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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
4 August 2016 (04.08.2016)

WIPO | PCT

(10) International Publication Number
WO 2016/123561 A1

(51) International Patent Classification:
H04L 12/715 (2013.01) *H04W 84/18* (2009.01)
H04W 56/00 (2009.01)

(21) International Application Number:
PCT/US2016/015802

(22) International Filing Date:
29 January 2016 (29.01.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

62/110,383	30 January 2015 (30.01.2015)	US
15/009,711	28 January 2016 (28.01.2016)	US
15/009,715	28 January 2016 (28.01.2016)	US

(71) Applicant: **SILVER SPRING NETWORKS, INC.** [US/US]; 555 Broadway Street, Redwood City, California 94063 (US).

(72) Inventor: **FLAMMER, George H., III**; 555 Broadway Street, Redwood City, California 94063 (US).

(74) Agents: **CAREY, John C.** et al.; Artegis Law Group, LLP, 7710 Cherry Park Drive, Suite T104, Houston, Texas 77095 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,

[Continued on next page]

WO 2016/123561 A1

(54) Title: TECHNIQUES FOR MANAGING HETEROGENEOUS NODES CONFIGURED TO SUPPORT A HOMOGENEOUS COMMUNICATION PROTOCOL

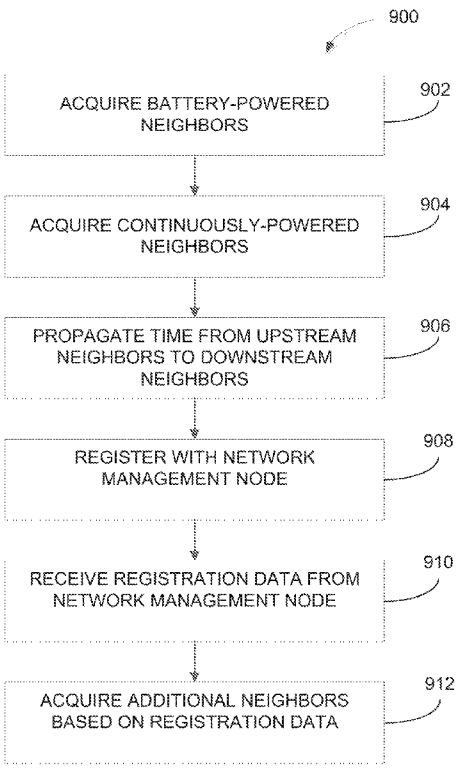


FIG. 9

(57) Abstract: A wireless mesh network includes heterogeneous types of nodes, including continuously-powered nodes and battery-powered nodes. The battery-powered nodes may reside in a sleeping state most of the time to conserve power. The various nodes in the network may communicate with one another by transmitting and receiving at scheduled times and on scheduled frequencies. The battery-powered nodes may become active during the scheduled transmit and receive times. Network management nodes may facilitate network formation by transmitting information that reflects the scheduled transmit and receive times across the network. Based on this data, the continuously-powered nodes and battery-powered nodes may establish communication links with one another.

SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

TECHNIQUES FOR MANAGING HETEROGENEOUS NODES CONFIGURED TO SUPPORT A HOMOGENEOUS COMMUNICATION PROTOCOL

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of United States provisional patent application titled "Techniques for Managing Heterogeneous Nodes Configured to Support a Homogeneous Communication Protocol," filed on January 30, 2015, and having serial number 62/110,383, United States patent application titled "Techniques for Managing Heterogeneous Nodes Configured to Support a Homogeneous Communication Protocol," filed on January 28, 2016, and having serial number 15/009,711, and of United States patent application titled "Techniques for Managing Heterogeneous Nodes Configured to Support a Homogeneous Communication Protocol," filed on January 28, 2016, and having serial number 15/009,715. The subject matter of each of these related applications is hereby incorporated herein by reference.

BACKGROUND OF THE INVENTION

Field of the Invention

[0002] Embodiments of the present invention relate generally to wireless network communications and, more specifically, to techniques for managing heterogeneous nodes configured to support a homogeneous communication protocol.

Description of the Related Art

[0003] A conventional network includes a plurality of nodes configured to communicate with one another across various communication links. In a first type of network, nodes within the network draw power from an electricity grid in order to perform various operations, such as communication operations and data processing operations. In a second type of network, nodes within the network draw power from batteries to support communication and data processing operations.

[0004] Continuously-powered nodes and battery-powered nodes generally rely on different communication protocols, because continuously-powered nodes and battery-powered nodes have different resource limitations. In particular, battery-powered nodes typically communicate via protocols that conserve power, while continuously-powered nodes typically do not. Consequently, continuously-powered nodes cannot

communicate with battery-powered nodes, and so conventional networks usually do not include both continuously-powered nodes and battery-powered nodes. Generally, different types of nodes do not employ similar communication protocols, and, therefore, networks typically do not include heterogeneous types of nodes. Such constraints place a substantial obstacle on the growing "Internet of Things" (IoT). Specifically, the IoT is meant to interconnect "things," such as nodes, of all different types. However, as set forth above, communication barriers impede the ability to interconnect heterogeneous nodes, such as continuously-powered nodes and battery-powered nodes. Such impediments jeopardize the ability of the IoT to continue growing.

[0005] As the foregoing illustrates, what is needed in the art is a technique for managing communication between different types of nodes that reside within a network.

SUMMARY OF THE INVENTION

[0006] One embodiment of the present invention sets forth a computer-implemented method for establishing a communication link between different network portions to generate a wireless mesh network, including identifying a first node within a first network portion, identifying a second node within a second network portion, where the first network portion and the second network portion include separate and distinct nodes, obtaining first network data that is associated with the first network portion and includes connectivity information related to the nodes included in the first network portion, and transmitting the first network data to the second node to enable the second node to communicate with the first node based on the first network data.

[0007] At least one advantage of the techniques described herein is that heterogeneous types of nodes, such as continuously-powered nodes and battery-powered nodes, may reside within the same wireless mesh network.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] So that the manner in which the above recited features of the present invention can be understood in detail, a more particular description of the invention, briefly summarized above, may be had by reference to embodiments, some of which are illustrated in the appended drawings. It is to be noted, however, that the appended drawings illustrate only typical embodiments of this invention and are therefore not to

be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

[0009] Figure 1 illustrates a network system configured to implement one or more aspects of the present invention;

[0010] Figure 2 illustrates a network interface configured to transmit and receive data within a mesh network, according to one embodiment of the present invention;

[0011] Figure 3 is a more detailed illustration of the network management node of Figure 1 configured to manage the operation of a portion of the network system, according to one embodiment of the present invention;

[0012] Figure 4 is a more detailed illustration of the server of Figure 1 configured to interoperate with the network management node, according to one embodiment of the present invention;

[0013] Figure 5 is a state diagram illustrating operating states associated with the different nodes of Figure 1, according to one embodiment of the present invention;

[0014] Figure 6 illustrates a time beacon for sharing time between the different nodes of Figure 1, according to one embodiment of the present invention;

[0015] Figures 7A-7B are conceptual diagrams illustrating a technique for performing data communication between the different nodes of Figure 1, according to one embodiment of the present invention;

[0016] Figure 8 illustrates a set of transmit parameters according to which a node transmits data within the network system of Figure 1, according to one embodiment of the present invention;

[0017] Figure 9 is a flow diagram of method steps for participating in a wireless mesh network, according to one embodiment of the present invention;

[0018] Figures 10A-10E are conceptual diagrams illustrating the formation of the wireless mesh network of Figure 1, according to one embodiment of the present invention;

[0019] Figures 11A-11D are conceptual diagrams illustrating network management nodes interoperating to couple together disjoint network portions to form the network system of Figure 1, according to one embodiment of the present invention; and

[0020] Figure 12 is a flow diagram of method steps for sharing connectivity information between disjoint network portions in order to form a network system, according to one embodiment of the present invention.

DETAILED DESCRIPTION

[0021] In the following description, numerous specific details are set forth to provide a more thorough understanding of the present invention. However, it will be apparent to one of skill in the art that the present invention may be practiced without one or more of these specific details. In other instances, well-known features have not been described in order to avoid obscuring the present invention.

System Overview

[0022] Figure 1 illustrates a network system 100 configured to implement one or more aspects of the present invention. As shown, network system 100 includes a wireless mesh network 102, which may include a source node 110, nodes 120, 130, and 140, and destination node 112. Nodes 120 represent continuously-powered nodes that are coupled to a power grid. Nodes 130 represent battery-powered nodes that draw power from a battery. Nodes 140 are network management nodes that are configured to manage the operation of portions 104 of wireless mesh network 102. Specifically, node 140-1 is configured to manage nodes 120-1, 130-1, and 130-2 within portion 104-1, and node 140-2 is configured to manage nodes 120-2, 130-3, and 130-4 within portion 104-2. Nodes 140-1 and 140-2 are also configured to communicate with one another in order to perform cross-mesh routing between portions 104-1 and 104-2.

[0023] Source node 110 is able to communicate with certain nodes 120 and 130 via communication links 132. Nodes 120, 130, and 140 communicate amongst themselves via communication links 134. Nodes 120 and 130 communicate with destination node 112 via communication links 136. Network system 100 may also include one or more access points (APs) 150, a network 152, and a server 160. A discovery protocol may be implemented to determine node adjacency to one or more

adjacent nodes. For example, node 120-1 may execute the discovery protocol to determine that nodes 110, 140-1, and 130-1 are adjacent to node 120-1. Furthermore, this node adjacency indicates that communication links 132 and 134 may be established. Any technically feasible discovery protocol may be implemented without departing from the scope and spirit of embodiments of the present invention.

[0024] The discovery protocol may also be implemented to determine the hopping sequences of adjacent nodes, i.e. the sequence of channels across which nodes periodically receive payload data. As is known in the art, a "channel" may correspond to a particular range of frequencies. Once adjacency is established between the source node 110 and at least one node 120, 130, or 140, source node 110 may generate payload data for delivery to the destination node 112, assuming a path is available. The payload data may comprise an Internet protocol (IP) packet, an Ethernet frame, or any other technically feasible unit of data. Similarly, any technically feasible addressing and forwarding techniques may be implemented to facilitate delivery of the payload data from the source node 110 to the destination node 112. For example, the payload data may include a header field configured to include a destination address, such as an IP address or Ethernet media access control (MAC) address.

[0025] Each node 120 and 130 may be configured to forward the payload data based on the destination address. Alternatively, the payload data may include a header field configured to include at least one switch label to define a predetermined path from source node 110 to destination node 112. A forwarding database may be maintained by each node 120 and 130 that indicates which communication link 132, 134, 136 should be used and in what priority to transmit the payload data for delivery to the destination node 112. The forwarding database may be supplied by one of nodes 140 in response to a registration procedure. The forwarding database may represent multiple paths to the destination address, and each of the multiple paths may include one or more cost values. Any technically feasible type of cost value may characterize a link or a path within network system 100. In one embodiment, each node within wireless mesh network 102 implements substantially identical functionality, and each node may act as a source node, destination node, intermediate node, or network management node.

[0026] In network system 100, access points 150-1 and 150-2 are configured to communicate with at least one node within wireless mesh network 102, such as intermediate node 120-1 or 130-3. Communication may include transmission of payload data, time-related data, or any other technically relevant data between an access point 150 and the at least one node within the wireless mesh network 102. For example, a communications link may be established between access point 150-1 and intermediate node 130-3 to facilitate transmission of payload data between portion 104-2 and network 152. Network 152 is coupled to the server 160 and configured to provide a communication path between wireless mesh network 102 and server 160. Network 152 may comprise any wired, optical, wireless, or hybrid network configured to transmit payload data between access points 150 and server 160.

[0027] In one embodiment, the server 160 represents a destination for payload data originating within the wireless mesh network 102 and a source of payload data destined for one or more nodes within wireless mesh network 102. For example, nodes within wireless mesh network 102 may perform measurements to generate measurement data, such as power consumption data. Server 160 may execute an application to collect the measurement data and report the measurement data. In one embodiment, server 160 queries nodes within wireless mesh network 102 for certain data. Each queried node replies with requested data, such as consumption data, system status and health data, and so forth. In an alternative embodiment, each node within wireless mesh network 102 autonomously reports certain data, which is collected by server 160 as the data becomes available via autonomous reporting.

[0028] The techniques described herein are sufficiently flexible to be utilized within any technically feasible network environment including, without limitation, a wide-area network (WAN) or a local-area network (LAN). Moreover, multiple network types may exist within a given network system 100. For example, communications between two nodes 130 or between a node 130 and a corresponding access point 150 may be via a radio-frequency local-area network (RF LAN), while communications between access points 150 and network 152 may be via a WAN such as a general packet radio service (GPRS). Each node 120, 130, and 140 within wireless mesh network 102 includes a network interface that enables the node to communicate wirelessly with other nodes, as described in greater detail below in conjunction with Figure 2.

[0029] Figure 2 illustrates a network interface configured to transmit and receive data within a mesh network, according to one embodiment of the present invention. Each node 110, 112, 120, 130, and 140 within wireless mesh network 102 of Figure 1 includes at least one instance of network interface 200. In continuously-powered nodes 120, network interface 200 may draw power from an external power source, such as the power grid, to perform the various functionalities described below. In battery-power nodes 130, network interface 200 may be coupled to a battery included within the node 130, and network interface 200 may draw power from the batter to perform the various functionalities described below. In the latter case, network interface 200 may become active only intermittently, and at scheduled times, in order to receive and/or transmit data. In this fashion, a battery-powered node 130 that includes network interface 200 may conserve power. This power-saving technique is described in greater detail below in conjunction with Figures 5-7B.

[0030] As shown, network interface 200 includes, without limitation, a microprocessor unit (MPU) 210, a digital signal processor (DSP) 214, digital to analog converters (DACs) 220, 221, analog to digital converters (ADCs) 222, 223, analog mixers 224, 225, 226, 227, a phase shifter 232, an oscillator 230, a power amplifier (PA) 242, a low noise amplifier (LNA) 240, an antenna switch 244, and an antenna 246. A memory 212 may be coupled to the MPU 210 for local program and data storage. Similarly, a memory 216 may be coupled to the DSP 214 for local program and data storage. Memory 212 and/or memory 216 may be used to buffer incoming data as well as store data structures such as, e.g., a forwarding database, and/or routing tables that include primary and secondary path information, path cost values, and so forth.

[0031] In one embodiment, the MPU 210 implements procedures for processing IP packets transmitted or received as payload data by the network interface 200. The procedures for processing the IP packets may include, without limitation, wireless routing, encryption, authentication, protocol translation, and routing between and among different wireless and wired network ports. In one embodiment, MPU 210 implements the techniques performed by the node that includes network interface 200, when MPU 210 executes a firmware program stored in memory.

[0032] MPU 214 is coupled to DAC 220 and DAC 221. Each DAC 220, 221 is configured to convert a stream of outbound digital values into a corresponding analog

signal. The outbound digital values are computed by the signal processing procedures for modulating one or more channels. MPU 214 is also coupled to ADC 222 and ADC 223. Each ADC 222, 223 is configured to sample and quantize an analog signal to generate a stream of inbound digital values. The inbound digital values are processed by the signal processing procedures to demodulate and extract payload data from the inbound digital values.

[0033] In one embodiment, MPU 210 and/or MPU 214 are configured to buffer incoming data within memory 212 and/or memory 216. The incoming data may be buffered in any technically feasible format, including, for example, raw soft bits from individual channels, demodulated bits, raw ADC samples, and so forth. MPU 210 and/or MPU 214 may buffer within memory 212 and/or memory 216 any portion of data received across the set of channels from which antenna 246 receives data, including all such data. MPU 210 and/or MPU 214 may then perform various operations with the buffered data, including demodulation operations, decoding operations, and so forth.

[0034] Persons having ordinary skill in the art will recognize that network interface 200 represents just one possible network interface that may be implemented within wireless mesh network 102 shown in Figure 1, and that any other technically feasible device for transmitting and receiving data may be incorporated within any of the nodes within wireless mesh network 102.

[0035] Figure 3 is a more detailed illustration of the network management node of Figure 1 configured to manage the operation of a portion of the network system, according to one embodiment of the present invention. As shown, network management node 140 includes a processor 300, input/output (I/O) devices 310, and a memory 320. Memory 320 includes a management application 322, network data 324, and node data 326.

[0036] Processor 300 may be any technically feasible hardware unit configured to execute applications and process data. Processor 300 could be, for example, a central processing unit (CPU), application specific integrated circuit (ASIC), microprocessor, and so forth, without limitation. I/O devices 310 may include any technically feasible devices configured to receive and/or provide data. I/O devices 310 could include, for example, input devices such as receivers, output devices such

as transmitters, or combination input/output devices such as transceivers. In one embodiment, I/O devices 310 include an instance of network interface 200. Memory 320 may be any technically feasible computer-readable and/or writable storage medium. Memory 320 could include, for example, random access memory (RAM), flash memory, read-only memory (ROM), and so forth, without limitation.

[0037] Management application 322 is a software program that, when executed by processor 300, causes node 140 to manage a network portion 104. In doing so, management application 322 causes node 140 to transmit network data 324 and node data 326 to other nodes within the network portion 104. Network data 324 generally includes routing information, including neighbor tables and forwarding databases, among other network-oriented data associated with network portion 104. Node data 326 generally includes operational parameters associated with the nodes in network portion 104. The operational parameters could include transmit parameters, receive parameters, and other parameters that influence node functionality across network portion 104.

[0038] A given node 120 or 130 may interact with a node 140 in order to obtain network data 324 and/or node data 326. Based on that data, the receiving node may then discover and acquire neighbors within network portion 104, perform routing operations, and generally participate in wireless mesh network 102. The process of obtaining the aforesaid data may occur in the context of a registration process initiated by a node 120 or 130 with node 140. Node 140 may provide network data 324 and node data 326 to each node within network portion 104 in order to facilitate the formation of that portion. This process is described in greater detail below in conjunction with Figures 11A-12. In one embodiment, node 140 receives network data 324 and node data 326 from server 160. An exemplary implementation of server 160 is described in greater details below in conjunction with Figure 4.

[0039] Figure 4 is a more detailed illustration of the server of Figure 1 configured to interoperate with the network management node, according to one embodiment of the present invention. As shown, server 160 includes a processor 400, input/output (I/O) devices 410, and a memory 420. Memory 420 includes a network application 422, network data 424, and node data 426.

[0040] Processor 400 may be any technically feasible hardware unit configured to execute applications and process data. Processor 400 could be, for example, CPU, ASIC, microprocessor, and so forth, without limitation. I/O devices 410 may include any technically feasible devices configured to receive and/or provide data. I/O devices 410 could include, for example, input devices such as receivers, output devices such as transmitters, or combination input/output devices such as transceivers. Memory 420 may be any technically feasible computer-readable and/or writable storage medium. Memory 420 could include, for example, RAM, flash memory, ROM, and so forth, without limitation.

[0041] Network application 324 is a software program that, when executed by processor 400, coordinates the functionality of server 160. In doing so, network application 422 causes server 160 to transmit network data 424 and node data 426 to nodes within wireless mesh network 102. Server 160 may transmit different network data 424 to each different network portion 104. For example, network data 324 shown in Figure 3 may reflect a portion of network data 424 that is associated with a particular network portion 104. Network data 424 generally includes routing information, including neighbor tables and forwarding tables, among other network-oriented data. Node data 426 may reflect operational parameters that include transmit parameters, receive parameters, and other parameters that influence node functionality across wireless mesh network 102.

[0042] A given node 140 may interact with server 160 in order to obtain network data 424 and/or node data 426. Node 140 may then manage the nodes within a network portion 104 by providing those nodes with the aforesaid network and/or node data, as discussed above on conjunction with Figure 3. In this manner, server 160 manages the operation of nodes 140, which, in turn, manage the operation of nodes 120 and 130 within a corresponding network portion 104. Nodes 120 and 130 are configured to operate according to a specific set of operating states, as described below in conjunction with Figure 5.

Node Functional Overview and Communication Protocol

[0043] Figure 5 is a state diagram illustrating operating states associated with the different nodes of Figure 1, according to one embodiment of the present invention. Any of nodes 110, 112, 120, 130, or 140 may implement state diagram 500. In some

embodiments, network management nodes 140 may implement a different state diagram that is not described herein. As shown, state diagram 500 includes five operating states: sleep state 510, sniff state 520, receive state 530, process states 540, and transmit state 550. Although several exemplary state transitions are shown, certain other state transitions have been omitted, for clarity. In practice, any of the states within state diagram 500 may transition to any of the other states under circumstances that may not be explicitly disclosed herein. As described in greater detail below, although continuously-powered nodes 120 and battery-powered nodes 130 alike may implement state diagram 500, those different types of nodes may spend differing amounts of time in each state. In particular, battery-powered nodes 130 may spend much more time in sleep state 510 compared to continuously-powered nodes 130. In this fashion, battery-powered nodes 130 may conserve power.

[0044] Sleep state 510 is a low-power operating mode during which some or most of the hardware within the node is powered down. The node may enter sleep state 510 periodically depending on the power resources and power consumption of the node. For example, a continuously-powered node 120 may enter sleep state 510 rarely or never, given that the node 120 has access to plentiful power resources. Conversely, a battery-powered node 130 may enter sleep state 510 often, given that the node 130 only has access to limited battery power resources. In practice, continuously-powered nodes 120 may not enter sleep state 510 at all under normal operating conditions, while battery-powered nodes 130 may reside in sleep state 510 most of the time.

[0045] Sniff state 520 is another low-power operating mode during which specific receiving hardware within the node is activated in order to "sniff" energy associated with radio signals. Other hardware within the node may be powered down, thereby conserving power. The node may enter sniff state 520 at a particular time, and on a specific channel or frequency, in order to detect radio signals intended for reception at that particular time and on that particular channel. When the active hardware detects at least a threshold level of energy, the remaining portions of hardware within the node may be powered up, in response, in order to receive the signal. Generally, nodes within wireless mesh network 102 transmit packets that include a preamble. Upon being sniffed by a receiving node, the preamble triggers the receiving node to

power on and receive the payload included in the packet. In some embodiments, continuously-powered nodes 120 may not enter sniff state 520, since those nodes are not power-limited. Instead, in such embodiments, continuously-powered nodes 120 may enter receive state 530.

[0046] Receive state 530 is an operating mode for receiving radio signals. As mentioned above, receive state 530 may be entered when a preamble is detected during sniff state 520. The node may also directly enter receive state 530 under various circumstances. While in receive state 530, the node generally powers on most or all of the receiving hardware within the node, including the receiving portions of network interface 200 shown in Figure 2.

[0047] Process state 540 is an operating mode whereby the node processes received data, processes data for subsequent transmission, and performs various other processing operations. The node includes a bit stream processor that is active during process state 540, as well as another processor for running protocol and executing software applications while in process state 540. In addition, the node may include other peripheral processors for performing cryptographic functions or other specialized operations.

[0048] Transmit state 550 is an operating mode within which the node transmits data to other nodes. The node may enter transmit state 550 at a specific time and then transmit on a particular channel in order to target another node. Specifically, and as described in greater detail below, the node is configured to compute a time and a channel for targeting a specific node that is based on the media access control (MAC) address of the specific node, the time the transmission will be received, and the number of channels used in wireless mesh network 102. Maintaining accurate time across the nodes within wireless mesh network 102 is therefore a priority, and so nodes 120, 130, and 140 are configured to exchange “time beacons” with one another that include the current time. An exemplary time beacon is described below in conjunction with Figure 6.

[0049] Figure 6 illustrates a time beacon for sharing time between the different nodes of Figure 1, according to one embodiment of the present invention. As shown, time beacon 600 includes a header 610, a time payload 620, authentication sections 630-1 through 630-N, and cyclic redundancy check 640. Time payload 620 includes

received time 622, current time 624, and time hops 626. Each authentication section 620 includes a short identifier (ID) 622 and a hash message identification code (HMAC) 624.

[0050] In practice, a continuously-powered node 120 receives a time signal from an external source and sets an internal clock based on that time. The external source could be any technically feasible service that is accessible via network 152 and provides reasonably accurate time. The continuously-powered node 120 then transmits instances of time beacon 600 to neighboring nodes. The neighboring nodes could be continuously-powered nodes 120 or battery-powered nodes 130. Upon receipt of an instance of time beacon 600, the receiving node likewise sets an internal clock, and may then transmit another instance of time beacon 600 to downstream neighboring nodes. In this fashion, time propagates across wireless mesh network 102. Generally, the nodes of wireless mesh network 102 share instances of time beacon 600 with one another according to a time beacon schedule. The time beacon schedule could be established by a network management node 140, for example, or be determined independently by each node.

[0051] In time beacon 600, received time 622 indicates the time as received by the node which transmitted time beacon 600. Current time 624 indicates the time when time beacon 600 was sent. Time hops 626 indicates the number of hops between the transmitting node and the external source of time. Time hops 626 generally reflects the accuracy of received time 622 and current time 624. In particular, time that has propagated many hops across the network is generally less accurate than time received closer to the external source of time, since each hop may introduce drift associated with an intermediate node.

[0052] Each authentication section 630 is generally associated with a different neighboring node and includes a short ID 622 that is associated with the corresponding neighboring node. When two nodes first discover each other, those two nodes may exchange short IDs 622 in preparation for subsequently exchanging time beacons 600. A node that receives time beacon 600 is configured to access a specific authentication section 630 based on the assigned short ID 622, and then perform various security checks with the corresponding HMAC 624. In this fashion, nodes may protect against network attacks that involve sourcing inaccurate time into wireless mesh network 102. The receiving node may also check CRC 640 for data

integrity purposes. In one embodiment, a given node may receive multiple different instances of time beacon 600, and then combine the data included therein in order to generate a time estimate. For example, the node could average the current time indicated by each such time beacon instance.

[0053] As mentioned above, the sharing of accurate time is important for synchronizing transmissions between nodes. In networks with many battery-powered nodes 130 that reside in sleep state 510 for long periods of time, targeting those nodes when sniffing or receiving is central to maintaining network connectivity. In other words, the ability for nodes within wireless mesh network 102 to properly and effectively communicate traffic across the network depends on accurate timing. Figures 7A-7B illustrate techniques for targeting transmissions based on time measurements.

[0054] Figures 7A-7B are conceptual diagrams illustrating a technique for performing data communication between the different nodes of Figure 1, according to one embodiment of the present invention. As shown in Figure 7A, a transmitting node 120-T transmits messages to a receiving node 130-R at different points in time along timeline 700. Node 130-R is a battery-powered node and therefore resides in sleeping mode 510 much of the time. However, node 130-R becomes active during receive window 710. Receive window 710 has a duration from t0 to t1. Node 120-T may be configured to schedule transmissions so that node 130-R receives transmissions during receive window 710.

[0055] Problems may arise when nodes 120-T and 130-R are not sufficiently time synchronized. In particular, if an internal clock associated with node 130-T is faster or slower than an internal clock associated with node 130-R, then node 120-T may transmit too early or too late, respectively, and miss receive window 710. For example, in Figure 7A, node 130-T could transmit packet 720 before time t0, and, consequently, node 130-R would be in sleep mode 510 when that packet arrived. Alternatively, node 120-T could transmit packet 740 after time t1, and node 130-R would have already returned to sleep mode 510 when that packet arrived. In either case, poor time synchronization between nodes could result in a communication disruption. If nodes 120-T and 130-R remain in relative synchronization, node 120-T could transmit packet 730 within receive window 710, and node 130-T would be able to receive that packet. Node 130-R is configured to transmit packets to node 130-T

based on a set of transmit parameters 750, as discussed in greater detail below in conjunction with Figure 8.

[0056] Referring now to Figure 7B, in practice, node 130-R enter sniff mode 520 during receive window 710. Again, if sufficient energy is detected while sniffing, then node 130-R may power up and receive transmissions. When operating in sniff mode 520, node 130-R performs a sequence of sniffs 760, as is shown. Each sniff reflects a short period of time when energy can be detected. Node 130-R may return to sleep mode 520 between sniffs. In this manner, power can be conserved. Node 130-T need only transmit messages with a preamble that is longer than the time between sniffs to trigger node 130-R to enter receive mode 530. Node 130-R may then receive and demodulate the actual payload of the message.

[0057] The width of receive window 710, the sniff rate with which sniffs 760 are generated, and the time between sniffs 760 (referred to hereinafter as the inter-sniff interval) are configurable parameters. Node 130-R may select these parameters based on battery life, current data rate, remaining storage space, or any other technically feasible parameter. In addition, node 130-R may acquire data from a network management node 140 that reflects these parameters. Node 120-T may likewise acquire these parameters (e.g., from node 140), in order to properly target transmissions for node 130-R. For example, node 120-T could scale the preamble of a packet based on the inter-sniff interval associated with node 130-R, thereby increasing the likelihood that transmissions are received by that node. Generally, node 120-T relies on transmit parameters 750 when generating transmissions, which includes both network-oriented parameters and node-specific parameters.

[0058] Figure 8 illustrates a set of transmit parameters according to which a node transmits data within the network system of Figure 1, according to one embodiment of the present invention. As shown, transmit parameters 750 includes gear data, channel data, receive bandwidth, zero channel, preamble length, receive window, receive schedule, duty cycle, ACK wait time, and beacon schedule. The transmit parameters shown in Figure 8 are provided for exemplary purposes only, and could include fewer parameters or additional parameters.

[0059] Transmit parameters 750 may be received, as a whole or in part, from the node 140 configured to manage the network portion 104 where node 120-T resides.

For example, node 140 could transmit network data 324 and node data 326 to node 120-T when node 120-T registers with node 140. The specific values of each transmit parameter associated with a given node may vary depending on the type of node. For example, a battery-powered node may implement different values for transmit parameters 750 compared to a continuously-powered node. Transmit parameters 750 may be received, as a whole or in part, from node 130-R when nodes 120-T and 120-R initially discover one another. For example, nodes 120-T could obtain receive window data from node 130-R indicating the width of receive window 710. Node 120-T could then target the center of that window, thereby improving the likelihood of successful communication. Likewise, node 120-T could obtain gear data from node 130-R, and then transmit to that node using the proper data rate.

[0060] Node 120-T is configured to compute a channel and time associated with receive window 710 based on the MAC address of node 130-R, the number of channels used in wireless mesh network 102, and a time value. Nodes 120-T and 130-R may share MAC addresses with one another upon initial discovery, or receive data from node 140 indicating the MAC addresses of nearby neighbors. In one embodiment, node 120-T performs a Jenkins hash with the MAC address of node 130-R, the number of available channels, and the time at which node 130-R should receive transmissions, in order to compute the channel on which to send the transmission.

[0061] As mentioned, transmit parameters 750 may be derived from network data 324 or node data 326 supplied by node 140 shown in Figure 3. Transmit parameters 750 may also be derived from network data 424 and node data 426 provided by server 160 shown in Figure 4. In general, during network formation, server 160, nodes 110, 112, 120, 130, and 140 may exchange any information discussed herein in order to form communication links with one another.

[0062] Referring generally to Figures 1-8, by implementing the techniques discussed herein, wireless mesh network 102 supports the coexistence of many different types of nodes within a single network. Specifically, wireless mesh network 102 relies on a single communication protocol governs network communications and that can be implemented by both continuously-powered nodes 120 and battery-powered nodes 130. Again, that protocol relies on accurate maintenance of time, as well as the sharing of data that allows the proper targeting of transmissions to occur. The

techniques described thus far are also described, in stepwise fashion, below in conjunction with Figure 9.

[0063] Figure 9 is a flow diagram of method steps for participating in a wireless mesh network, according to one embodiment of the present invention. Although the method steps are described in conjunction with the systems of Figures 1-8, persons skilled in the art will understand that any system configured to perform the method steps, in any order, is within the scope of the present invention.

[0064] As shown, a method 900 begins at step 902, where a node within wireless mesh network 102 acquires one or more battery-powered neighbors. The node could be a continuously-powered node 120 or a battery-powered node 130. The node may be preprogrammed to include MAC addresses associated with the battery-powered neighbors, or may acquire those MAC addresses by performing a discovery process with the neighbors. Once the node obtains the neighbor MAC addresses, the node may transmit data thereto using the targeting approach described above in conjunction with Figures 7A-7B.

[0065] At step 904, the node acquires one or more continuously-powered neighbors. Similar to step 902, the node may be preprogrammed to include MAC addresses associated with the continuously-powered neighbors, or may acquire those MAC addresses by performing a discovery process. Once the node obtains the neighbor MAC addresses, the node may transmit data to those neighbors using the targeting approach described above in conjunction with Figures 7A-7B.

[0066] At step 906, the node propagates time across the network. In doing so, the node may receive a first time beacon from an upstream neighbor. The upstream neighbor could be, for example, any of the neighboring nodes acquired in steps 902 and 904. The node may then generate a second time beacon based on the first time beacon, and transmit the second time beacon to one or more downstream neighbors. In this manner, the node propagates time from upstream neighbors to downstream neighbors.

[0067] At step 908, the node registers with one or more network management nodes 140. The network management node 140 with which the node registers could manage, for example, the network portion 104 within which the node resides, or could

also manage another network portion 104. When the node registers with the network management node 140, the network management node 140 then transmits various registration data to the node, including network data 324 and node data 326.

[0068] At step 910, the node receives the registration data from the network management node 140. The registration data generally includes network-oriented data, such as the number of channels used within wireless mesh network 102, among others. The registration data also includes node-oriented data that may assist the node with acquiring neighbors. For example, the node data 326 could include MAC addresses of potential neighbors of the node, routing tables, forwarding tables, and so forth.

[0069] At step 912, the node acquires additional neighbors based on the registration data. For example, the node could parse a set of MAC addresses from the registration data, as well as the number of channels used within wireless mesh network 102, and then compute the channel and time when one or more neighboring nodes may receive data. When the receive window of one such neighbor opens, then the node may transmit a discovery packet to that node, thereby acquiring a new neighbor.

[0070] Referring generally to Figures 1-9, any of the nodes discussed herein may perform any of the techniques described thus far. Accordingly, those nodes may interoperate regardless of whether those nodes are battery-powered or continuously-powered. As a general matter, the techniques discussed above may be implemented by any heterogeneous collection of nodes, beyond those with different types of power sources. When a collection of such heterogeneous nodes initializes and first comes online, each such node may implement the method 900 in order to form a wireless mesh network, such as wireless mesh network 102 of Figure 1. An example formation of a network is described below in conjunction with Figures 10A-10E.

Network Functional Overview and Exemplary Mesh Formation

[0071] Figures 10A-10E are conceptual diagrams illustrating the formation of the wireless mesh network of Figure 1, according to one embodiment of the present invention. In one embodiment, the techniques described herein may be implemented in order to form network portions 104-1 and 104-2 or, alternatively, wireless mesh network 102 as a whole. As shown in Figure 10A, a network 1000 includes nodes

120-A and 120-B, nodes 130-A through 130-E, and node 140-A. Node 140-A is a network management node that is configured to manage some or all of network 1000. Nodes 120-A and 120-B are continuously-powered nodes that are configured to source time into network 1000, as described in greater detail below. Nodes 130-A through 130-E are battery-powered nodes. Figure 10A illustrates the initial state of network 1000, before time has been sourced into network 100 and before any communication links have been established.

[0072] In Figure 10B, nodes 120-A and 120-B receive a time signal from an external source. In doing so, nodes 120-A and 120-B may access the source of time via a network, such as network 152 of Figure 1. Generally, continuously-powered nodes 120 source time into network 1000, rather than battery-powered nodes 130, because sourcing time requires power that battery-powered nodes 130 typically hold in short supply. Continuously-powered nodes 130 may not have a comparatively strict energy budget, and therefore can assume responsibility for sourcing time. Once nodes 120-A and 120-B have set internal clocks based on the received time, those nodes may begin establishing communication links with neighboring nodes.

[0073] In Figure 10C, nodes 120-A and 120-B have begun establishing communication links with neighbors. Once those links are established, nodes 120-A and 120-B may begin propagating time across those links to the newly-acquired neighboring nodes. In particular, nodes 120-A and 120-B may transmit instances of time beacons 600 to those neighbors. Then, based on the time indicated in those beacons, the neighboring nodes may in turn acquire additional neighbors.

[0074] In Figure 10D, the nodes within network 1000 have acquired significantly more neighbors, and the time signal thus propagates throughout almost the entire network, the exception being node 130-E. With time shared across so many nodes, each such node may accurately target transmissions to any of the other nodes. Accurate targeting of this nature depends on the current time, the MAC address of the targeted node, and the number of channels used within network 1000, as previously described. A given node in network 1000 may acquire network data 324 and/or node data 326 that includes such information from network management node 140-A.

[0075] In Figure 10E, node 130-E has finally joined network 1000. Since node 130-E is substantially farther away from the other nodes of network 1000, nodes 130-D

and/or 130-E may perform various techniques in order to establish a communication link with one another. For example, node 130-D may periodically reduce a data rate when sending discovery packets to node 130-E. A lower data rate may permit transmissions from node 130-D to 130-E to cover the distance between those nodes more effectively. In addition, node 130-E may periodically widen a receive window to capture additional packets. Since node 130-E is isolated from an accurate source of time while the other nodes of network 1000 form communication links, node 130-E may not be synchronized with those other nodes. Consequently, node 130-E may not enter sniff mode 520 when node 130-D sends discovery packets. To mitigate this potential issue, node 130-E could widen the receive window until a discovery packet is received.

[0076] Referring generally to Figures 10A-10E, the nodes of network 1000 may obtain various data that facilitates the formation of that network. As mentioned, each node may obtain an accurate measure of time as well as MAC addresses of potential neighbors. This data allows a given node to target discovery packets to neighbors quickly. Each node may also be supplied with a network beacon schedule that indicates when local advertisement beacons will be transmitted by neighboring nodes. Based on this schedule, a given node may periodically widen the receive window until capturing a local advertisement beacon from a neighboring node. The advertisement beacon would include the MAC address of the neighboring node. Based on that MAC address, the node may then transmit discovery packets to the neighboring node to establish a communication link. In this fashion, nodes may populate neighbor tables.

[0077] In addition, node 140-A may perform various techniques to assist with the formation and growth of network 1000. For example, node 140-A, once coupled to one or more other nodes, may broadcast network data 324 and/or node data 326 to those nodes. Such data may include MAC addresses, data rates, channel information, and other types of data that may allow receiving nodes to acquire additional neighbors. As a general matter, nodes 140 may perform a wide variety of network management operations that facilitate both the formation of a network and the ongoing ability of the network to function in a robust manner. In addition, nodes 140 may interoperate with one another to join together previously disjoint portions of an overarching network. This particular functionality of nodes 140 is described in greater detail below in conjunction with Figures 11A-11D.

[0078] Figures 11A-11D are conceptual diagrams illustrating network management nodes interoperating to couple together disjoint network portions to form the network system of Figure 1, according to one embodiment of the present invention. The techniques described below may be implemented to join together network portions 104-1 and 104-2, which, as shown in Figure 1, may initially be disjoint from one another.

[0079] As shown in Figure 11A, each of network portions 104-1 and 104-2 of Figure 1 are coupled to a network management node 140-3. Node 140-3 allows nodes 140-1 and 140-2 to exchange various data within one another, including network data 324 and node data 326. Network data 324 may include connectivity information, while node data 326 may include MAC addresses. This data permits nodes 130-1 and 130-4 to establish communication link 1100, as described in greater detail below.

[0080] In Figure 11B, nodes 140-1 and 140-2 exchange data 1110-1 and 1110-2 via node 140-3. Again, data 1110 includes network-oriented data and node-oriented data. Nodes 140-1 and 140-2 may exchange data 1110 when one of nodes 130-1 or 130-4 indicates that communication link 1100 is needed. In Figure 11C, node 140-1 transmits data 1110-2 to node 130-1, and node 140-2 transmits data 1110-1 to node 130-4.

[0081] In Figure 11D, node 130-1 transmits discovery packet 1120 to node 130-4 based on data 1110-2, thereby establishing communication link 1100. In doing so, node 130-1 may rely on at least two different techniques. First, node 130-1 may parse connectivity information from data 1110-2 that reflects the topology of network portion 104-2. Then, node 130-1 may transmit discovery packet 1120 through network portion 104-1 to an access point 150 (not shown here) and then across network portion 104-2 to node 130-4. Alternatively, node 130-1 may parse the MAC address and data rate associated with node 130-4 and target that node directly. Node 130-1 may rely on either such technique based on the distance between nodes 130-1 and 130-4.

[0082] Referring generally to Figures 11A-11C, these figures illustrate one exemplary scenario where network management nodes 140 may couple together previously disjoint network portions 104. In some embodiments, the functionality of node 140-3 may be performed by either of nodes 140-1 and 140-2. In further embodiments,

nodes 140-1 and 140-2 may share responsibility for managing network portions 104-1 and 1042. The general approach discussed above is also described in stepwise fashion below in conjunction with Figure 12.

[0083] Figure 12 is a flow diagram of method steps for sharing connectivity information between disjoint network portions in order to form a network system, according to one embodiment of the present invention. Although the method steps are described in conjunction with the systems of Figures 1-8 and 10A-11D, persons skilled in the art will understand that any system configured to perform the method steps, in any order, is within the scope of the present invention.

[0084] As shown, a method 1200 begins at step 1202, where node 140-3 identifies node 130-1 within network portion 104-1. Node 140-3 may identify node 130-1 based on communicating with node 140-1. At step 1202, node 140-2 identifies node 130-4 within network portion 104-2. Similarly, node 140-3 may identify node 130-4 based on communication with node 140-2. At step 1206, node 140-3 determines that nodes 130-1 and 130-4 should be coupled together. Node 140-3 may make this determination based on node 130-1 requesting to be coupled to node 130-4, or node 130-4 requesting to be coupled to node 130-1, among other possibilities.

[0085] At step 1208, node 140-3 acquires data 1110-1 associated with network portion 104-1 and provides that data to node 130-4 via node 140-2. Data 1110-1 may include connectivity information associated with network portion 104-1, as well as the MAC address and data rate of node 130-1. At step 1210, node 140-3 acquires data 1110-2 associated with network portion 104-2 and provides that data to node 130-1 via node 140-1. Data 1110-2 may include connectivity information associated with network portion 104-2, as well as the MAC address and data rate of node 130-4.

[0086] Based on data 1110, nodes 130-1 and 130-4 may then establish communication link 1100 with one another. Nodes 130-1 and 130-4 may communicate directly, via targeted transmissions, or indirectly by traversing network portions 104-1 and 104-2. By implementing this approach, a given network management node 140, configured to manage one or more network portions 104, is capable of coupling nodes within those network portions to other nodes within other network portions 104.

[0087] In sum, a wireless mesh network includes heterogeneous types of nodes, including continuously-powered nodes and battery-powered nodes. The battery-powered nodes may reside in a sleeping state most of the time to conserve power. The various nodes in the network may communicate with one another by transmitting and receiving at scheduled times and on scheduled frequencies. The battery-powered nodes may become active during the scheduled transmit and receive times. Network management nodes may facilitate network formation by transmitting information that reflects the scheduled transmit and receive times across the network. Based on this data, the continuously-powered nodes and battery-powered nodes may establish communication links with one another.

[0088] At least one advantage of the techniques described herein is that heterogeneous types of nodes, such as continuously-powered nodes and battery-powered nodes, may reside within the same wireless mesh network. Accordingly, networks such as the "Internet of Things" may continue to grow, despite having a heterogeneous mixture of device types. When each such device implements the homogeneous communication protocol described herein, those devices may establish communication links with one another. In addition, the functionality of the network management nodes permits disjoint network portions to be coupled with one another, further improving the degree to which heterogeneous devices can be connected.

[0089] The descriptions of the various embodiments have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments.

[0090] Aspects of the present embodiments may be embodied as a system, method or computer program product. Accordingly, aspects of the present disclosure may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present disclosure may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

[0091] Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

[0092] Aspects of the present disclosure are described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the disclosure. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, enable the implementation of the functions/acts specified in the flowchart and/or block diagram block or blocks. Such processors may be, without limitation, general purpose processors, special-purpose processors, application-specific processors, or field-programmable processors.

[0093] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more

executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

[0094] While the preceding is directed to embodiments of the present disclosure, other and further embodiments of the disclosure may be devised without departing from the basic scope thereof, and the scope thereof is determined by the claims that follow.

The Invention Claimed Is:

1. A non-transitory computer readable medium storing instructions that, when executed by a processor, causes the processor to establish a communication link between different network portions to generate a wireless mesh network, by performing the steps of:

identifying a first node within a first network portion;

identifying a second node within a second network portion, wherein the first network portion and the second network portion include separate and distinct nodes;

obtaining first network data that is associated with the first network portion and includes connectivity information related to the nodes included in the first network portion; and

transmitting the first network data to the second node to enable the second node to communicate with the first node based on the first network data.

2. The non-transitory computer readable medium of claim 1, further comprising the steps of:

obtaining second network data that is associated with the second network portion and includes connectivity information related to the nodes included in the second network portion; and

transmitting the second network data to the first node to enable the first node to communicate with the second node based on the second network data.

3. The non-transitory computer readable medium of claim 2, wherein transmitting the second network data to the first node comprises transmitting the second network data during formation of the second network portion.

4. The non-transitory computer readable medium of claim 3, wherein the second network data includes one or more addresses corresponding to one or more neighbor nodes that reside proximate to the second node and are included in the second network portion, and wherein the second node transmits data to the one or more neighbor nodes based on the second network data.

5. The non-transitory computer readable medium of claim 1, further comprising the step of determining that the second node generates data for transmission to the first node prior to transmitting the first network data to the second node.

6. The non-transitory computer readable medium of claim 1, further comprising the step of causing a time-sourcing node within the second network portion to transmit a first time beacon that indicates a time value to the second node.

7. The non-transitory computer readable medium of claim 6, wherein the time-sourcing node comprises a continuously-powered node, and where causing the continuously-powered node to transmit the first time beacon comprises providing the continuously-powered node with an address associated with the second node.

8. The non-transitory computer readable medium of claim 1, further comprising the steps of:

configuring a first set of continuously-powered nodes included in the first network portion to communicate with a first set of battery-powered nodes included in the first network portion according to a unified communication protocol; and

configuring a second set of continuously-powered nodes included in the second network portion to communicate with a second set of battery-powered nodes included in the second network portion according to the unified communication protocol.

9. The non-transitory computer readable medium of claim 8, wherein configuring the first set of continuously-powered nodes to communicate with the first set of battery-powered nodes comprises causing the first set of continuously-powered nodes to schedule data transmissions to the first set of battery-powered nodes based on a first set of media access control addresses associated with the first set of battery-powered nodes and based on a current time value.

10. The non-transitory computer readable medium of claim 8, wherein configuring the second set of continuously-powered nodes to communicate with the second set of

battery-powered nodes comprises causing the second set of continuously-powered nodes to schedule data transmissions to the second set of battery-powered nodes based on a second set of media access control addresses associated with the second set of battery-powered nodes and based on a current time value.

11. The non-transitory computer readable medium of claim 8, wherein the first network data includes first media access control addresses associated with the nodes included in the first network portion, and wherein the second node, upon receipt of the first network data, schedules data transmissions to the nodes included in the first network portion based on the first media access control addresses and based on a current time value.

12. A system for coupling together different network portions to generate a wireless mesh network, comprising:

a primary network management node, including:

a memory that includes a network management application, and

a processor that, in conjunction with executing the network management application, performs the steps of:

identifying a first node within a first network portion,

identifying a second node within a second network portion, wherein the first network portion and the second network portion include separate and distinct nodes,

obtaining first network data that is associated with the first network portion and includes connectivity information related to the nodes included in the first network portion, and

transmitting the first network data to the second node to enable the second node to communicate with the first node based on the first network data.

13. The system of claim 12, further comprising:

a first network management node that resides within the first network portion, including:

a first memory that includes a first network management application, and

a first processor that, in conjunction with executing the first network management application, transmits the first network data to the primary network management node.

14. The system of claim 13, wherein the first processor, in conjunction with executing the first network management application, configures a first set of continuously-powered nodes included in the first network portion to communicate with a first set of battery-powered nodes included in the first network portion according to a unified communication protocol by causing the first set of continuously-powered nodes to schedule data transmissions to the first set of battery-powered nodes based on a first set of media access control addresses associated with the first set of battery-powered nodes and based on a current time value.

15. The system of claim 14, wherein the first set of battery-powered nodes and the second set of battery-powered nodes reside in a sleeping state for a fraction of time that exceeds a fraction of time that the first set of continuously-powered nodes and the second set of continuously-powered nodes reside in the sleeping state.

16. The system of claim 12, wherein the processor, in conjunction with executing the network management application, performs the additional steps of:

obtaining second network data that is associated with the second network portion and includes connectivity information related to the nodes included in the second network portion; and

transmitting the second network data to the first node to enable the first node to communicate with the second node based on the second network data.

17. The system of claim 16, further comprising:

a second network management node that resides within the second network portion, including:

a second memory that includes a second network management application, and

a second processor that, in conjunction with executing the second network management application, transmits the second network data to the primary network management node.

18. The system of claim 17, wherein the second processor, in conjunction with executing the second network management application, configures a second set of continuously-powered nodes included in the second network portion to communicate with a second set of battery-powered nodes included in the second network portion according to a unified communication protocol by causing the second set of continuously-powered nodes to schedule data transmissions to the second set of battery-powered nodes based on a second set of media access control addresses associated with the second set of battery-powered nodes and based on a current time value.

19. A computer-implemented method for establishing a communication link between different network portions to generate a wireless mesh network, the method comprising:

identifying a first node within a first network portion;

identifying a second node within a second network portion, wherein the first network portion and the second network portion include separate and distinct nodes;

obtaining first network data that is associated with the first network portion and includes connectivity information related to the nodes included in the first network portion; and

transmitting the first network data to the second node to enable the second node to communicate with the first node based on the first network data.

20. The computer-implemented method of claim 19, wherein the first network data includes first media access control addresses associated with the nodes included in the first network portion, and wherein the second node, upon receipt of the first network data, schedules data transmissions to the nodes included in the first network portion based on the first media access control addresses and based on a current time value.

21. The computer-implemented method of claim 20, further comprising:

obtaining second network data that is associated with the second network portion and includes connectivity information related to the nodes included in the second network portion; and

transmitting the second network data to the first node to enable the first node to communicate with the second node based on the second network data.

22. The computer-implemented method of claim 21, wherein the second network data includes first media access control addresses associated with the nodes included in the second network portion, and wherein the first node, upon receipt of the second network data, schedules data transmissions to the nodes included in the second network portion based on the second media access control addresses and based on a current time value.

23. A non-transitory computer readable medium including instructions that, when executed by a processor, cause the processor to perform the steps of:

obtaining registration data associated with a network;

transmitting a first packet to a first battery-powered node participating within the network based on the registration data and according to a first communication protocol; and

transmitting a second packet to a first continuously-powered node participating within the network based on the registration data and according to the first communication protocol,

wherein the first communication protocol governs at least a portion of communications across a plurality of nodes participating within the network, and the plurality of nodes includes one or more battery-powered nodes and one or more continuously-powered nodes.

24. The non-transitory computer-readable medium of claim 23, wherein the registration data comprises data indicating a number of available channels within the network and a measure of time maintained within the network, and data indicating a media access control (MAC) address associated with the first battery-powered node and a MAC address associated with the continuously-powered node.

25. The non-transitory computer-readable medium of claim 24, further comprising executing a hash function based on the number of available channels within the network, the measure of time, and the MAC address of the first battery-powered node to determine a first channel on which to transmit the first packet to the battery-powered node and a first time at which to transmit the first packet

26. The non-transitory computer-readable medium of claim 25, wherein the battery-powered node is receptive to communications at the first time and on the first channel and inactive for a period of time following the first time.

27. The non-transitory computer-readable medium of claim 24, further comprising executing a hash function based on the number of available channels within the network, the measure of time, and the MAC address of the first continuously-powered node to determine a first channel on which to transmit the second packet to the continuously-powered node and a first time at which to transmit the second packet

28. The non-transitory computer-readable medium of claim 23, further comprising transitioning to a first state for a first interval of time in order to periodically monitor a first channel with a first frequency for a response from either the battery-powered node or the continuously-powered node.

29. The non-transitory computer-readable medium of claim 28, further comprising, upon determining that the response was not received during the first interval of time, transitioning to the first state for a second interval of time, wherein the second interval of time exceeds the first interval of time.

30. The non-transitory computer-readable medium of claim 28, further comprising, upon determining that the response was not received during the first interval of time, periodically monitoring the first channel with a second frequency, wherein the second frequency exceeds the first frequency.

31. The non-transitory computer-readable medium of claim 23, wherein at least one of the first packet and the second packet comprises a time beacon that includes a measure of time maintained within the network.

32. The non-transitory computer-readable medium of claim 23, wherein each of the first packet and the second packet is transmitted across the network based on transmit parameters that includes at least one of gear data, channel data, receive bandwidth, zero channel, preamble length, receive window, receive schedule, duty cycle, acknowledge wait time, and beacon schedule.

33. The non-transitory computer-readable medium of claim 32, wherein the first packet is transmitted to the battery-powered node based on a first set of values for the transmit parameters, the first set of values being unique to battery-powered nodes in the network, and wherein the second packet is transmitted to the continuously-powered node based on a second set of values for the transmit parameters, the second set of values being unique to continuously-powered nodes in the network.

34. The non-transitory computer readable medium of claim 33, wherein at least one of the battery-powered node and the continuously-powered node also participates in at least one other network.

35. A computing device, comprising:
a memory storing a software application; and
a processor coupled to the memory that, when executing the software application, is configured to:
obtain registration data associated with a network,
transmit a first packet to a first battery-powered node participating within the network based on the registration data and according to a first communication protocol, and
transmit a second packet to a first continuously-powered node participating within the network based on the registration data and according to the first communication protocol,
wherein the first communication protocol governs at least a portion of communications across a plurality of nodes participating within

the network, and the plurality of nodes includes one or more battery-powered nodes and one or more continuously-powered nodes.

36. The computing device of claim 35, wherein the registration data comprises data indicating a number of available channels within the network and a measure of time maintained within the network, and data indicating a media access control (MAC) address associated with the first battery-powered node and a MAC address associated with the continuously-powered node.

37. The computing device of claim 36, further comprising executing a Jenkins hash with different portions of the registration data to determine a first channel on which to transmit the first packet to the battery-powered node at a first time, and a second channel on which to transmit the second packet to the continuously-powered node at a second time.

38. The computing device of claim 35, wherein the processor is further configured to transition to a first state for a first interval of time in order to periodically monitor a first channel with a first frequency for a response from either the battery-powered node or the continuously-powered node.

39. The computing device of claim 38, wherein the processor is further configured to, upon determining that the response was not received during the first interval of time, transition to the first state for a second interval of time, wherein the second interval of time exceeds the first interval of time.

40. The computing device of claim 38, wherein the processor is further configured to, upon determining that the response was not received during the first interval of time, periodically monitoring the first channel with a second frequency, wherein the second frequency exceeds the first frequency.

41. The computing device of claim 35, wherein each of the first packet and the second packet is transmitted across the network based on transmit parameters that includes at least one of gear data, channel data, receive bandwidth, zero channel,

preamble length, receive window, receive schedule, duty cycle, acknowledge wait time, and beacon schedule.

42. The computing device of claim 41, wherein the processor transmits the first packet to the battery-powered node based on a first set of values for the transmit parameters, the first set of values being unique to battery-powered nodes in the network, and transmits the second packet to the continuously-powered node based on a second set of values for the transmit parameters, the second set of values being unique to continuously-powered nodes in the network.

43. A computer-implemented method for configuring a node to participate in a network, the method comprising:

obtaining registration data associated with a network;

transmitting a first packet to a first battery-powered node participating within the network based on the registration data and according to a first communication protocol; and

transmitting a second packet to a first continuously-powered node participating within the network based on the registration data and according to the first communication protocol,

wherein the first communication protocol governs at least a portion of communications across a plurality of nodes participating within the network, and the plurality of nodes includes one or more battery-powered nodes and one or more continuously-powered nodes.

44. The computer-implemented method of claim 43, wherein obtaining the registration data comprises registering to participate in the network with a network management entity configured to manage communications between nodes within the network, and wherein the network management entity is also configured to manage communications between nodes that reside within the network and nodes that reside within at least one other network.

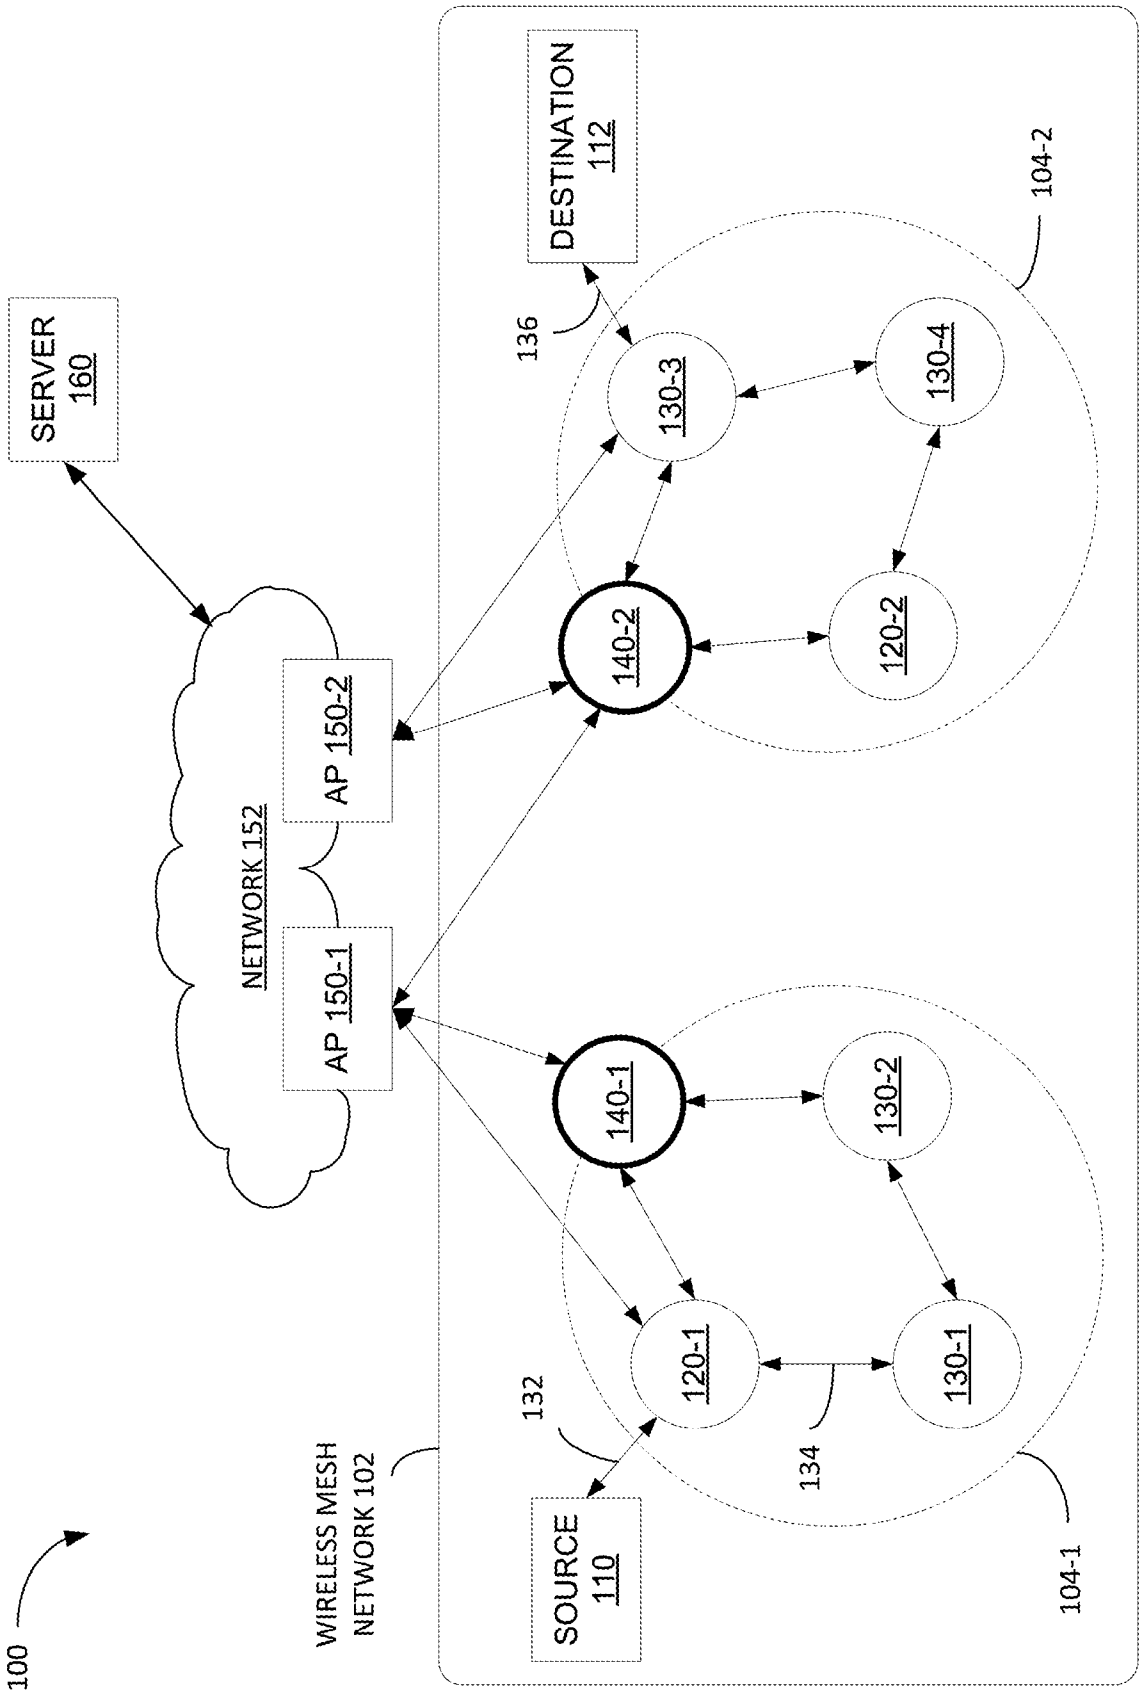

100
SERVER 160
NETWORK 152
AP 150-1
AP 150-2
WIRELESS MESH NETWORK 102
SOURCE 110
DESTINATION 112
132
134
136
104-1
104-2
120-1
130-1
130-2
140-1
120-2
130-3
130-4
140-2

FIG. 1

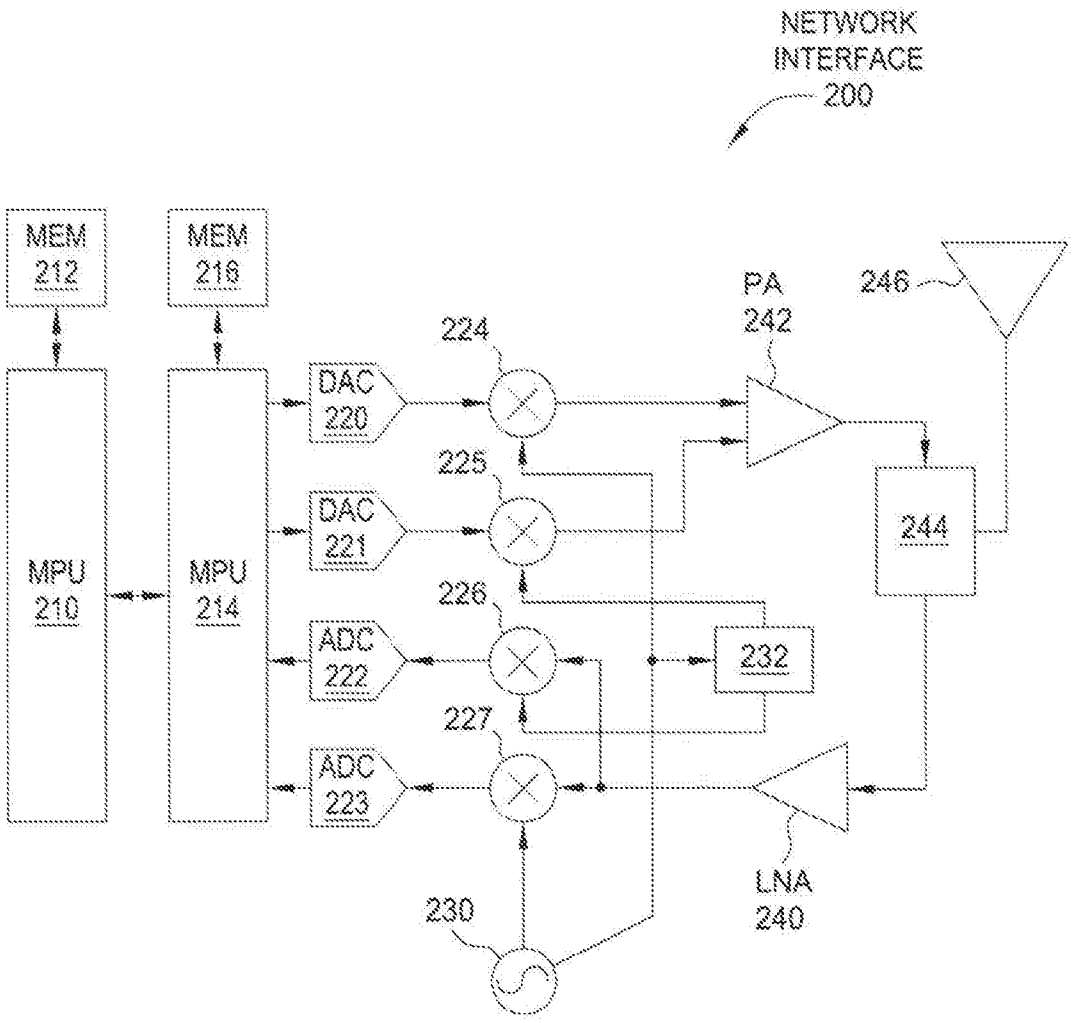

NETWORK
INTERFACE
200
MEM
212
MEM
216
MPU
210
MPU
214
DAC
220
DAC
221
ADC
222
ADC
223
224
225
226
227
230
PA
242
246
244
232
LNA
240

FIG. 2

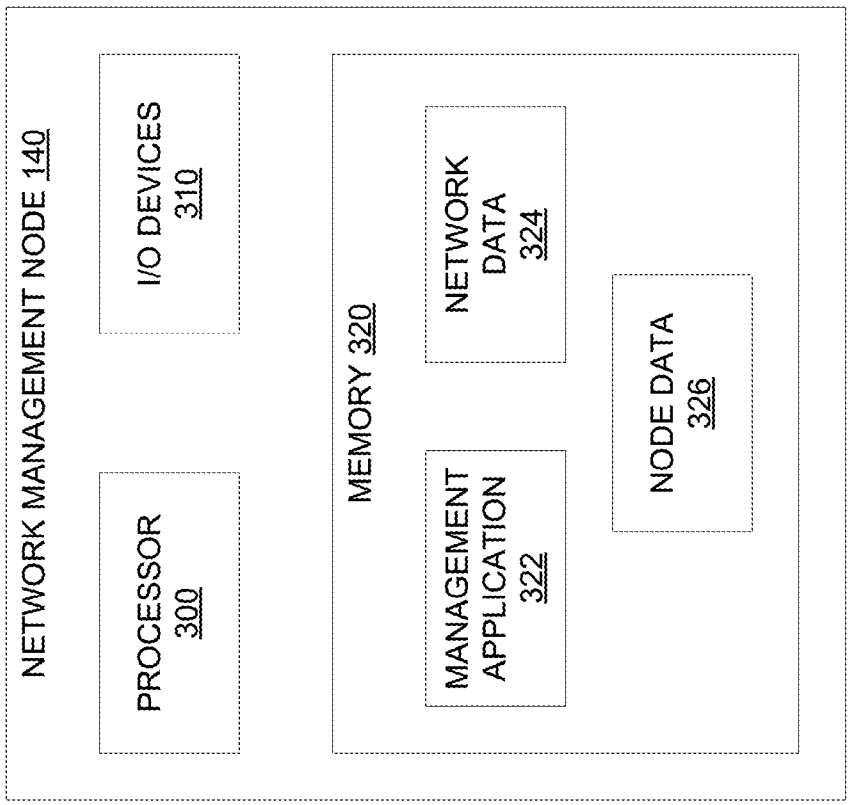

NETWORK MANAGEMENT NODE 140
PROCESSOR 300
I/O DEVICES 310
MEMORY 320
MANAGEMENT APPLICATION 322
NETWORK DATA 324
NODE DATA 326

FIG. 3

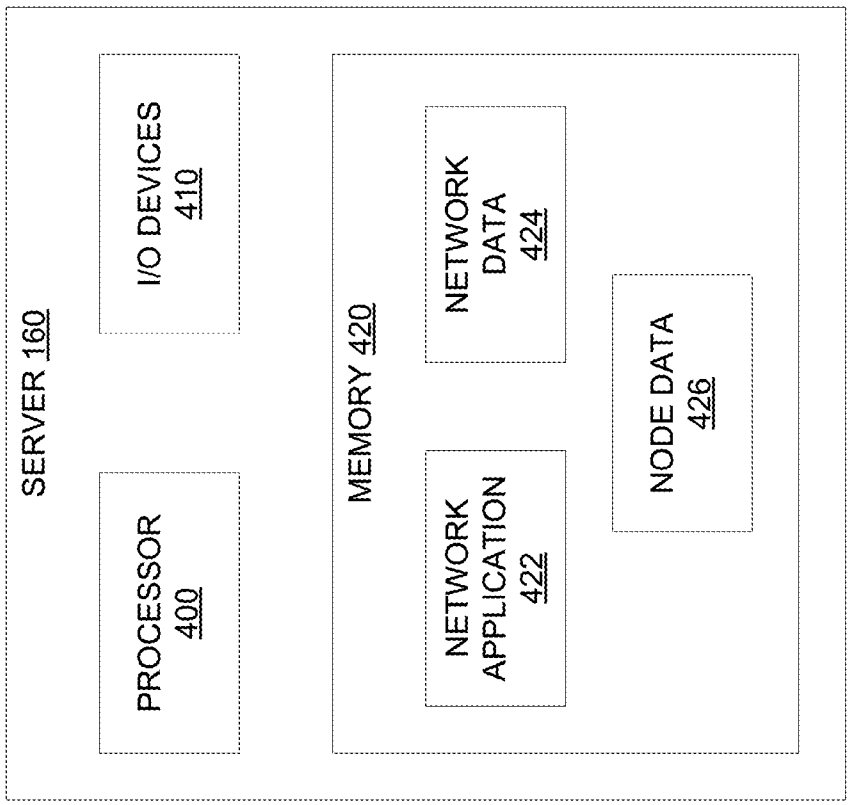

SERVER 160
PROCESSOR 400
I/O DEVICES 410
MEMORY 420
NETWORK APPLICATION 422
NETWORK DATA 424
NODE DATA 426

FIG. 4

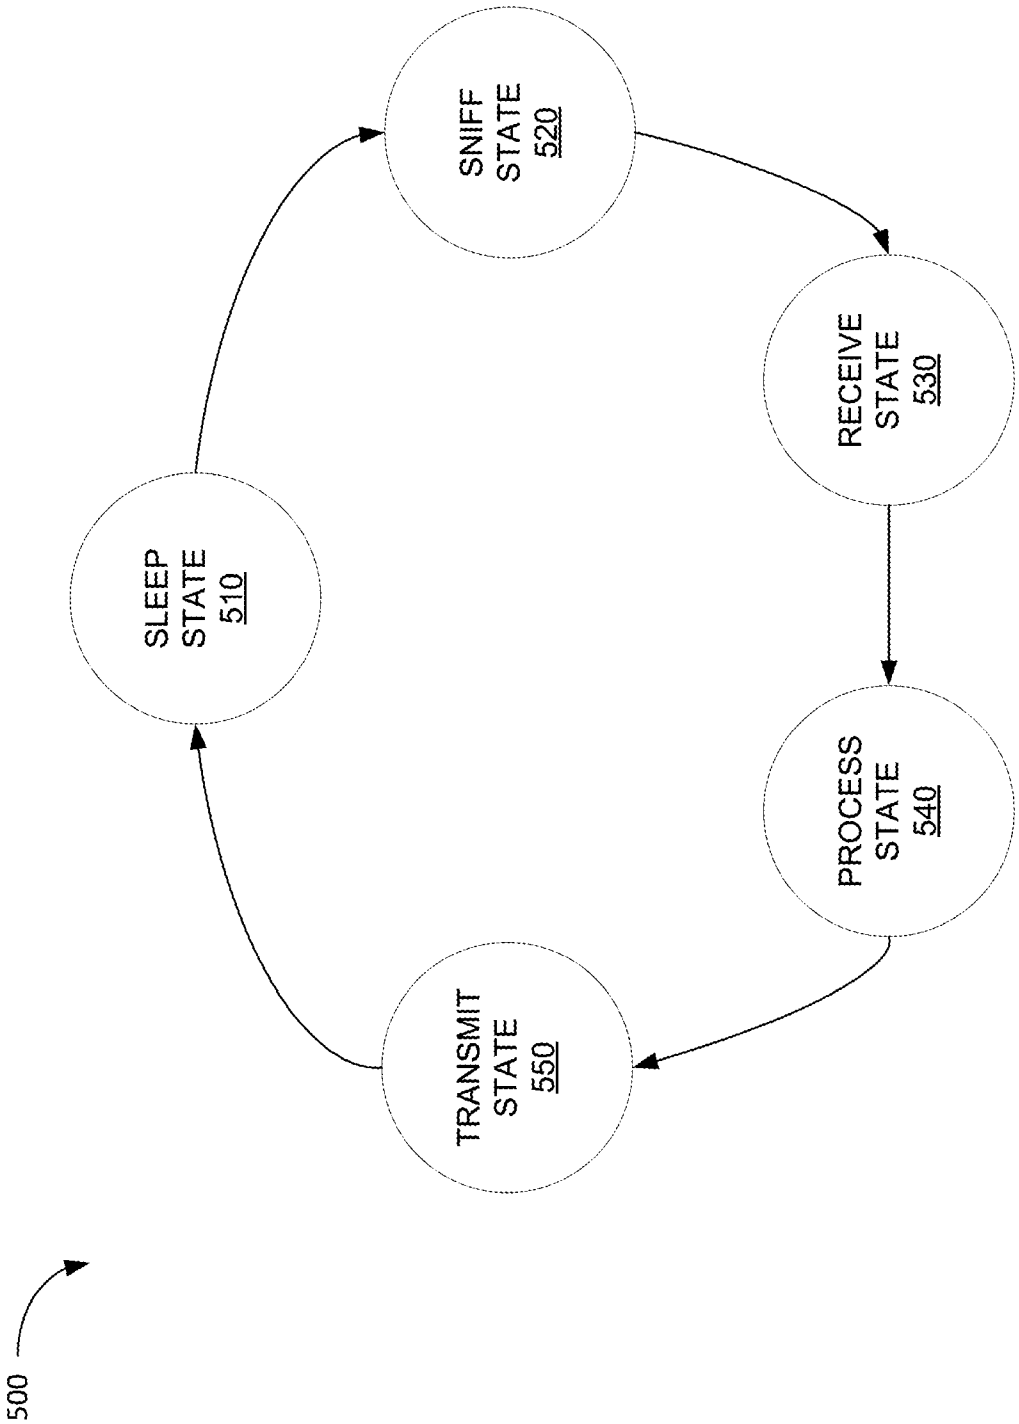

SNIFF STATE 520
RECEIVE STATE 530
PROCESS STATE 540
TRANSMIT STATE 550
SLEEP STATE 510
500

FIG. 5

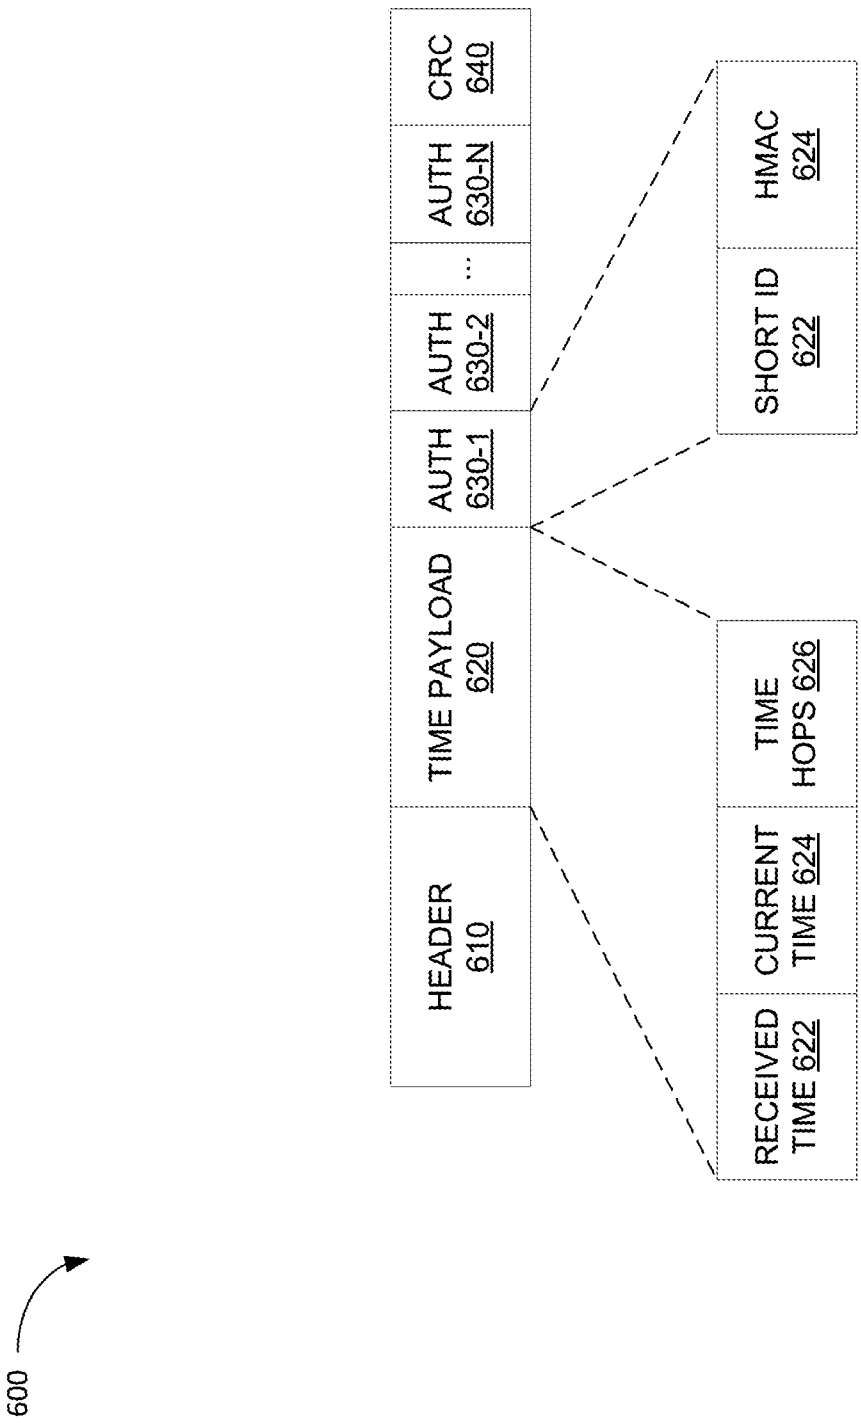

600
HEADER 610
TIME PAYLOAD 620
AUTH 630-1
AUTH 630-2
...
AUTH 630-N
CRC 640
RECEIVED TIME 622
CURRENT TIME 624
TIME HOPS 626
SHORT ID 622
HMAC 624

FIG. 6

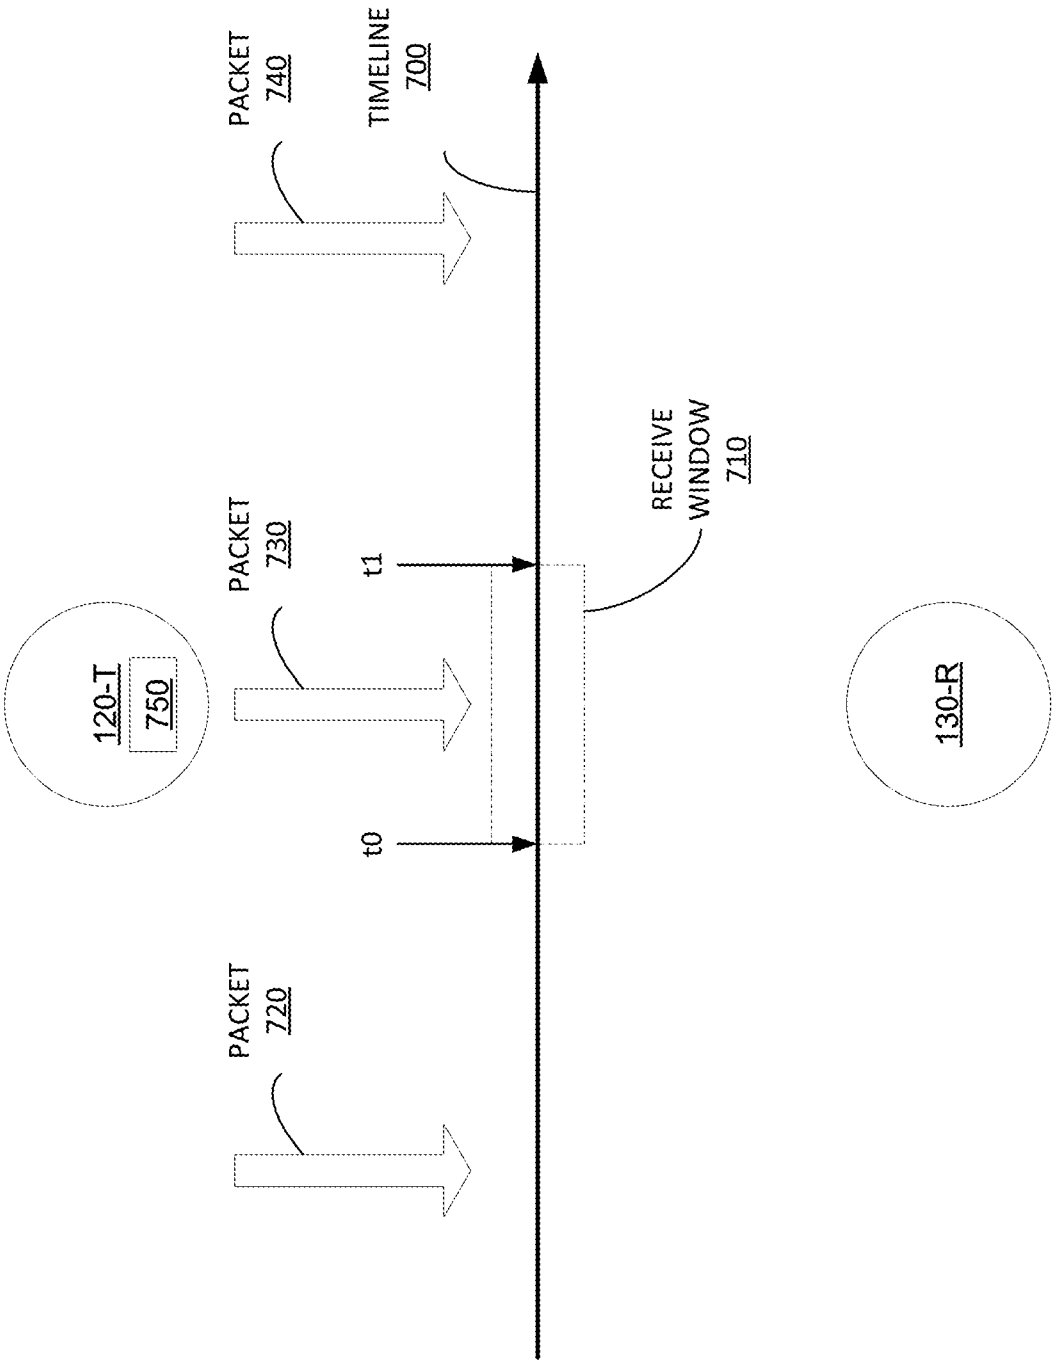

PACKET 740
TIMELINE 700
RECEIVE WINDOW 710
PACKET 730
t1
120-T 750
t0
130-R
PACKET 720

FIG. 7A

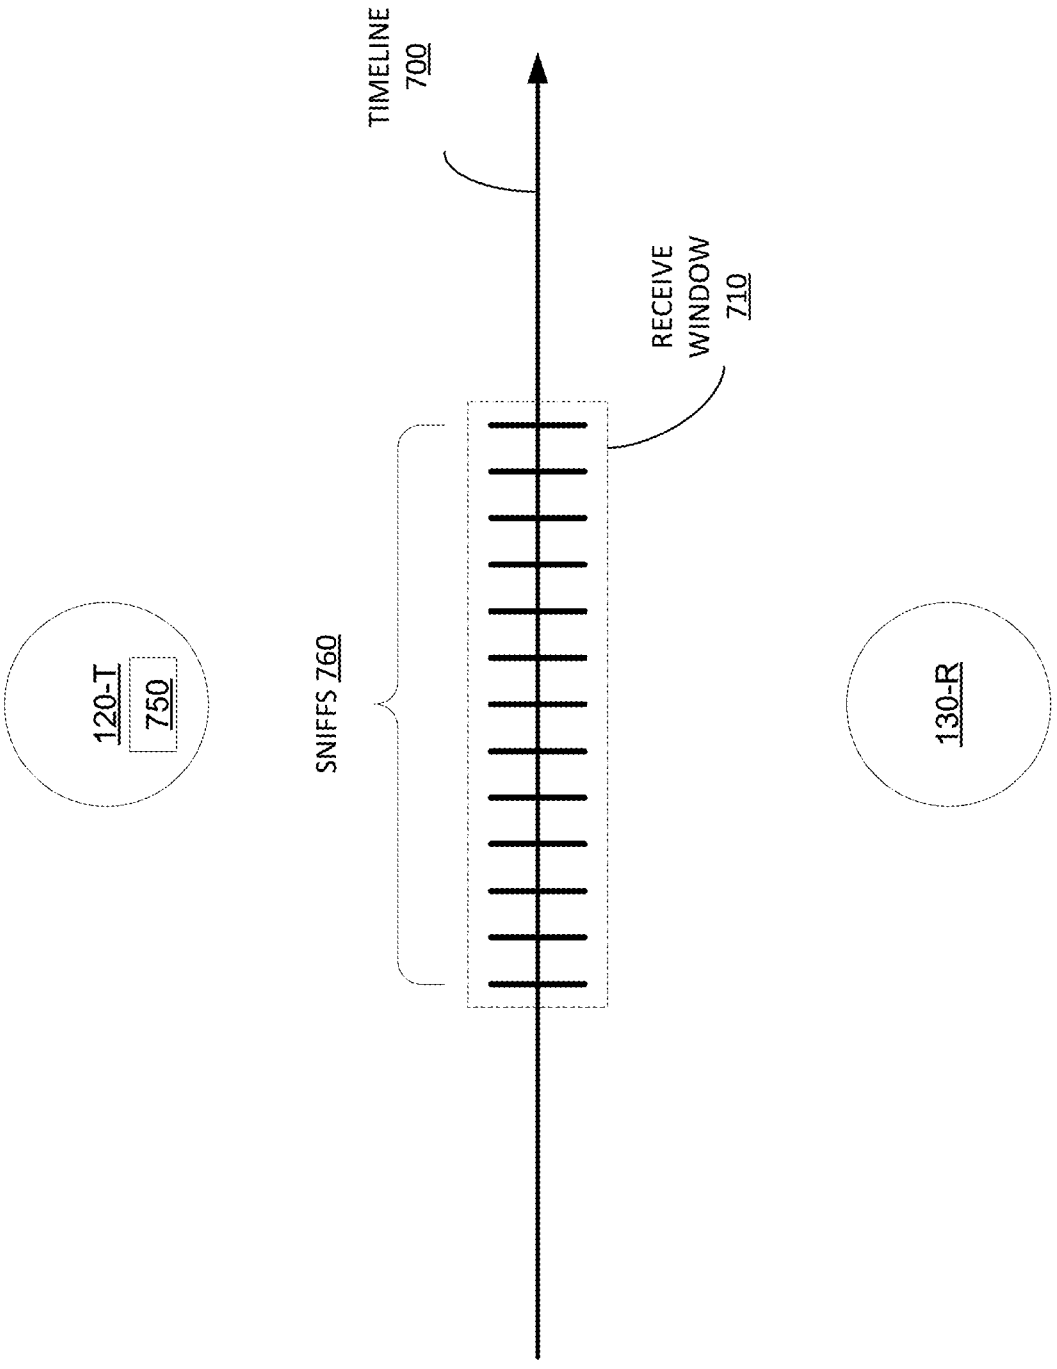


FIG. 7B

TRANSMIT PARAMETERS
750

- GEAR DATA
- CHANNEL DATA
- RECEIVE BANDWIDTH
- ZERO CHANNEL
- PREAMBLE LENGTH
- RECEIVE WINDOW
- RECEIVE SCHEDULE
- DUTY CYCLE
- ACK WAIT TIME
- BEACON SCHEDULE

FIG. 8

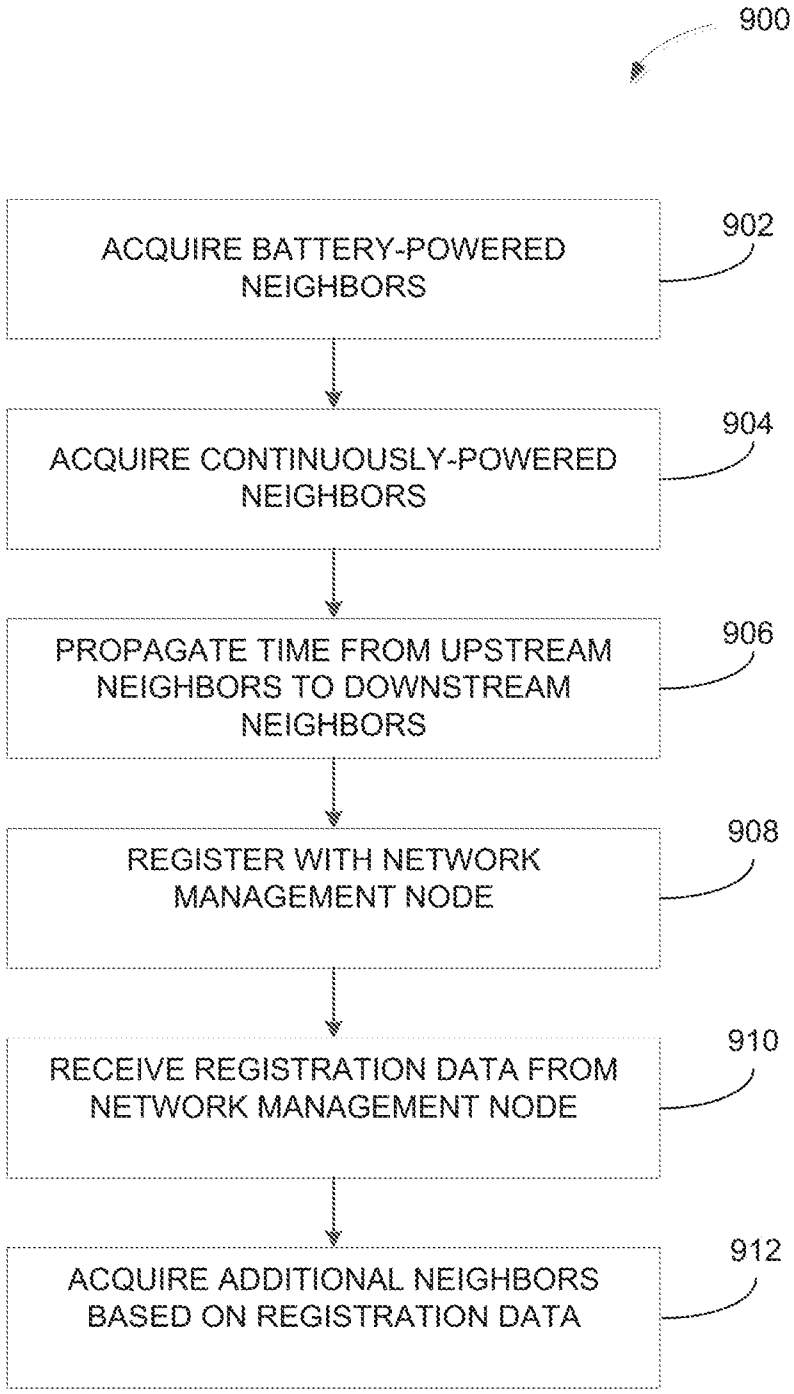

900
ACQUIRE BATTERY-POWERED NEIGHBORS
902
ACQUIRE CONTINUOUSLY-POWERED NEIGHBORS
904
PROPAGATE TIME FROM UPSTREAM NEIGHBORS TO DOWNSTREAM NEIGHBORS
906
REGISTER WITH NETWORK MANAGEMENT NODE
908
RECEIVE REGISTRATION DATA FROM NETWORK MANAGEMENT NODE
910
ACQUIRE ADDITIONAL NEIGHBORS BASED ON REGISTRATION DATA
912

FIG. 9

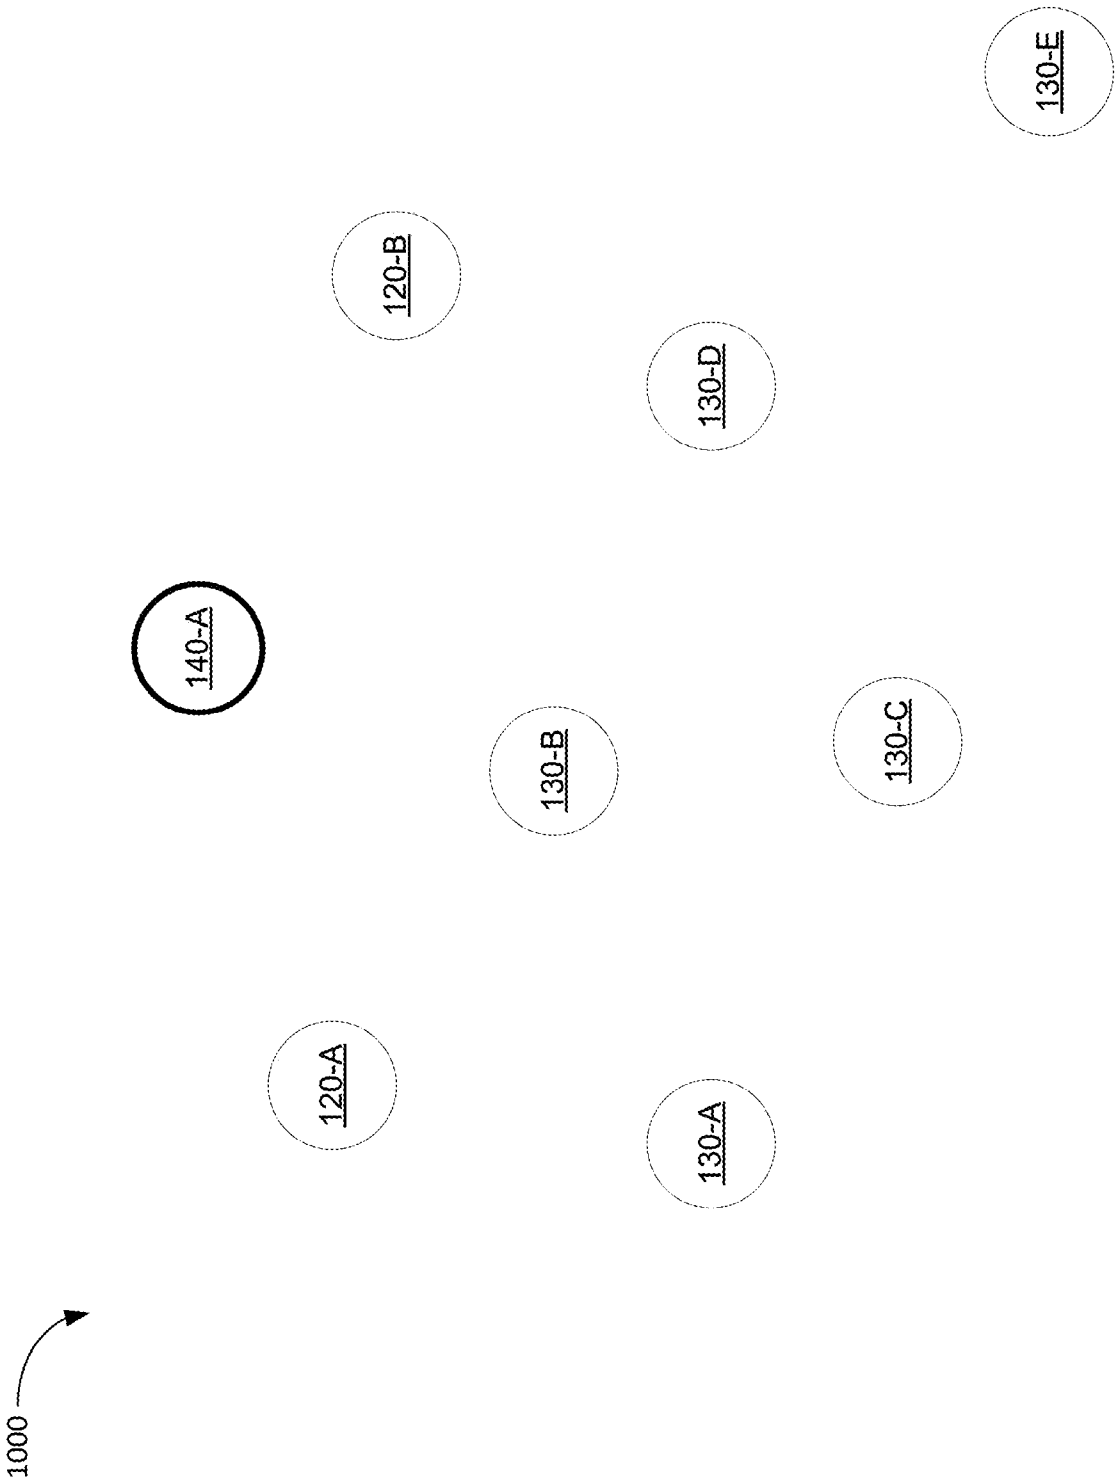


FIG. 10A

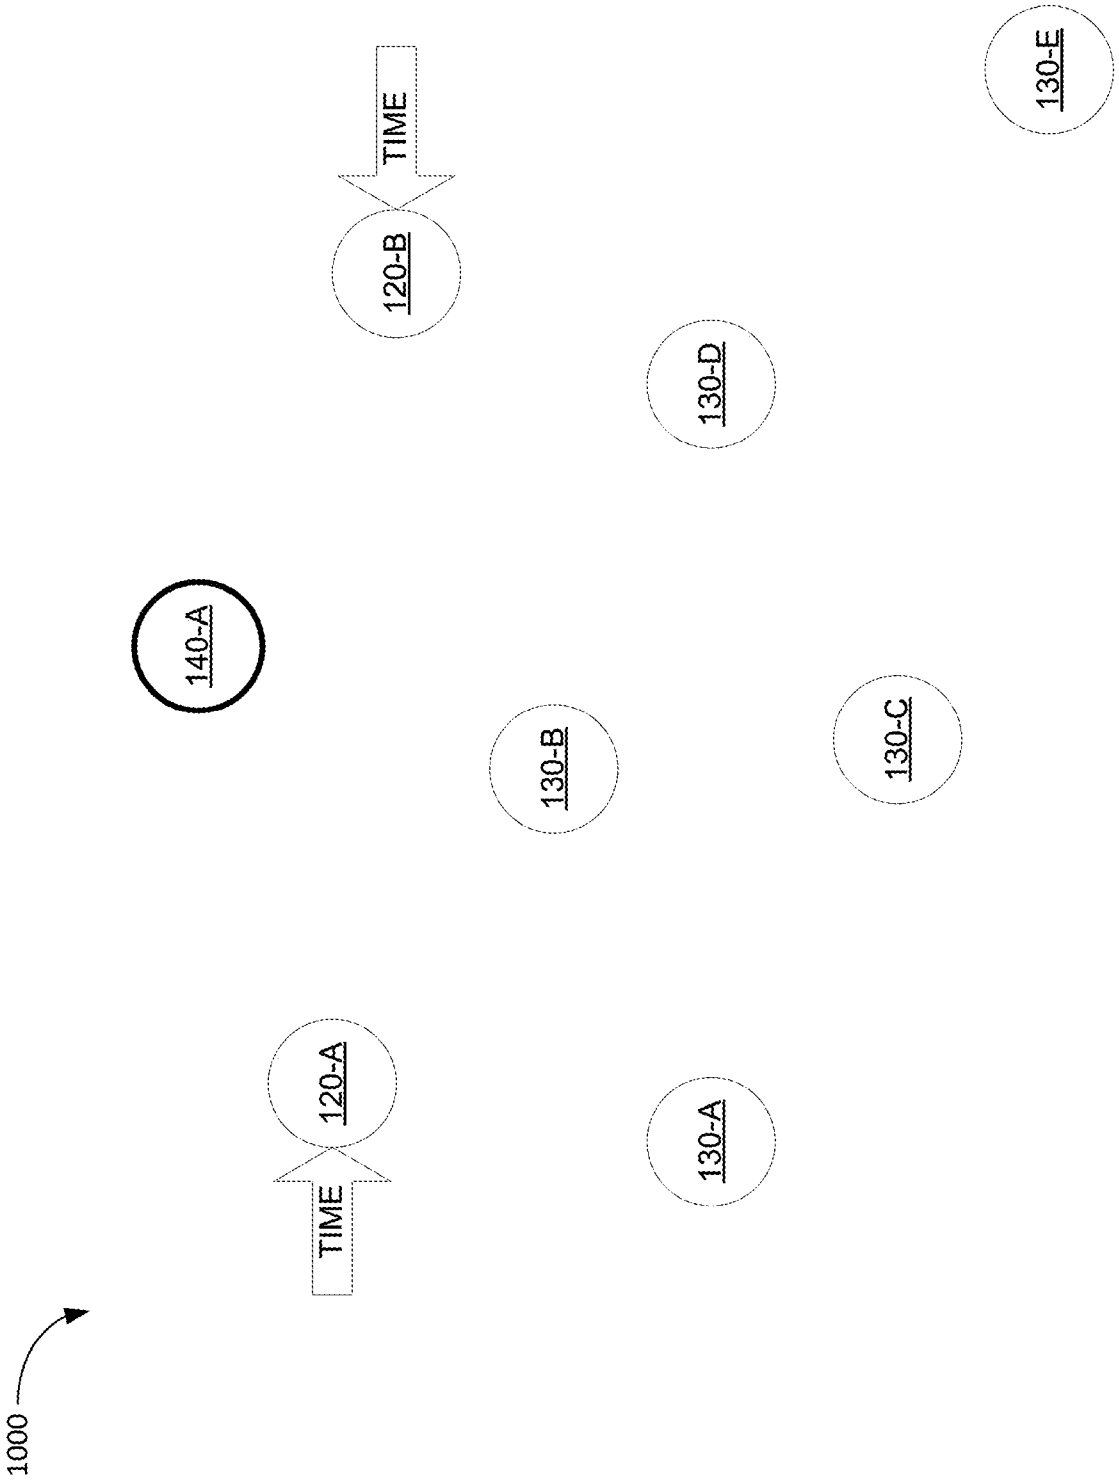


FIG. 10B

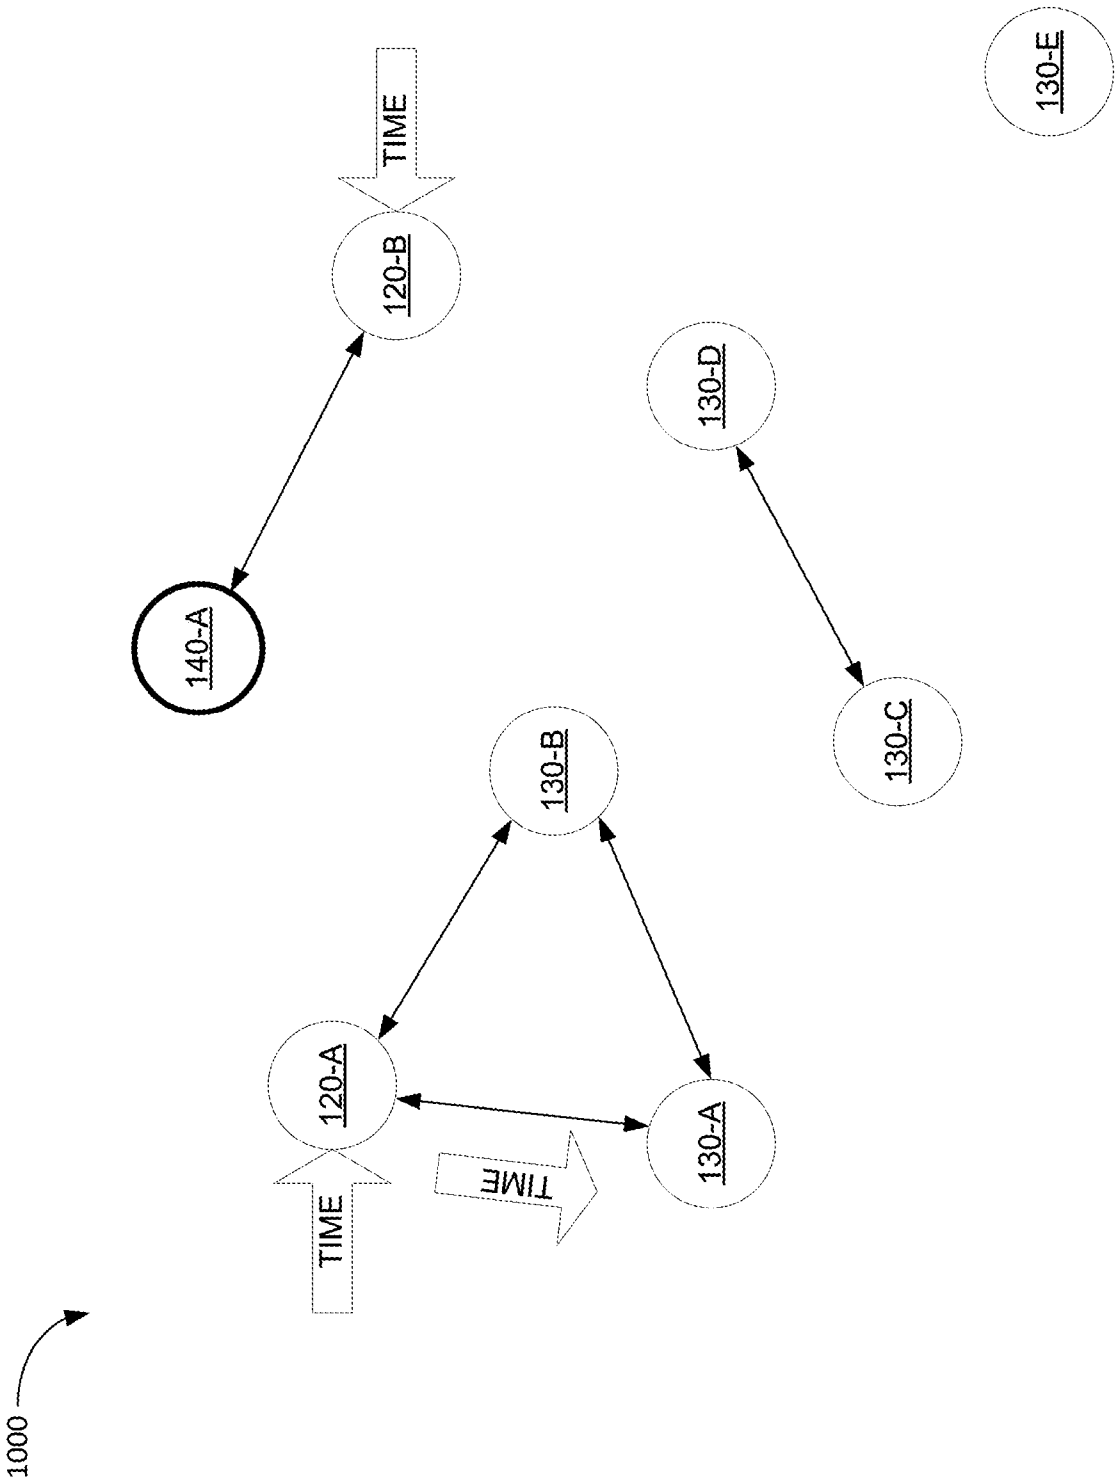
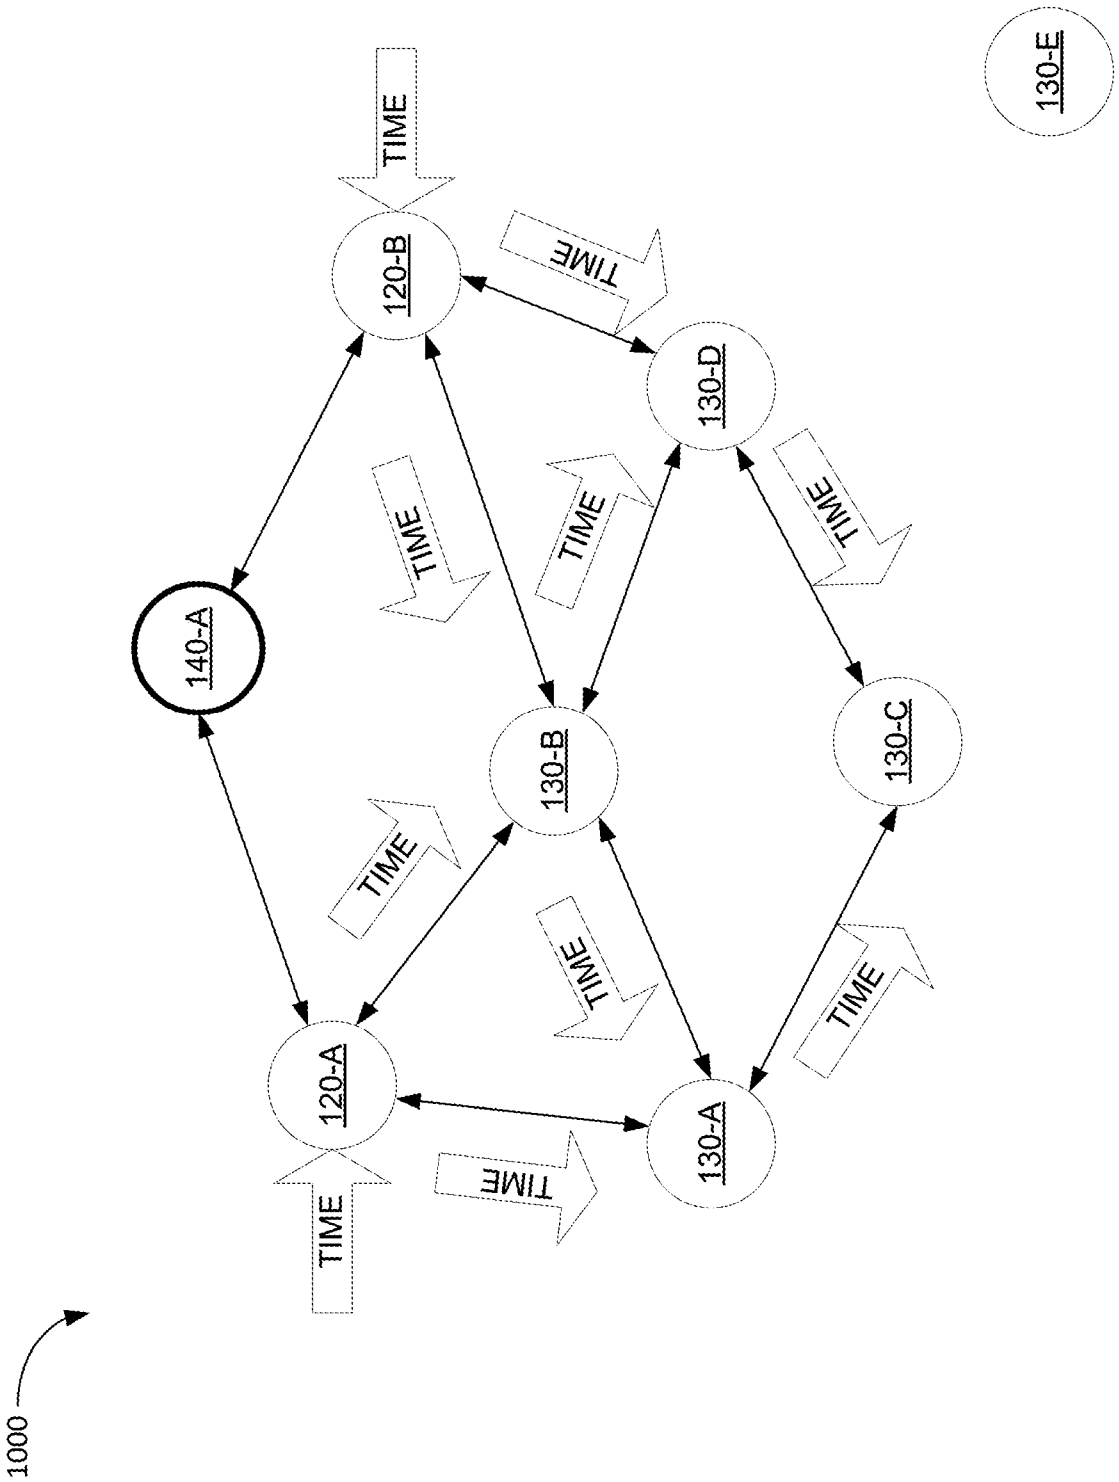


FIG. 10C


1000
140-A
120-A
120-B
130-A
130-B
130-C
130-D
130-E
TIME
TIME
TIME
TIME
TIME
TIME
TIME
TIME
TIME

FI10D1C

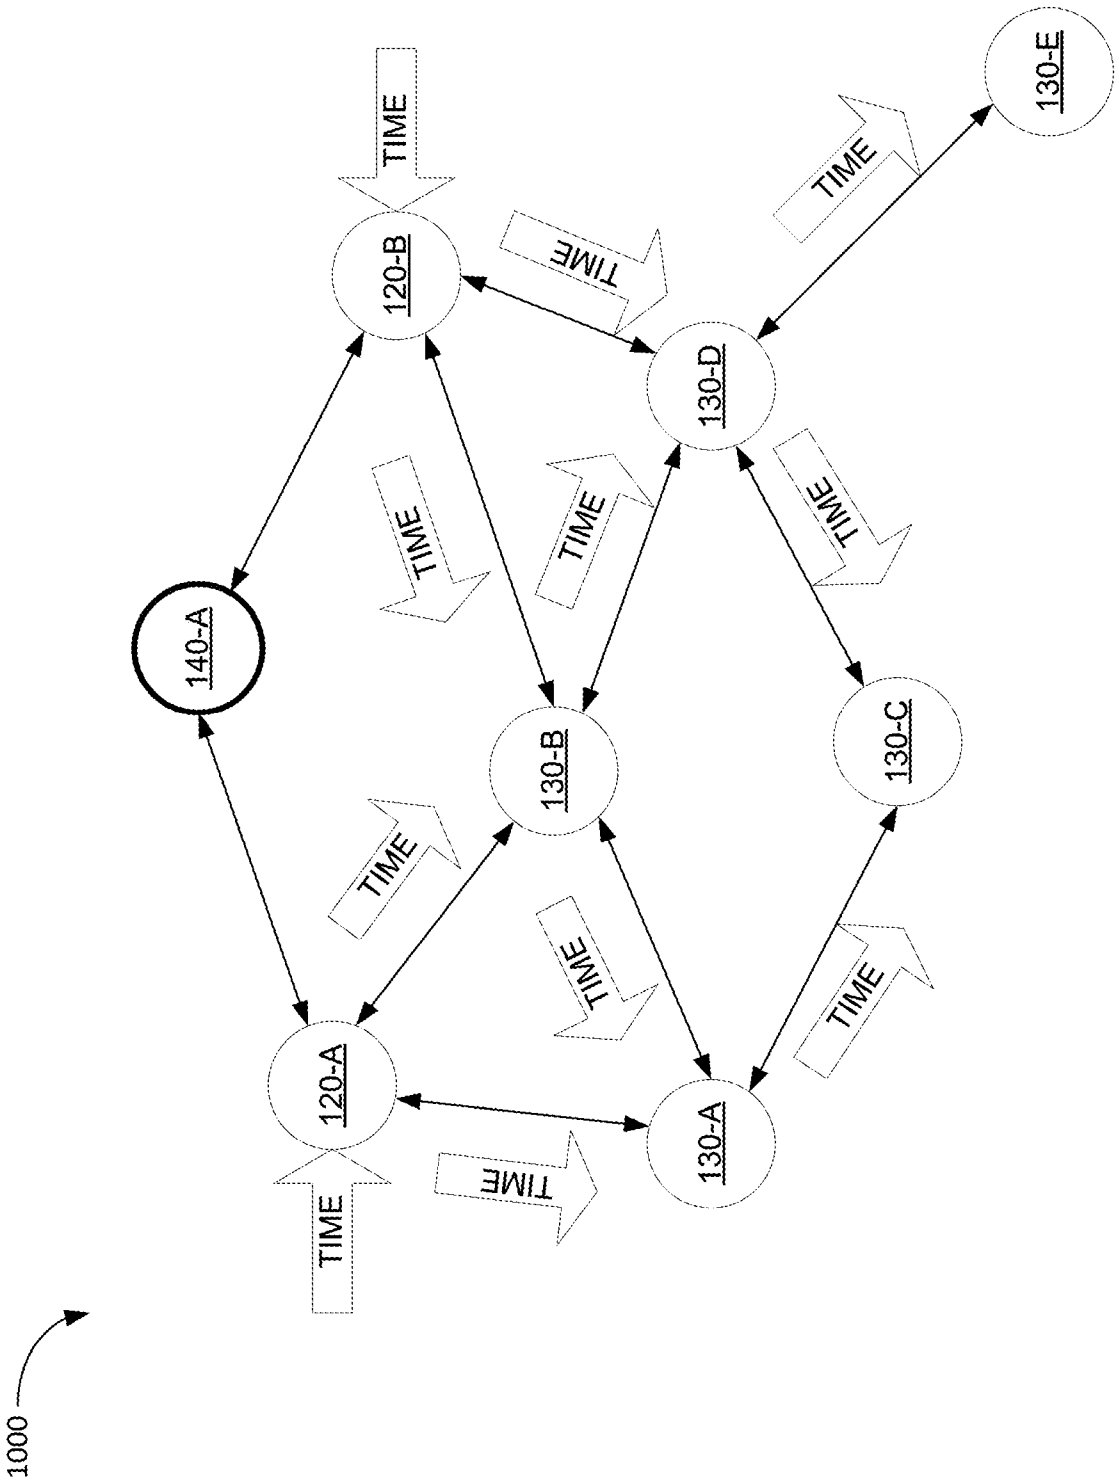

1000
140-A
120-A
120-B
130-A
130-B
130-C
130-D
130-E
TIME
TIME
TIME
TIME
TIME
TIME
TIME
TIME
TIME
TIME

FIG. 10E

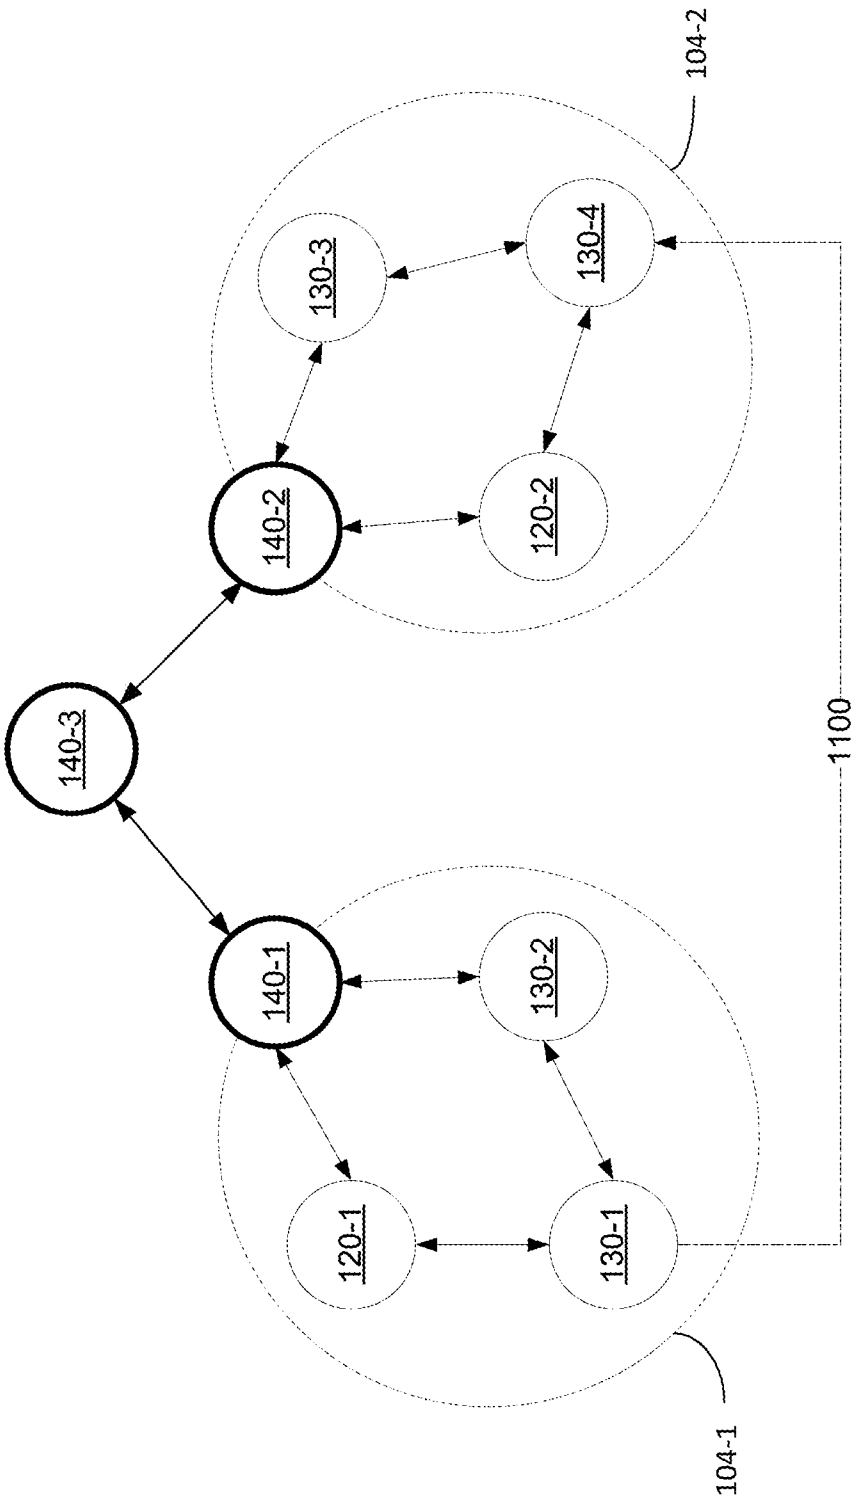


FIG. 11A

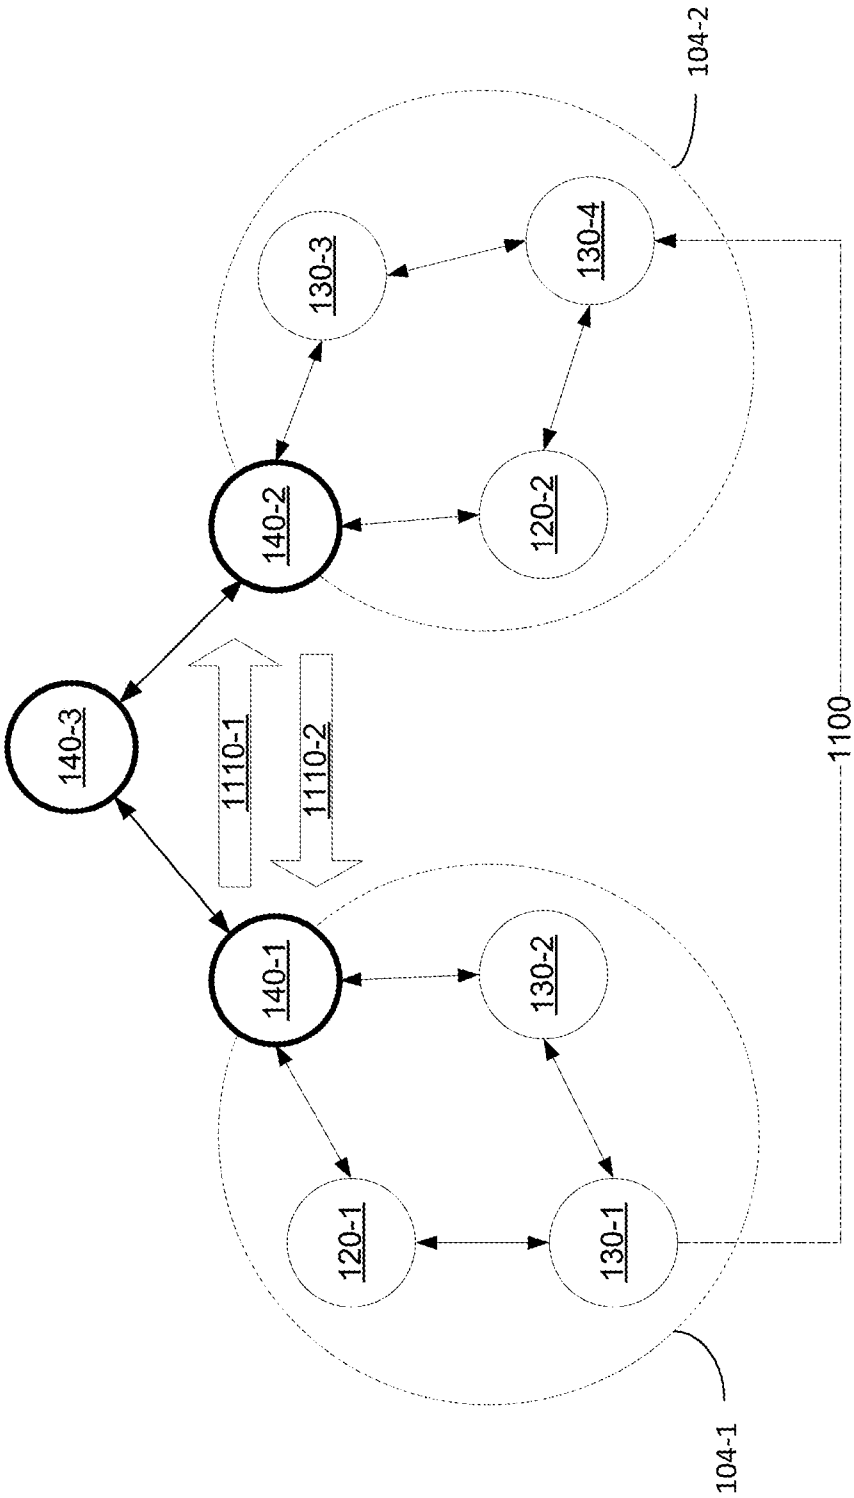

104-2
130-3
130-4
140-2
120-2
140-3
1110-1
1110-2
1100
140-1
130-2
120-1
130-1
104-1

FIG. 11B

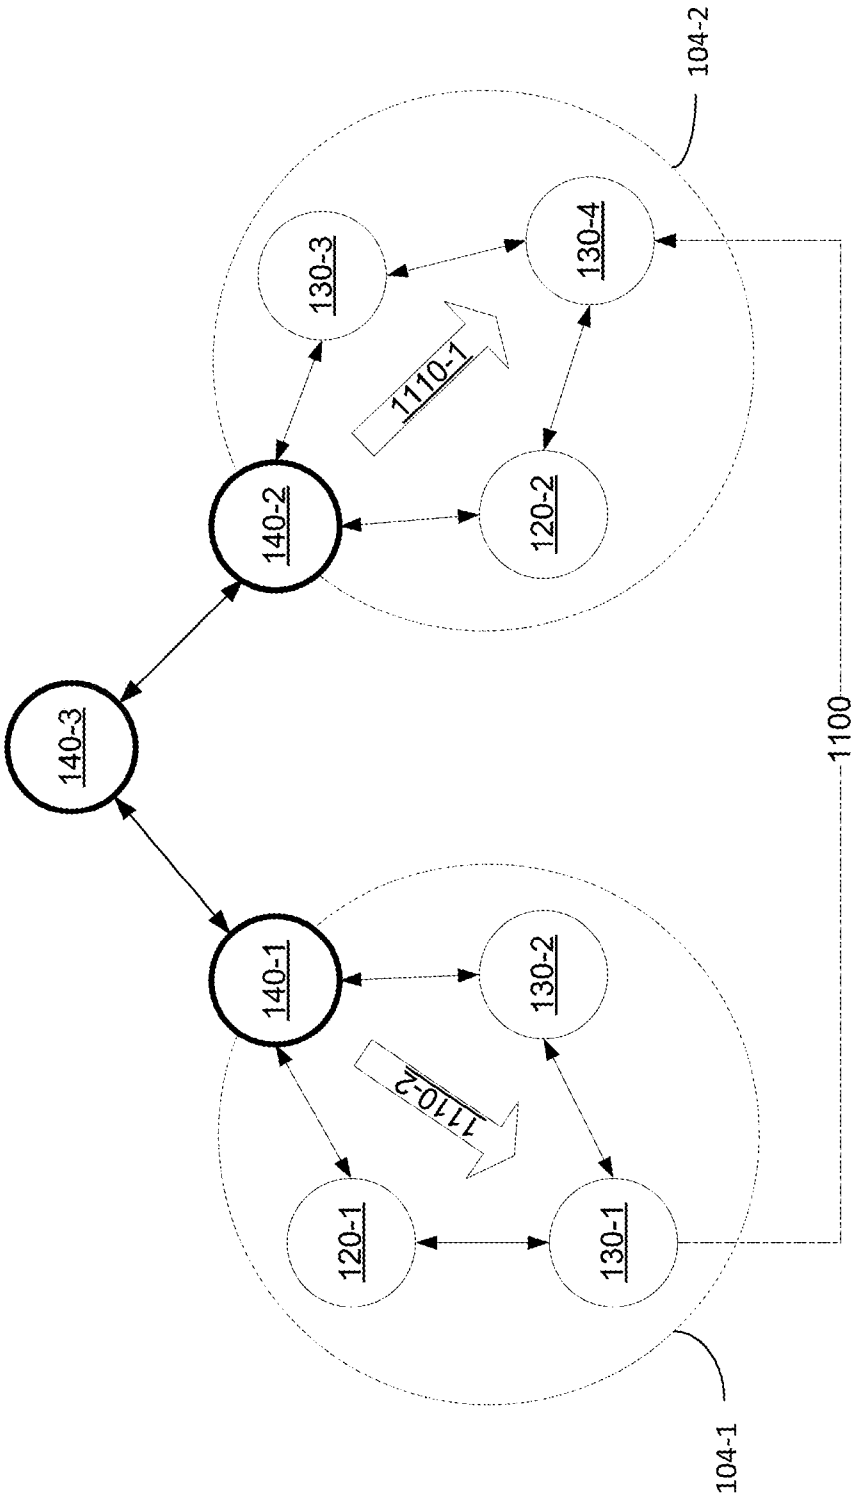

104-2
130-3
130-4
1110-1
140-2
120-2
140-3
1100
140-1
130-2
1110-2
120-1
130-1
104-1

FIG. 11C

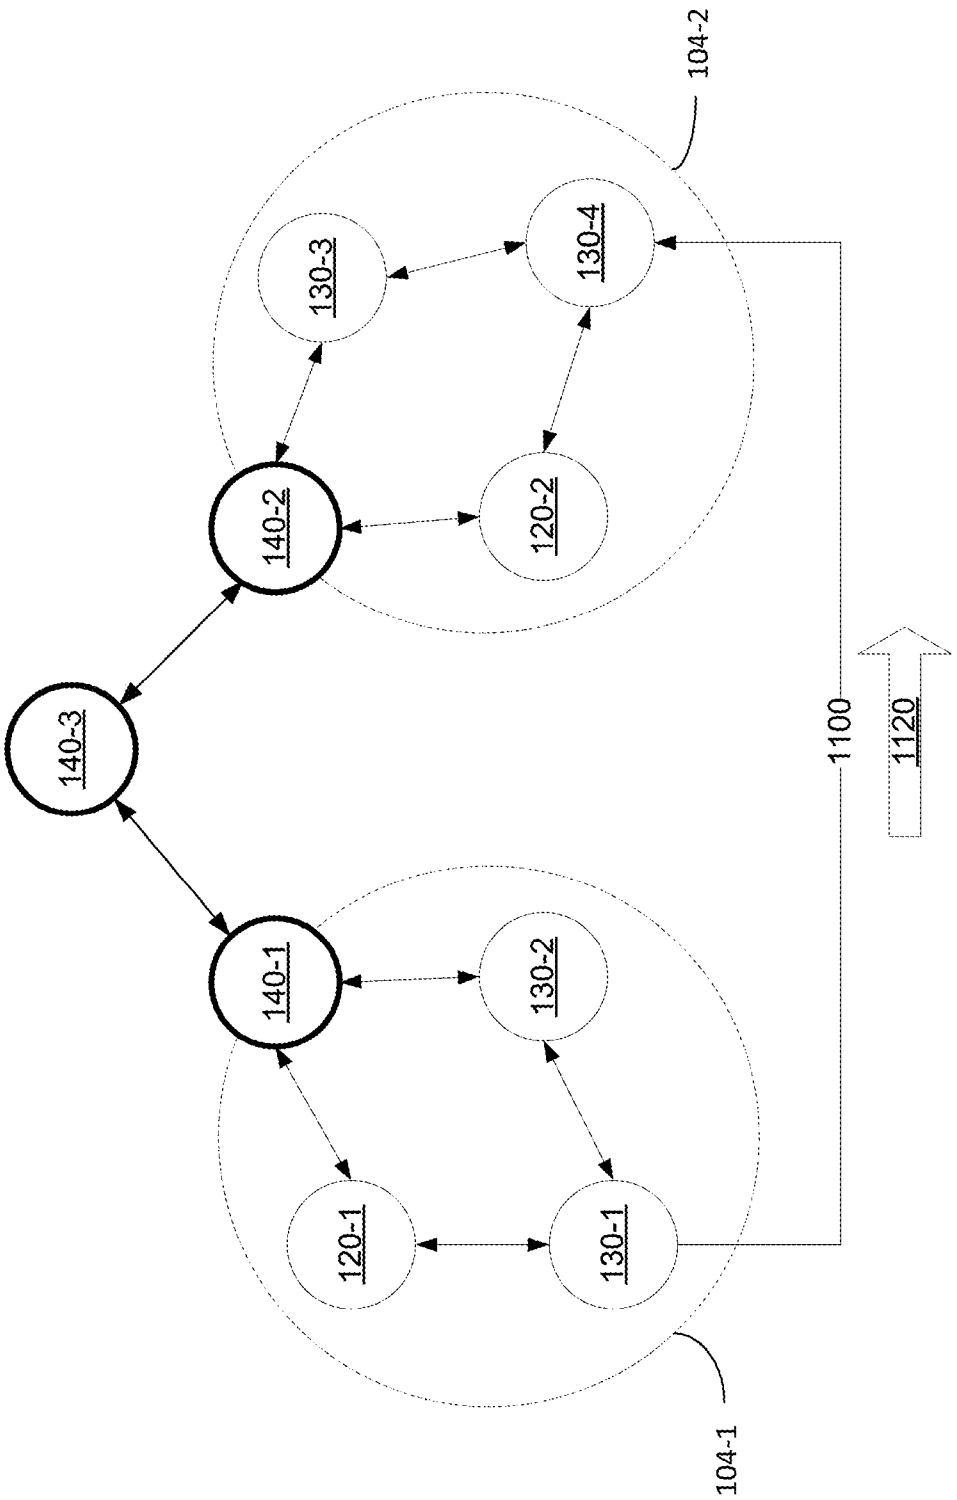

104-2
130-3
130-4
140-2
120-2
140-3
1100
1120
140-1
130-2
120-1
130-1
104-1

FIG. 11D

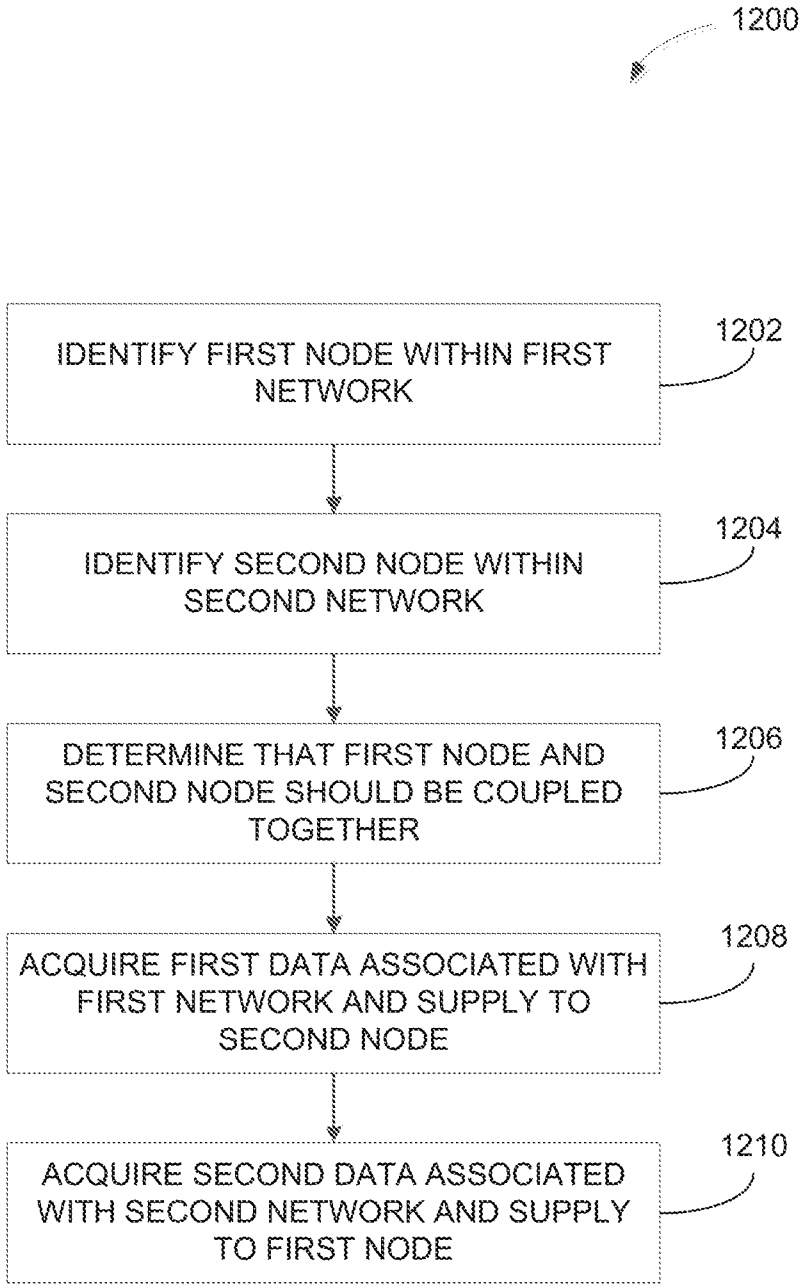

1200
IDENTIFY FIRST NODE WITHIN FIRST NETWORK
1202
IDENTIFY SECOND NODE WITHIN SECOND NETWORK
1204
DETERMINE THAT FIRST NODE AND SECOND NODE SHOULD BE COUPLED TOGETHER
1206
ACQUIRE FIRST DATA ASSOCIATED WITH FIRST NETWORK AND SUPPLY TO SECOND NODE
1208
ACQUIRE SECOND DATA ASSOCIATED WITH SECOND NETWORK AND SUPPLY TO FIRST NODE
1210

FIG. 12

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/015802

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/715 H04W56/00
ADD. H04W84/18

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)
H04L H04W

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, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 2 461 547 A1 (JUNIPER NETWORKS INC [US]) 6 June 2012 (2012-06-06)	1-22
A	page 5, paragraph 29 - page 12, paragraph 80 page 12, paragraph 86 - page 13, paragraph 92 page 16, paragraph 117 - page 17, paragraph 125 page 17, paragraph 128 - page 19, paragraph 141 figures 1, 3, 5, 7	23-44

	-/--	

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

[X] 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	Date of mailing of the international search report
28 June 2016	06/07/2016

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Koch, György

4

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/015802

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2009/154343 A1 (FITCH JEFFERY D [US] ET AL) 18 June 2009 (2009-06-18)	23-44
A	page 1, paragraph 4 - page 1, paragraph 6 page 2, paragraph 18 - page 3, paragraph 25 page 5, paragraph 40 - page 5, paragraph 43 figures 1, 2, 8	1-22

Y	US 2007/214254 A1 (AGUINIK ANATOLY [US]) 13 September 2007 (2007-09-13)	23-44
A	page 2, paragraph 14 - page 2, paragraph 17 figure 1	1-22

A	US 2005/050225 A1 (TATMAN LANCE A [US]) 3 March 2005 (2005-03-03) page 4, paragraph 34 - page 6, paragraph 50 page 8, paragraph 60 - page 8, paragraph 65 figures 3, 5	1-44

A	YICK J ET AL: "Wireless sensor network survey", COMPUTER NETWORKS, ELSEVIER SCIENCE PUBLISHERS B.V., AMSTERDAM, NL, vol. 52, no. 12, 22 August 2008 (2008-08-22), pages 2292-2330, XP022825257, ISSN: 1389-1286, DOI: 10.1016/J.COMNET.2008.04.002 [retrieved on 2008-04-14] page 2294, section "3. Types of sensor networks" page 2303, section "6. Network services" page 2309, section "7. Communication protocol"	1-44

4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/015802

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.

☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.

☒ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2016/015802

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 2461547 A1	06-06-2012	CN 102571557 A	11-07-2012
		EP 2461547 A1	06-06-2012
		US 2012144066 A1	07-06-2012
		US 2014229581 A1	14-08-2014
US 2009154343 A1	18-06-2009	US 2009154343 A1	18-06-2009
		WO 2009075889 A2	18-06-2009
US 2007214254 A1	13-09-2007	NONE	
US 2005050225 A1	03-03-2005	EP 1511248 A2	02-03-2005
		JP 2005080291 A	24-03-2005
		US 2005050225 A1	03-03-2005

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-22

   Performing data communication between two networks comprising separate and distinct nodes.

   ---

2. claims: 23-44

   Performing time synchronization in a wireless mesh network comprising battery-powered nodes and continuously-powered nodes.
