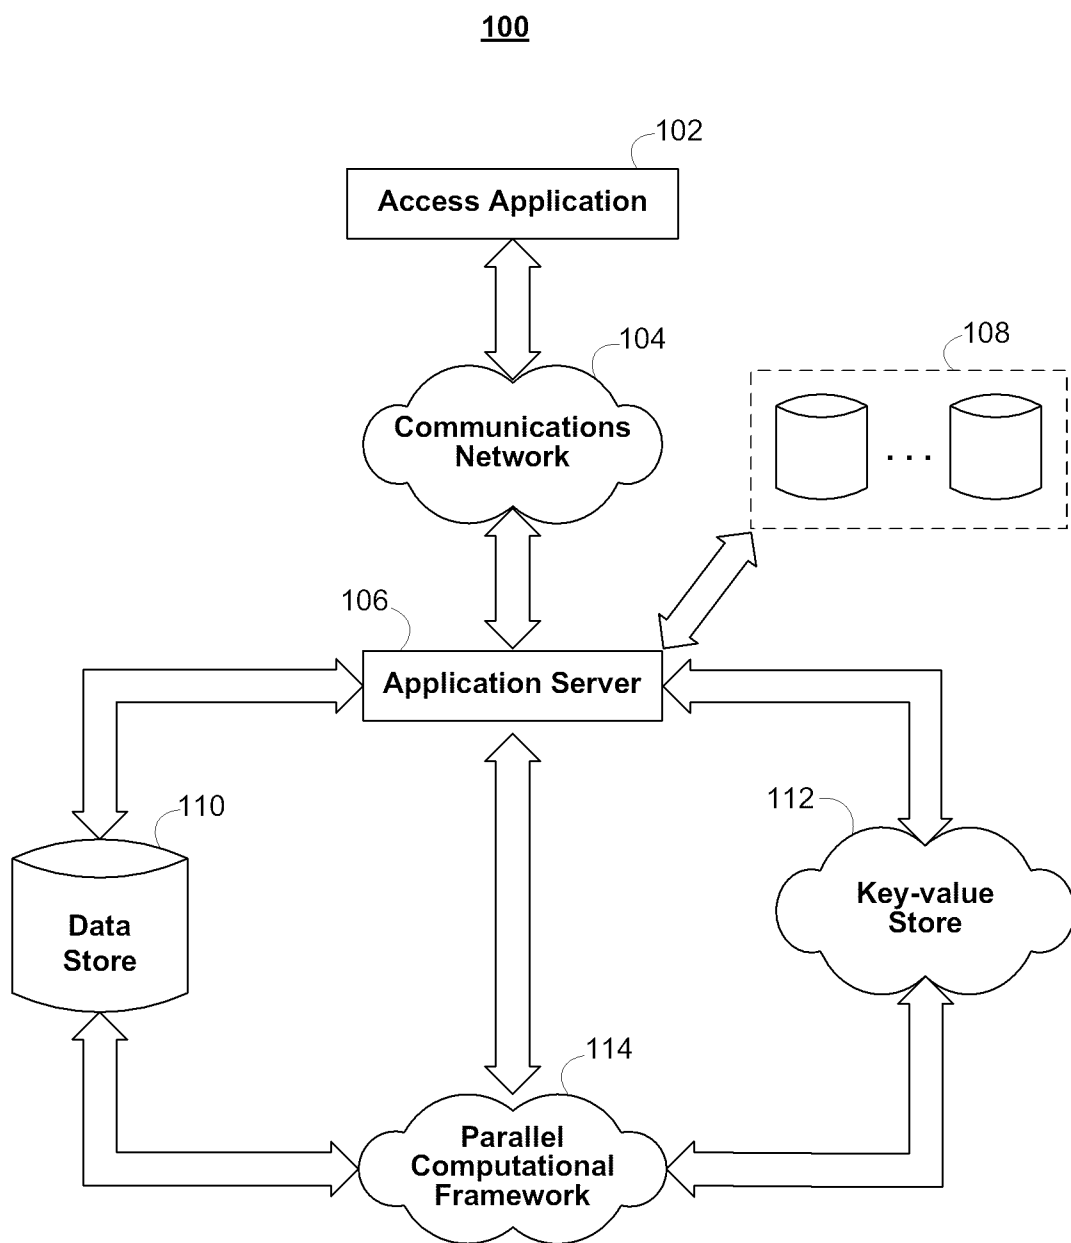
27. The system of claim 26 further comprising:
means for accessing third-party ratings data; and
means for determining the network connectivity indication based, at least in part, on the third-party ratings information.

28. The system of claim 21 wherein the means for identifying paths to the second node from the first node within a network community comprise means for accessing data from a social networking service.

29. The system of claim 21 wherein the means for identifying paths to the second node from the first node within a network community comprises means for retrieving a pre-stored identification of the paths to the second node from the first node from a table in a database.

5 30. The system of claim 21 further comprising means for automatically making at least one network-based decision based, at least in part, on the network connectivity indication.

1/10

**FIG. 1**

200

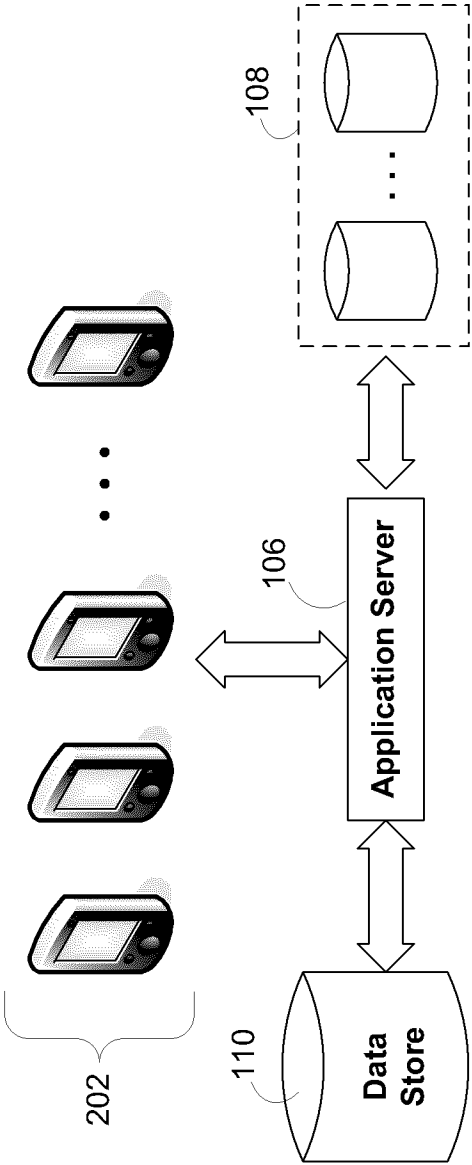
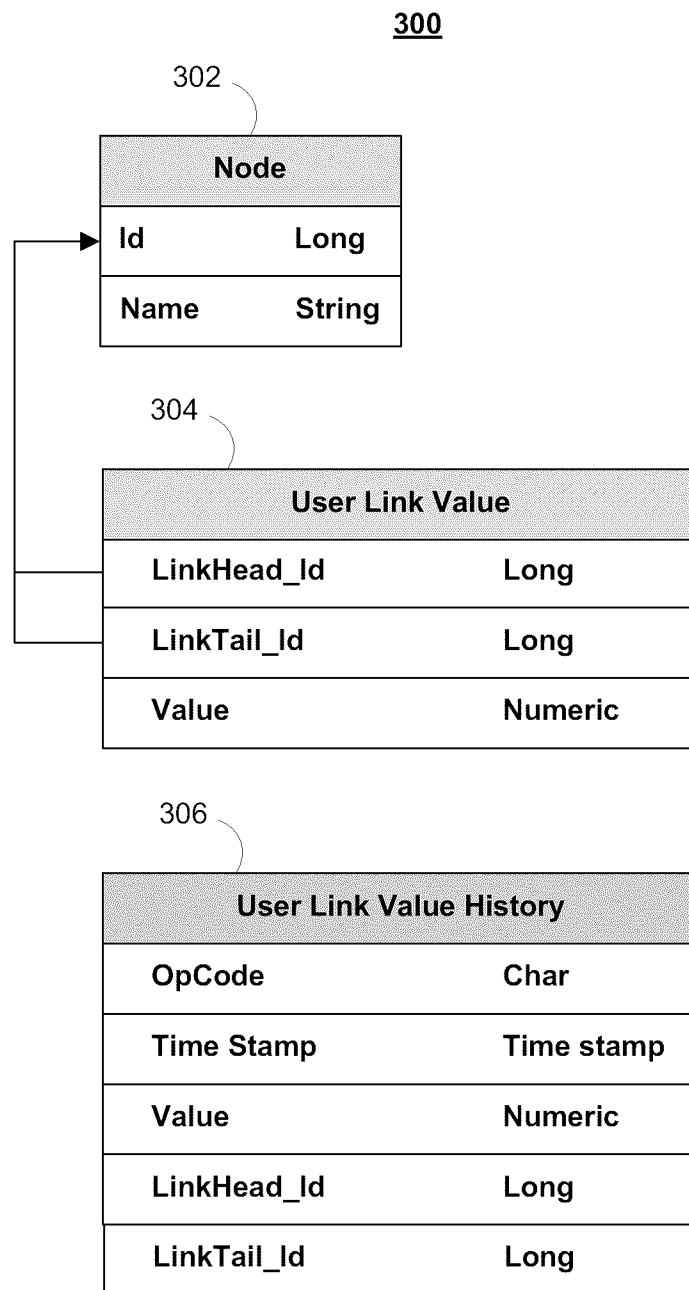


FIG. 2

3/10

**FIG. 3A**

4/10

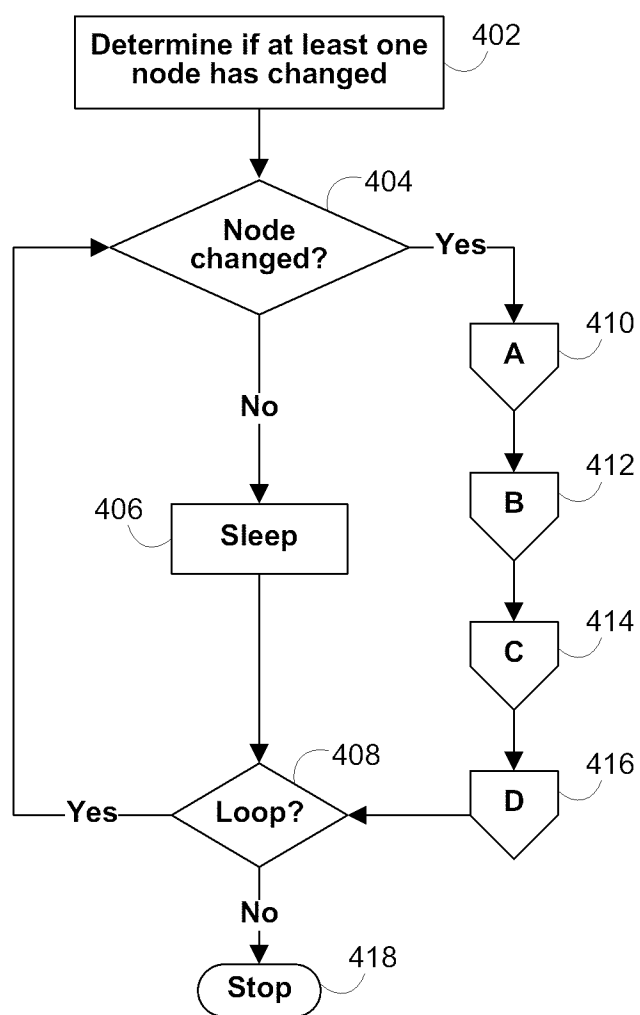
310

312

Node Table	
314	RowId 64-bit Integer
316	"info:inlinks" List
318	"info:outlinks" List
320	"info:oldHeadIdBuckets" List
322	"bucket:" + target node id List

FIG. 3B

5/10

400**FIG. 4A**

6/10

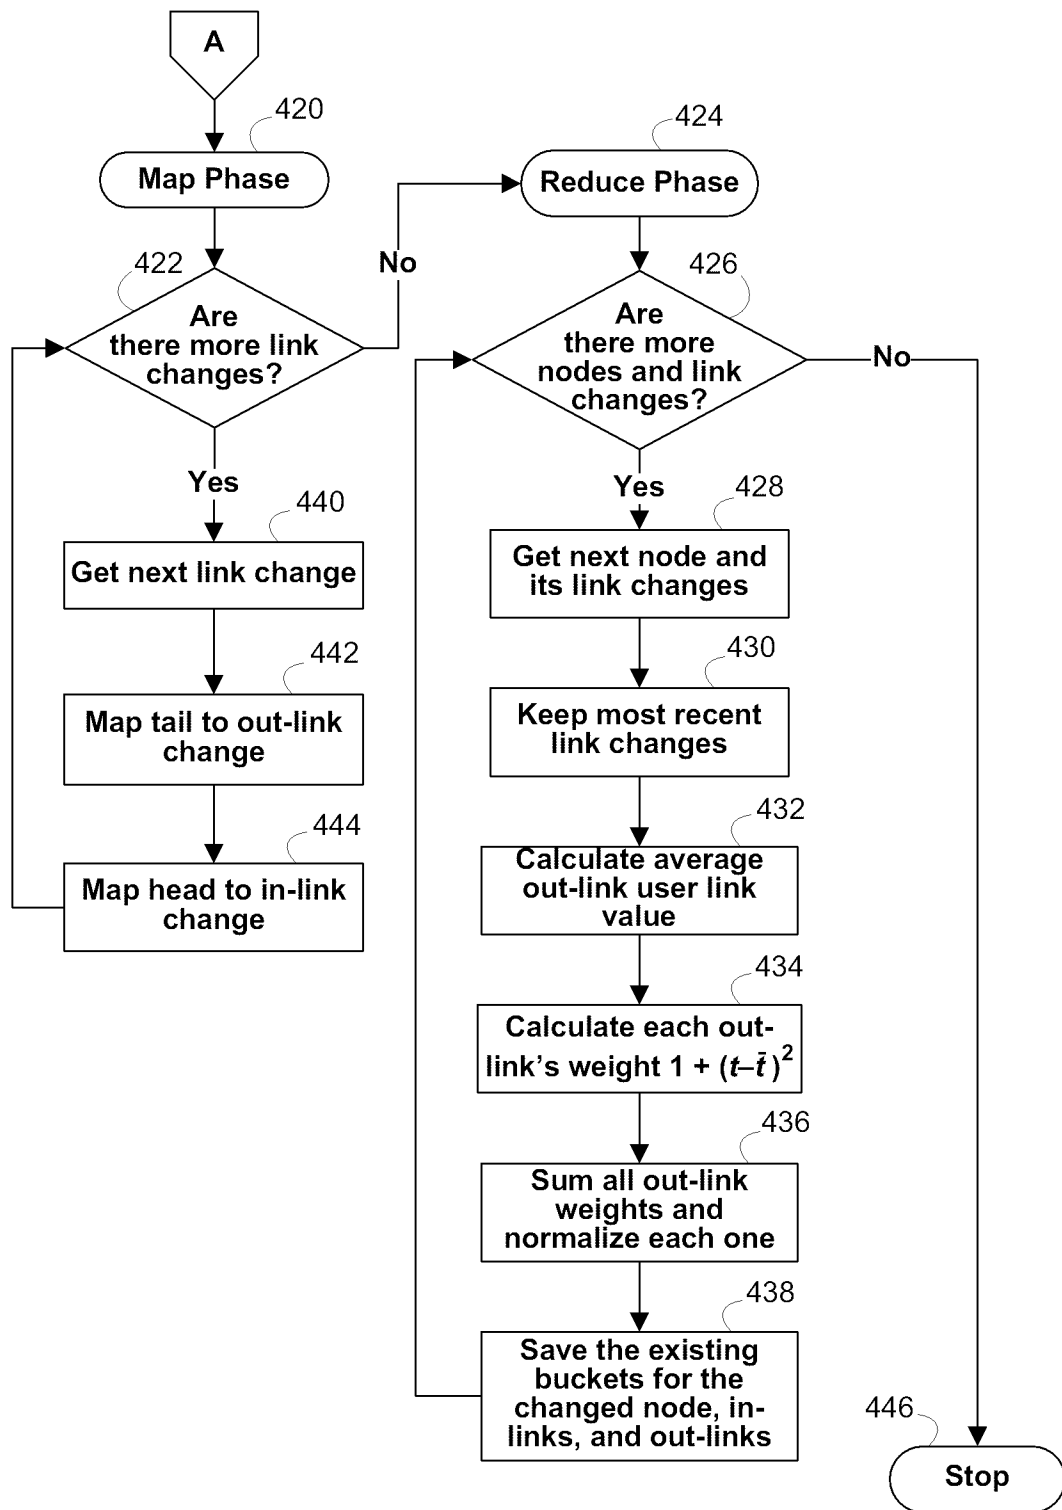


FIG. 4B

7/10

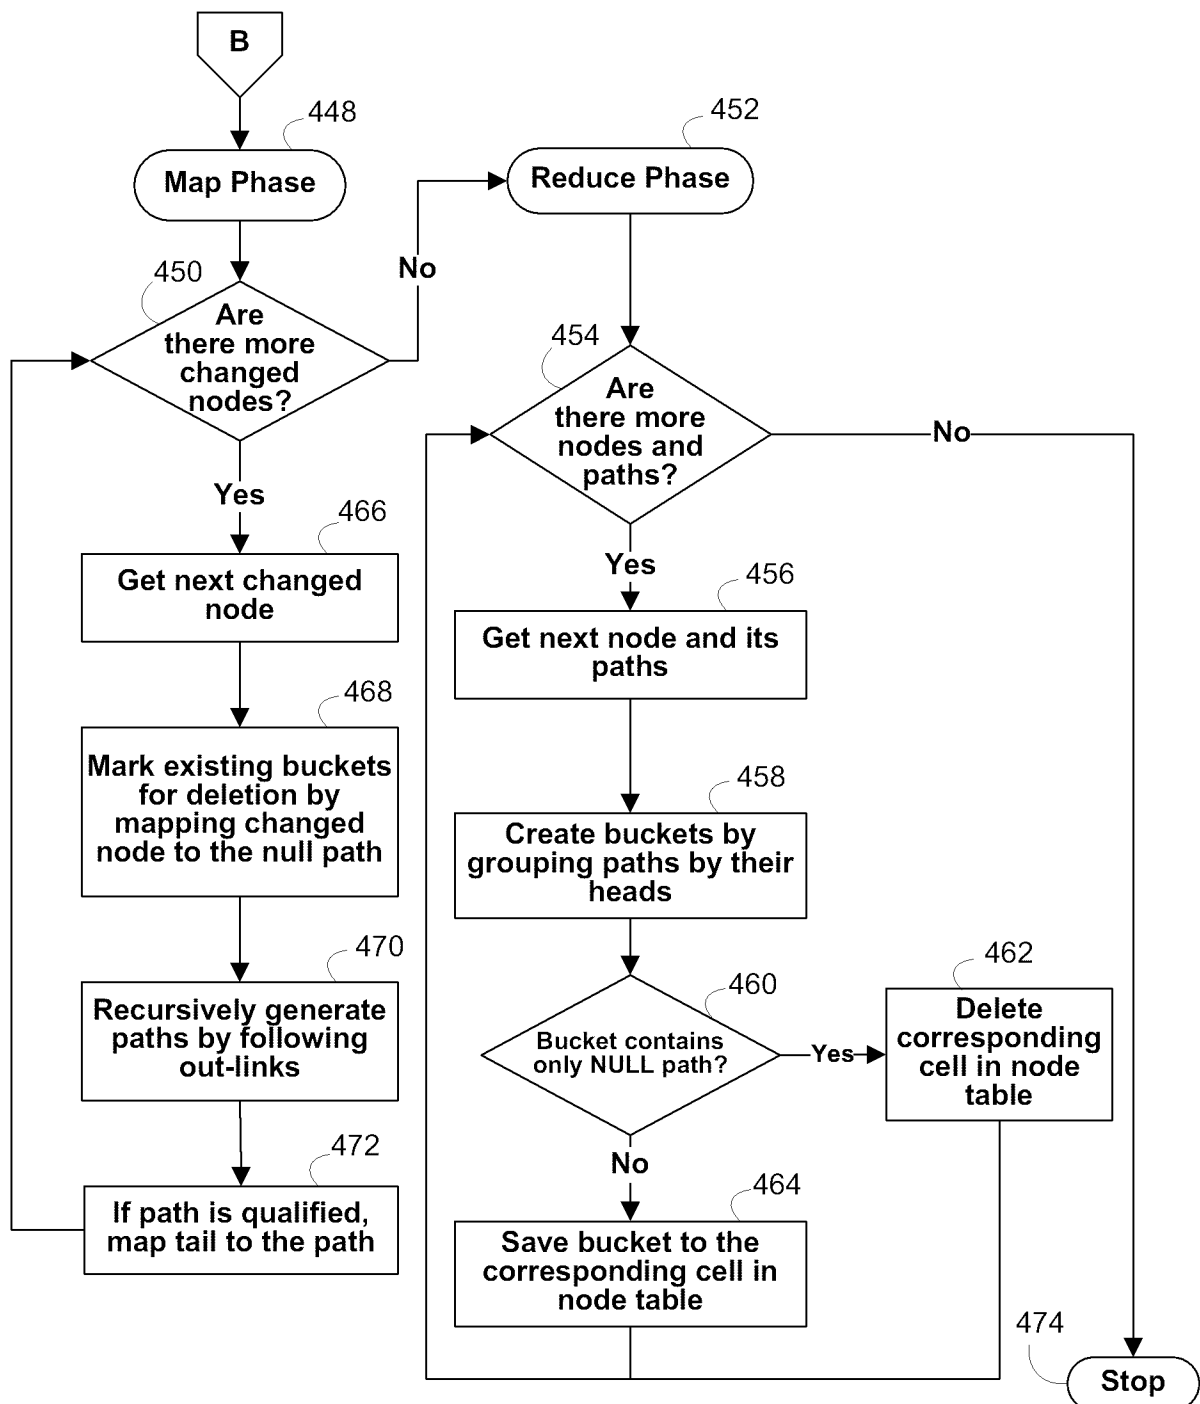


FIG. 4C

8/10

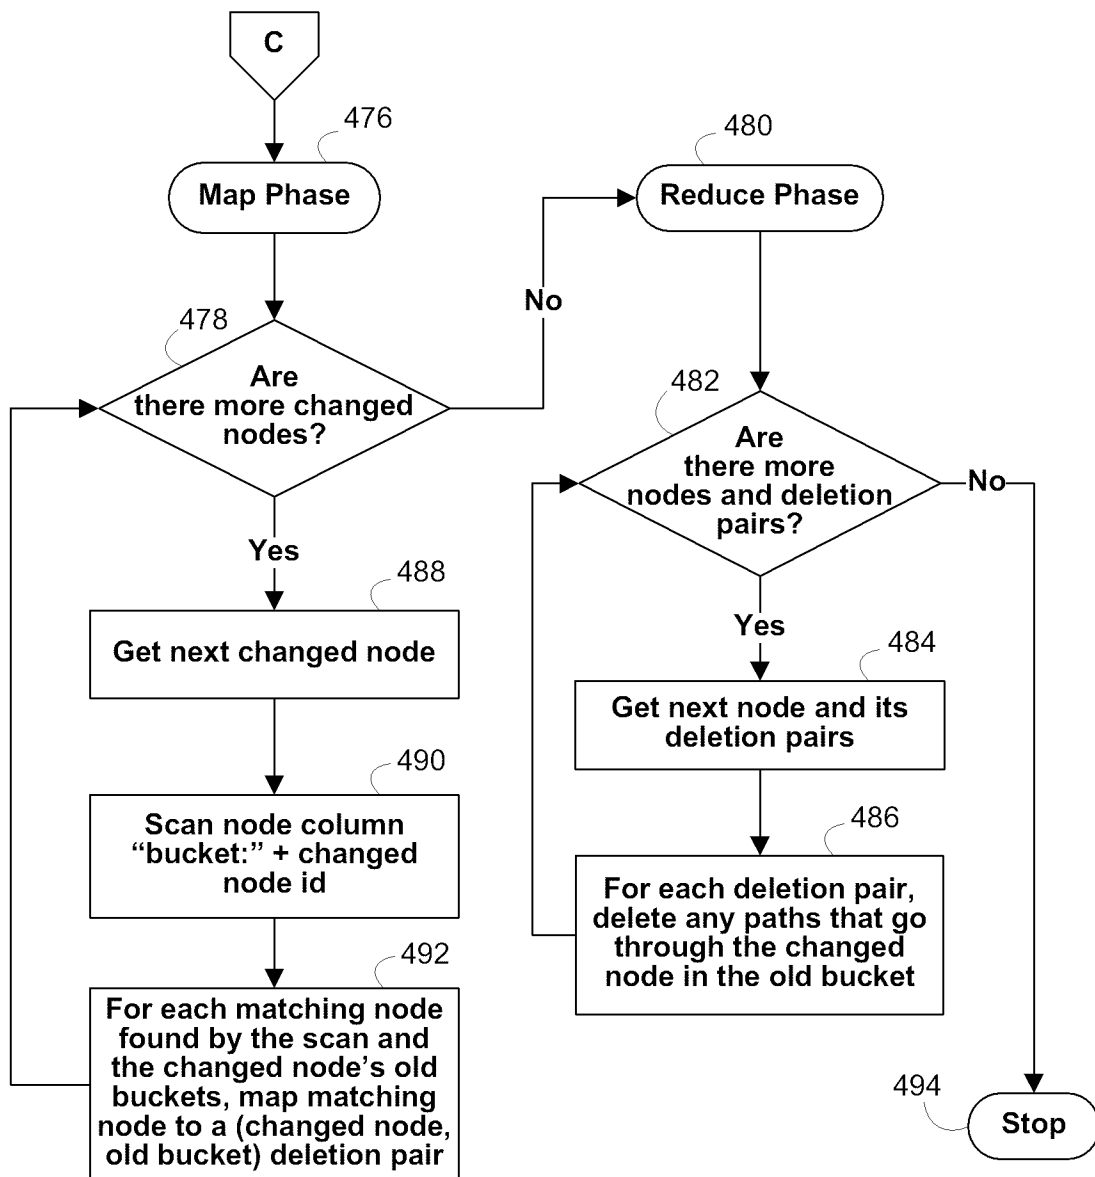


FIG. 4D

9/10

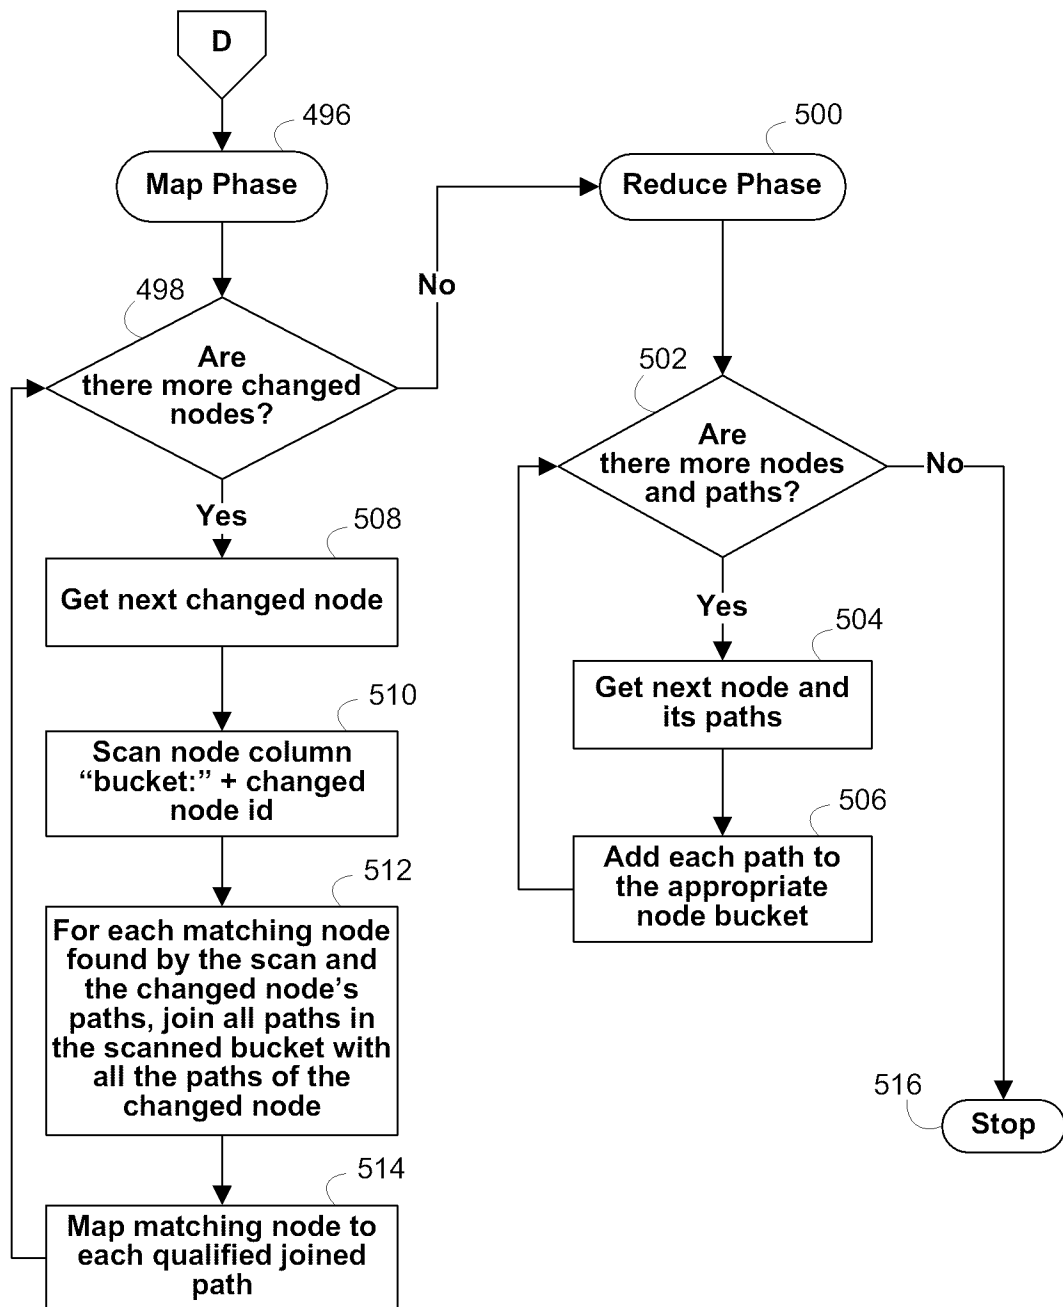


FIG. 4E

10/10

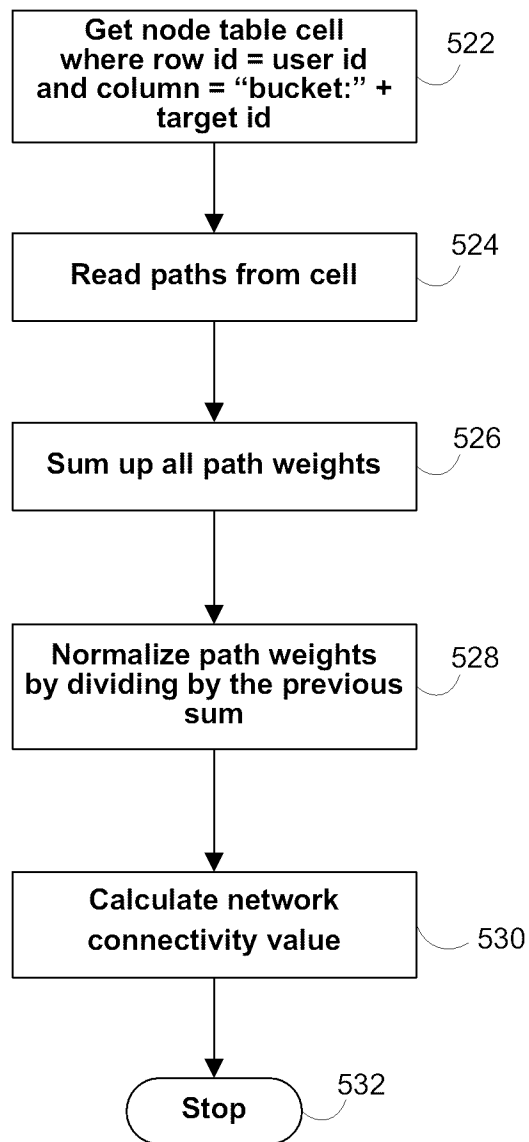
520

FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/CA2010/001531

A. CLASSIFICATION OF SUBJECT MATTER

IPC: **H04L 12/26** (2006.01) , **H04L 12/56** (2006.01)

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)

IPC: H04L 12/26 (2006.01), H04L 12/56 (2006.01)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic database(s) consulted during the international search (name of database(s) and, where practicable, search terms used)

Databases: EPOQUE (Epodoc, English Full-text), Canadian Patents Database, IEEEExplore, Google

Keywords: social graph/network, data analytics, determine/calculate connectivity/trust score/value, network/community, trust/reputation/status/membership/influence/competence, key-value store, map/reduce parallel computational framework/paradigm, link/path weight(s), threshold, normalized path weight, third-party ratings

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y	MUI, L. et al., "A Computational Model of Trust and Reputation", Proceedings of the 35th Annual Hawaii International Conference on System Sciences, HICSS '02, vol. 7, pages 2431-2439, 7-10 January 2002 (07-01-2002) *abstract; sections 1, 5-7*	1,4-6,8,10-11,14-16,18,20-21,24-26,28,30 2-3,7,9,12-13,17,19,22-23,27,29
X Y	SAFAEI, M. et al., "Social Graph Generation & Forecasting Using Social Network Mining", Proceedings of the 33rd Annual IEEE International Computer Software and Applications Conference, COMPSAC '09, pages 31-35, 20-24 July 2009 (20-07-2009) *whole document*	1,5-6,8,10-11,15-16,18,20-21,25-26,28,30 2-3,7,9,12-13,17,19,22-23,27,29
Y	US 2004/0122803 A1 (DOM et al.) 24 June 2004 (24-06-2004) *abstract; paragraphs [0045], [0070]-[0082]*	2-3,9,12-13,19,22-23,29
Y	HUYNH, T.D. et al., "An Integrated Trust and Reputation Model for Open Multi-Agent Systems", Journal of Autonomous Agents and Multi-Agent Systems, vol. 13, issue 2, September 2006 (09-2006) *abstract; sections 1, 3, 4; Appendix A*	7,17,27

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :	"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
"A" document defining the general state of the art which is not considered to be of particular relevance	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"&" document member of the same patent family
"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	

Date of the actual completion of the international search

9 December 2010 (09-12-2010)

Date of mailing of the international search report

10 January 2011 (10-01-2011)

Name and mailing address of the ISA/CA
Canadian Intellectual Property Office
Place du Portage I, C114 - 1st Floor, Box PCT
50 Victoria Street
Gatineau, Quebec K1A 0C9
Facsimile No.: 001-819-953-2476

Authorized officer

Daniela Savin (819) 934-4890

INTERNATIONAL SEARCH REPORT

International application No.
PCT/CA2010/001531

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	CHAKRABORTY, S. et al., "TrustBAC - Integrating Trust Relationships into the RBAC Model for Access Control in Open Systems", Proceedings of the eleventh ACM symposium on Access Control Models and Technologies, SACMAT '06, pages 49-58, 7-9 June 2006 (07-06-2006) *whole document*	1-30
A	GAN, Z. et al., "A Novel Reputation Computing Model for Mobile Agent-Based E-Commerce Systems", Proceedings of the International Conference on Information Security and Assurance, ISA 2008, pages 253-260, 24-26 April 2008 (24-04-2008) *whole document*	1-30
A	US 2009/0106822 A1 (OBASANJO et al.) 23 April 2009 (23-04-2009) *whole document*	1-30
A	US 7 512 612 B1 (AKELLA et al.) 31 March 2009 (31-03-2009) *whole document*	1-30
A	GOLBECK, J. et al., "Inferring Trust Relationships in Web-based Social Networks", Journal of ACM Transactions on Internet Technology (TOIT), vol. 6, issue 4, November 2006 (11-2006) *whole document*	1-30
A	US 2005/0096987 A1 (MIYAUCHI) 5 May 2005 (05-05-2005) *whole document*	1-30
A	WO 2009/109009 A1 (KOSHY et al.) 11 September 2009 (11-09-2009) *whole document*	1-30

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/CA2010/001531

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
US2004122803A1	24 June 2004 (24-06-2004)	None	
US2009106822A1	23 April 2009 (23-04-2009)	CN101836230A EP2203890A2 WO2009055241A2 WO2009055241A3	15 September 2010 (15-09-2010) 07 July 2010 (07-07-2010) 30 April 2009 (30-04-2009) 16 July 2009 (16-07-2009)
US7512612B1	31 March 2009 (31-03-2009)	US7512612B1 US7539697B1	31 March 2009 (31-03-2009) 26 May 2009 (26-05-2009)
US2005096987A1	05 May 2005 (05-05-2005)	CN1619567A EP1528498A2 EP1528498A3 JP2005135071A KR20050040805A US7707068B2	25 May 2005 (25-05-2005) 04 May 2005 (04-05-2005) 02 August 2006 (02-08-2006) 26 May 2005 (26-05-2005) 03 May 2005 (03-05-2005) 27 April 2010 (27-04-2010)
WO2009109009A1	11 September 2009 (11-09-2009)	AU2009221644A1	11 September 2009 (11-09-2009)