

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
19 April 2007 (19.04.2007)

PCT

(10) International Publication Number
WO 2007/044985 A2

(51) International Patent Classification:
H04Q 7/00 (2006.01)

(21) International Application Number:
PCT/US2006/040499

(22) International Filing Date: 13 October 2006 (13.10.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/727,025 13 October 2005 (13.10.2005) US
11/331,789 14 January 2006 (14.01.2006) US

(71) Applicant (*for all designated States except US*):
TRAPEZE NETWORKS, INC. [US/US]; 5753 W.
Las Positas Blvd., Pleasanton, CA 94588 (US).

(72) Inventor: **MATTA, Sudheer.**

(74) Agent: **AHMANN, William, F.**; Perkins Coie LLP, 101
Jefferson Drive, Menlo Park, CA 94025 (US).

(81) Designated States (*unless otherwise indicated, for every
kind of national protection available*): AE, AG, AL, AM,

AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN,
CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI,
GB, GD, GE, GH, GM, HN, HR, HU, ID, IL, IN, IS, JP,
KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT,
LU, LV, LY, MA, MD, 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, 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, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,
RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished
upon receipt of that report*

*For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.*

(54) Title: SYSTEM AND METHOD FOR WIRELESS NETWORK MONITORING

(57) Abstract: A technique for wireless network monitoring involves scanning channels using clients instead of access points. An example of a method according to the technique may include, for example, receiving from a wireless access point a command to perform a channel scanning function, listening on a channel associated with the channel scanning function, and sending RP data found on the channel to the wireless access point. Another example of a method according to the technique may include, for example, scanning a first channel, switching from the first channel to a second channel, sending data on the second channel to an access point, switching from the second channel to the first channel, and resuming scanning on the first channel. A system according to the technique may include one or more scanning clients, proxy clients, multi-channel clients, or other clients that are capable of scanning channels in lieu of an access point.



WO 2007/044985 A2

SYSTEM AND METHOD FOR WIRELESS NETWORK MONITORING

BACKGROUND

Frame requests proved a way to sense frames transmitted on air on a channel. This can be valuable in a wireless network because it introduces a way to understand who is communicating on a channel, including learning, for example, Received Signal Strength Indication (RSSI) or other values associated with a client.

The 802.11k standard introduces Frame Request in 802.11k D3.0, section 7.3.2.21.7. However, the standard does not help to identify or locate a particular rogue device. For example, with the 802.11k, as it is proposed today, a device cannot query all trusted stations to go and look for a particular rogue device. In addition, the full frame report may or may not include a specific device, based upon the length of the report and what else might be happening on the channel. Moreover, the querying station has to digest a large report, when its needs may be for a single device. The standard also is relatively useless at identifying any disassociation or deauthentication storms or at identifying any CTS storm blocking a particular channel.

These are but a subset of the problems that may exist with the 802.11k standard, as it is proposed today, that are intended to characterize weaknesses in the prior art by way of example. The foregoing examples of the related art and limitations related therewith are intended to be illustrative and not exclusive. Other limitations of the related art will become apparent to those of skill in the art upon a reading of the specification and a study of the drawings.

SUMMARY

The following embodiments and aspects thereof are described and illustrated in conjunction with systems, tools, and methods that are meant to be exemplary and illustrative, not limiting in scope. In various embodiments, one or more of the above-described problems have been reduced or eliminated, while other embodiments are directed to other improvements.

A technique for wireless network monitoring involves scanning channels using clients instead of access points. An example of a method according to the technique may include, for example, receiving from a wireless access point a command to perform a channel scanning function, listening on a channel associated with the channel scanning function, and sending RF data found on the channel to the wireless access point. The command may or may not be sent from a switch to the wireless access point. In an embodiment, data may or may not be filtered

using one or more of channel number, mac address, frame type, and frame subtype. Reports generated from the data may be used to determine what countermeasures to use in response.

Another example of a method according to the technique may include, for example, scanning a first channel, switching from the first channel to a second channel, sending data on the second channel to an access point, switching from the second channel to the first channel, and resuming scanning on the first channel. In this way, continuous or nearly continuous scanning of a channel may be possible. Or, all channels could be scanned in this way.

Another example of a method according to the technique may include, for example, sending a command from an access point on a first channel to a proxy client, switching the proxy client to a second channel, forwarding the command from the proxy client to a client on the second channel, generating a report that is responsive to the command, switching the proxy client to the first channel, and forwarding the report from the proxy client to the access point.

A system according to the technique may include one or more scanning clients, proxy clients, multi-channel clients, or other clients that are capable of scanning channels in lieu of an access point.

The proposed system can offer, among other advantages, clients that can scan channels in lieu of an access point scanning the channels. These and other advantages of the present invention will become apparent to those skilled in the art upon a reading of the following descriptions and a study of the several figures of the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the invention are illustrated in the figures. However, the embodiments and figures are illustrative rather than limiting; they provide examples of the invention.

FIG. 1 depicts a system including a wireless access domain.

FIG. 2 depicts a computer system for use in the system of FIG. 1.

FIG. 3 depicts an access area with multiple scanning clients.

FIGS. 4-8 depict flowcharts of examples of methods for wireless network monitoring.

FIG. 9 depicts a conceptual diagram of a report.

FIG. 10 depicts a conceptual diagram of an example of a frame request frame.

FIG. 11 depicts a conceptual diagram of an example of a frame report frame.

In the following description, several specific details are presented to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that the invention can be practiced without one or more of the specific details, or in combination with other components, etc. In other instances, well-known implementations or operations are not shown or described in detail to avoid obscuring aspects of various embodiments, of the invention.

FIG. 1 depicts a system 100 including a wireless access domain. The system 100 includes a computer system 102, a network 104, and a wireless access domain 106. The system 100 may or may not include multiple wireless access domains. The computer system 102 may be practically any type of device that is capable of communicating with a communications network, such as, by way of example but not limitation, a workstation. The network 104 may be practically any type of communications network, such as, by way of example but not limitation, the Internet. The term "Internet" as used herein refers to a network of networks which uses certain protocols, such as the TCP/IP protocol, and possibly other protocols such as the hypertext transfer protocol (HTTP) for hypertext markup language (HTML) documents that make up the World Wide Web (the web). The physical connections of the Internet and the protocols and communication procedures of the Internet are well known to those of skill in the art.

In a non-limiting embodiment, the computer system 102 may be running a program such as, by way of example but not limitation, ethereal, to decode, by way of example but not limitation, IEEE 802.11 standard packets encapsulated in TZSP that are received from the wireless access domain 106. In a non-limiting embodiment, the computer system 102 is connected to a wireless backbone network (not shown), either directly or indirectly through a wireless network.

In a non-limiting embodiment, the network 104 provides a Layer 2 path for Layer 3 traffic, preserving IP addresses, sessions, and other wired Layer 3 attributes as users roam throughout the wireless access domain 106. The network may or may not include a wireless backbone network, or be connected directly or indirectly to a wireless backbone network. Communications between the computer system 102 and the wireless access domain 106 are, therefore, Layer 3 traffic tunneled through Layer 2. Advantageously, by tunneling Layer 3 traffic at Layer 2, users stay connected with the same IP address and keep the same security and Quality of Service (QoS) policies from the wired network while they roam the wireless side. Since Layer 3 attributes are maintained, mobile devices that are connected to the wireless access domain 106 can retain persistent identities.

The seven layers of the Open System Interconnection (OSI) model, of which Layers 2 and 3 are a part, are well-known to those of skill in the relevant art, and are, therefore, not described herein in any substantial detail. It should be noted, however, that Layer 3 is known as the "Network Layer" because it provides switching and routing technologies, creating logical paths, known as virtual circuits, for transmitting data from node to node. Routing and forwarding are functions of this layer, as well as addressing, internetworking, error handling, congestion control and packet sequencing. Layer 2 is known as the "Data Link Layer" because at Layer 2 data packets are encoded and decoded into bits; and Layer 2 furnishes transmission protocol knowledge and management and handles errors in the physical layer, flow control and frame synchronization. The data link layer is divided into two sublayers: The Media Access Control (MAC) layer and the Logical Link Control (LLC) layer. The MAC sublayer controls how a computer on the network gains access to the data and permission to transmit it. The LLC layer controls frame synchronization, flow control, and error checking.

In non-limiting embodiments, the wireless access domain 106 may be referred to as, by way of example but not limitation, a Local Area Network (LAN), virtual LAN (VLAN), and/or wireless LAN (WLAN). The wireless access domain 106 gives each user a persistent identity that can be tracked and managed, no matter where they roam. The wireless access domain 106 may have one or more associated snoop filters. In an embodiment, the wireless access domain 106 may include one or more radios.

In the example of FIG. 1, the wireless access domain 106 includes access areas 108-1 to 108-N (hereinafter collectively referred to as access areas 108). The access areas 108 have characteristics that depend upon, among other things, a radio profile. A radio profile is a group of parameters such as, by way of example but not limitation, beacon interval, fragmentation threshold, and security policies. In an embodiment, the parameters may be configurable in common across a set of radios in one or more access areas 108. In another embodiment, a few parameters, such as the radio name and channel number, must be set separately for each radio. An example of the implementation of a wireless access domain, provided by way of example but not limitation, includes a Trapeze Networks "identity-aware" Mobility Domain™.

In the example of FIG. 1, the following elements are associated with each of the access areas 108: Wireless exchange switches 110-1 to 110-N (hereinafter collectively referred to as wireless exchange switches 110), networks 112-1 to 112-N (hereinafter collectively referred to as networks 112), and access points 114-1 to 114-N (hereinafter collectively referred to as access points 114).

In an embodiment, the wireless exchange switches 110 swap topology data and client information that details each user's identity, location, authentication state, VLAN membership, permissions, roaming history, bandwidth consumption, and/or other attributes assigned by, by way of example but not limitation, an Authentication, Authorization, and Accounting (AAA) backend (not shown). In an embodiment, the wireless exchange switches 110 provide forwarding, queuing, tunneling, and/or some security services for the information the wireless exchange switches 110 receive from their associated access points 114. In another embodiment, the wireless exchange switches 110 coordinate, provide power to, and/or manage the configuration of the associated access points 114. An implementation of a wireless exchange switch, provided by way of example but not limitation, includes a Trapeze Networks Mobility Exchange™ switch. The Trapeze Networks Mobility Exchange™ switches may, in another implementation, be coordinated by means of the Trapeze Access Point Access (TAPA) protocol.

In an embodiment, the networks 112 are simply wired connections from the wireless exchange switches 110 to the access points 114. The networks 112 may or may not be part of a larger network. In a non-limiting embodiment, the networks 112 provides a Layer 2 path for Layer 3 traffic, preserving IP addresses, sessions, and other wired Layer 3 attributes as users roam throughout the wireless access domain 106. Advantageously, by tunneling Layer 3 traffic at Layer 2, users stay connected with the same IP address and keep the same security and Quality of Service (QoS) policies from the wired network while they roam the wireless side.

In a non-limiting embodiment, the access points 114 are hardware units that act as a communication hub by linking wireless mobile 802.11 stations such as PCs to a wired backbone network. In an embodiment, the access points 114 connect users to other users within the network and, in another embodiment, can serve as the point of interconnection between a WLAN and a fixed wire network. The number of users and size of a network help to determine how many access points are desirable for a given implementation. An implementation of an access point, provided by way of example but not limitation, includes a Trapeze Networks Mobility System™ Mobility Point™ (MP™) access point.

The access points 114 are stations that transmit and receive data (and may therefore be referred to as transceivers) using one or more radio transmitters. For example, an access point may have two associated radios, one which is configured for IEEE 802.11a standard transmissions, and the other which is configured for IEEE 802.11b standard transmissions. In a non-limiting embodiment, an access point transmits and receives information as radio frequency (RF) signals to and from a wireless client over a 10/100BASE-T Ethernet connection. The

access points 114 transmit and receive information to and from their associated wireless exchange switches 110. Connection to a second wireless exchange switch provides redundancy.

A station, as used herein, may be referred to as a device with a media access control (MAC) address and a physical layer (PHY) interface to the wireless medium that comply with the IEEE 802.11 standard. As such, in a non-limiting embodiment, the access points 114 are stations. Similarly, the wireless client 116 may be implemented as a station. In alternative embodiments, a station may comply with a different standard than IEEE 802.11, and may have different interfaces to a wireless or other medium.

In operation, a wireless client 116 can roam from one of the access areas 108 to another of the access areas 108. For example, in the example of FIG. 1 the wireless client 116 moves from the access area 108-1 to the access area 108-N. In an embodiment, the wireless client 116 can maintain a single IP address and associated data sessions. The ability of the wireless client 116 to roam across the access areas 108 while maintaining a single IP address and associated data sessions may be referred to as subnet mobility. Advantageously, the system 100 may be implemented using identity-based networking, which is a technique that enforces network authorization attributes to the wireless client 116 based on client identity rather than the port or device through which the wireless client 116 connects to the network. This technique enables both a single persistent login and passport free roaming which permits the introduction of services such as voice to a wireless LAN.

FIG. 2 depicts a computer system 200 for use in the system 100 (FIG. 1). The computer system 200 may be a conventional computer system that can be used as a client computer system, such as a wireless client or a workstation, or a server computer system. The computer system 200 includes a computer 202, I/O devices 204, and a display device 206. The computer 202 includes a processor 208, a communications interface 210, memory 212, display controller 214, non-volatile storage 216, and I/O controller 218. The computer 202 may be coupled to or include the I/O devices 204 and display device 206.

The computer 202 interfaces to external systems through the communications interface 210, which may include a modem or network interface. It will be appreciated that the communications interface 210 can be considered to be part of the computer system 200 or a part of the computer 202. The communications interface 210 can be an analog modem, ISDN modem, cable modem, token ring interface, satellite transmission interface (e.g. "direct PC"), or other interfaces for coupling a computer system to other computer systems.

The processor 208 may be, for example, a conventional microprocessor such as an Intel Pentium microprocessor or Motorola power PC microprocessor. The memory 212 is coupled to

the processor 208 by a bus 220. The memory 212 can be Dynamic Random Access Memory (DRAM) and can also include Static RAM (SRAM). The bus 220 couples the processor 208 to the memory 212, also to the non-volatile storage 216, to the display controller 214, and to the I/O controller 218.

5 The I/O devices 204 can include a keyboard, disk drives, printers, a scanner, and other input and output devices, including a mouse or other pointing device. The display controller 214 may control in the conventional manner a display on the display device 206, which can be, for example, a cathode ray tube (CRT) or liquid crystal display (LCD). The display controller 214 and the I/O controller 218 can be implemented with conventional well known technology.

10 The non-volatile storage 216 is often a magnetic hard disk, an optical disk, or another form of storage for large amounts of data. Some of this data is often written, by a direct memory access process, into memory 212 during execution of software in the computer 202. One of skill in the art will immediately recognize that the terms "machine-readable medium" or "computer-readable medium" includes any type of storage device that is accessible by the processor 208 and
15 also encompasses a carrier wave that encodes a data signal.

 The computer system 200 is one example of many possible computer systems which have different architectures. For example, personal computers based on an Intel microprocessor often have multiple buses, one of which can be an I/O bus for the peripherals and one that directly connects the processor 208 and the memory 212 (often referred to as a memory bus).
20 The buses are connected together through bridge components that perform any necessary translation due to differing bus protocols.

 Network computers are another type of computer system that can be used in conjunction with the teachings provided herein. Network computers do not usually include a hard disk or other mass storage, and the executable programs are loaded from a network
25 connection into the memory 212 for execution by the processor 208. A Web TV system, which is known in the art, is also considered to be a computer system, but it may lack some of the features shown in FIG. 2, such as certain input or output devices. A typical computer system will usually include at least a processor, memory, and a bus coupling the memory to the processor.

30 In addition, the computer system 200 is controlled by operating system software which includes a file management system, such as a disk operating system, which is part of the operating system software. One example of operating system software with its associated file management system software is the family of operating systems known as Windows® from Microsoft Corporation of Redmond, Washington, and their associated file management systems.

Another example of operating system software with its associated file management system software is the Linux operating system and its associated file management system. The file management system is typically stored in the non-volatile storage 216 and causes the processor 208 to execute the various acts required by the operating system to input and output data and to store data in memory, including storing files on the non-volatile storage 216.

Some portions of the detailed description are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

The present invention, in some embodiments, also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language, and various embodiments may thus be implemented using a variety of programming languages.

FIG. 3 depicts an access area 300 with multiple scanning clients. In the example of FIG. 3, the access area includes a switch 302, Access Points (APs) 304-1 to 304-N (hereinafter collectively referred to as APs 304), clients 306, and (for illustrative purposes) a rogue device 308. Advantageously, using multiple clients for scanning channels provides significantly better location detection than would be achieved using a single AP. When multiple clients, for example, look for a particular mac address, it becomes easier to triangulate thanks to the increased number of points of reference. The enhanced location detection can be useful not only for detecting and locating a rogue device 308, but also for interfering clients, or just to acquire information about where a client is.

Advantageously, by finding out which clients are on a particular set of channels, it may be possible to do load balancing. For example, if more clients are on one channel than another, the clients could be instructed to change channels.

Rogues may try to spoof an AP. In an embodiment, when the clients 306 look for a rogue channel, they also know the signal strength, which is passed to APs 304 and to the switch 302. The switch 302 knows where the APs 304 are (all of which may be passing RSSI information). There will be a large discrepancy between signal strengths of the APs 304 and the rogue device 308.

FIG. 4 depicts a flowchart 400 of a method for wireless network monitoring. This method and other methods are depicted as serially arranged modules. However, modules of the methods may be reordered, or arranged for parallel execution as appropriate. FIG. 4 is intended to illustrate a first example of operation of a system such as that depicted in FIG. 3, using the techniques described herein. The flowchart 400 continues from a determination that a trigger has occurred to the conclusion of activity that is a result of the trigger. When describing the flowchart 400, for illustrative purposes only, reference is made to components depicted in FIG. 3. However, it should be noted that the method of FIG. 4 is not intended to be limited to the components depicted in FIG. 3, and may be applicable to other systems and configurations.

In the example of FIG. 4, the flowchart 400 starts at decision point 402 where it is determined whether a trigger has occurred. In the context of FIG. 4, a trigger is any event that stimulates the system to take the steps necessary to generate a report regarding a rogue device, a client, a channel, or some other aspect or component of a wireless network. The trigger may include, but is not limited to a random trigger, an unscheduled trigger, or a scheduled trigger. The trigger may or may not be in response to analysis of prior reports, instructions from an administrator, messages from another switch, an apparent network outage, messages from a client, messages from an AP, or practically any other stimulus that is deemed to be potentially indicative of a problem, or useful in helping to detect a problem, in a wireless network. The commands should be appropriate for the trigger. For example, in response to an apparent network outage on a particular channel, a switch may send a command to listen for beacons on that channel.

In the example of FIG. 4, if it is determined that a trigger has not occurred (402-N), then the flowchart 400 repeats decision point 402 until a trigger does occur. It may be noted that the loop need not be an active check for a triggering condition; rather the loop could simply represent the time lag until a trigger actually occurs. The stimulus for a trigger may originate within a switch, or higher up.

In the example of FIG. 4, when it is determined that a trigger has occurred (402-Y), the flowchart 400 continues at module 404 with the switch 302 sending to one or more of the APs 304 a command to find out more about a device, such as the rogue device 308 or one or more of the clients 306, or channel. The command could be sent to a single AP or multiple APs. When sending to multiple APs, the command may be sent to each AP for a given switch, all of the APs of a mobile domain, or some subset of the APs of a mobile domain. In some implementations, it may even be desirable to send the command to APs of multiple mobile domains.

In the example of FIG. 4, the flowchart 400 continues at module 406 with sending the command to clients associated with and/or within range of the AP. The clients may include all clients 306, or a subset of clients 306 within the access area of the relevant APs 304. The command may include, for example, instructions for the client to go to a particular channel and listen for a specified time. Each of the clients may be given a slightly different command. For example, a first client may be instructed to go to channel 1, and a second client may be instructed to go to a channel 2. In a specific embodiment, an AP may have multiple dedicated clients. In this specific embodiment, it may be unnecessary to send the command to clients other than the dedicated clients, since the dedicated clients perform all of the scanning routines. Of course,

other clients could serve a backup function, such as when the dedicated clients are all busy or otherwise unavailable.

In the example of FIG. 4, the flowchart 400 continues at module 408 with one or more clients 306 going to a channel specified in the command and listening (or scanning) for a period of time specified in the command. In alternative embodiments, the command need not specify a channel. By way of example but not limitation, each client may be pre-configured to scan a particular channel, obviating the need for a command to specify the channel. In alternative embodiments, the command need not specify a period of time. By way of example but not limitation, each client (or all of the clients 306) may be pre-configured to scan for a set amount of time, obviating the need for a command to specify a period of time.

In the example of FIG. 4, the flowchart 400 continues at module 410 with the clients going back to the AP's channel to send raw data. APs typically operate on a single channel, so it may be assumed that clients will also operate on a single channel, which would have the benefit of more inexpensive client devices. Thus, the clients must typically switch between channels to respectively scan and to send data (assuming they are scanning a channel other than the one on which the AP is operating). In alternative embodiments, the clients may operate on multiple channels. In such a case, it may not be necessary for the client to switch from a scanned channel to the AP's channel; rather, the client may scan on a channel, and send on another channel without switching.

In order to reduce the amount of data sent from the clients 306 to the APs 304, it may be desirable to give the clients 306 the ability to generate reports from the raw data collected. Alternatively, the clients 306 could be given the ability to pre-process some of the raw data.

In the example of FIG. 4, the flowchart 400 continues at module 412 with the AP reporting to the switch 302. The APs 304 may be given the task of analyzing raw data from the clients 306, depending upon the devices and/or the implementation. For example, the APs 304 could generate a report similar to that depicted in FIG. 9 from the raw data received from the clients 306. One of skill in the relevant art would recognize that such a report can be used to diagnose certain problems, and perform countermeasures in response if desired or available.

Alternatively, the APs 304 could be configured to perform some pre-processing of raw data and allow the switch 302 (or a device that is higher up) generate the desired reports. Or, the APs 304 could simply pass the raw data on to the switch 302 (or a device that is higher up) for processing. This is an implementation decision. When pre-processing data at the APs 304, it may be desirable to remove some raw data if it can be omitted. It is believed that clients 306, when properly configured, can obtain as much information for the reports as an AP, so the AP

should not be required to add additional data when pre-processing the raw data. Nevertheless, if it is determined that the AP can add data to the reports, then the APs 304 may be so configured.

In the example of FIG. 4, the flowchart 400 ends at module 414 with the switch performing countermeasures. The switch 302 may or may not analyze the reports. For example, the analysis could be done at the AP or, in the alternative, higher up. The switch 302 (or some other device) may decide upon and set into action countermeasures to respond to the rogue device 308, if the rogue device 308 is detected. The countermeasures could be in response to some other stimulus than the detection of the rogue device 308, as well.

Countermeasure techniques are known to those of skill in the relevant arts, and generally include measures that attack rogue devices within an access area, shore up the defenses of the network against such devices, or improve the network in some other way in response to a network problem. However, in the example of FIG. 4, these measures are intended to include any procedure that is a response to an analyzed report, including but not limited to sending a notification to an administrator, shutting down an access point, rebooting, or gathering additional information. As the reports include raw data, the problems detected may include anything that can be detected using the raw data. The number of problems or attacks that can be detected is practically impossible to list in its entirety, but includes and is not limited to spoofed network detection, CTS storms, rogue devices, etc.

Advantageously, these techniques obviate the need for the APs to scan any of the channels. It may be desirable for APs to scan channels if client scanning is not deemed sufficient, but the need for the APs to scan channels is at least reduced. Consequentially, service disruptions at the AP can be reduced.

FIG. 5 depicts a flowchart 500 of a method for wireless network monitoring. FIG. 5 is intended to illustrate a second example of operation of a system such as that depicted in FIG. 3, using the techniques described herein, such as client sensors. In the example of FIG. 5, it is assumed that the client sensors are dedicated wireless sensors, rather than active clients spending time on channels to perform scanning. This is not required, but since dedicated wireless sensors may be considered extremely useful, considering their relatively low cost compared to an AP, and relatively simple deployment within a wireless network, it is appropriate to illustrate this advantage with reference to at least one flowchart—in this case FIG. 5. When describing the flowchart 500, for illustrative purposes only, reference is made to components depicted in FIG. 3. However, it should be noted that the method of FIG. 5 is not intended to be limited to the components depicted in FIG. 3, and may be applicable to other systems and configurations.

In the example of FIG. 5, the flowchart 500 starts at module 502 with using the client sensors 306 to scan a plurality of channels. For example, 802.11b may operate on 11 channels (in the United States, at least). So, 11 of the clients 306 may operate on the 11 channels (or a smaller number of clients 306 may operate on a smaller number of channels). Advantageously, this allows each of the client sensors 306 to generate nearly continuous reports. It should be noted that, in this example, a trigger may not be necessary since the channels can be scanned continuously and, if a threat is detected, countermeasures can be executed. In some cases, it may be desirable to include a trigger as well.

It should be noted that promiscuous listening might encounter thousands of nodes for a given period of time. Since large reports increase network usage, it may be desirable to provide more accurate detection so that reports can be smaller. Therefore, continuous reports may not be a desirable outcome. The balance between continuous, complete reports and deterministic rogue detection (or, for example, security detection, location detection, etc.) is an implementation decision that may be impacted one way or another depending upon the embodiment or configuration that is used. Advantageously, the techniques provided herein allow for either a broader (e.g., complete) report than ever before in 802.11, or for a more focused report than ever before in 802.11.

In the example of FIG. 5, the flowchart 500 continues at module 504 with a client sensor switches to the AP's channel. If a client sensor is scanning on the AP's channel then, of course, the client sensor need not change channels at this step. In other cases, assuming that the client sensor does not have a mechanism for scanning on one channel and transmitting on another, the client sensor may need to change to the AP's channel prior to communicating with the AP.

In the example of FIG. 5, the flowchart 500 continues at module 506 with the client sensor sends raw data to the AP. There is no technical reason why the client sensors 306 would be unable to scan for any RF activity that is within range. So, the amount of raw data can theoretically include all 802.11 RF activity in an embodiment in which 802.11 RF activity is the activity of interest. (In embodiments that do not use 802.11, other RF activity may be included.) In an embodiment, all of the raw data is sent from the client sensor to the AP.

In another embodiment, the raw data may be pre-processed to reduce the amount of data that is sent to the AP. In such an embodiment, the amount of data sent would depend upon the desired implementation of the system. In another embodiment, the client sensor may actually generate a report for sending to the AP. In such an embodiment, the AP (or switch, or other component that is higher up) may not need to perform any report generating functions. Indeed,

the client sensors could even be programmed with countermeasure procedures, which could obviate the need for sending a report at all.

In the example of FIG. 5, the flowchart 500 continues at module 508 with the client sensor switches to another channel, and module 510 with the client sensor resumes scanning, and at module 504, which was described previously. (In the example of FIG. 5, the flowchart 500 also continues from module 506 to module 512, which is described later.) In an embodiment, at module 508, the client sensor switches back to the channel that it was scanning previously.

With appropriate configurations, every channel can be scanned continuously. A minimalist example (in which there are, for illustrative purposes, 11 channels) of such a configuration would include an AP and 11 clients. The AP scans its own channel. Periodically, each of the 10 clients switches to the AP's channel to send a report. When the clients switch to the AP's channel to send a report, they cannot continue to scan the previously scanned channel. However, when (or just before) the clients switch to the AP's channel, the 11th client switches to the previously scanned channel so that scanning can be essentially continuous. As each of the 10 clients, in turn, finish their reports, the 11th client switches back to the AP's channel to deliver the report for the previously scanned channel. This can result in continuous scanning of all channels using one client per channel.

Referring once again to FIG. 5, in an alternative embodiment, at module 508 the client sensor may switch to a channel that it was not immediately previously scanning. For example, it may be desirable to have client sensors rotate through channels so that scanning is continuous. If raw data proves to be too much to handle with only 12 (or 11) client sensors, the number could be increased. For example, 21 client sensors could be used to scan 11 channels, two for each channel, except the AP's channel, which may only require one client sensor because the client sensor does not need to change channels to send data. When one of the client sensors on a given channel is sending data to the AP, the other is scanning. When the reporting client sensor switches back, the scanning client sensor can switch to the AP's channel to deliver its data.

In the example of FIG. 5, the flowchart 500 continues at module 512 with the access point generates a report using the raw data, at module 514 with the access point reports to the switch, and at module 516 with the switch performs countermeasures.

Even with continuous scanning of channels, it may be desirable to implement additional redundancy. For example, it may be desirable to deploy, e.g., four clients (or client sensors) at four edges of an access area, and perhaps one client near the AP. This may facilitate detection of relatively distant rogue devices more quickly. It may also be desirable to place a client (or client sensor) approximately equidistant from two adjacent APs, and have the client report to both APs.

FIG. 6 depicts a flowchart 600 of a method for wireless network monitoring. FIG. 6 is intended to illustrate a third example of operation of a system such as that depicted in FIG. 3, in which the amount of data that is sent from the clients 306 to the APs 304 is intelligently reduced. In the example of FIG. 6, it is assumed that the clients 306 have report generating capabilities. This is not required, and this capability could be replaced with data pre-processing capabilities. When describing the flowchart 600, for illustrative purposes only, reference is made to components depicted in FIG. 3. However, it should be noted that the method of FIG. 6 is not intended to be limited to the components depicted in FIG. 3, and may be applicable to other systems and configurations.

In the example of FIG. 6, the flowchart 600 starts at decision point 602 where it is determined whether a trigger has occurred. In the example of FIG. 6, if it is determined that a trigger has not occurred (602-N), then the flowchart 600 repeats decision point 602 until a trigger does occur. When a trigger occurs (602-Y), the flowchart 600 continues at module 604 with the switch sending to one or more of the access points 304 a command.

In the example of FIG. 6, the flowchart 600 continues at module 606 with the AP sending the command to clients 306 associated with the AP. The clients to which the AP sends the command may vary depending upon implementation and/or the command itself. For example, if the command is a directive to listen on a particular channel for a given period of time, then only clients that are configured to listen on that channel would respond to the command. In an embodiment, only those clients that are configured to respond to a given command receive the command. Other commands may include listening for a particular mac address. In such a case, it may be desirable for all of the clients 306 to receive the command.

In the example of FIG. 6, the flowchart 600 continues at module 608 with the clients generating a report in accordance with the command. The amount of raw data that a client retains would depend primarily upon implementation. Very large storage capacities have the benefit of providing better reports in some cases, but at greater expense. Large storage capacities will also tend to increase physical size, and may or may not increase access times as well. Of course, generating reports from a large pool of data also takes longer. One of skill in the relevant art should be able to weigh the advantages and disadvantages to come up with a reasonable compromise between storage capacity and expense (or, e.g., access time).

In the example of FIG. 6, the flowchart 600 continues at module 610 with the client switches to the AP's channel and reports, at module 612 with the APs report to the switch, and at module 614 with the switch performs countermeasures.

In an embodiment, the APs have to switch channels in order to send commands to the various clients (or client sensors). This can result in a service disruption similar to that caused by active scan. To remedy this problem, it may be desirable to implement a proxy client.

FIG. 7 depicts a flowchart 700 of a method for providing commands to client sensors.

5 FIG. 7 is intended to illustrate a fourth example of operation of a system such as that depicted in FIG. 3, in which APs can transmit commands to client sensors without a service disruption.

In the example of FIG. 7, the flowchart 700 starts at module 702 with a switch sending a command to an AP. This may be, for example, in response to a trigger (not shown).

10 In the example of FIG. 7, the flowchart 700 continues at module 704 with an AP sending the command to a proxy client. The proxy client may be a client sensor that is on the same channel as the AP. The proxy client may or may not also be responsible for scanning on the AP's channel.

15 In the example of FIG. 7, the flowchart 700 continues at module 706 with the proxy client switching to another channel. The channel to which the proxy client switches may or may not be dependent upon the command. For example, if the command is to scan a particular channel, the proxy client may switch to that channel. As another example, if the command is to scan for a particular type of frame on any channel, then the proxy client may switch to each other channel in turn (or other proxy clients could be used to switch to the channels more quickly).
20 Alternatively, the proxy client could just switch to each channel in turn regardless of the command, and the APs may or may not respond to the command depending upon the command's relevance to them.

In the example of FIG. 7, the flowchart 700 continues at module 708 with the proxy client sending the command to a client. Since the proxy client is a "proxy", the proxy client by definition passes the command on to a client. It should be noted, however, that a client could
25 wait on the AP's channel until given a command, then switch to an appropriate channel or channels to carry out the command.

In the example of FIG. 7, the flowchart 700 continues at module 710 with the client generating a report. Since, in an embodiment, the client does not leave the channel, the client can amass continuous RF data on the channel. This data can be used to generate reports without
30 any additional scanning. For example, if a command is sent periodically to check whether any suspicious activity is occurring on a channel, recent data may be as valuable as data that is acquired shortly after the command is received. So, the client could use a slightly dated report, and perhaps respond slightly faster to the command.

In the example of FIG. 7, the flowchart 700 continues at module 712 with the client sending the report to the proxy client and at module 714 with the proxy client switching to the AP's channel. In an alternative embodiment, the client and the proxy could switch places. For instance, the client could generate a report and switch to the AP's channel, and the proxy remain on the channel and begin scanning. An advantage of this embodiment is that the report need not be transmitted from client to client.

In the example of FIG. 7, the flowchart 700 continues at module 716 with the proxy client reporting to the AP, and at module 718 with the AP reporting to the switch. It may be noted that since the client generates the report, the AP only needs to forward the report on to the switch. Of course, the AP could perform additional processing, depending upon the implementation and/or configuration of the AP.

FIG. 8 depicts a flowchart 800 of a method for providing commands to multi-channel clients. FIG. 8 is intended to illustrate a fifth example of operation of a system such as that depicted in FIG. 3, in which continuous scanning of channels is possible without switching between channels.

In the example of FIG. 8, the flowchart 800 begins at module 802 with a switch sending a command to an AP and at module 804 with the AP sending the command to a multi-channel client. A multi-channel client may include a two or more sensors that can be set to different channels. The multi-channel client may further include a link between the two or more sensors. The link allows the multi-channel client to receive data on one channel, and transmit the data to another channel.

Alternatively or in addition, the multi-channel client can receive data on one channel, process the data into, for example, a report, and transmit the processed data to another channel. Alternatively or in addition, the multi-channel client can receive data on multiple channels, process the data into, for example, a report, and transmit the processed data to another channel (or one of the channels on which the data was received). In an embodiment, the multi-channel client could include one sensor for each channel. Thus, by way of example but not limitation, a multi-channel client with 14 channels could operate on all 14 channels of the 802.11b standard in Japan (in the US, 802.11b operates on only 11 channels at 2.4 GHz).

In the example of FIG. 8, the flowchart 800 continues at module 806 with the multi-channel client generating a report for the relevant channel. For example, if the command requires listening on all channels, the report could include data from each of the channels (in an embodiment or implementation wherein the multi-channel client includes a sensor for each

channel). A command that only asks for a report on a single channel may only need data from that single channel.

In the example of FIG. 8, the flowchart 800 continues at module 808 with the multi-channel client reporting to the AP. It may be noted that it was not necessary for the multi-channel client to switch between channels. Accordingly, the multi-channel client can continue to scan on, potentially, all channels while reporting to the AP.

In the example of FIG. 8, the flowchart 800 continues at module 810 with the AP reporting to the switch and at module 812 with the switch performing countermeasures.

FIG. 9 depicts a conceptual diagram of a report. In the example of FIG. 9, the report includes zero or more mac addresses and type, subtype, and number of frames associated with the mac addresses. FIG. 9 is intended to illustrate a report that can be generated using data acquired by clients. Those of skill in the relevant art would recognize that this report is similar to, for example, a report generated at an AP after the AP performed an active scan. Advantageously, using the techniques described herein, the report can be generated using clients and without active scan (and the associated service disruption).

FIG. 10 depicts a conceptual diagram of an example of a frame request frame 1000. In the example of FIG. 10, the frame request includes several fields, most of which would be self-explanatory to one of skill in the relevant art. The fields are, therefore, only briefly described, with the exception of field 1010, for which additional explanation is provided. It should be noted that FIG. 10 is intended to illustrate one of many possible frame requests. Other frame request embodiments could easily be created and used without deviating from the teachings provided herein.

In an embodiment, the frame request frame 1000 is part of a measurement request report mechanism. The frame request frame 1000 can be sent to clients in a wireless network. Clients go to the channel indicated in the Channel Number field. The clients then get frames and send data back, such as mac address and an associated RSSI, (plus, if desired, average and last RSSI).

The frame request frame 1000 facilitates looking for frames of a specific type or subtype. Requesting frames of a particular type allows recognition of security attacks, such as, by way of example but not limitation, CTS attack. You can ask stations to look for disassociate frames or disassociate frames from a particular address. Advantageously, in an embodiment, the mac address and frame type/subtype choices are distinct, thereby providing more flexibility in detecting, e.g., CTS storms, spoofed networks, and other attacks. For example, the AP could

send a command to a client to scan a channel for a particular mac address and a particular type of frame, allowing for relatively specific reports. In general, it helps to know what kinds of frames, for example, interfering clients are using.

Field 1002 is Channel Number, which indicates the channel number for which the measurement request applies. Channel Number is defined within a Regulatory Class.

Field 1004 is Regulatory Class, which indicates the frequency band for which the measurement request applies.

Field 1006 is Randomization Interval, which specifies the upper bound of the random delay to be used prior to making the measurement in units of TU.

Field 1008 is Measurement Duration, which is set to the preferred duration of the requested measurement, expressed in TUs. If the Duration Mandatory bit is set to 1 in the Measurement Request Mode field this is interpreted as a mandatory measurement duration. If the Duration Mandatory bit is set to 0 this shall be interpreted as a target measurement duration.

Field 1010 is Frame Request Parameter Set, which is split, for illustrative purposes, into six subfields, Match Type bit, Match Subtype bit, Match Mac bit, Type, Subtype, and Reserved.

The Match Type bit indicates whether the type fields indicated in the Frame Request Parameter Set 1010 should match for the frames counted for frame report generation. If the bit is set to 1, only frames that match the type should be counted towards the generation of a frame report.

The Match Subtype bit is only valid when the match type bit is set to 1. This bit indicates whether the subtype field indicated in the Frame Request Parameter Set 1010 should match for the frames counted for frame report generation. If the bit is valid and is set to 1, only frames that match the subtype should be counted towards the generation of a frame report.

The Match Mac bit indicates whether the mac address included in the frame request frame 1000 should match for the frames counted for the frame report generation. If the bit is set to 1, only frames that match the Mac Address field 1012 should be counted towards the generation of a frame report. When this bit set to 1, the mac address field 1012 is mandatory.

The Type field is used to indicate the type of packets that would be counted towards the frame report generation. The Type field is only used when the Match Type bit is set to 1.

The Subtype field is used to indicate the subtype of packets that would be counted towards the frame report generation. The Subtype field is only used when the Match Subtype bit is valid and is set to 1.

The Reserved field is set to zeros on transmit and should be ignored on receive.

Field 1012 is Mac Address, which is included in the frame request 1000 if the match mac address field in the frame request parameter set is set to 1. If this field is included, only frames from this mac address are counted towards the frame report generated in response to this frame request.

FIG. 11 depicts a conceptual diagram of an example of a frame report frame 1100. In the example of FIG. 11, the frame report frame 1100 includes several fields, most of which would be self-explanatory to one of skill in the relevant art. The fields are, therefore, only briefly described, with the exception of field 1110, for which additional explanation is provided. It should be noted that FIG. 11 is intended to illustrate one of many possible frame report frames. Other frame report embodiments could easily be created and used without deviating from the teachings provided herein.

Field 1102 is Channel Number, which indicates the channel number for which the measurement request applies. Channel Number is defined within a Regulatory Class.

Field 1104 is Regulatory Class, which indicates the frequency band for which the measurement request applies.

Field 1106 is Actual Measurement Start Time, which is set to the value of the measuring STA's TSF timer at the time the measurement started.

Field 1108 is Measurement Duration, which is set equal to the duration over which the Frame Report was measured, expressed in TUs.

Field 1110 is Frame Report Entry, which includes the fields Transmit Address, BSSID, Phy Type (Phy), Average RCPI (Avg), RSNI Last (RL), RCPI, Antenna ID (Ant ID), and Frame Report Parameter Set. The Transmit Address field contains the Transmit Address from the frames being reported. The BSSID field contains the BSSID from the frames being reported. PHY Type indicates the physical medium type for the frame(s) being reported. Valid entries are coded according to the value of dot11PHYType. Average RCPI indicates the average value for the received channel power of all the frames counted for this report. Average RCPI is reported in dBm, as defined in the RCPI measurement clause for the PHY Type. RSNI indicates the received signal to noise indication of the received frame in dBm. This field is the RSNI value for the most recently received frame. Last RCPI indicates the received channel power of the most recently counted frame in this Frame Report entry. Last RCPI is reported in dBm, as defined in the RCPI measurement clause for the PHY Type. The Antenna ID field contains the

identifying number for the antenna used to receive the most recently counted frame in this Frame Report entry.

The Frame Report Parameter has four subfields: Type, Subtype, Reserved, and Number of Frames. The Type field indicates the type of the packets counted towards this frame report parameter set. The Subtype field is used to indicate the subtype of packets that were counted towards this frame report parameter set. Number of Frames is a count of the frames of the type and subtype mentioned in the corresponding frame report parameter set with the indicated Transmit Address and BSSID during the measurement duration. The value 255 indicates a count of 255 or more.

In many of the examples provided herein, multicast packets are assumed to be 802.11-compatible. 802.11-compatible is intended to mean the multicast packet can be sent in accordance with, by way of example but not limitation, 802.11a, 802.11b, 802.11g, or other current or future 802.11 standards. It is to be understood that other wireless implementations other than 802.11 will likely have problems that can be reduced using the techniques described herein. Therefore, although the 802.11 standard is ubiquitous, the teachings provided herein are not limited to the 802.11 standard.

Frame report request methodology of, for example, 802.11k can be improved with techniques described herein. The techniques should be applicable to other 802.11 standards and to wireless network techniques in general.

As used herein, a rogue device is a wireless Access-Point or Client device that is violating policies or hampering legal wireless access in a network. The rogue device is often assumed to be harmful for an enterprise network. An interfering device, on the other hand, is a wireless Access-Point or Client device that is coexisting with a legal enterprise network without causing any intentional damage to it. A known device is a wireless Access-Point or Client device that is part of the legal enterprise network wireless installation. A client sensor is a wireless client device that scans 802.11 (or other) channels and provides information to an AP or set of APs.

As used herein, a wireless network refers to any type of wireless network, including but not limited to a structured network or an ad hoc network. Data on a wireless network is often encrypted. However, data may also be sent in the clear, if desired. With encrypted data, a rogue device will have a very difficult time learning any information (such as passwords, etc.) from clients before countermeasures are taken to deal with the rogue. The rogue may be able to confuse the client, and perhaps obtain some encrypted data, but the risk is minimal (even less than for some wired networks).

Active scan involves sending an AP to other channels (for a short time). One well-known problem with active scanning is service disruption. In addition to the short period of time where an AP is not on its primary channel, the switching can cause annoying problems in other respects. For example, when voice is sent over a wireless channel, active scan can cause an audible click when an AP scans another channel.

As used herein, the term "embodiment" means an embodiment that serves to illustrate by way of example but not limitation.

It will be appreciated to those skilled in the art that the preceding examples and embodiments are exemplary and not limiting to the scope of the present invention. It is intended that all permutations, enhancements, equivalents, and improvements thereto that are apparent to those skilled in the art upon a reading of the specification and a study of the drawings are included within the true spirit and scope of the present invention. It is therefore intended that the following appended claims include all such modifications, permutations and equivalents as fall within the true spirit and scope of the present invention.

What is claimed is:

1. A method comprising:

receiving from a wireless access point a command to perform a channel scanning
5 function;

listening on a channel associated with the channel scanning function;

sending RF data found on the channel to the wireless access point.

2. The method of claim 1, further comprising sending the command from a switch to the
wireless access point.

3. The method of claim 1, further comprising sending the command from the wireless
10 access point to a wireless client associated with the wireless access point.

4. The method of claim 1, further comprising listening on the channel for a period of time
that is specified in the command.

5. The method of claim 1, further comprising filtering data using one or more of the
15 categories consisting of channel number, mac address, frame type, and frame subtype.

6. The method of claim 1, further comprising generating a report from the RF data.

7. The method of claim 1, further comprising generating a report from the RF data, wherein
the sending the RF data found on the channel includes sending the report.

8. The method of claim 1, further comprising performing countermeasures in response to
20 the RF data.

9. The method of claim 1, wherein the RF data is 802.11 RF data.

10. A method comprising:

scanning a first channel of a plurality of wireless channels;

switching from the first channel to a second channel of the plurality of wireless channels,
wherein an access point is on the second channel;

5 sending data on the second channel to the access point;

switching from the second channel to the first channel;

resuming scanning on the first channel.

11. The method of claim 10, further comprising scanning each of the plurality of wireless channels.

10 12. The method of claim 10, further comprising receiving a command from the access point, wherein the data sent to the access point is in accordance with the command.

13. The method of claim 10, wherein

switching from the first channel to the second channel includes a first client switching
from the first channel to the second channel;

15 switching from the second channel to the first channel includes a second client switching
from the second channel to the first channel.

14. The method of claim 10, further comprising generating a report using the data, wherein
the sending the data on the second channel includes sending the report.

15. The method of claim 10, further comprising:

20 generating a report using the data;

performing countermeasures in response to the report.

16. The method of claim 10, further comprising using a proxy client to forward a command
from the access point to a client, and to provide the data to the access point from the client.

17. A method comprising:

5 sending a command from an access point on a first channel to a proxy client;
 switching the proxy client to a second channel;
 forwarding the command from the proxy client to a client on the second channel;
 generating a report that is responsive to the command;
 switching the proxy client to the first channel;
 forwarding the report from the proxy client to the access point.

18. The method of claim 17, wherein generating the report includes generating the report at the client, further comprising transmitting the report from the client to the proxy client.

10 19. The method of claim 17, further comprising, after forwarding the command from the proxy client to the client, swapping the roles of the client and the proxy client such that the proxy client becomes the client, and the client becomes the proxy client, wherein, after the swapping of roles, the generating a report includes generating the report at the proxy client.

20. The method of claim 17, wherein the report includes a frame report frame.




FIG. 1

200 →




FIG. 2

300 →

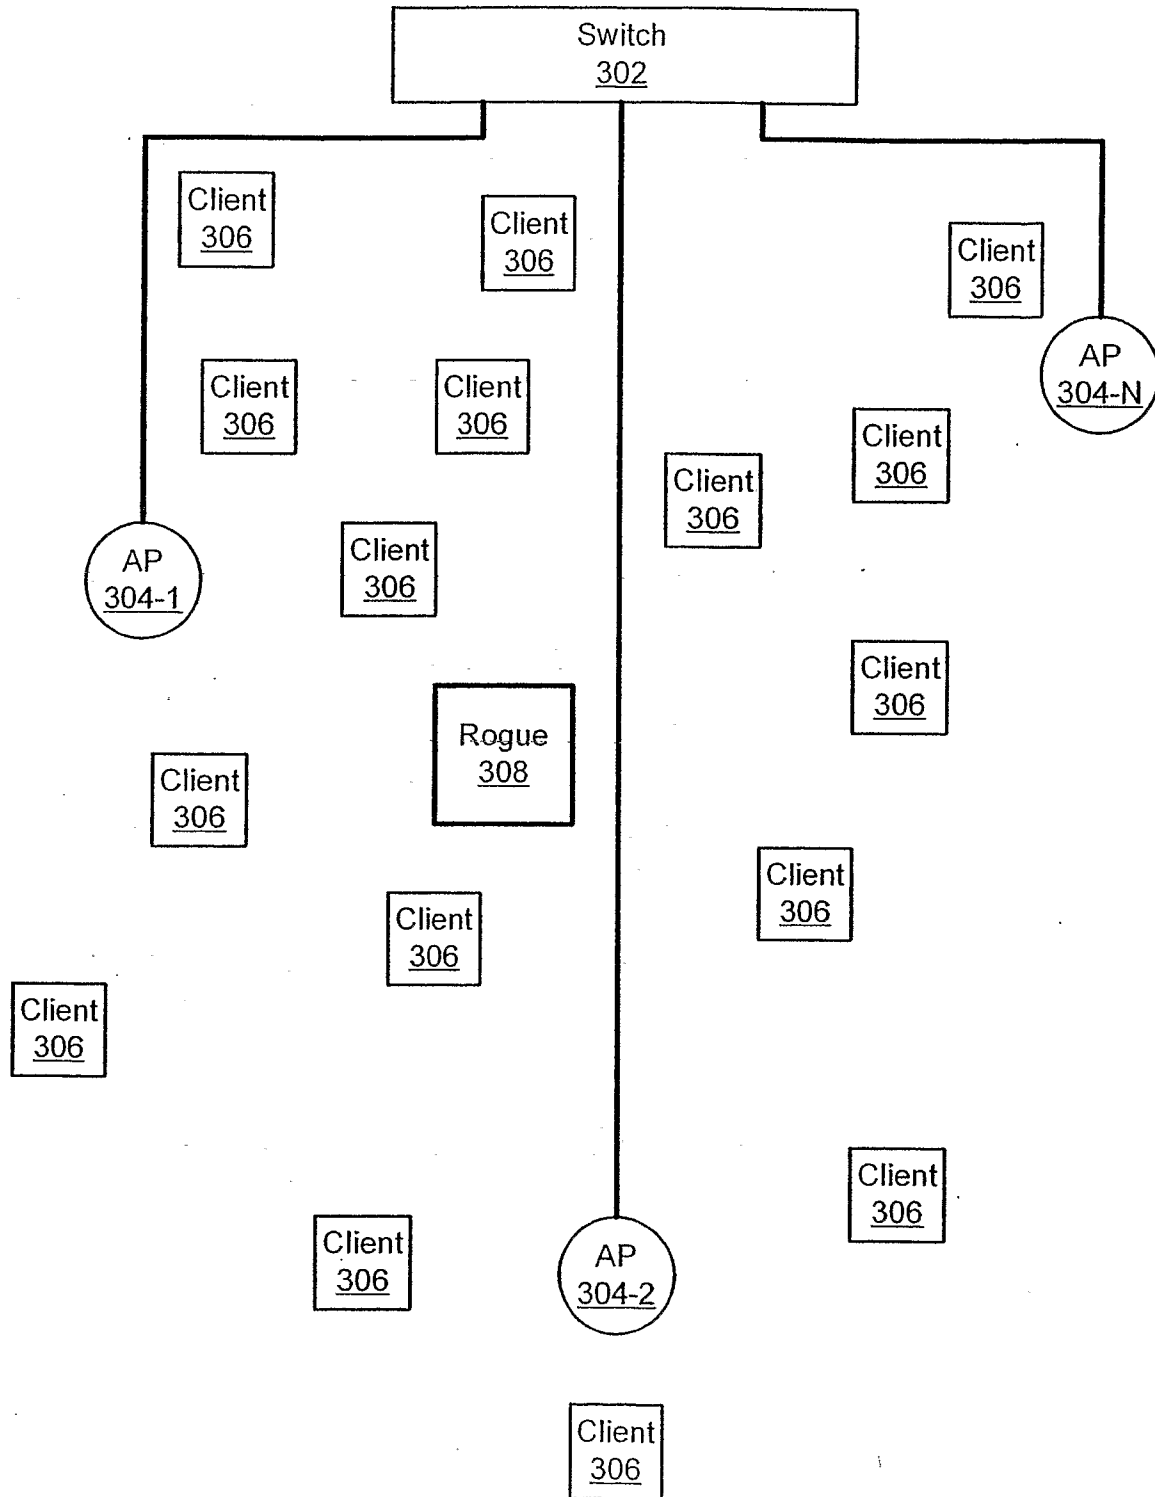


FIG. 3

400 →

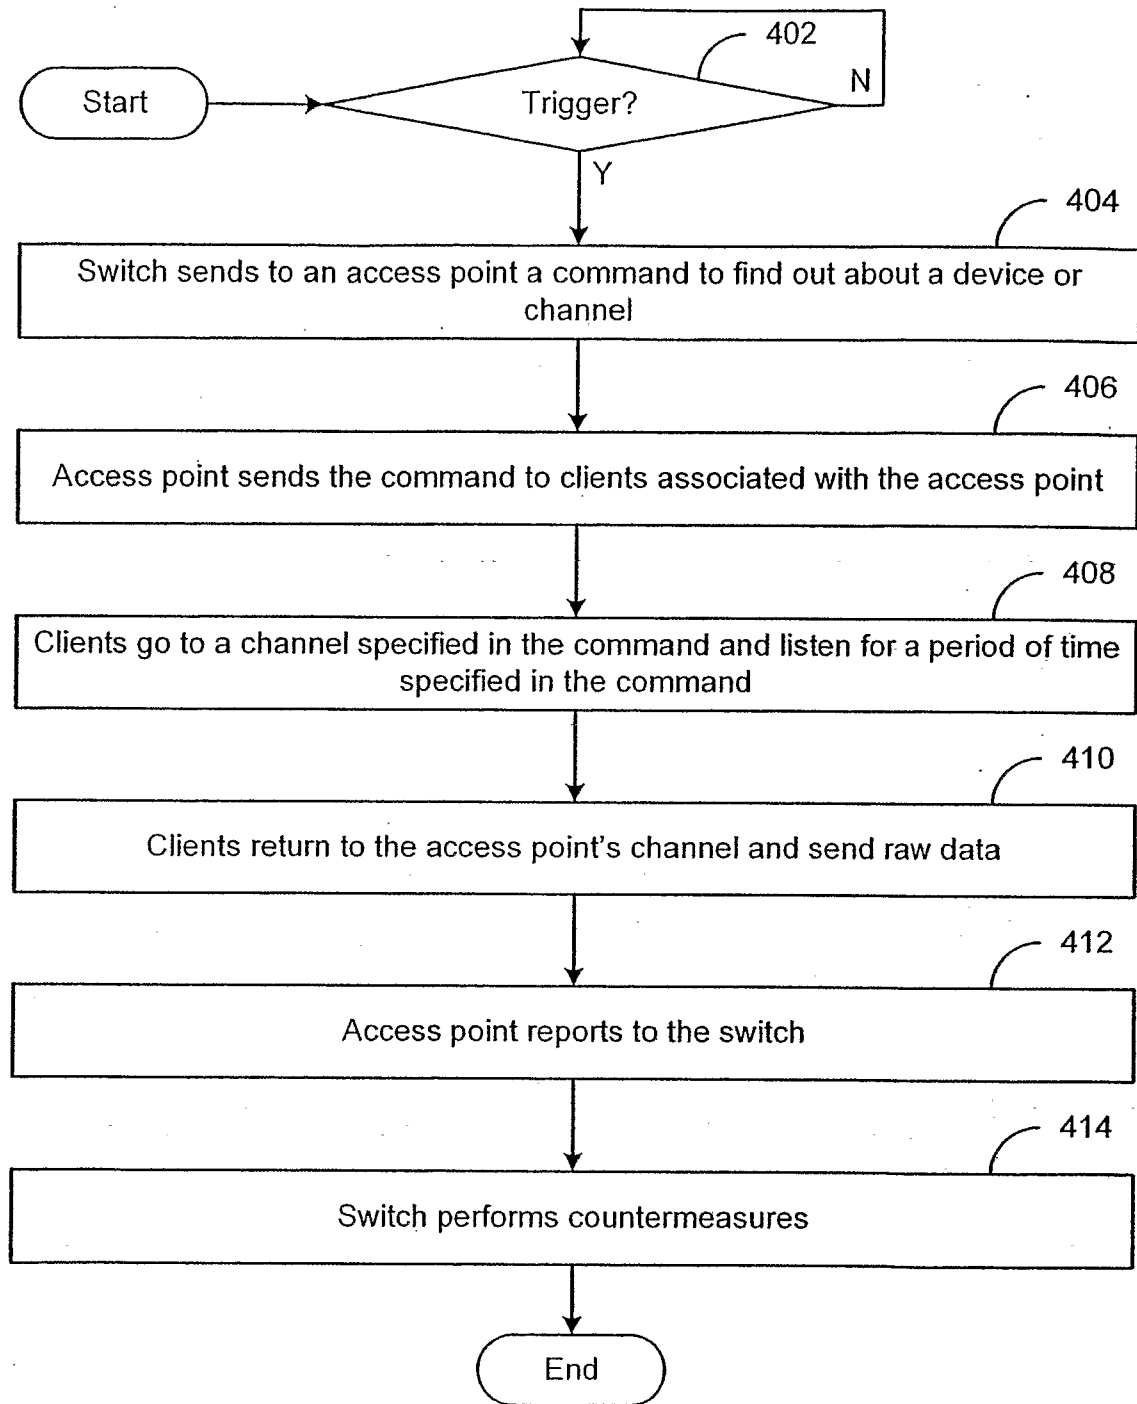


FIG. 4

500 →

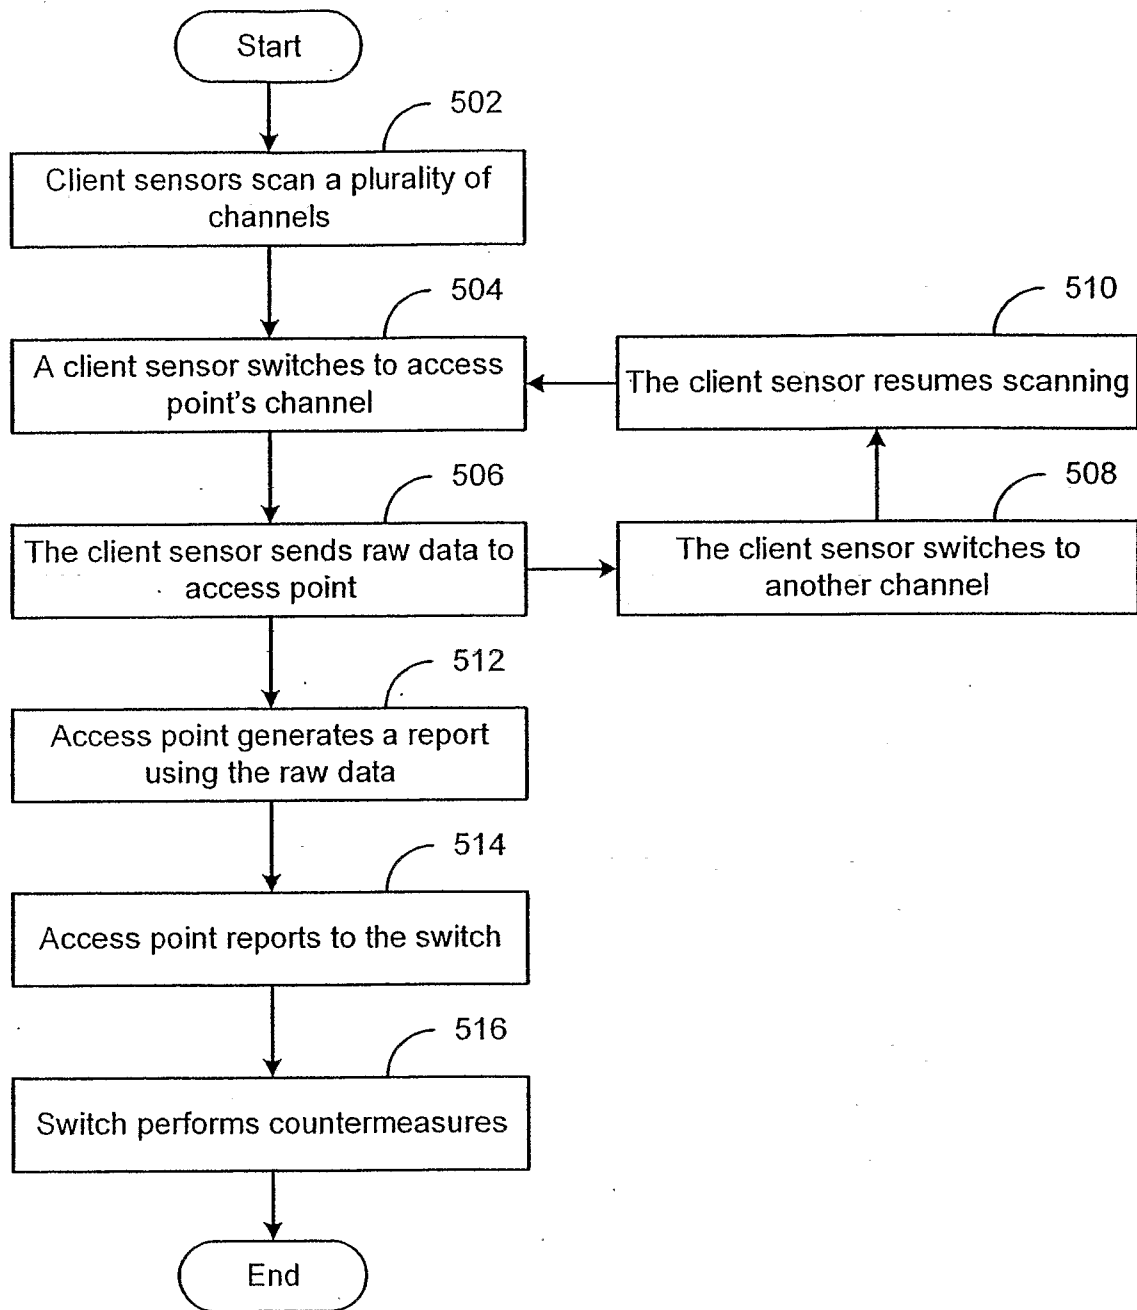


FIG. 5

600 →

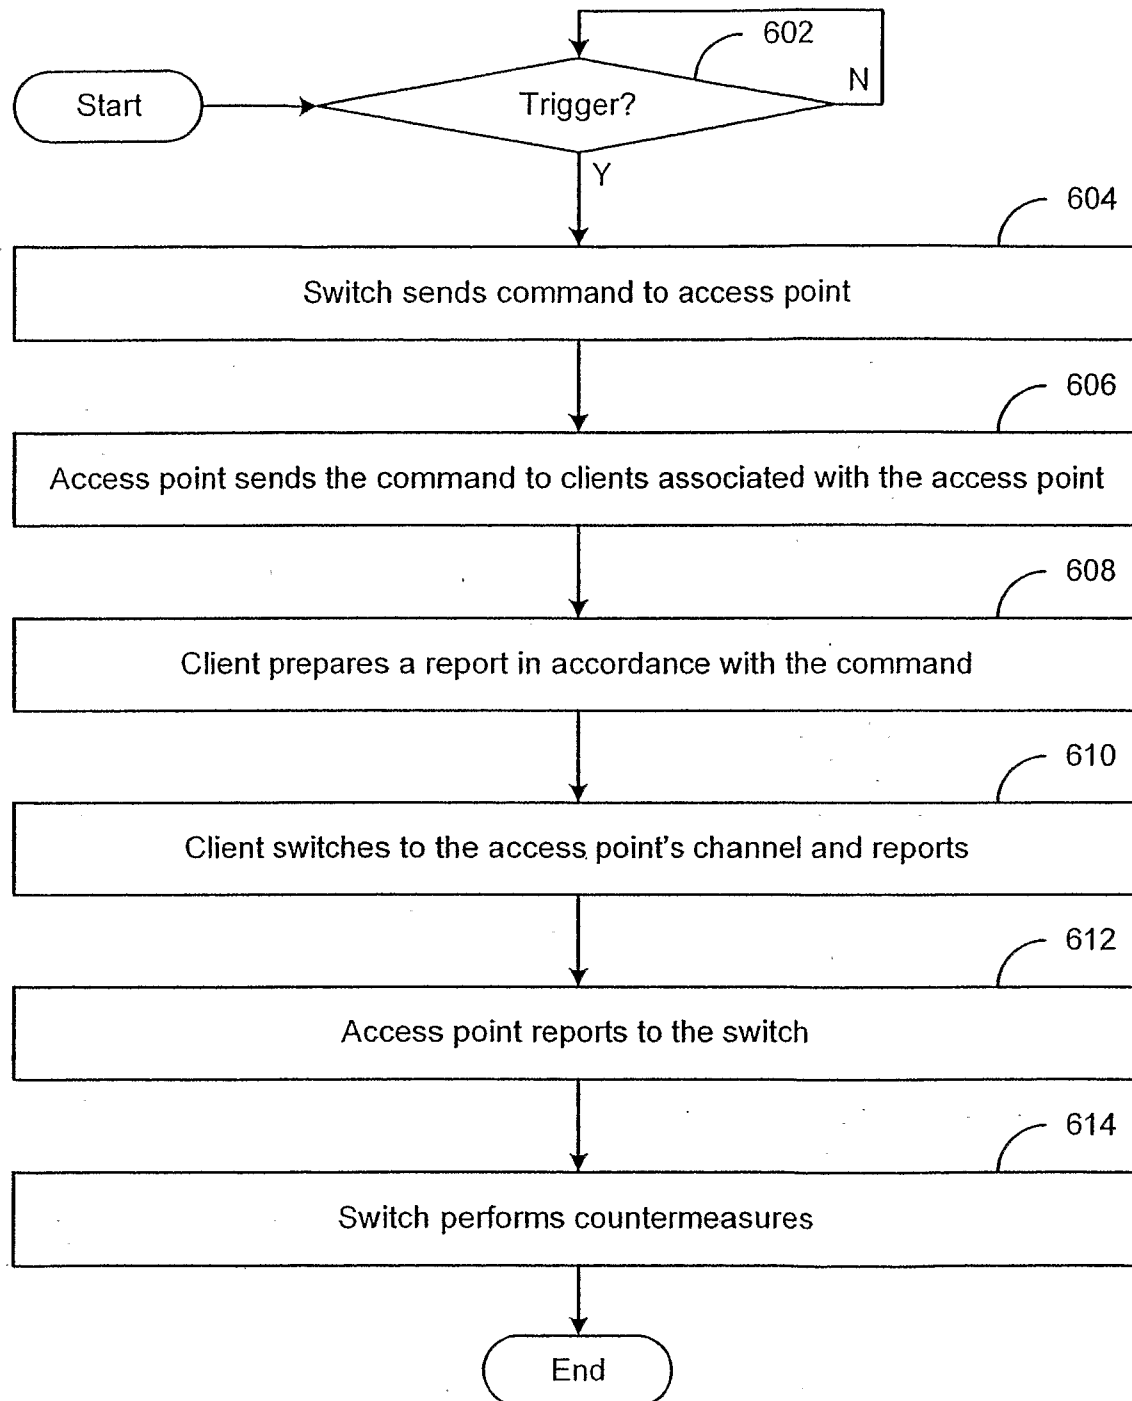


FIG. 6

700 →

7/10

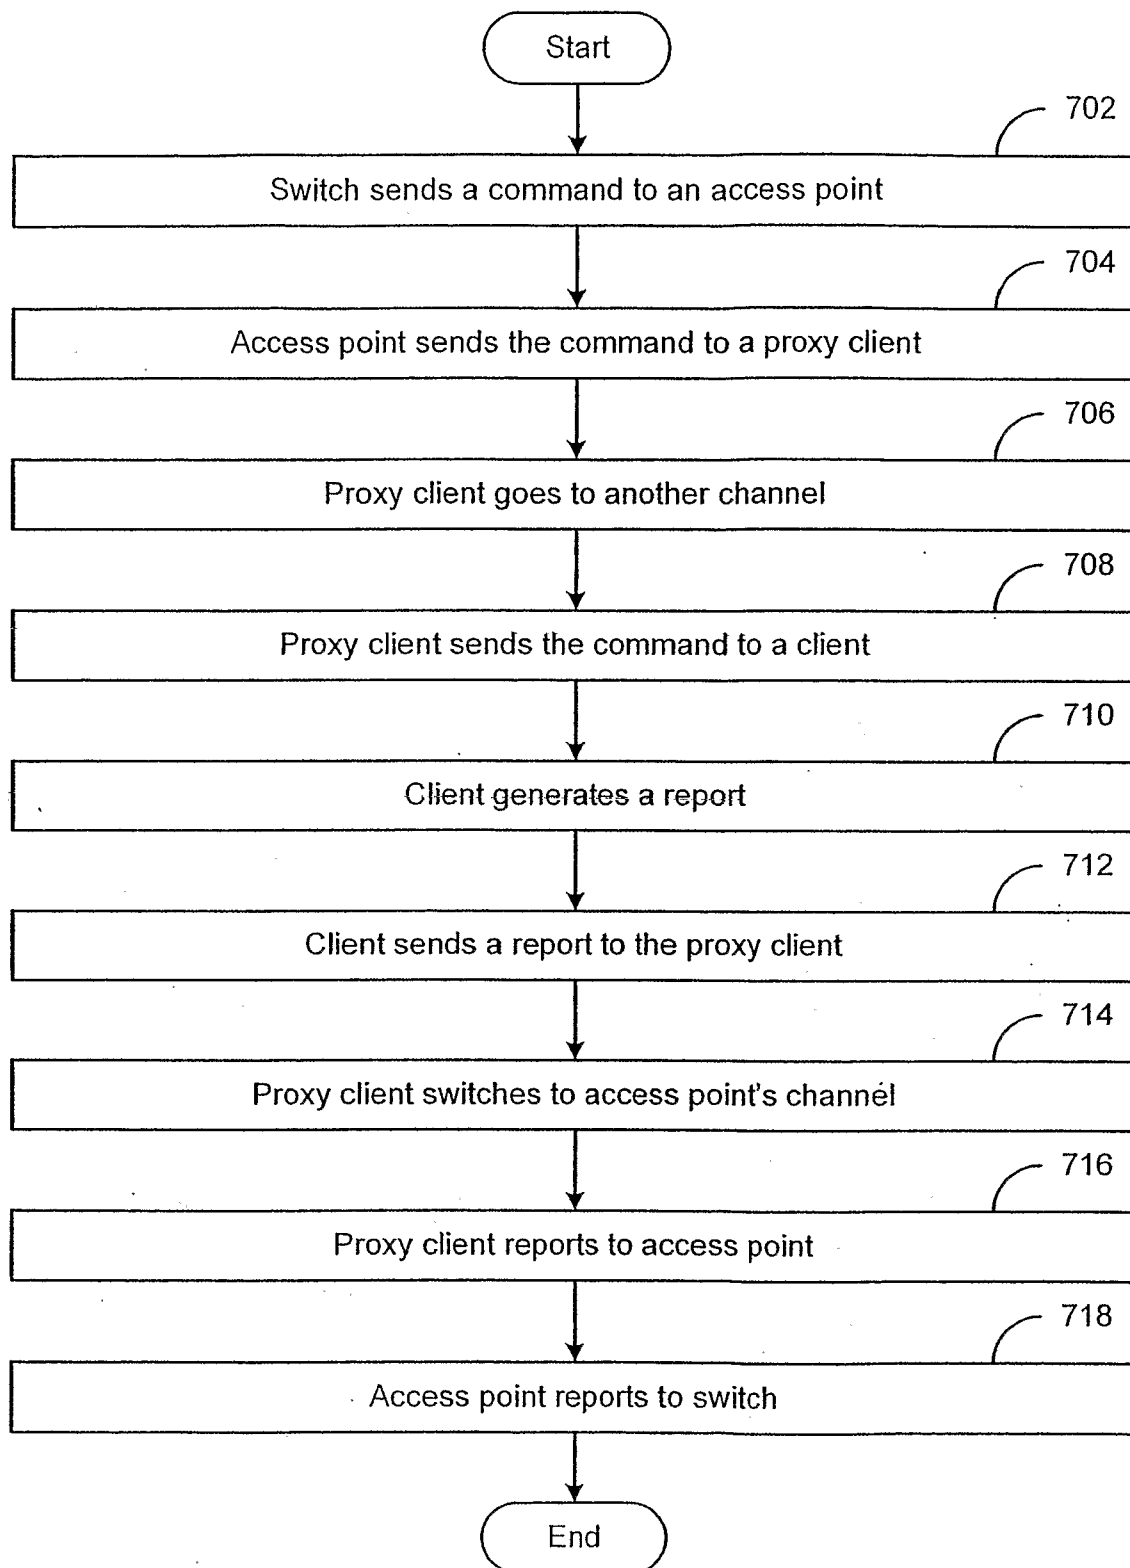


FIG. 7

800 →

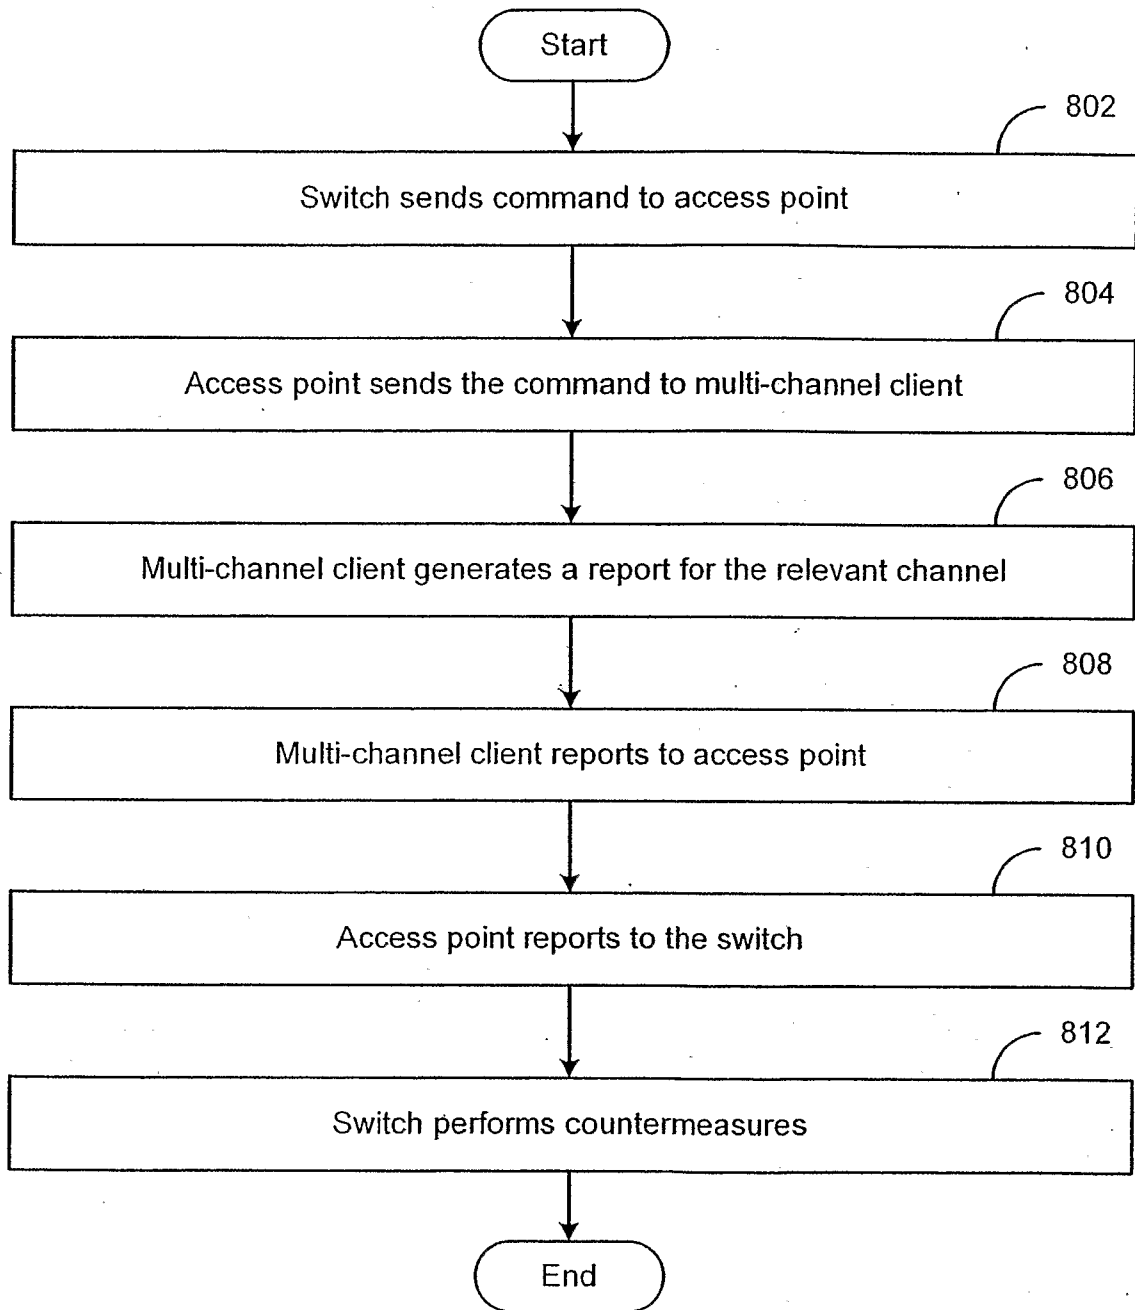


FIG. 8

Mac Address	Type	Subtype	# of Frames
Mac Address 1	Type 1	Subtype 1	2
Mac Address 1	Type 1	Subtype 2	6
Mac Address 1	Type 2	Subtype 1	3
Mac Address 1	Type 3	Subtype 4	6
Mac Address 2	Type 2	Subtype 3	8
.	.	.	.
.	.	.	.
.	.	.	.

FIG. 9

1000 →




FIG. 10

1100 →




FIG. 11