

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

(19) World Intellectual Property Organization

International Bureau



(10) International Publication Number

WO 2017/058547 A1

(43) International Publication Date

6 April 2017 (06.04.2017)

(51) International Patent Classification:

H04L 12/721 (2013.01) *H04W 40/02* (2009.01)
H04L 12/753 (2013.01) *H04W 84/18* (2009.01)

(21) International Application Number:

PCT/US2016/052444

(22) International Filing Date:

19 September 2016 (19.09.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/236,406 2 October 2015 (02.10.2015) US

(71) Applicant: **LANDIS+GYR INNOVATIONS, INC.**, [US/US]; 30000 Mill Creek Avenue, Suite 100, Alpharetta, Georgia 30022 (US).

(72) Inventors: **HOLCOMBE, Michael Sean**; 174 Spring Drive, Roswell, Georgia 30075 (US). **CALVERT, Christopher Sean**; 10405 Amberside Ct., Roswell, Georgia 30076 (US).

(74) Agents: **HOLMES, Brenda, Ozaki et al.**; Kilpatrick Townsend & Stockton LLP, Suite 2800 1100 Peachtree Street, Atlanta, Georgia 30309 (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, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

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



WO 2017/058547 A1

(54) Title: DETERMINING NETWORK RANK FOR COMMUNICATION TO NEIGHBORING NODES

(57) Abstract: The determination and advertisement of a node's rank in an RF mesh network reflects sustained network changes. A node's advertised rank may be based on the node's induced rank and a previous rank candidate. Induced rank is the rank of the node at a particular point in time. Rank candidate is based on a running average of the node's past rank. The advertised rank may be based on a weighted sum of the induced rank and a previous rank candidate. Induced rank may consider a Link Quality Indicator (LQI), which may be based on a running average LQI.

DETERMINING NETWORK RANK FOR COMMUNICATION TO NEIGHBORING NODES

[01] This application claims priority to U.S. Provisional Patent Application No. 62/236,406 filed on October 2, 2015, entitled "Maintaining Performance and Stability in a Mesh Network," which is incorporated herein in its entirety by reference.

Background

[02] Network rank is a quantification of the quality of a node's connection to a root node. A node's network rank may be advertised to its neighbors and used by the neighbors to evaluate the node as a potential parent. A node's rank is based on the number of hops required to reach the root node and other factors, such as the quality of the links between nodes.

[03] When a node's advertised rank fluctuates in real time, a network may become unstable. Brief aberrations in network quality may cause a node to select a parent based on aberrant data or may impact its advertised rank. Frequent changes in advertised rank may cause nodes to frequently change parents which may lead to instability in the network.

[04] It is desirable to have a system that incorporates gradual changes to the rank advertised by a node. Such a system would minimize the impact of temporary fluctuations in rank to provide for greater network stability, while still being responsive to persistent changes in rank.

Summary

[05] The present invention provides a method for calculating and advertising a rank that reflects changes in the network but ignores transitory anomalies. The advertised rank is based on the induced rank of the node and the previously advertised rank. The induced rank is the node's rank at a point in time and is based on the parent node's rank and the

quality of the link between the node and the parent node. A weighted running average or rank candidate is used to determine the advertised rank. In one example, the rank candidate is the weighted sum of the previous rank candidate value and the induced rank and is calculated each time the node advertises its rank. In this manner, only a sustained move to a new induced rank will result in a new advertised rank. This will result in a more stable network.

[06] Once the rank candidate is calculated, the difference between the rank candidate and a neighboring rank is determined. A neighboring rank is a rank value that is one step in rank value higher or lower than the previously advertised rank. When the difference between the rank candidate and the closest neighboring rank is within a rank threshold, the advertised rank is set to the neighboring rank value. When the difference between the rank candidate and the neighboring rank is greater than the rank threshold, the advertised rank is set to the previous advertised rank. The advertised rank is communicated to the other nodes.

[07] A Link Quality Indicator (LQI) may reflect the quality of the link between the node and the parent node. The value of the LQI may result in an LQI adjustment to the induced rank. The value of the LQI may be calculated using a running average to account for differences in the number of packets received by the node.

[08] These and other features, aspects, and advantages of the present disclosure are better understood when the following Detailed Description is read with reference to the accompanying drawings.

Brief Description of the Drawings

[09] Figure 1 is a block diagram illustrating a portion of an exemplary RF mesh network.

[010] Figure 2 is a block diagram illustrating the selection of a parent node.

[011] Figure 3 is a flowchart illustrating an exemplary method of adjusting the advertised rank.

[012] Figure 4 is a graph showing gradual changes in advertised rank.

[013] Figure 5 is a flowchart illustrating an exemplary method of adjusting the advertised rank using an upper threshold and a lower threshold.

[014] Figure 6 is a block diagram illustrating an exemplary node.

[015] Detailed Description

[016] Aspects of the invention support determining the rank advertised by a node so that the advertised rank reflects persistent changes in the network. A node's advertised rank may be based on the node's induced rank and a previous rank candidate. Induced rank is the rank of the node at a particular point in time. Rank candidate is based on a running average of the node's past rank. In one implementation, the advertised rank is based on a weighted sum of the induced rank and a previous rank candidate. When determining its induced rank, the node may consider the quality of its connection to its parent using a Link Quality Indicator (LQI). The LQI may be based on the average LQI over a time period, as well as a running average LQI. In one implementation, the LQI is a weighted sum of the average LQI over a time period and a previous running average LQI.

Exemplary Operating Environment

[017] Figure 1 is a diagram illustrating a portion of an exemplary RF mesh network. The nodes of the network may be configured to transmit and receive communications, as well as relaying communications for other nodes. The network may utilize the RPL routing protocol so that each node advertises its network rank to neighboring nodes.

[018] The network includes a root node 20. The network also includes a number of additional nodes, illustrated here as nodes 21 – 31. The root node 20 may communicate with additional nodes outside of the network depicted in Figure 1. Although Figure 1 depicts a single root node and eleven additional nodes for illustrative purposes, the RF mesh network can include any number of root nodes and any number of additional nodes. Figure 1 illustrates a number of parent-child relationships. For example, root node 20 is the parent node for nodes 21, 22, and 23. A node may be both a parent and a child. For example, node 22 is a child of root node 20 and is also a parent of node 26. A parent node is a node that is between the root node and the child node and that communicates directly with the child node. The child node will periodically re-evaluate its parent node and may select a new parent node if the new parent node provides a better quality connection to the root node. In one non-limiting example, root node 20 is a collector for an AMI network and nodes 21-31 include metering devices for measuring resource consumption at different premises.

[019] Figure 2 illustrates the selection of a parent node by node 250. The portion of the network shown in Figure 2 includes root node 200 and nodes 210, 220, 240, 250, 260. Node 250 selects a parent node to minimize its own rank. In the following examples, rank is expressed in the form of binary numbers, or multiples of 256. Rank may also be expressed in the form of hexadecimal numbers, or multiples of 100. Other numbering systems may also be used.

[020] To select a parent node, node 250 considers the advertised rank of its neighboring nodes and the link quality between it and each neighboring node. Node 250 maintains LQI information for its neighboring nodes based on its communications with the nodes. The LQI may result in one or more rank values being added to the rank calculation as an LQI

adjustment. In one implementation, when the LQI is below a threshold, indicating that the link quality is good, then the LQI adjustment is zero, when the LQI is between a first value and the threshold, indicating that the link quality is fair, one rank value may be added, and when the LQI is above the first value, indicating that the link quality is poor, two rank values may be added.

[021] Returning to Figure 2, assume the advertised rank of node 210 is 256, the advertised rank of node 220 is 256, the advertised rank of node 240 is 512, and the advertised rank of node 260 is 512. Also assume that the LQI between nodes 210 and 250 is above the first value, the LQI between nodes 240 and 250 is between the first value and the threshold, the LQI between nodes 220 and 250 is above the first value, and the LQI between nodes 250 and 260 is below the threshold. Node 250 may consider nodes 210, 220, 240, and 260 as potential parent nodes and select the one that provides node 250 with the lowest rank.

[022] To compare the potential parent nodes, node 250 may determine the impact of the selection of the different potential parent nodes on its rank. If node 250 selects node 210 as its parent, then its rank is 1024 (advertised rank of node 210 + 2 rank values for LQI adjustment + 1 rank value for one hop). If node 250 selects node 240 as its parent, then its rank is 1024 (advertised rank of node 240 + 1 rank value for LQI adjustment + 1 rank value for one hop). If node 250 selects node 220 as its parent, then its rank is 1024 (advertised rank of node 220 + 2 rank values for LQI adjustment + 1 rank value for one hop). If node 250 selects node 260 as its parent, then its rank is 768 (advertised rank of Node 260 + no LQI adjustment + 1 rank value for one hop). Since node 250 is seeking to minimize its own rank, it selects node 260 as its parent.

[023] In another example, assume the advertised rank of node 210 is 256, the advertised rank of node 220 is 512, the advertised rank of node 240 is 512, and the advertised rank of

node 260 is 768. Also assume that the LQI between nodes 210 and 250 is above the first value, the LQI between nodes 240 and 250 is between the first value and the threshold, the LQI between nodes 220 and 250 is below the threshold, and the LQI between nodes 250 and 260 is below the threshold.

[024] In this example, if node 250 selects node 210 as its parent, then its rank is 1024 (advertised rank of node 210 + 2 rank values for LQI adjustment + 1 rank value for one hop). If node 250 selects node 240 as its parent, then its rank is 1024 (advertised rank of node 240 + 1 rank value for LQI adjustment + 1 rank value for one hop). If node 250 selects node 220 as its parent, then its rank is 768 (advertised rank of node 220 + no LQI adjustment + 1 rank value for one hop). If node 250 selects node 260 as its parent, then its rank is 1024 (advertised rank of node 260 + no LQI adjustment + 1 rank value for one hop). Since node 250 is seeking to minimize its own rank, it selects node 220 as its parent.

[025] A node selects only one parent node. If there are multiple potential parent nodes that result in the same rank for the node, then the node may give one of the factors used to determine its rank more weight. One option is to select the potential parent having the fewest number of hops to the root. Another option is to select the potential parent with the best LQI. Other options may use multiple factors or different factors to select one of the nodes as its parent when there is no difference in rank for the node.

[026] In another example, assume the advertised rank of node 210 is 256, the advertised rank of node 220 is 256, the advertised rank of node 240 is 512, and the advertised rank of node 260 is 512. Also assume that the LQI between nodes 210 and 250 is between the threshold and the first value, the LQI between nodes 240 and 250 is below the threshold, the LQI between nodes 220 and 250 is above the first value, and the LQI between nodes 250 and 260 is between the threshold and the first value.

[027] In this example, if node 250 selects node 210 as its parent, then its rank is 768 (advertised rank of node 210 + 1 rank values for LQI adjustment + 1 rank value for one hop). If node 250 selects node 240 as its parent, then its rank is 768 (advertised rank of node 240 + no LQI adjustment + 1 rank value for one hop). If node 250 selects node 220 as its parent, then its rank is 1024 (advertised rank of node 220 + 2 rank values for LQI adjustment + 1 rank value for one hop). If node 250 selects node 260 as its parent, then its rank is 1024 (advertised rank of node 260 + 1 rank value for LQI adjustment + 1 rank value for one hop). In this example, nodes 210 and 240 result in the same rank for Node 250. If the criteria for selecting a parent node includes minimizing the number of hops to the root, then node 250 selects node 210 as its parent. The node may determine the number of hops to the root using a routing metric, such as the Hop Count Object defined by the RPL specification. If the criteria for selecting a parent node includes the best LQI between the parent and child, then node 250 selects node 240 as its parent.

[028] A node's rank is dynamic and may change as the rank of its parent or the LQI between it and its parent changes. The term induced rank as used herein, refers to the node's rank at a particular time. It is based on the advertised rank of the parent node, plus one rank value for the hop from the parent node to the child, plus any LQI adjustment for the link between the parent node and the child node.

[029] A node advertises its rank to its neighboring nodes. To improve network stability, a node may advertise a rank referred to herein as its advertised rank that may be different than its induced rank. In one example, the advertised rank is calculated and sent out in a DIO (DODAG Information Object) message as defined by the RPL specification. The frequency of the transmission of the DIO message is variable. In one example, the DIO is sent approximately once per hour, but it may be sent more frequently after a network

change. In a mesh network, the DIO may be broadcast to the other nodes in the network.

In one exemplary system, the frequency of the calculation of the induced rank is independent of the frequency of the transmission of the advertised rank. When the network or the particular node is initialized, the advertised rank is set to be equal to the induced rank. Subsequently, the advertised rank is determined as described below.

[030] Exemplary Method for Determining Advertised Rank

Rank Candidate = ($\frac{3}{4}$ * previous Rank Candidate) + ($\frac{1}{4}$ * current Induced Rank).

[031] The rank candidate is determined each time the node advertises its rank.

[032] In 306, the difference between the rank candidate and a neighboring rank is determined. A neighboring rank is a rank value that is one step in rank value higher or lower than the previously advertised rank. For example, if the previously advertised rank is 768, the neighboring rank values are 512 and 1024. The difference between the rank candidate and the closest neighboring rank is compared to a threshold, referred to herein as a rank threshold, in 308. The rank threshold is configurable, and may vary according to the installation. The rank threshold may be the same for all nodes throughout the network. In one example, the rank threshold is set to approximately 10% of one rank value.

[033] When the difference between the rank candidate and the neighboring rank is within, or less than, the rank threshold, the process continues to 310 and the advertised rank is set to the neighboring rank value. In one example, the rank candidate value is also set to the advertised rank value, for the purpose of calculating the rank candidate in the next iteration.

[034] When the difference between the rank candidate and the neighboring rank is greater than the rank threshold, the process continues to 312 and the advertised rank is set to the previous advertised rank.

[035] After the advertised rank has been set at 310 or 312, the process continues to 314, where the advertised rank is communicated to the other nodes.

[036] Figure 4 illustrates how the advertised rank may change over time and how the advertised rank relates to the induced rank. The y-axis represents the rank value. The x-axis represents sequential time periods at which the advertised rank is determined. In one implementation the advertised rank is determined whenever the node transmits a DIO message. Although the divisions along the x-axis are shown as being equally spaced, the time between communications may vary. For the purposes of illustration, one rank value is 256 and the rank threshold is 20.

[037] At time 1 and time 2, the advertised rank, the induced rank, and the rank candidate are all 512. At time 3, the induced rank is 768. The rank candidate may be calculated using the equation provided above. If so, then the rank candidate is 576 (.75(previous rank candidate of 512) + .25(induced rank of 768)). The rank candidate is compared to the neighboring rank, 768. Since the difference between the rank candidate (576) and the neighboring rank (768) is greater than the rank threshold (192 > 20), the advertised rank is set to be the previous advertised rank or 512.

[038] At time 14, the advertised rank is 512, the rank candidate is 744, and the induced rank is 768. At time 15, the induced rank is 768. The rank candidate may be calculated using the equation provided above. If so, then the rank candidate is 750 (.75(previous rank candidate of 744) + .25(induced rank of 768)). Since the difference between the rank

candidate (750) and the neighboring rank (768) is within the rank threshold ($18 < 20$), the advertised rank is set to a new rank value of 768.

[039] At time 18, the advertised rank, the rank candidate, and the induced rank are all 768. At time 19, the induced rank is 512. The rank candidate may be calculated using the equation provided above. If so, then the rank candidate is 704 (.75(previous rank candidate of 768) + .25(induced rank of 512)). Since the difference between the rank candidate (704) and the neighboring rank (512) is greater than the rank threshold ($192 > 20$), the advertised rank is set to the previous advertised rank of 768.

[040] At time 26, the advertised rank is 768, the rank candidate is 520, and the induced rank is 512. At time 27, the induced rank is 512. The rank candidate may be calculated using the equation provided above. If so, then the rank candidate is 518 (.75(previous rank candidate of 520) + .25(induced rank of 512)). Since the difference between the rank candidate (518) and the neighboring rank (512) is within the rank threshold ($6 < 20$), the advertised rank is set to a new rank value of 512.

[041] Figure 4 illustrates that the advertised rank changes more slowly than the induced rank. Although Figure 4 used the same rank threshold for a neighboring rank above the previous advertised rank or below the previous advertised rank, some implementations may use different rank thresholds. This is illustrated in Figure 5. In this case, an upper threshold may be used to determine when to increase the advertised rank and a lower threshold may be used to determine when to decrease the advertised rank. In 502, this method starts with the rank candidate determined in step 304 in Figure 3. In 504, the node evaluates whether the rank candidate is above or below the previous advertised rank. When the rank candidate is above the previous advertised rank, the node proceeds to 506 and determines the difference between the rank candidate and the upper neighboring rank. At 508, the

node compares the difference calculated in 506 to the upper threshold. When the difference is within the upper threshold, the node proceeds to step 310 in Figure 3. When the difference is outside of the upper threshold, the node proceeds to step 312 in Figure 3. In this example, the advertised rank is only changed when the rank candidate is within the upper threshold.

[042] When the rank candidate is below the previous advertised rank in 504, the node proceeds to 510 and determines the difference between the rank candidate and the lower neighboring rank. At 512, the node compares the difference from 510 to the lower threshold. If the difference is within the lower threshold, the node proceeds to step 310 in Figure 3. If the difference is outside of the lower threshold, the node proceeds to step 312 in Figure 3. The advertised rank is only changed when the rank candidate is within the lower threshold.

Link Quality Indicator

[043] Link Quality Indicator (LQI) is an indicator of the RF signal quality between a node and a neighboring node. Each packet that a node receives from a neighboring node has an associated LQI. The LQI may be based on the signal to noise ratio (SNR) and the packet success rate. In one system each time a frame is received from a neighboring node, the SNR is recorded. The node includes a translation table that is based upon the RF characteristics of the node's transceiver hardware. The translation table is used to convert the SNR to LQI so that a "good" SNR value is translated to a "good" LQI value for RPL. The LQI scale is consistent in RPL across hardware platforms, whereas the translation tables are adapted for the different hardware platforms used by the nodes. Some systems use packet success rate. One way that packet success rate may be used is to adjust the LQI based on packet success rate. For example, if the SNR value is "good", but the packet success rate is bad, then the

translation table provides a “good” LQI value, but the packet success rate adjusts the LQI value so that it is not quite as “good”. In this manner, the LQI reflects both SNR and other factors that affect packet success rate.

[044] Since a node is likely to receive more packets from certain nodes, *e.g.*, its parent, there may be significant variations in the number of packets used to calculate LQI values for different nodes. To account for the different number of packets received from different nodes, a running average calculation may be used. An average LQI is determined over a predetermined time period by adding the LQI values for the packets received from a node during the time period and dividing the total by the number of received packets. A running average LQI is determined by a weighted sum of the average LQI and the previous running average LQI. In one example, the following is used:

$$\text{LQI(running average)} = (\frac{3}{4} * \text{LQI(previous running average)}) + (\frac{1}{4} * \text{LQI(average)}).$$

[045] If the node does not receive any packets from one of its neighboring nodes during the time period, then the average LQI for the previous time period may be used. The time period used to calculate the average LQI may be independent of when the node communicates its advertised rank.

[046] Figure 6 illustrates an exemplary node that may be used in the network of Figure 1. In one example, a node can include a meter located at a premises, which measures the consumption of a resource such as gas, water, or electric power. Such a meter can be part of an RF network used for AMI. Other examples of nodes include a router, collector or collection point, host computer, hub, or other electronic device that is attached to a

network and is capable of sending, receiving, or forwarding information over a communication channel.

[047] A node 600 may include a network interface 604 and an antenna 602 so that it may communicate with like nodes and/or other devices in the mesh network. A network interface 604 may communicate with the microcontroller 612 using the bus 606, which enables the node to function like a computer, carrying out computer and command functions to provide implementations of the present invention described herein. The microcontroller 612 may include a processor 608 and a storage medium, such as memory 610. The processor 608 can be any suitable processing device or group of devices configured to execute instructions stored in the memory.

[048] The memory 610 can be any suitable non-transitory computer-readable medium for providing computer-executable program instructions, such as firmware or other operating instructions for controlling one or more hardware components of a node, to the processor 608. Examples of a non-transitory computer-readable medium may include, but are not limited to, a memory chip, ROM, RAM, an ASIC, a configured processor, electronic storage, optical storage, magnetic storage, or other storage device from which a computer processor can read instructions. The instructions may include processor-specific instructions generated by a compiler and/or an interpreter from code written in any suitable computer-programming language. The memory may be located internal to the node or accessible by the node via a network, for example. A node may also include a crystal oscillator (i.e. a clock) to provide timekeeping and an energy storage device (i.e. a battery) to provide back-up power. Some nodes may be powered only by an energy storage device. Some nodes include additional components, such as a utility meter that measures the consumption of a resource or service.

General Considerations

[049] The foregoing is provided for purposes of illustrating, describing, and explaining aspects of the present invention and is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Further modifications and adaptations of these embodiments will be apparent to those skilled in the art and may be made without departing from the scope and spirit of the invention. For example, different weighting factors may be used to determine the rank candidate or the running average LQI. Additional factors may also be used to determine induced rank or LQI. Different arrangements of the components described above, as well as components and steps not shown or described are possible. Similarly, some features and sub-combinations are useful and may be employed without reference to other features and sub-combinations. Embodiments of the invention have been described for illustrative and not restrictive purposes, and alternative embodiments will become apparent to readers of this application.

CLAIMS

What is claimed is:

1. A method for communicating an advertised rank of a node in a mesh network, comprising:
 - receiving an advertised rank of a parent node of the node;
 - determining an induced rank of the node, wherein the induced rank is based on the advertised rank of the parent node and a Link Quality Indicator (LQI) adjustment for a link between the node and the parent node;
 - determining a rank candidate for the node based on the induced rank and a previous rank candidate for the node, wherein the previous rank candidate was determined in connection with a previous advertised rank of the node;
 - comparing the rank candidate and a neighboring rank, wherein the neighboring rank is an adjacent rank value to the previous advertised rank of the node;
 - when the rank candidate is within a rank threshold of the neighboring rank, then determining that the advertised rank is the neighboring rank;
 - when the rank candidate is outside the rank threshold of the neighboring rank, then determining that the advertised rank is the previous advertised rank; and
 - communicating the advertised rank to other nodes in the mesh network.

2. The method of claim 1, wherein the LQI adjustment is equal to at least one rank value, and determining an induced rank comprises summing the advertised rank of the parent node, the at least one rank value for the LQI adjustment, and one rank value.

3. The method of claim 1, wherein the LQI for the link between the parent node and the node is a weighted sum of an average LQI for a time period and a running average LQI for a previous time period.
4. The method of claim 1, wherein the rank threshold has one value when the neighboring rank is above the previous advertised rank and another value when the neighboring rank is below the previous advertised rank.
5. The method of claim 1, wherein the rank candidate is a weighted sum of the induced rank and the previous rank candidate.
6. The method of claim 1, wherein communicating the advertised rank to other nodes in the mesh network comprises including the advertised rank in a DODAG Information Object (DIO) message.

7. A method for communicating an advertised rank of a node in a mesh network, comprising:
 - determining an induced rank of the node, wherein the induced rank is based on a rank of a parent node of the node and a Link Quality Indicator (LQI) for a link between the node and the parent node;
 - determining a rank candidate for the node based on the induced rank and a previous rank candidate for the node, wherein the previous rank candidate was determined in connection with a previous advertised rank of the node;
 - comparing the rank candidate and a neighboring rank, wherein the neighboring rank and the previous advertised rank are adjacent rank values;
 - when the rank candidate is within a rank threshold of the neighboring rank, then determining that the advertised rank is the neighboring rank;
 - when the rank candidate is outside the rank threshold of the neighboring rank, then determining that the advertised rank is the previous advertised rank; and
 - communicating the advertised rank to other nodes in the mesh network.

8. The method of claim 7, wherein determining the rank candidate for the node based on the induced rank and the previous rank candidate for the node, further comprises:
 - determining that the rank candidate is equal to a weighted sum of the previous rank candidate and the induced rank, wherein a weight associated with the previous rank candidate is greater than the weight associated with the induced rank.

9. The method of claim 7, wherein communicating the advertised rank to other nodes in the mesh network comprises transmitting a DIO (DODAG Information Object) message.
10. The method of claim 7, wherein the rank threshold has one value when the neighboring rank is above the previous advertised rank and another value when the neighboring rank is below the previous advertised rank.
11. The method of claim 7, wherein determining the induced rank of the node is performed more often than determining the rank candidate.
12. The method of claim 7, wherein communicating the advertised rank to other nodes in the mesh network comprises broadcasting the advertised rank in the mesh network.
13. The method of claim 7, wherein the node determines the LQI for the link between the node and the parent node based on signal to noise ratio and packet success rate for communications with the parent node.

14. A method for determining an advertised rank of a node in a mesh network,

comprising:

receiving a plurality of packets from a parent node of the node, wherein each packet has an associated Link Quality Indicator (LQI);

determining an induced rank of the node, wherein the induced rank is based on a rank of the parent node and an LQI adjustment based on a running average LQI;

determining a rank candidate for the node based on the induced rank and a previous rank candidate for the node, wherein the previous rank candidate was determined in connection with a previous advertised rank of the node;

comparing the rank candidate and a neighboring rank, wherein the neighboring rank and the previous advertised rank are adjacent rank values;

when the rank candidate is within a rank threshold of the neighboring rank, then determining that the advertised rank is the neighboring rank;

when the rank candidate is outside the rank threshold of the neighboring rank, then determining that the advertised rank is the previous advertised rank; and

communicating the advertised rank to other nodes in the mesh network.

15. The method of claim 14, wherein the plurality of packets are received during a time period, an average LQI is determined using the LQIs for the received packets, and a previous running average LQI was determined during a prior time period, and the running average LQI is a weighted sum of the average LQI and the previous running average LQI.

16. The method of claim 14, wherein the LQI adjustment is zero when the running average LQI is below a threshold value, the LQI adjustment is one rank value when the running average LQI is between the threshold and a first value, and the LQI adjustment is two rank values when the running average LQI is above the first value.

17. The method of claim 14, wherein determining an induced rank comprises summing the advertised rank of the parent node, the LQI adjustment, and one rank value.

18. The method of claim 14, wherein the rank threshold has one value when the neighboring rank is above the previous advertised rank and another value when the neighboring rank is below the previous advertised rank.

19. The method of claim 14, further comprising:
translating a signal to noise ratio for a frame received from the parent node to the LQI using a translation table.

20. The method of claim 14, wherein the LQI is based in part on a packet success rate.

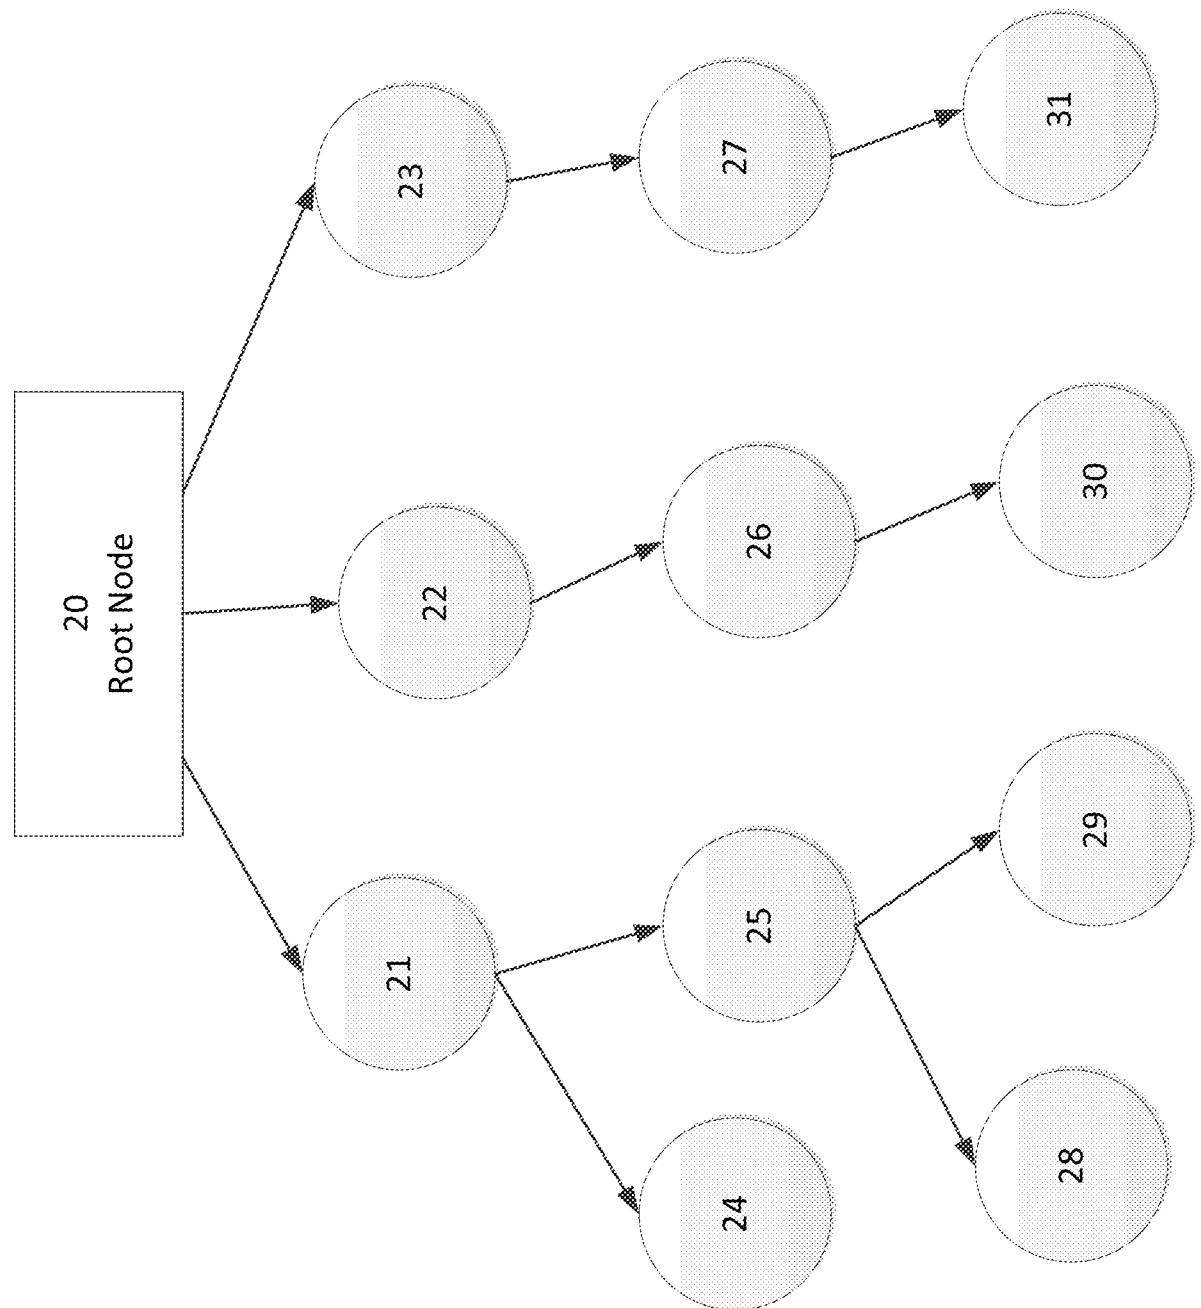


Figure 1

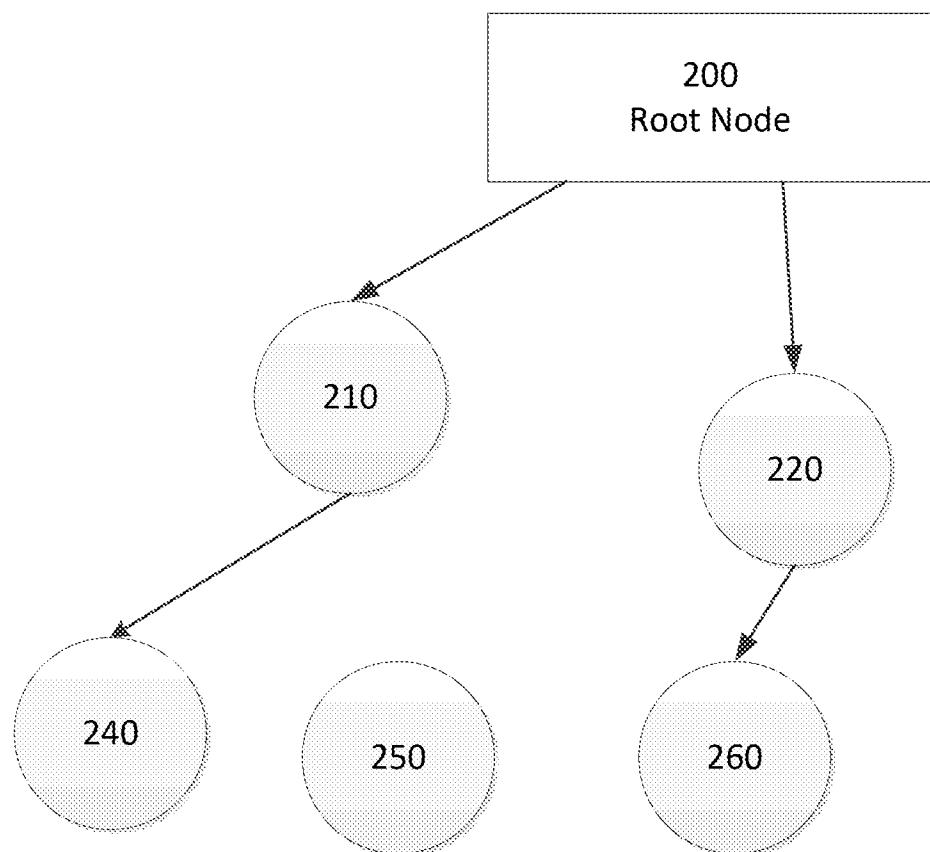


Figure 2

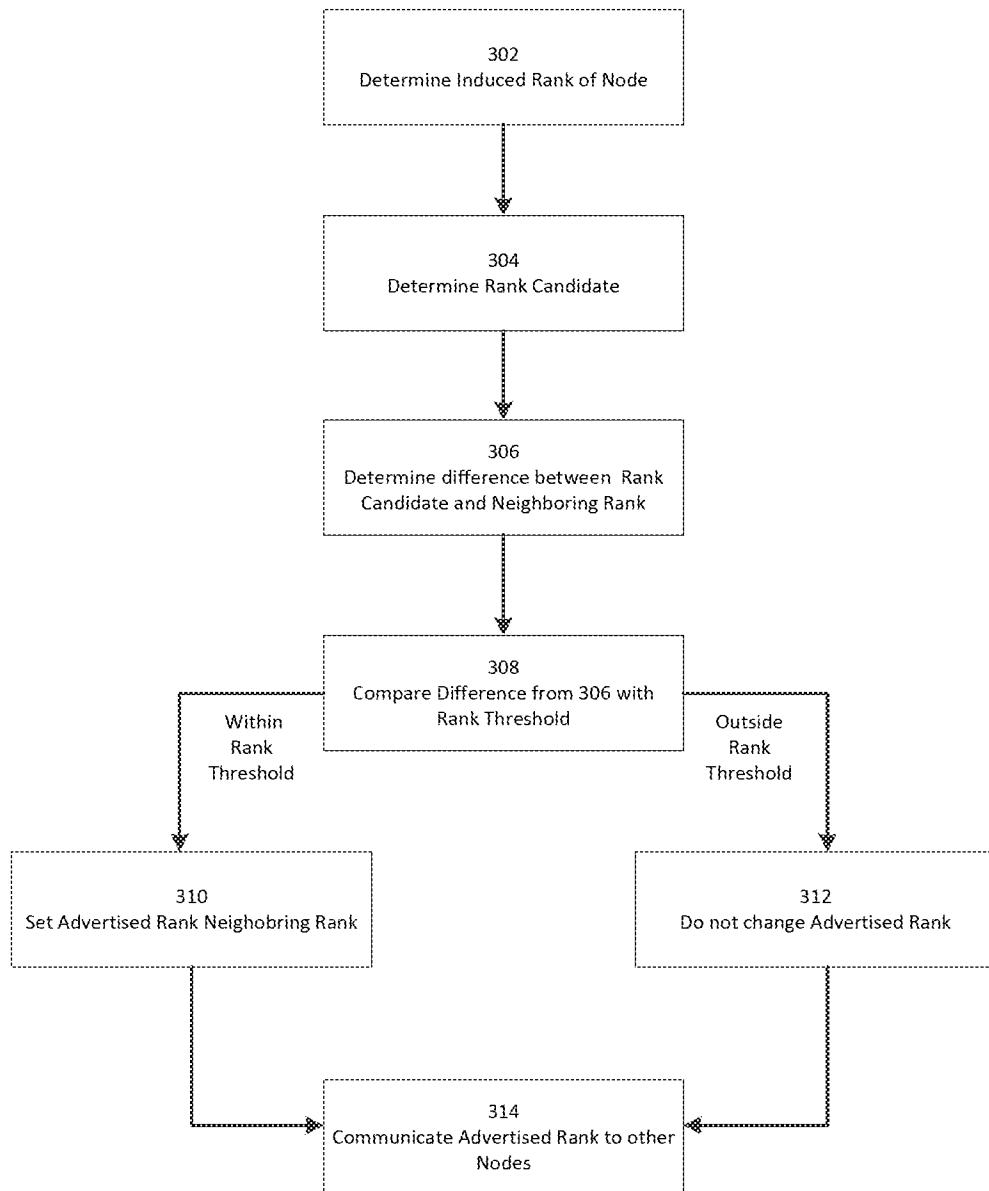


Figure 3

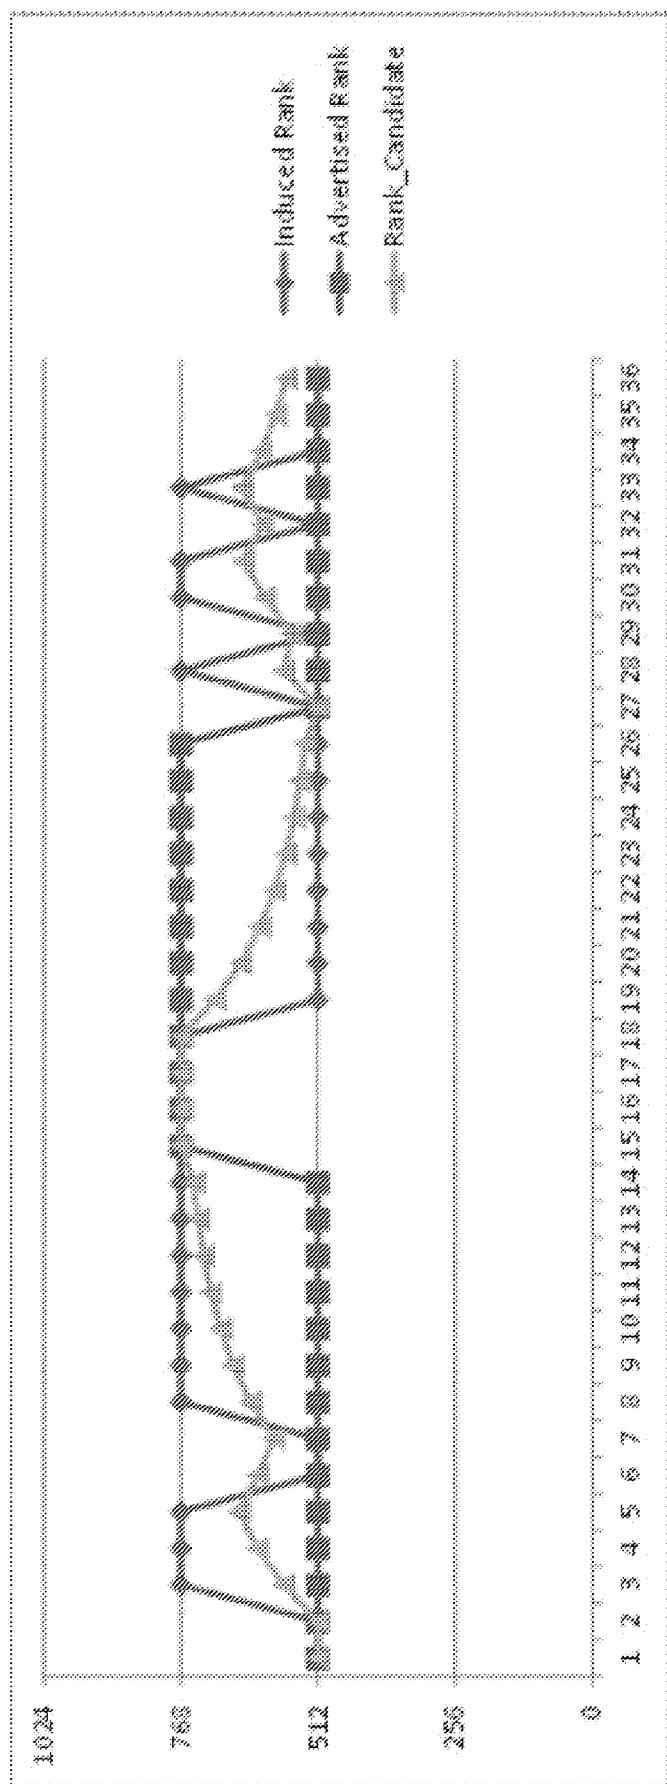


Figure 4

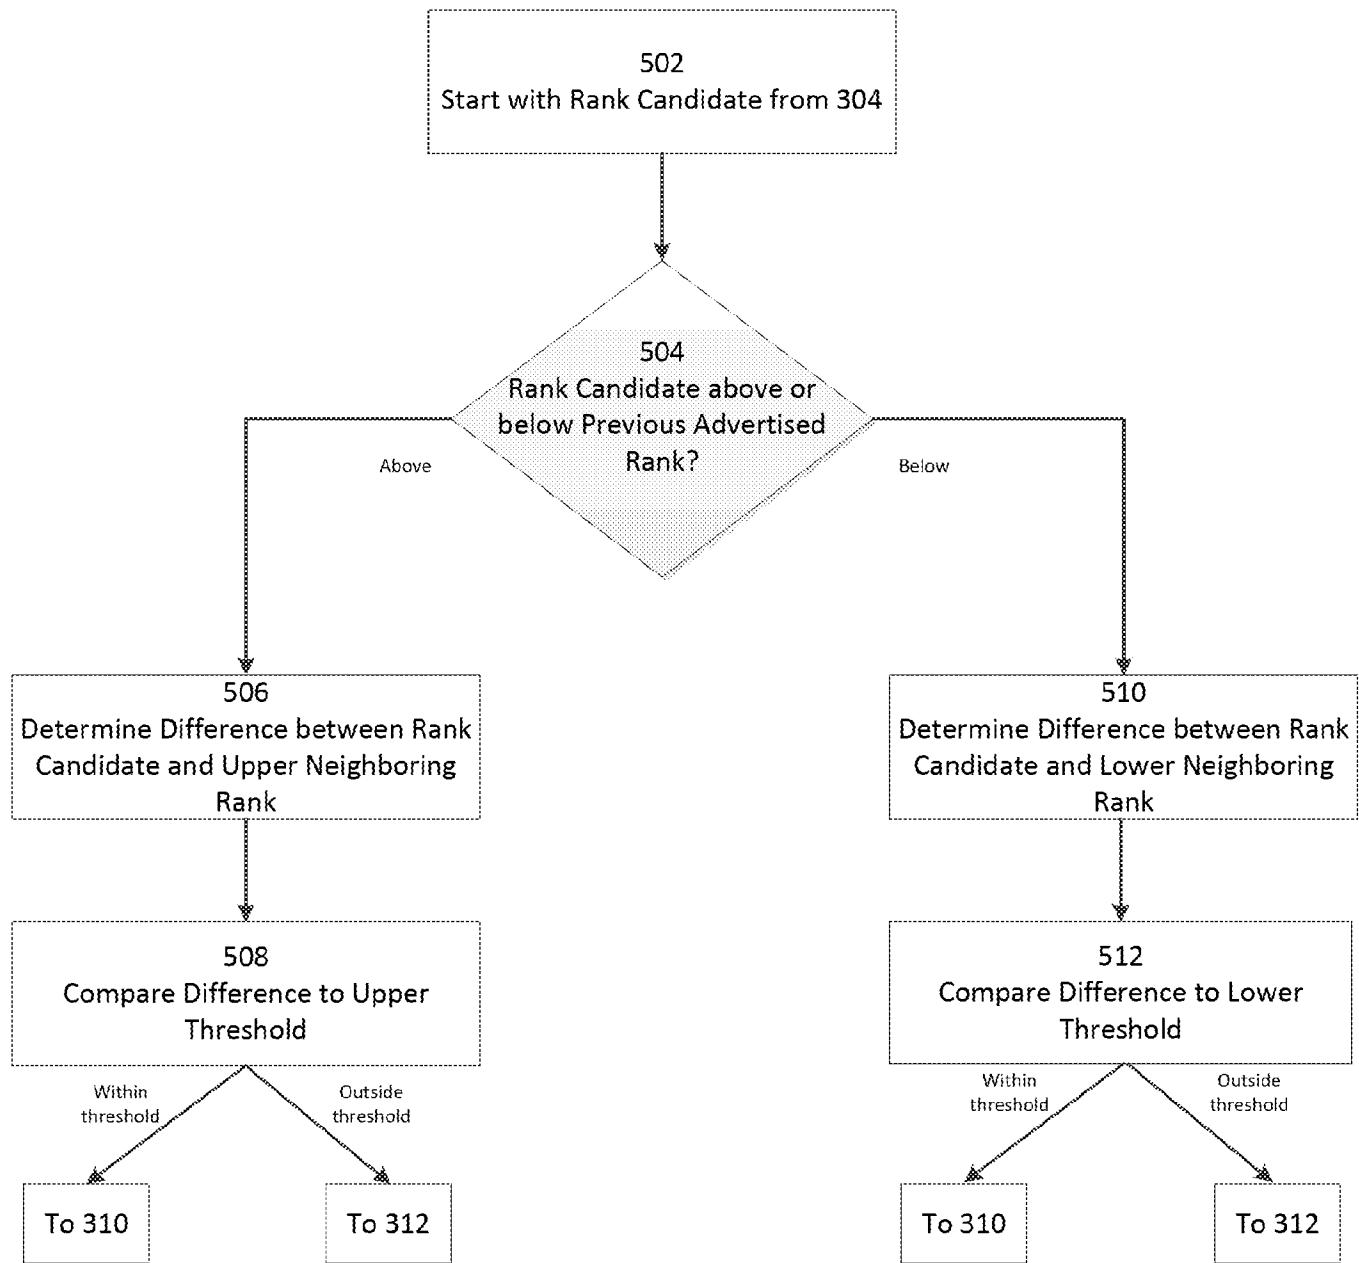


Figure 5

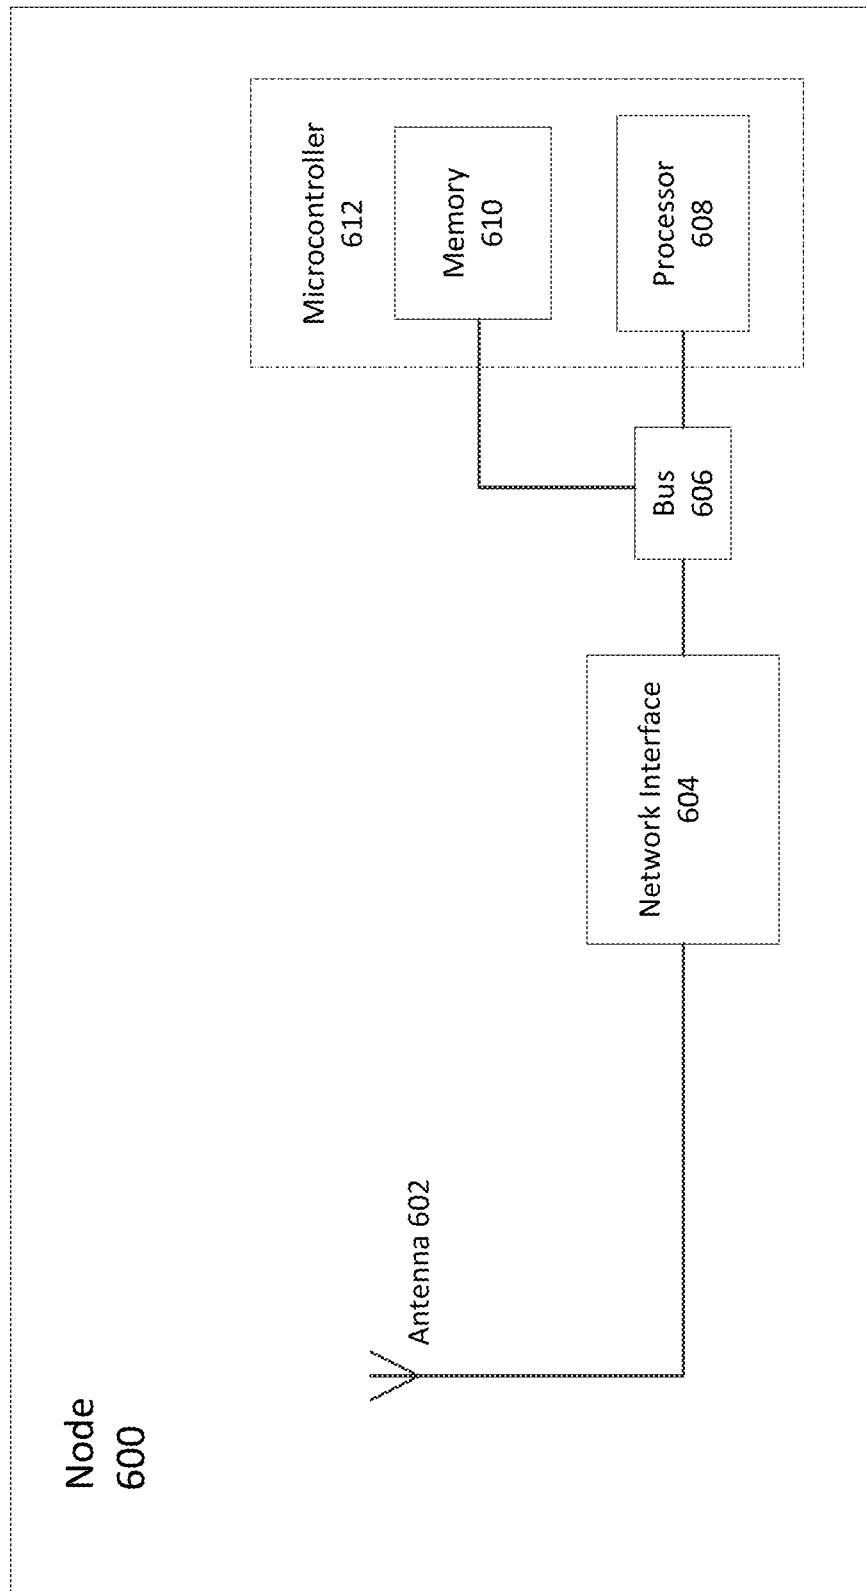


Figure 6

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/052444

A. CLASSIFICATION OF SUBJECT MATTER
 INV. H04L12/721 H04L12/753
 ADD. H04W40/02 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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	GNAWALI UNIVERSITY OF HOUSTON P LEVIS STANFORD UNIVERSITY 0: "The Minimum Rank with Hysteresis Objective Function; rfc6719.txt", THE MINIMUM RANK WITH HYSTERESIS OBJECTIVE FUNCTION; RFC6719.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARD, INTERNET SOCIETY (ISOC) 4, RUE DES FALAISES CH- 1205 GENEVA, SWITZERLAND, 11 September 2012 (2012-09-11), pages 1-13, XP015086433, [retrieved on 2012-09-11] chapter 1; page 2 - page 3 chapter 2; page 3 chapter 3; page 4 chapter 3.5.;	1-13
Y	chapter 1; page 2 - page 3 chapter 2; page 3 chapter 3; page 4 chapter 3.5.;	14-20 -/-

Further documents are listed in the continuation of Box C.

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
20 December 2016	04/01/2017
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Bos, Jürgen

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/052444

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	<p>page 8 chapter 3.3.; page 7; table 1 chapter 3.2.2.; page 6 chapter 3.4. abstract chapter 4 & WINTER T ET AL: "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks; rfc6550.txt", RPL: IPV6 ROUTING PROTOCOL FOR LOW-POWER AND LOSSY NETWORKS; RFC6550.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARD, INTERNET SOCIETY (ISOC) 4, RUE DES FALAISES CH- 1205 GENEVA, SWITZERLAND, 26 March 2012 (2012-03-26), pages 1-157, XP015081473, [retrieved on 2012-03-26] chapter 9.10; page 89 & JP VASSEUR ET AL: "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks; rfc6551.txt", ROUTING METRICS USED FOR PATH CALCULATION IN LOW-POWER AND LOSSY NETWORKS; RFC6551.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARD, INTERNET SOCIETY (ISOC) 4, RUE DES FALAISES CH- 1205 GENEVA, SWITZERLAND, 26 March 2012 (2012-03-26), pages 1-30, XP015081474, [retrieved on 2012-03-26] chapter 4.3.1.; page 19 - page 20 chapter 2.1.; page 9</p> <p>-----</p> <p>CACCIAPUOTI ANGELA SARA ET AL: "Link quality estimators for multi-hop mesh network", 2014 EURO MED TELCO CONFERENCE (EMTC), AEIT, 12 November 2014 (2014-11-12), pages 1-6, XP032714841, DOI: 10.1109/EMTC.2014.6996637 [retrieved on 2014-12-23]</p>	14-20
A	<p>chapter II.A and II.B; page 1</p> <p>-----</p>	1-13