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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
3 December 2009 (03.12.2009)

PCT

(10) International Publication Number
WO 2009/145930 A2

(51) International Patent Classification:
H04W 72/10 (2009.01)

(21) International Application Number:
PCT/US2009/031031

(22) International Filing Date:
15 January 2009 (15.01.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/030,584 13 February 2008 (13.02.2008) US

(71) Applicant *(for all designated States except US)*: **CISCO TECHNOLOGY, INC.** [US/US]; 170 West Tasman Drive, San Jose, CA 95134 (US).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: **KRUYS, Johannes, Petrus** [NL/NL]; 5 Tiendhof, NL-3481GT Hermelen (NL). **GONG, Xiaohong** [US/US]; 1352 Highland Court, Milpitas, CA 95025 (US). **RAJU, Jyoti** [IN/US]; 5953 Killarney Circle, San Jose, CA 95138 (US). **MEIL, Christopher, G.** [US/US]; 975 Alder Run Way, Akron, OH 44333 (US).

(74) Agent: **SHOWALTER, Barton, E.**; Baker Botts, L.L.P., 2001 Ross Avenue, Suite 600, Dallas, TX 75201 (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, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:
— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:
— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

WO 2009/145930 A2

(54) **Title**: COORDINATED CHANNEL CHANGE IN MESH NETWORKS

(57) **Abstract**: A coordinated channel change system. In particular implementations, a method includes receiving a prepare-to-change message, wherein the prepare-to-change message indicates instructions to prepare to change channels and includes a designated channel, and forwarding the prepare-to-change message to one or more child nodes. The method also includes receiving a ready-to-change message from the one or more child nodes, and transmitting a change-to-channel message to the one or more child nodes, wherein the change-to-channel message indicates instructions to switch to the designated channel. The method also includes receiving an acknowledgement message from the one or more child nodes, and changing to the designated channel.

Coordinated Channel Change in Mesh Networks

TECHNICAL FIELD

The present disclosure relates generally to wireless mesh networks.

BACKGROUND

Market adoption of wireless LAN (WLAN) technology has exploded, as users from a wide range of backgrounds and vertical industries have brought this technology into their homes, offices, and increasingly into the public air space. This inflection point has highlighted not only the limitations of earlier-generation systems, but also the changing role that WLAN technology now plays in people's work and lifestyles across the globe. Indeed, WLANs are rapidly changing from convenience networks to business-critical networks. Increasingly users are depending on WLANs to improve the timeliness and productivity of their communications and applications, and in doing so, require greater visibility, security, management, and performance from their networks.

Dynamic Frequency Selection (DFS) requirements give radar priority in wireless mesh network channels operating in the 5GHz range. When a given wireless mesh node detects a radar signal on a given DFS channel, the wireless mesh node is required to move off that channel. In a mesh network not all mesh nodes may see particular radar signals, and therefore any detection of a radar signal needs to be signaled throughout the mesh network. Further, rather than stopping the use of a DFS channel when a mesh node detects radar, a DFS mechanism should trigger the mesh node to change channels to one that is known to be free of radar signals. One of the main advantages of mesh networks is their ability to re-configure themselves in response to a loss of nodes, interference, changing traffic conditions, etc. Such changes may require changing the frequency of operation of one or more links in the network. In networks with a single backhaul radio, a whole sub-tree rooted at a gateway node may have to change frequency. Changing channels, however, presents a risk in terms of loss of connectivity when errors occur.

DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a wireless mesh network according to one implementation of the present invention.

Figure 2A illustrates the logical relationship between mesh access points and controller 20 relative to wireless clients, according to one possible implementation of the invention.

Figure 2B illustrates an example hardware system 100, which may be used to implement a controller 20.

Figure 3 illustrates for didactic purposes a hardware system, which may be used to implement a wireless mesh access point in a wireless mesh network.

Figure 4 is an example hierarchy of the mesh access points.

Figure 5 illustrates an example process flow associated with a bottom-up channel change.

Figure 6 illustrates an example process flow implemented at a child node and associated with a bottom-up channel change.

Figure 7 is an example hierarchy of the mesh access points.

Figure 8 illustrates an example process flow associated with a top-down channel change.

DESCRIPTION OF EXAMPLE EMBODIMENTS

A. Overview

Particular implementations facilitate channel changes in wireless mesh networks. According to particular implementations, when a given mesh node detects an event (such as a Dynamic Frequency Selection (DFS) event), the mesh node signals the root node or a controller. The root node or controller re-computes one or more channel assignments for the mesh network and assigns the new channel(s) to the appropriate mesh nodes. In one implementation, the wireless mesh achieves the channel change using a protocol where a given parent node and child node exchange messages before changing channels. While the channel assignment messages initially flow from the root node down to the leaf nodes, execution of the channel change protocol begins at the leaf nodes and then propagates up toward the root node. To ensure continuous connectivity, a number of messages may be exchanged between each parent-child pair before they switch channels.

In another implementation, after the root node selects one or more channels for the mesh network, the root node initiates a top-down channel change protocol. In a particular implementation, after the root node receives a DFS event, the root node transmits channel change messages to its child nodes multiple times before switching to a designated channel. Each child node in turn transmits the channel change message to its child nodes multiple times. This process continues down to the leaf nodes until all of the mesh nodes have switched to the designated channel. The channel change protocols described herein may be applied to various types of networks, such as mesh networks, to facilitate continuous connectivity.

B. Example Wireless Mesh Network System Architecture

B.1. Network Topology

Figure 1 illustrates a wireless mesh network according to one implementation of the present invention. In one implementation, the wireless mesh network includes a wireless mesh controller 20, a root access point 21, and a plurality of child wireless mesh access points. In one implementation, the mesh access points are logically arranged in a hierarchy for purposes of routing traffic to the root access point (RAP), and on to a network. In one implementation, this hierarchy can be dynamically configured and shifted through the exchange of wireless management messages between wireless mesh access points, or statically configured.

In one implementation, a hierarchical architectural overlay is imposed on the mesh network of routing nodes to create a downstream direction towards leaf routing nodes 35, and an upstream direction toward the root access point 21. For example, in the hierarchical mesh network illustrated in Figure 1, first hop mesh access point 31 is the parent of intermediate mesh access point 33. In addition, intermediate mesh access points 33 and 34 are the parent to leaf mesh access point 35. In one implementation, this hierarchical relationship is used in routing packets between wireless clients 60, or between wireless clients 60 and network 30. Of course, a variety of wireless mesh network configurations are possible, including non-hierarchical configurations, and hierarchical configurations with fewer or greater number of hierarchical tree structures.

The mesh access points in the mesh network, in one implementation, generally include one radio, operating in a first frequency band, and associated wireless communication functionality to communicate with other mesh access points to thereby implement the wireless backbone, as discussed more fully below. All or a subset of the mesh access points, in one implementation, also include an additional radio, operating in a second, non-interfering frequency band, and other wireless communication functionality to establish and maintain wireless connections with mobile stations, such as wireless client 60. For example, in 802.11 wireless networks, the backbone radios on the wireless routing nodes may transmit wireless packets between each other using the 802.11a protocol on the 5 GHz band, while the second radio on each mesh access point may interact with wireless clients on the 2.4 GHz band (802.11b/g). Of course, this relation can also be reversed with backhaul traffic using the 802.11b/g frequency band, and client traffic using the 802.11a band. In addition, the mesh access points may include only a single radio for client and backhaul traffic, or additional radios.

In one implementation, some wireless mesh networks can include a controller and a plurality of mesh access points that are configured into one or more routing and control hierarchies based on automatic neighbor and route discovery protocols. In some environments, individual mesh access points automatically discover their neighbors and configure hierarchical routing configurations by selecting parent nodes based on a variety of factors. Mesh access points, in some systems, connect to a wireless controller through one or more parents nodes in the routing hierarchy.

B.2. Central Controller

Figure 2A illustrates the logical relationship between mesh access points and controller 20 relative to wireless clients, according to one possible implementation of the invention. In one implementation, the mesh access points, in connection with the controller 20, implement a hierarchical processing scheme for management of wireless connections with clients 60. For example, each mesh access point may be configured to autonomously implement time-critical link layer functions (such as transmitting acknowledgements), while encapsulating and forwarding wireless management frames (e.g., association requests, etc.) and other client traffic to controller 20 for processing. The encapsulated frames may traverse one or more intermediate mesh access points in the mesh hierarchy as indicated by Figure 2A.

In other implementations, the controller 20 may be implemented as a wireless domain management server (WDMS). If the controller 20 is implemented as a WDMS, the functionality implemented by the mesh access points may comprise the full range of processing functions for wireless data frames as well wireless management frames (e.g., association requests, etc.) and other client traffic. Of course, a variety of other mesh routing and control schemes can be used in connection with the real-time transport protocol described herein.

Figure 2B illustrates an example hardware system 100, which may be used to implement a controller 20. As Figure 2B shows, in one implementation, the central controller 20 includes a network interface 102. Controller 20, in one implementation, further comprises a processor 106, a memory 108, one or more software modules stored in memory 108, including instructions for performing the functions described herein, and a system bus 110 operably connecting these components. The central control elements may optionally include an administrative port 112 allowing for administrative access for such purposes as configuration and diagnostic access.

B.3. Wireless Mesh Access Point

Figure 3 illustrates for didactic purposes a hardware system, which may be used to implement a wireless mesh access point in a wireless mesh network. In one implementation, the wireless mesh access point 300 comprises a processor 308, a read-only memory (ROM) 309, and an electronically erasable read-only memory (EEPROM) 310. The wireless mesh access point 300 may also include one or more of the following: a memory 312, a network interface 314 (e.g., an 802.3 interface) for communication with a LAN, a cache 316 for storing WLAN information, and a persistent memory 318. The wireless mesh access point 300 may also include a backhaul wireless network interface 320 having an antenna 321. Backhaul wireless network interface 320 is configured to transmit and receive messages to/from one or more other wireless mesh access points in a mesh network. The wireless mesh access point 300 may also include a client wireless network interface 322 (e.g., an IEEE 802.11 WLAN interface) having an antenna 323. Client wireless network interface 322 is configured for wireless communication with one or more wireless clients 60. The backhaul wireless network interface 320 and client wireless network interface 322 may be radio interfaces. The wireless mesh access point 300 may also include input/output (I/O) ports 324 and a system bus 330 interconnecting these components.

In some implementations, wireless mesh access point use one or more of the following standards: WiFi/802.11, WiMax/802.16, 2G, 3G, or 4G Wireless, Bluetooth/802.15, Zigbee, or any other suitable wireless communication standards. In one implementation, wireless mesh access point may have a separate access radio, and associated interface components, for communicating with a wireless client or other portable computer. The wireless mesh access points may also include software modules, including Dynamic Host Configuration Protocol (DHCP) clients, transparent bridging, Lightweight Access Point Protocol (LWAPP), Cisco® Discovery Protocol (CDP) modules, wireless access point modules, Simple Network Management Protocol (SNMP) functionality, etc., and device drivers (e.g., network and WLAN interface drivers) stored in persistent memory 318 (e.g., a hard disk drive, flash memory, EEPROM, etc.). At start up, one or more of these software components are loaded into system memory 312 and then accessed and executed by processor 310. In one implementation, the wireless mesh access point 300 includes software or firmware modules for recognizing the reception of network management information and for storing such information in memory (e.g., EEPROM 310).

C. Example Hierarchy of Mesh Access Points

Figure 4 is an example hierarchy of the mesh access points to explain operation of certain embodiments of the invention. The hierarchy includes root access point 21, mesh access point 1 31 (labeled Node 1), mesh access point 2 31 (labeled Node 2), mesh access point 3 33 (labeled Node 3), mesh access point 4 33 (labeled Node 4), and mesh access point 5 33 (labeled Node 5).

In operation, generally, when one or more of the mesh access points (e.g., mesh access point 3 33) detects an event, the mesh access point signals the root access point 21 or controller 20. In one implementation, the event may be a DFS event (e.g., detection of radar). In one implementation, DFS detection events flow upward from child nodes to the parent nodes and eventually reach the top or root node (in implementations, where controller 20 is separate from the root node). In one implementation, each parent node may apply filtering rules to received events so as to improve the detection process in terms of accuracy and reducing false alarms. For example, in one implementation, the controller 20 may initiate a channel change if two or more mesh access points detect a DFS event within a threshold window of time.

While an event may be a DFS event, the event may be any type of event such as a reconfiguration, a node failure, a command issued by a network administrator, etc. The root access point 21 then recomputes one or more channel assignments for the mesh network and assigns the new channel(s) to the mesh access points. In one implementation, all mesh access points operate on the same channel. In other implementations, the mesh access points operate on different channels. For example, a give mesh access point may include a parent radio tuned to a first channel, and a child radio tuned to a second channel. A channel change in a mesh network means that a given node may have to change either the parent channel, the child channel or both. In particular implementations, the wireless mesh network achieves a channel change using a protocol that is implemented between parent nodes and child nodes. The channel change protocols described herein may apply to both single channel and multi-channel networks. In one implementation, the channel change can be for one channel for the entire mesh network or for one or more channels of many channels.

In one implementation, the root access point 21 or controller 20 initiates the channel change by transmitting channel assignment messages that propagate down to the leaf nodes of the mesh network. In one implementation, the leaf nodes execute a channel change, which execution propagates up towards the root access point 21. The messages used to execute a change between two mesh nodes can belong to the class of adjacency protocols executed between mesh nodes. Examples of the channel change protocol are described in more detail below. For ease of illustration, the following descriptions address scenarios where one channel is being changed throughout the mesh network.

D. Channel Change Protocol

As described herein, a channel change may be initiated in response to an event such as detection of radar signals in a DFS band. As Figure 4 shows, a mesh access point 33 (Node 3) may detect an event and transmit a notification toward the root access point. Figure 4 shows an example message flow across a portion of the mesh network that may result as part of execution of a bottom-up channel change protocol.

D.1. Bottom-up Channel Change Protocol

Figure 5 illustrates an example process flow associated with a bottom-up channel change. Referring to both Figures 4 and 5, a mesh access point 31 (Node 2) receives a prepare-to-change message from the root access point 21 (502), where the prepare-to-change message includes instructions to prepare to change channels and includes a designated channel. If mesh access point 2 31 is a leaf node (504), mesh access point 2 31 transmits a ready-to-change message (506) to its parent node, where the ready-to-change message is a channel change notification that indicates that mesh access point 2 31 is ready to change channels. If mesh access point 2 31 is not a leaf node, mesh access point 2 31 forwards the prepare-to-change message to its child nodes 33 (508). As Figure 4 illustrates, the prepare-to-change message, as a result of these decisional steps, propagates down to the leaf node 33 (Node 5) (Figure 4, Ref. Nos. 1 & 2).

For ease of illustration, the remainder of this process flow is described with respect to child node 5 33. The same process flow applies to all of mesh access point's 31 child nodes.

Next, mesh access point 2 31 determines if it has received a ready-to-change message from child node 5 33 (510), where the ready-to-change message indicates that the child node 5 33 is ready to change channels (see Figure 4, Ref. No. 3). If mesh access point 2 31 does not receive a ready-to-change message and there is a time out (512), mesh access point 2 31 aborts the process (514). In one implementation, mesh access point 2 31 sends an abort message to the child node 5 33. In one implementation, if the channel change fails and mesh access point 2 31 aborts the process, the child node 5 33 may perform a default recovery procedure (e.g., switch to a bootstrap channel or scan all channels for a new parent). If mesh access point 2 31 receives a ready-to-change message, the mesh access point sends a message to child node 5 33 to switch to the designated channel (516) (Figure 4, Ref. No. 4).

Mesh access point 2 31 determines if it has received an acknowledgement message from child node 5 33 (518) (Figure 4, Ref. No. 5). If mesh access point 2 31 does not receive an acknowledgement message from child node 5 33 and there is a time out (520), mesh access point 2 31 aborts the process (514). If mesh access point 2 31 receives an acknowledgement message, mesh access point 2 31 changes to the designated channel with respect to the child node 5 33 (522).

Next, mesh access point 2 31 determines if it is successful in changing the channel (i.e., there is successful connectivity between mesh access point 2 31 and the child node 5 33) (524). After switching to the new channel, the child node 5 33 polls (probes) mesh access point 2 31 to see if it has also switched. If mesh access point 2 31 is not successful in changing the channel, mesh access point 2 31 aborts the process (514). If mesh access point 2 31 is successful, mesh access point 2 31 transmits a ready-to-change message to its parent (526) (Figure 4, Ref. No. 6), which initiates a similar process and message flow. Mesh access point 31 then follows a channel changing process similar to child node 5 33.

Figure 6 illustrates an example process flow implemented at a child node and associated with a bottom-up channel change. As Figure 6 shows, child node 5 33 receives a message from its parent node 33 to change to a designated channel (602). Child node 5 33 then transmits a response to the parent node 33 (604). The child node 5 33 then changes to the designated channel (606). In one embodiment, after changing to the designated channel, child node 5 33 sends a probe to the parent. Child node 5 33 then determines if the parent node has responded to the probe (608). If not, child node 5 33 initiates a failure recover process (610). Various failure recovery processes are possible. For example, a failure may cause both parent and child nodes to revert back to the old channel. If that fails, the child node may revert to a back-up configuration (e.g., a bootstrap channel) or do an exhaustive channel scan to identify potential parent nodes. If the parent has responded to the probe, child node 5 33 continues with its normal processes.

D.2. Top-Down Channel Change

Figure 7 illustrates an example message flow in an example hierarchy of the mesh access points according to a top-down channel change protocol. In one embodiment, when one or more of the mesh access points (e.g., mesh access point 4 33) detects an event such as a DFS event, the mesh access point signals the root access point 21. As described in more detail below in connection with Figure 8, after the root access point 21 recomputes one or more channel assignments for the mesh network, the root access point 21 implements a top-down channel change mechanism.

Figure 8 illustrates an example process flow associated with a top-down channel change. Referring to both Figures 7 and 8, mesh access point 2 31 receives a prepare-to-change message from the root access point 21 multiple times (802). If mesh access point 2 31 is a leaf node (804), mesh access point 2 31 changes channels (806). If mesh access point 2 31 is not a leaf node, mesh access point 2 31 transmits a channel change message to its child nodes multiple times (808). In one implementation, mesh access point 2 31 transmits a channel change for a predefined number of times (e.g., 5 times) in predefined intervals (e.g., a beacon interval). Mesh access point 2 31 then changes the channel to a designated channel (810). In one implementation, if the channel change fails, child node 5 33 may perform a recovery procedure (e.g., default to a suitable default channel or perform a neighbor scan to find a parent).

The present invention has been explained with reference to specific embodiments. For example, while embodiments of the present invention have been described as operating in connection with IEEE 802.11 networks, the present invention can be used in connection with any suitable wireless network environment. Other embodiments will be evident to those of ordinary skill in the art. It is therefore not intended that the present invention be limited, except as indicated by the appended claims.

CLAIMS

What is claimed is:

1. A method comprising:

receiving a prepare-to-change message, wherein the prepare-to-change message indicates instructions to prepare to change channels and includes a designated channel;

forwarding the prepare-to-change message to one or more child nodes;

receiving a ready-to-change message from the one or more child nodes;

transmitting a change-to-channel message to the one or more child nodes, wherein the change-to-channel message indicates instructions to switch to the designated channel;

receiving an acknowledgement message from the one or more child nodes; and

changing to the designated channel.

2. The method of claim 1 further comprising:

determining if there is successful connectivity with the one or more child nodes; and

if there is successful connectivity, transmitting a ready-to-change message to a parent node.

3. The method of claim 2 further comprising:

receiving a change-to-channel message from the parent node, wherein the change-to-channel message indicates instructions to switch to the designated channel;

transmitting a response to the parent node; and

changing to the designated channel.

4. The method of claim 3 further comprising:

determining if the parent node has responded to a probe; and

if the parent node has not responded to the probe, initiating a failure recovery process.

5. The method of claim 4 wherein the failure recover process comprises performing a neighbor scan to find a parent.

6. The method of claim 1 further comprising detecting an event and transmitting an event notification to a root node.

7. An apparatus comprising:

one or more processors;

a memory;

one or more wireless network interfaces; and

logic encoded in one or more tangible media for execution and when executed operable to cause the one or more processors to:

receive a prepare-to-change message, wherein the prepare-to-change message indicates instructions to prepare to change channels and includes a designated channel;

forward the prepare-to-change message to one or more child nodes;

receive a ready-to-change message from the one or more child nodes;

transmit a change-to-channel message to the one or more child nodes, wherein the change-to-channel message indicates instructions to switch to the designated channel;

receive an acknowledgement message from the one or more child nodes; and

change to the designated channel.

8. The apparatus of claim 7 wherein the logic is further operable to cause the one or more processors to:

determine if there is successful connectivity with the one or more child nodes; and

if there is successful connectivity, transmit a ready-to-change message to a parent node.

9. The apparatus of claim 7 wherein the logic is further operable to cause the one or more processors to:

receive a change-to-channel message from the parent node, wherein the change-to-channel message indicates instructions to switch to the designated channel;

transmit a response to the parent node; and

change to the designated channel.

10. The apparatus of claim 7 wherein the logic is further operable to cause the one or more processors to:

determine if the parent node has responded to a probe; and

if the parent node has not responded to the probe, initiate a failure recovery process.

11. The apparatus of claim 10 wherein the failure recover process comprises performing a neighbor scan to find a parent.

12. The apparatus of claim 7 wherein the logic is further operable to cause the one or more processors to detect an event and transmitting an event notification to a root node.

13. A method comprising:

receiving a channel change message from a parent node;

transmitting the change message to one or more child nodes a plurality of times; and

switching to the designated channel.

14. The method of claim 13 further comprising entering a failure recovery mode when a loss of connectivity to a parent is detected.

15. The method of claim 14 wherein the failure recover mode comprises performing a neighbor scan to find a parent.

16. Logic encoded in one or more tangible media for execution and when executed operable to:

receive a channel change message from a parent node;

transmit the change message to one or more child nodes a plurality of times; and

switch to the designated channel.

17. The logic of claim 16 further comprising entering a failure recovery mode when a loss of connectivity to a parent is detected.

18. The logic of claim 17 wherein the failure recover mode comprises performing a neighbor scan to find a parent.

19. An apparatus comprising:

one or more processors;

a memory;

one or more wireless network interfaces; and

logic encoded in one or more tangible media for execution and when executed operable to cause the one or more processors to:

receive a channel change message from a parent node;

transmit the change message to one or more child nodes a plurality of times; and

switch to the designated channel.

20. The apparatus of claim 19 wherein the logic is further operable to cause the one or more processors to enter a failure recovery mode when a loss of connectivity to a parent is detected.

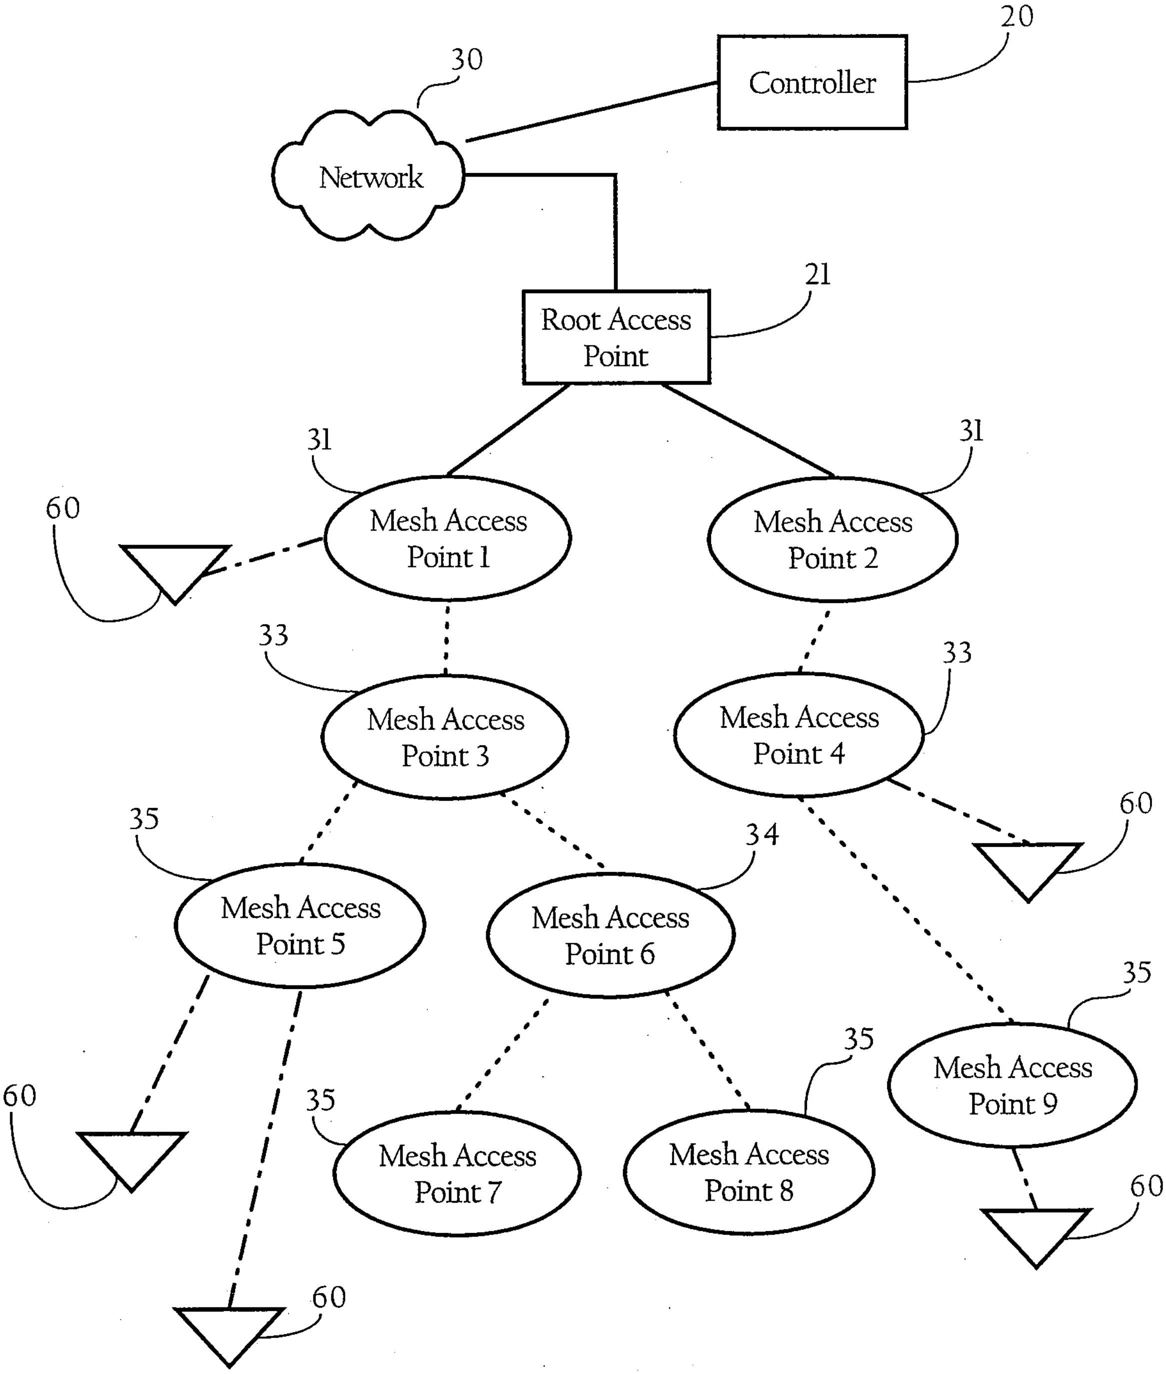
20
Controller
30
Network
21
Root Access Point
31
Mesh Access Point 1
31
Mesh Access Point 2
60
33
Mesh Access Point 3
33
Mesh Access Point 4
60
35
Mesh Access Point 5
34
Mesh Access Point 6
35
Mesh Access Point 9
35
35
Mesh Access Point 7
Mesh Access Point 8
60
60
60
60

Fig._1

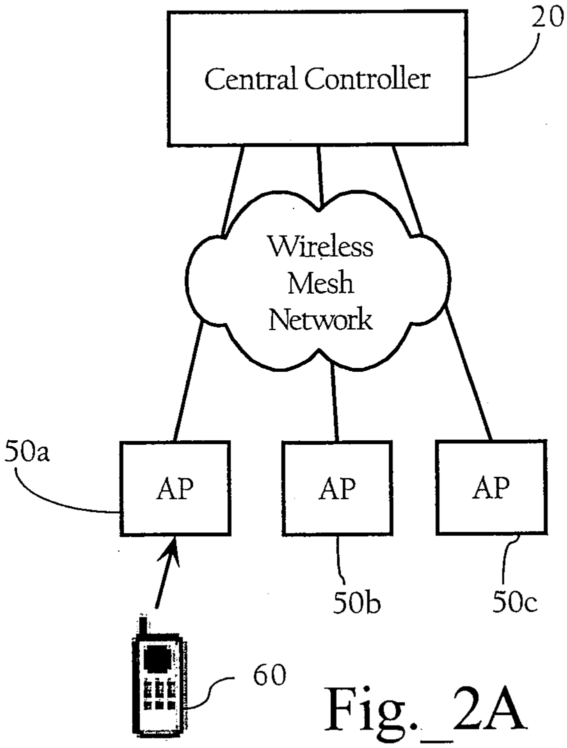
20
Central Controller
Wireless
Mesh
Network
50a
AP
AP
AP
50b
50c
60

Fig._2A

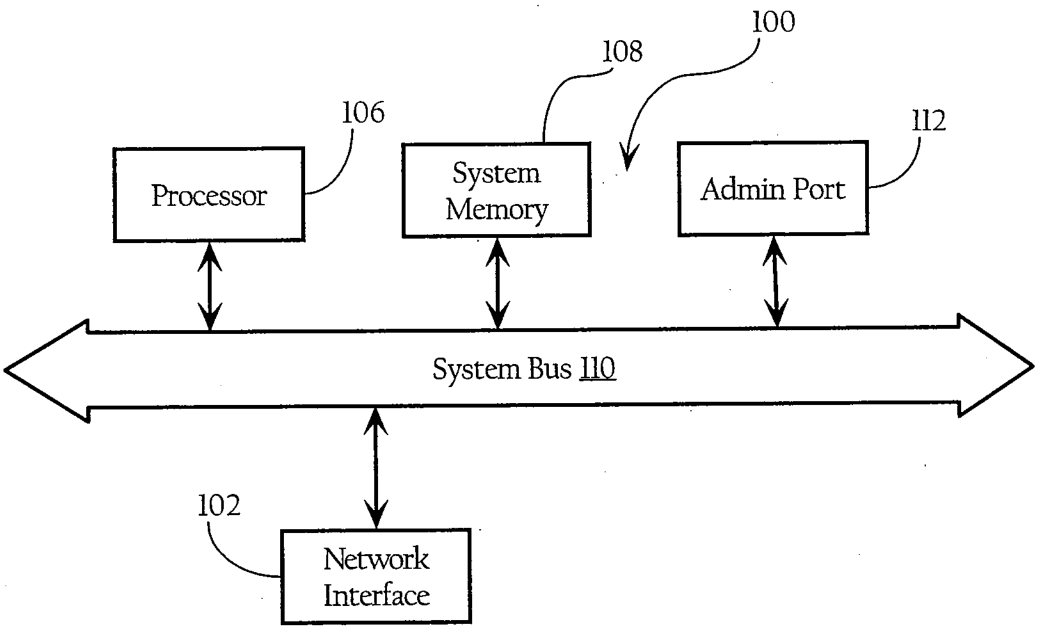
100
108
106
112
Processor
System
Memory
Admin Port
System Bus 110
102
Network
Interface

Fig._2B

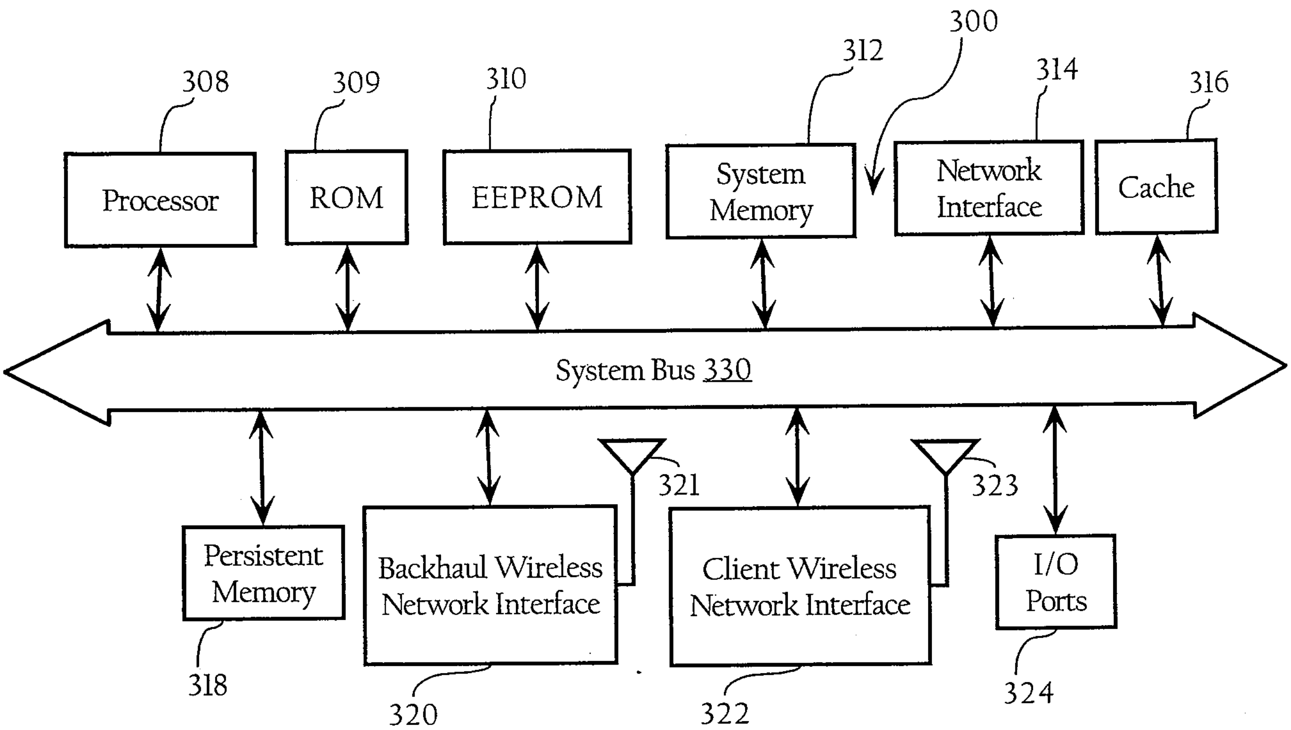
300
308
309
310
312
314
316
Processor
ROM
EEPROM
System Memory
Network Interface
Cache
System Bus 330
321
323
Persistent Memory
Backhaul Wireless Network Interface
Client Wireless Network Interface
I/O Ports
318
320
322
324

Fig._3

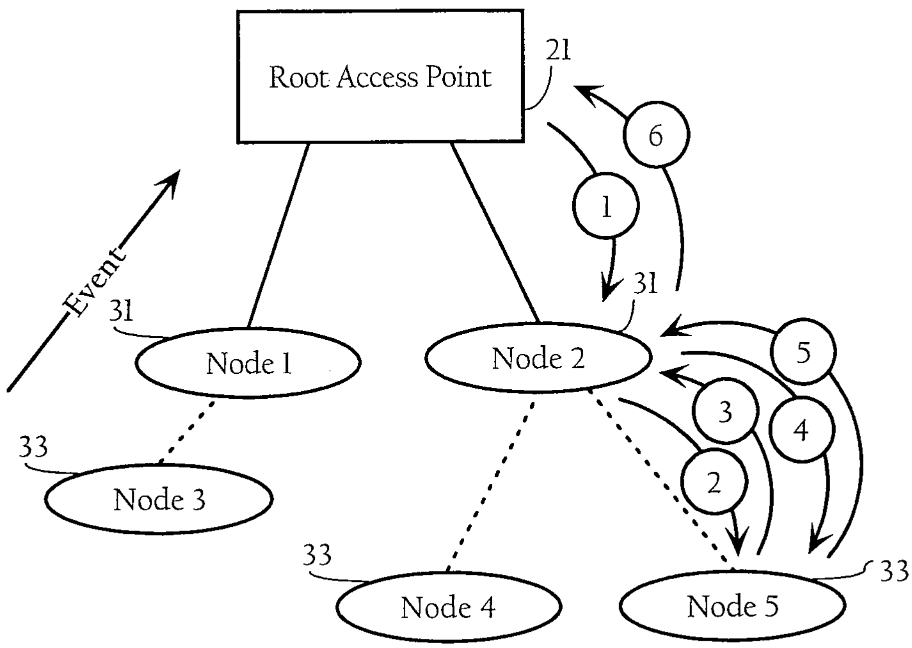
Root Access Point
21
6
1
31
Event
31
Node 1
Node 2
5
3
4
33
2
Node 3
33
33
Node 4
Node 5

Fig._4

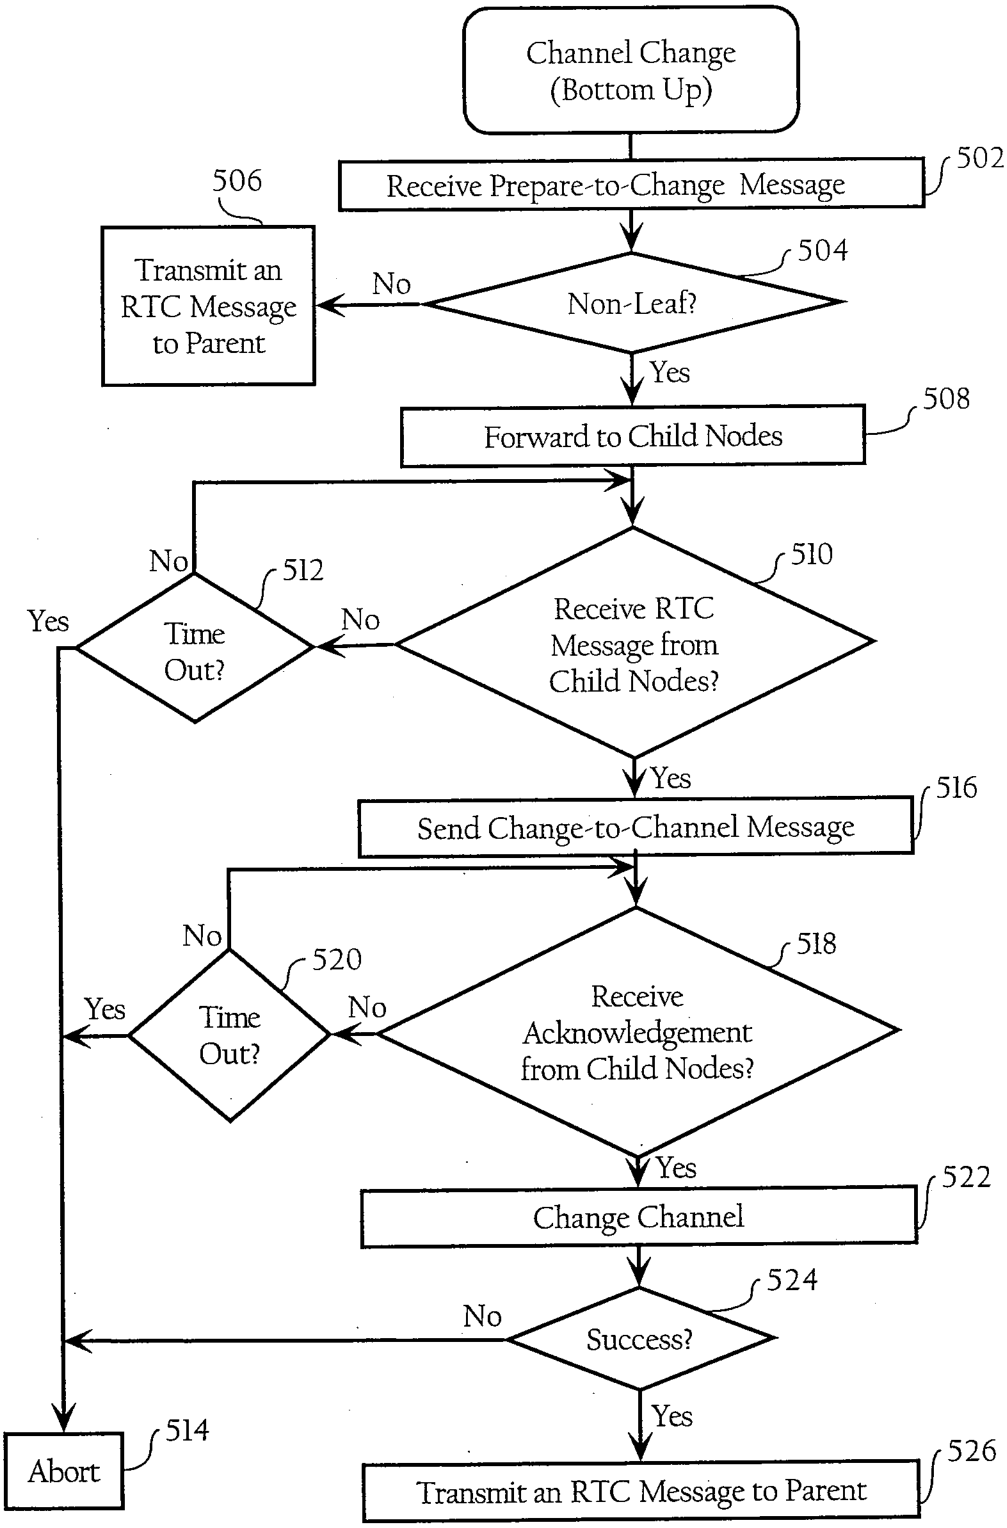
Channel Change
(Bottom Up)
Receive Prepare-to-Change Message
502
506
Transmit an RTC Message to Parent
No
Non-Leaf?
504
Yes
Forward to Child Nodes
508
No
512
Yes
Time Out?
No
Receive RTC Message from Child Nodes?
510
Yes
Send Change-to-Channel Message
516
No
520
Yes
Time Out?
No
Receive Acknowledgement from Child Nodes?
518
Yes
Change Channel
522
No
Success?
524
Yes
Abort
514
Transmit an RTC Message to Parent
526

Fig._5

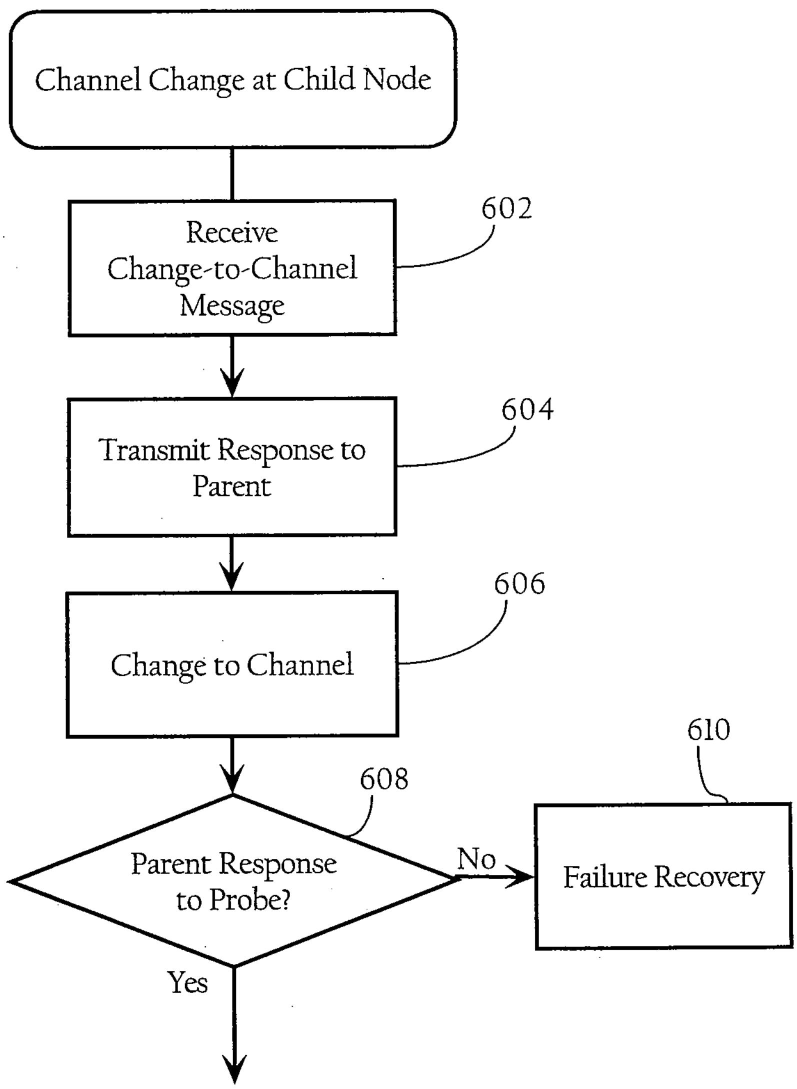
Channel Change at Child Node
Receive Change-to-Channel Message
602
Transmit Response to Parent
604
Change to Channel
606
Parent Response to Probe?
608
No
Yes
Failure Recovery
610

Fig._6

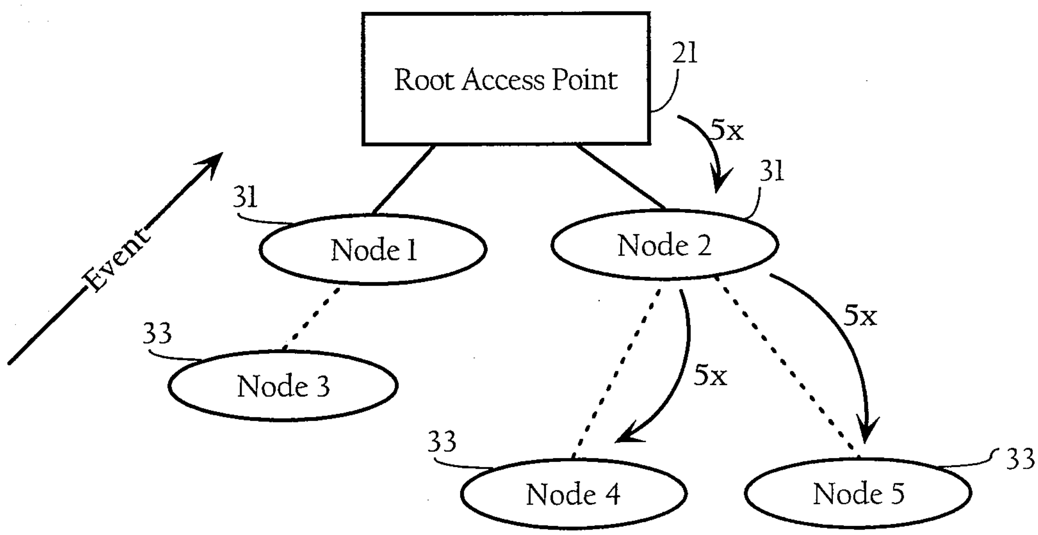
Root Access Point
21
5x
31
31
Event
Node 1
Node 2
33
5x
Node 3
5x
33
33
Node 4
Node 5

Fig._7

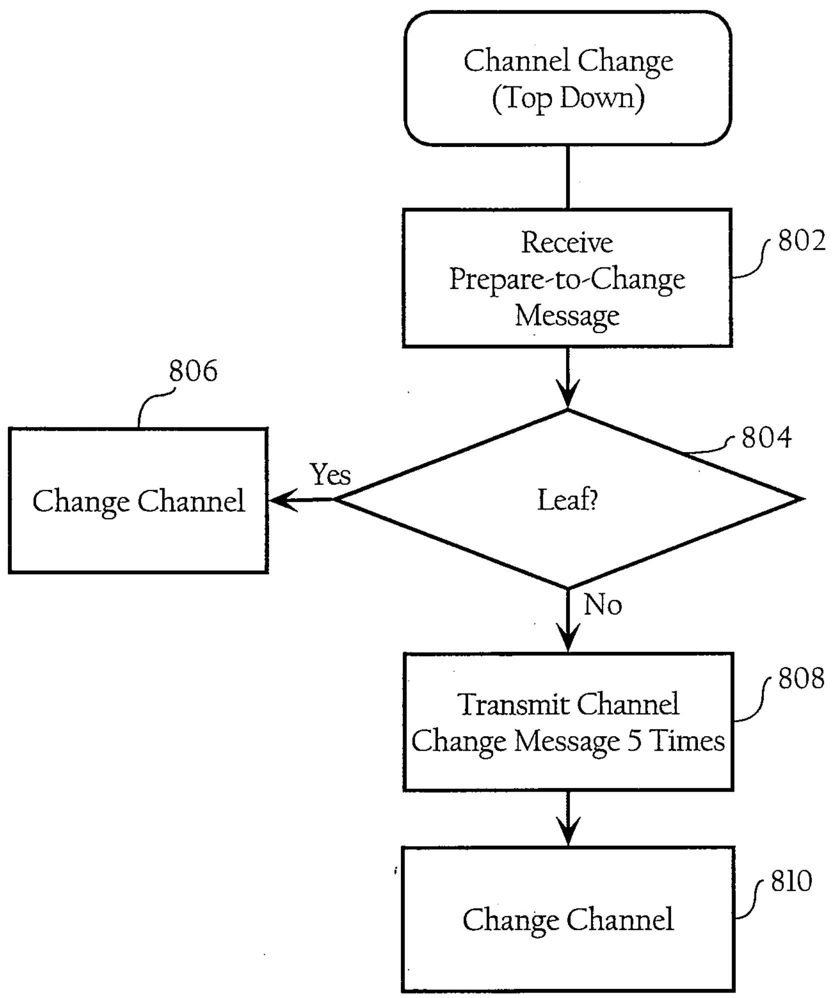
Channel Change
(Top Down)
Receive
Prepare-to-Change
Message
802
804
Leaf?
Yes
806
Change Channel
No
Transmit Channel
Change Message 5 Times
808
Change Channel
810

Fig._8