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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
13 September 2007 (13.09.2007)

PCT

(10) International Publication Number
WO 2007/103794 A2

(51) International Patent Classification:
H04L 12/12 (2006.01) *H04L 12/56* (2006.01)

(21) International Application Number:
PCT/US2007/063190

(22) International Filing Date: 2 March 2007 (02.03.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/779,235 3 March 2006 (03.03.2006) US
60/779,824 7 March 2006 (07.03.2006) US

(71) Applicant *(for all designated States except US)*: **QUALCOMM INCORPORATED** [US/US]; Attn: International Ip Administration, 5775 Morehouse Drive, San Diego, California 92121 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **MEYLAN, Arnaud** [CH/US]; 3103 Natchez Avenue, San Diego, California 92117 (US). **DESHPANDE, Manoj M.** [US/US]; 11688 Castile Way, San Diego, California 92128 (US). **NANDA, Sanjiv** [US/US]; 16808 Daza Drive, Ramona, California 92065 (US).

(74) Agent: **OGROD, Gregory D.**; Attn: International Ip Administration, 5775 Morehouse Drive, San Diego, California 92121 (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, 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, 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, MT, 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).

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

[Continued on next page]

(54) Title: STANDBY TIME IMPROVEMENTS FOR STATIONS IN A WIRELESS NETWORK

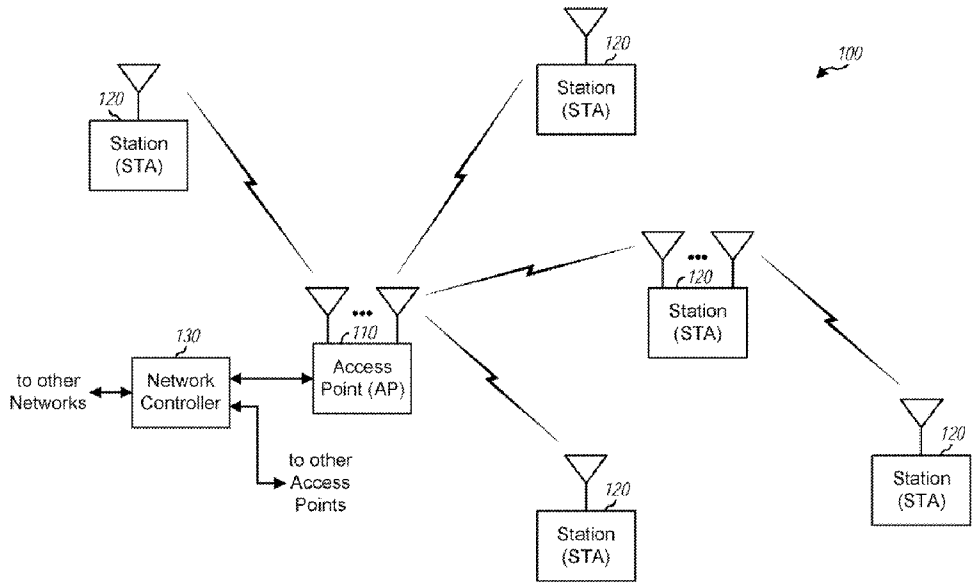

(57) Abstract: Techniques to improve the standby time of a station in a wireless network are described. An access point may advertise or convey a maximum listen interval and/or an association timeout supported by that access point. A station may operate in a power-save mode and may wake up every listen interval to receive a beacon and any potential traffic for the station. The station may select a suitable listen interval based on the maximum listen interval. The station may be dormant for a longer duration than the listen interval and may become active at least once in every association timeout in order to keep the association with the access point alive. The access point may also send broadcast and multicast traffic that might be of interest to stations in the power-save mode less frequently and using a special indication message.

WO 2007/103794 A2

— *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*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

STANDBY TIME IMPROVEMENTS FOR STATIONS IN A WIRELESS NETWORK

[0001] The present application claims priority to provisional U.S. Application Serial No. 60/779,235, entitled "STANDBY TIME IMPROVEMENTS FOR WLAN," filed March 3, 2006, and provisional U.S. Application Serial No. 60/779,824, entitled "STANDBY TIME IMPROVEMENTS FOR WLAN," filed March 7, 2006, both assigned to the assignee hereof and incorporated herein by reference.

BACKGROUND

I. Field

[0002] The present disclosure relates generally to communication, and more specifically to techniques for improving standby time of a station in a wireless network.

II. Background

[0003] Wireless networks are widely deployed to provide various communication services such as voice, video, packet data, broadcast, messaging, etc. These wireless networks may be capable of supporting communication for multiple users by sharing the available network resources. Examples of such networks include wireless local area networks (WLANs), wireless metropolitan area networks (WMANs), wireless wide area networks (WWANs), and wireless personal area networks (WPANs). The terms "network" and "system" are often used interchangeably.

[0004] A wireless network may include any number of access points (APs) and any number of stations (STAs). An access point may act as a coordinator for communication with the stations. A station may actively communicate with an access point, may be idle, or may be powered down at any given moment depending on the data requirements of the station.

[0005] Standby time is an important selling point for portable devices that are battery powered. Current WLAN portable devices tend to have poor standby time performance in comparison to cellular phones. For example, the standby time for currently available WLAN Voice-over-IP (VoIP) phones typically ranges between 40

and 80 hours on batteries similar to those used in cellular phones. In comparison, cellular phones may be able to achieve 400 hours of standby time on similar batteries.

[0006] IEEE 802.11 is a family of standards developed by The Institute of Electrical and Electronics Engineers (IEEE) for WLANs. IEEE 802.11 defines a method for a station to sleep and thus save power. However, the efficiency of the method is limited for stations desiring very low power consumption because of signaling limitations in the standard as well as limited support by the access points and/or stations.

[0007] There is therefore a need in the art for techniques to improve the standby time of a station in a wireless network.

SUMMARY

[0008] Various techniques to improve the standby time of a station in a wireless network are described herein. In an aspect, an access point conveys a maximum listen interval supported by that access point, e.g., via a beacon or a unicast frame. The maximal listen interval for a given access point indicates the maximal time interval for which a given station may sleep when associated with that access point. The listen interval for a given station indicates how often that station may wake up to receive the beacon and potential traffic. A station may operate in a power-save mode and may wake up periodically to receive the beacon and any potential traffic for the station. The station may receive the maximum listen interval from an access point and may efficiently select a suitable listen interval for that station based on the maximum listen interval supported by the access point, without having to repeatedly negotiate different listen intervals with the access point. The station may also discover the maximum listen interval in other manners, as described below.

[0009] In another aspect, an access point conveys its association timeout to stations within its coverage. The association timeout is a duration of time in which the access point will keep an association for a station even when the station shows no activity during this time duration. The station may be dormant for longer than its listen interval in order to extend battery life. The station may obtain the association timeout of the access point and may ensure to be active at least once in every association timeout in order to keep the association with the access point alive.

[0010] In yet another aspect, an access point sends broadcast and multicast traffic that might be of interest to stations in the power-save mode (or PS-stations) in a manner such that these stations may achieve improved power saving. Such broadcast and

multicast traffic may include traffic associated with managing network connectivity, network monitoring, etc. The access point may send (i) a Delivery Traffic Indication Message (DTIM) to indicate regular broadcast and multicast traffic being sent by the access point and (ii) a slow DTIM to indicate broadcast and multicast traffic of potential interest to the PS-stations. The slow DTIM may be sent at a slower rate than the DTIM. The PS-stations may listen for the slow DTIM and may sleep through the DTIM.

[0011] Various aspects and features of the disclosure are described in further detail below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] FIG. 1 shows a wireless network with an access point and multiple stations.

[0013] FIG. 2 shows an example transmission timeline for the access point.

[0014] FIG. 3 shows a design of a beacon frame.

[0015] FIG. 4 shows a process for negotiating a listen interval.

[0016] FIG. 5 shows an apparatus for negotiating a listen interval.

[0017] FIG. 6 shows a process for renegotiating the listen interval on the fly.

[0018] FIG. 7 shows an apparatus for renegotiating the listen interval on the fly.

[0019] FIG. 8 shows a process for avoiding association time out.

[0020] FIG. 9 shows an apparatus for avoiding association time out.

[0021] FIG. 10 shows a process for receiving broadcast and multicast traffic in the power-save mode.

[0022] FIG. 11 shows an apparatus for receiving broadcast and multicast traffic in the power-save mode.

[0023] FIG. 12 shows a block diagram of an access point and a station.

DETAILED DESCRIPTION

[0024] The power saving techniques described herein may be used for various wireless networks such as WLANs, WMANs, WWANs, WPANs, etc. A WLAN may implement a radio technology such as any defined by IEEE 802.11, Hiperlan, etc. A WWAN may be a cellular network such as a Code Division Multiple Access (CDMA) network, a Time Division Multiple Access (TDMA) network, a Frequency Division Multiple Access (FDMA) network, an Orthogonal FDMA (OFDMA) network, a Single-Carrier FDMA (SC-FDMA) network, etc. A WMAN may implement a radio

technology such as any defined by IEEE 802.16 such as 802.16e, which is commonly referred to as WiMAX, or IEEE 802.20. A WPAN may implement a radio technology such as Bluetooth. For clarity, the techniques are described below for an IEEE 802.11 WLAN.

[0025] **FIG. 1** shows a wireless network 100 with an access point (AP) 110 and multiple stations (STAs) 120. In general, a wireless network may include any number of access points and any number of stations. A station is a device that can communicate with another station via a wireless medium. The terms "wireless medium" and "channel" are often used interchangeably. A station may communicate with an access point or peer-to-peer with another station. A station may also be called, and may contain some or all of the functionality of, a terminal, a mobile station, a user equipment, a subscriber unit, etc. A station may be a cellular phone, a handheld device, a wireless device, a personal digital assistant (PDA), a laptop computer, a wireless modem, a cordless phone, etc. An access point is a station that provides access to distribution services via the wireless medium for stations associated with that access point. An access point may also be called, and may contain some or all of the functionality of, a base station, a base transceiver station (BTS), a Node B, an evolved Node B (eNode B), etc.

[0026] For a centralized network, a network controller 130 couples to the access points and provides coordination and control for these access points. Network controller 130 may be a single network entity or a collection of network entities. For a distributed network, the access points may communicate with one another as needed without the uses of network controller 130.

[0027] Wireless network 100 may implement the IEEE 802.11 family of standards. For example, wireless network 100 may implement IEEE 802.11, 802.11a, 802.11b, 802.11e and/or 802.11g, which are existing IEEE 802.11 standards. Wireless network 100 may also implement IEEE 802.11n and/or 802.11s, which are IEEE 802.11 standards being formed. IEEE 802.11, 802.11a, 802.11b, 802.11g and 802.11n cover different radio technologies and have different capabilities. IEEE 802.11e covers quality of service (QoS) enhancements for a Medium Access Control (MAC) layer. In IEEE 802.11e, a station that supports QoS facility is referred to as QSTA, and an access point that supports QoS facility is referred to as QAP. QoS facility refers to mechanisms used to provide parameterized and prioritized QoS.

[0028] A station may communicate with an access point for one or more flows. A flow is a higher layer data stream that is sent via a link. A flow may utilize Transmission Control Protocol (TCP), User Datagram Protocol (UDP), or some other protocol at a transport layer. A flow may also be referred to as a data stream, a traffic stream, etc. A flow may carry any type of data such as voice, video, packet data, etc. A flow may be for a particular traffic class and may have certain requirements on data rate, latency or delay, etc. A flow may be (a) periodic and sent at regular intervals or (b) non-periodic and sent sporadically, e.g., whenever there is data to send. A periodic flow is a flow in which data is sent at regular intervals. For example, a flow for VoIP may send a data frame every 10 or 20 milliseconds (ms). As used herein, a frame is a unit of transmission and may be a data frame, a null frame, a control frame, or some other type of frame. A frame may also be referred to as a packet, a data block, a data unit, a protocol data unit (PDU), a service data unit (SDU), a MAC PDU (MPDU), etc. A call for a station may have one or more flows for one or more traffic types.

[0029] **FIG. 2** shows an example transmission timeline 200 for access point 110 in wireless network 100. In general, each access point in a wireless network may maintain a separate timeline for all transmissions covered by that access point. The transmission timeline for access point 110 is described below. Access point 110 periodically transmits a beacon on the downlink. This beacon carries a preamble and an access point identifier (AP ID) that allows the stations to detect and identify the access point. The time interval between the start of two consecutive beacons is called a target beacon transmit time (TBTT) or a beacon interval. The beacon interval may be fixed or variable and may be set to a suitable duration, e.g., 100 ms.

[0030] Each beacon interval may include any number of service periods for any number of stations. A service period is a contiguous time duration during which an access point may transmit one or more downlink frames to a station and/or may grant one or more transmission opportunities (TXOPs) to the same station. A TXOP is an allocation of time for transmission on a link. A service period may be scheduled or unscheduled. A given station may have any number of service periods within a given beacon interval.

[0031] A station typically performs association procedures to associate with an access point when the station is first powered up or moves into a new WLAN coverage area. Association refers to the mapping of a station to an access point, which enables

the station to receive distribution service. The association allows the distribution service to know which access point to contact for the station. The station attempts to disassociate whenever it leaves the network. The station performs reassociation procedures to "move" a current association from one access point to another access point. The association, disassociation, and reassociation procedures are described in IEEE 802.11 documents.

[0032] A station typically performs negotiation with an access point for various features or attributes such as security, Internet Protocol (IP) address, QoS, flows, power management, etc. The negotiation typically entails exchanging request and response frames between the station and the access point until pertinent parameter values are agreed upon between the station and the access point. Thereafter, the station operates in accordance with the states or context defined by the parameters negotiated with the access point.

[0033] IEEE 802.11 defines a power-save (PS) mode for stations desiring to conserve battery power. A station that desires to operate in the power-save mode indicates this intention to an access point by setting a "PS-mode" bit to 1 in a MAC header of a transmission sent to the access point. A station that is in the power-save mode is referred to as a PS-station. In response, the access point recognizes that the station will be sleeping and will wake up only at agreed upon times to receive traffic. The access point then buffers any incoming traffic data for the station and delivers the data to the station when the station is awake.

[0034] A station that is in the power-save mode may choose to wake up to receive a Traffic Indication Map (TIM) and/or a Delivery Traffic Indication Message (DTIM). The TIM is a bitmap that is present in every beacon transmitted by an access point. The TIM in a given beacon indicates to the station whether there is pending unicast traffic for that station in the upcoming beacon interval. At the time of association, the station and the access point negotiate a listen interval that indicates how often the station will wake up to listen for beacon and hence receive the TIM. The listen interval is typically multiple times the beacon interval, as shown in FIG. 2. For example, if the station has a listen interval of five, then the station may wake up at every fifth beacon to decode the TIM and receive potential traffic for that station.

[0035] The DTIM is a bitmap that indicates whether broadcast and multicast traffic is being delivered in the upcoming beacon interval. The DTIM is sent at an interval that

is selected by the access point. The DTIM interval is typically multiple times the beacon interval and is fixed for a Basic Service Set (BSS), which is a network of stations associated to the access point. A station that is willing to receive broadcast or multicast traffic would decode the DTIM independently of the listen interval for that station.

[0036] An access point may select a DTIM interval based on a tradeoff between latency, buffer size requirements, and power saving. Similarly, a station that is in the power-save mode may select a listen interval as well as whether or not to wake for the DTIM based on a tradeoff between latency, buffer size requirements, and power saving.

[0037] In general, a longer listen interval provides more power saving for a station in the power-save mode but results in more delay, which may be tolerable for some types of traffic. Hence, the station may request a large listen interval at the time of association with an access point if the station favors power saving. However, a larger listen interval results in larger buffer size requirements for the access point to store potential incoming traffic for all stations supported by that access point. Supporting a large listen interval is thus a constraint for the access point because the buffers used for storing potential incoming traffic should be sized according to the amount of data that might be received during the listen interval for all stations supported by the access point.

[0038] IEEE 802.11 does not impose a requirement on the maximum listen interval that an access point needs to support. An access point may support listen intervals within a particular range, e.g., from 1 to 20 times the beacon interval, or possibly more. The supported listen interval range may be dependent on various factors such as the capability of the access point, the number of stations being served by the access point, the number of stations in the power-save mode, etc. Different access points from different vendors may support different ranges of listen intervals. Furthermore, the maximum listen interval supported by a given access point may change over time, e.g., depending on the number of stations that are in the power-save mode with that access point. Conventionally, a station has no easy way of knowing the maximum listen interval supported by an access point.

[0039] In an aspect, an access point conveys the maximum listen interval supported by that access point, and a station uses this information to more efficiently select a suitable listen interval. In one design, the access point advertises the maximum listen

interval in the beacon. A beacon frame includes various information elements carrying various types of information. An information element may be defined for maximum listen interval and may be included in a beacon frame sent by the access point.

[0040] **FIG. 3** shows a design of a beacon frame 300 that may be transmitted by an access point. Beacon frame 300 includes a Timestamp field that indicates the timing of the access point, a Beacon Interval field that indicates the time duration between beacons, a Capability Information field that indicates the requested or advertised capabilities of the access point, a Service Set Identity (SSID) field that carries an identifier for the WLAN, and other information elements defined by IEEE 802.11. In the design shown in FIG. 3, beacon frame 300 includes a Maximum Listen Interval information element 310. Information element 310 includes an Element ID field 312 that is set to a unique value assigned to information element 310, a Length field 314 that indicates the length of subsequent field 316, and field 316 that carries the maximum listen interval supported by the access point.

[0041] A station may listen for a beacon frame upon power up or moving into a new WLAN coverage area. The station may then determine the maximum listen interval supported by the access point. If the station desires to maximize its power saving, then the station may select the maximum listen interval advertised by the access point. The station may also select a shorter listen interval based on its traffic requirements. In any case, the station is able to select and include a suitable listen interval in the first association request sent to the access point.

[0042] In another design, an access point conveys the maximum listen interval that it supports in a unicast frame sent to a station. An information element may be defined for maximum listen interval and may be included in a frame sent by the access point to the station. In one scheme, the maximum listen interval is conveyed during system access. A station may send a probe request to an access point. The access point may include the maximum listen interval supported by that access point in a probe response sent to the station, e.g., if the "PS-mode" bit is set to 1. In another scheme, the maximum listen interval is conveyed during association. A station may select a listen interval desired by the station and include the selected listen interval in the first association request sent to the access point. If the selected service interval is not supported by the access point, then the access point may send an association response

that includes the maximum listen interval supported by the access point. The station may then select a suitable listen interval and include it in the next association request.

[0043] An access point may also convey the maximum listen interval supported by that access point in other manners. A station may determine the maximum listen interval based on a beacon, a probe response, an association response, or some other transmission. The station may then select a suitable listen interval without or with little guesswork and may save power in the association procedure.

[0044] **FIG. 4** shows a process 400 for negotiating a listen interval. A station initially determines a maximum listen interval supported by an access point (block 412). The station may obtain the maximum listen interval from a beacon, a control frame sent in a probe response or an association response, or some other transmission sent by the access point. The station then selects a listen interval based on the maximum listen interval (block 414). For example, the selected listen interval may be equal to the maximum listen interval if the station desires to maximize power saving and its traffic can tolerate the delay. The station then sends the selected listen interval to the access point (block 416).

[0045] **FIG. 5** shows a design of an apparatus 500 for negotiating a listen interval. Apparatus 500 includes means for determining a maximum listen interval supported by an access point (module 512), means for selecting a listen interval based on the maximum listen interval (module 514), and means for sending the selected listen interval to the access point (module 516). Modules 512 to 516 may comprise processors, electronics devices, hardware devices, electronics components, logical circuits, memories, etc., or any combination thereof.

[0046] A station may also determine the maximum listen interval supported by an access point in other manners. The station may send one or more requests to determine the maximum listen interval supported by the access point. If a large listen interval is desired, then the station may try one or more listen intervals at association time until the access point accepts one of the listen intervals. The station may request the largest listen interval in the first association request sent to the access point. If the requested listen interval is too large, then the access point may simply respond with an error status code in an association response, e.g., a status code of 51 for "Association denied because the listen interval is too large". The status code in the response does not convey to the station the largest listen interval supported by the access point. Hence,

the station may then request a smaller listen interval in the next association request sent to the access point. The station may request progressively smaller listen intervals until a requested listen interval is within the range supported by the access point.

[0047] The station may also send the requests for listen intervals in other orders. For example, the station may send a request for a listen interval of N. If that listen interval is supported, then the station may send a request for a larger listen interval, e.g., N+1. Otherwise, the station may send a request for a smaller listen interval, e.g., N–1. The station may repeat sending requests until the maximum listen interval is discovered and may then use it.

[0048] In general, a station may determine the maximum listen interval supported by an access point in a heuristic manner. The station may send multiple requests for multiple listen interval values until the station receives a response accepting one of the requests and another response denying another one of the requests. The station may send one request for one listen interval value at a time. The station may start with a request for a largest listen interval value and conclude with a request for a smallest listen interval value, until the response accepting one of the requests is received. The station may also start with a request for a smallest listen interval value and conclude with a request for a largest listen interval value, until the response denying one of the requests is received. The station may also start with a request for a middle listen interval value, send a request for a larger listen interval value if a response accepting the request is received, and send a request for a smaller listen interval value if a response denying the request is received. The station may also send the requests in other orders. The station may determine a suitable listen interval for use based on the received responses.

[0049] A station typically negotiates an appropriate listen interval when it associates with an access point. The station thereafter uses the negotiated listen interval for the entire duration of the association with the access point. The negotiated listen interval may be inadequate or undesirable for various reasons, and the station may desire to select a more suitable listen interval. In this case, the station would disassociate with the access point and then reassociate with the same access point. The station may negotiate a more suitable listen interval during the reassociation with the access point. The current IEEE 802.11 standard does not provide a mechanism for updating the listen interval while a station is associated with an access point.

[0050] In another aspect, a station renegotiates the listen interval on the fly without having to disassociate and reassociate with an access point. This capability may provide certain advantages, as described below.

[0051] **FIG. 6** shows a process 600 for a station to renegotiate the listen interval on the fly. The station initially establishes an association with an access point (block 612). The station negotiates a first listen interval and establishes states or context (e.g., for security, IP address, QoS, power management, etc.) during establishment of the association (block 614). The station thereafter receives data based on the first listen interval and the established states (block 616).

[0052] The station thereafter determines that the first listen interval is insufficient for whatever reason. The station then renegotiates with the access point for a second listen interval without disassociating with the access point (block 618). The second listen interval may be shorter or longer than the first listen interval, depending on the requirements of the station. The station may select the second listen interval with or without knowledge of the maximum listen interval supported by the access point. For the renegotiation, the station may send to the access point a control frame with the second listen interval. The access point may grant or deny the request by the station. If the request is granted, then the access point may send a response (e.g., an acknowledgement) indicating that the request is granted. The station receives the response and thereafter receives data based on the second listen interval (block 620). If the access point denies the second listen interval, then the station and the access point may continue negotiation until a suitable listen interval is selected and accepted.

[0053] In one design, the states or context (e.g., for security, IP address, QoS, power management, etc.) are retained for the station, and only the listen interval is changed during the renegotiation. In this design, the station thereafter communicates with the access point based on the states/context established earlier during association establishment (block 622). In another design, one or more parameters may also be renegotiated and modified either in the control frame carrying the second listen interval sent by the station or in one or more subsequent frames.

[0054] **FIG. 7** shows a design of an apparatus 700 for renegotiating the listen interval on the fly. Apparatus 700 includes means for establishing an association with an access point (module 712), means for negotiating a first listen interval and establishing states or context during establishment of the association (module 714),

means for receiving data based on the first listen interval and the established states (module 716), means for renegotiating with the access point for a second listen interval without disassociating with the access point (module 718), means for receiving data based on the second listen interval (module 720), and means for communicating with the access point based on the states/context established earlier during association establishment (module 722). Modules 712 to 722 may comprise processors, electronics devices, hardware devices, electronics components, logical circuits, memories, etc., or any combination thereof.

[0055] In one design, new frames that are not currently in IEEE 802.11 are defined and used for sending a request for a new listen interval and a response to this request. In another design, new information elements are defined and included in existing frames to request for a new listen interval and to respond to the request.

[0056] A station may be dormant for an extended duration that is longer than the maximum listen interval supported by an access point. In this case, entities operating above Layer 2 in the network may buffer traffic for the station and synchronize delivery from a page buffering function (PBF) to the station. The traffic buffering may thus occur upstream of the access point. The PBF may be synchronized with the station and may send traffic for the station such that the traffic reaches the access point shortly before the station wakes up at the start of the next listen interval and decodes the TIM in the beacon. The access point does not need to be aware of the extended dormancy by the station. The access point may operate in a normal manner as if the station is waking up every listen interval to decode the TIM. However, the station may wake up at a multiple of the listen interval, which is synchronized with the PBF. The longer actual listen interval may allow the station to save more power. The PBF may ensure that traffic is sent when the station is awake to receive it.

[0057] When a station wakes up at a rate that is less frequent than the listen interval and an access point does not have this knowledge, there is a risk that the access point may disassociate the station if the access point does not detect any activity from the station for an extended period of time. If the access point disassociates the station, then the states/context for the station may be lost. The station may need to perform reassociation procedures in order to re-establish the states/context with the access point. This is undesirable since the reassociation procedures consume time and power.

[0058] In yet another aspect, an access point conveys its association timeout to the stations, and a station may use this information to avoid being timed out by the access point. In one design, the access point advertises its association timeout in the beacon. Referring to the design shown in FIG. 3, beacon frame 300 includes an Association Timeout information element 320 that includes the association timeout for the access point. The access point may set this information element to the current value of the association timeout used by the access point. In another design, the access point conveys the association timeout in a unicast frame, e.g., for a probe response or an association response.

[0059] A station that is dormant for longer than the negotiated listen interval with an access point may obtain the association timeout of the access point. The station may then ensure to be active at least once in every association timeout in order to keep the association with the access point alive. Correspondingly, the access point may ensure to keep the station associated for at least the advertised association timeout duration, even when the station shows no activity.

[0060] **FIG. 8** shows a process 800 for a station to avoid association time out, e.g., due to extended dormancy. The station initially establishes an association with an access point (block 812). The station determines an association timeout for the access point, e.g., based on a beacon or a frame sent by the access point (block 814). The station may negotiate with the access point for a listen interval during establishment of the association. The station may sleep for a period longer than the listen interval, possibly without informing the access point (block 816). The station becomes active at least once in every association timeout to keep the association with the access point alive (block 818).

[0061] **FIG. 9** shows a design of an apparatus 900 for avoiding association time out. Apparatus 900 includes means for establishing an association with an access point (module 912), means for determining an association timeout for the access point, e.g., based on a beacon or a frame sent by the access point (module 914), means for sleeping for a period longer than a listen interval negotiated during establishment of the association, possibly without informing the access point (module 916), and means for becoming active at least once in every association timeout to keep the association with the access point alive (module 918). Modules 912 to 918 may comprise processors,

electronics devices, hardware devices, electronics components, logical circuits, memories, etc., or any combination thereof.

[0062] An access point may advertise the maximum listen interval and/or the association timeout supported by that access point in the beacon, as described above. The maximum listen interval and/or the association timeout may be received by, and applied to, all stations within the coverage of the access point. The access point may also convey the maximum listen interval and/or the association timeout in unicast frames sent to specific stations. The access point may use different maximum listen intervals and/or different association timeouts for different stations, different access categories, etc. For example, a longer listen interval and/or a longer association timeout may be used for a station with higher priority. Conversely, a shorter listen interval and/or a shorter association timeout may be used for a station with lower priority. Individual listen interval and/or association timeout for a specific station may be conveyed in the beacon or unicast frames.

[0063] The current IEEE 802.11 standard supports power save for a station and groups traffic into two categories:

- Broadcast and multicast traffic indicated by the DTIM, and
- Unicast traffic sent in directed frames to the station after the presence of the traffic is indicated in the beacon's TIM at every listen interval.

[0064] A station may desire to receive broadcast or multicast traffic (e.g., audio, streaming video, etc.) from an application layer. The station may then wake up a sufficient amount of time in order to receive these traffic streams. The station is likely not a candidate for deep sleep and significant power-save operation because the DTIM may be sent often, e.g., every beacon.

[0065] Conventionally, the broadcast and multicast traffic indicated by the DTIM includes (1) broadcast and multicast traffic from the application layer, which is referred to herein as "application" broadcast and multicast traffic, and (2) broadcast and multicast traffic associated with managing network connectivity, network monitoring, etc., which is referred to herein as "network" broadcast and multicast traffic. Some examples of network broadcast and multicast traffic include Address Resolution Protocol (ARP) traffic, Dynamic Host Configuration Protocol (DHCP) traffic, topology updates, and other such types of traffic. ARP is used to map MAC addresses to IP addresses. DHCP is used for dynamic IP configuration. A station may desire to receive

network broadcast and multicast traffic even when the station intends to be idle and operating in deep sleep.

[0066] Conventionally, both application and network broadcast and multicast traffic are included together and sent using the DTIM mechanism. An idle station that desires to save power will likely not be interested in receiving the application broadcast and multicast traffic. Otherwise, the station may have its display, keypad, and processor turned on, and therefore may not be saving much power anyway. However, the idle terminal may desire to ensure that its connectivity for Layer 2 and above is operational. For example, the station may desire to respond to ARP requests, possible DHCP messages, etc. These messages are also broadcast traffic and are thus sent using the DTIM.

[0067] The DTIM indicates both (a) network broadcast and multicast traffic used to maintain Layer 2 and above connectivity and (b) application broadcast and multicast traffic. Hence, a station that is interested in receiving only network broadcast and multicast traffic would need to wake up for each DTIM in order to receive potential network broadcast and multicast traffic. The DTIM is typically sent in every beacon (or every few beacons) in order to reduce the delay of the application broadcast and multicast traffic. In this case, the station may need to wake up at each beacon for the DTIM, which may severely impact power save performance.

[0068] In yet another aspect, an access point sends network broadcast and multicast traffic in a manner to improve power saving for stations in the power-save mode. A new service access point (SAP) may be made available for a new broadcast and multicast traffic class that is associated with stations in the power-save mode. This traffic class may be referred to as Power Save (PS) broadcast and multicast traffic and may include network broadcast and multicast traffic and/or other broadcast and multicast traffic that might be of interest to stations in the power-save mode.

[0069] A new DTIM may also be made available and may be referred to as a SlowDTIM or slow DTIM. The PS broadcast and multicast traffic may be sent using the SlowDTIM. The SlowDTIM may be sent at every SlowDTIM interval, which is a predetermined number of beacon intervals. The SlowDTIM interval may be larger than the DTIM interval and may be selected based on a tradeoff between power saving and messaging delay. For example, the SlowDTIM may be sent in every 2, 3, 4, or some other multiple of the DTIM.

[0070] Traffic relevant to maintaining Layer 2 and above connectivity and/or other types of PS broadcast and multicast traffic may be sent using the SlowDTIM. The PS broadcast and multicast traffic may also be copied and sent using the DTIM so that stations that do not receive the SlowDTIM can also receive this traffic. Application broadcast and multicast traffic is not sent using the SlowDTIM and is instead sent using the DTIM.

[0071] A station in the power-save mode may be able to maintain Layer 2 and above connectivity by listening to the SlowDTIM and receiving network broadcast and multicast traffic used to maintain network connectivity. The station may sleep through the DTIM, which may be sent at a shorter interval or a faster rate. The SlowDTIM may improve power saving for the station.

[0072] In one design, the SlowDTIM is sent every N DTIM, where N may be any integer greater than one. In this design, a station in the power-save mode may wake up at each listen interval as well as for each SlowDTIM. In another design, the SlowDTIM is sent in a manner to reduce the number of times that power-save stations need to wake up. For example, the SlowDTIM may be sent in each listen interval for a station or a group of stations having the same listen interval. A station may then receive the SlowDTIM and the TIM from the same beacon in each listen interval.

[0073] A station may desire deep sleep and may select a listen interval that is much longer than the SlowDTIM interval. In one design, to avoid requiring the station to wake up at each SlowDTIM interval, the network broadcast and multicast traffic may be sent directly to the station in unicast frames. The station may then receive the network broadcast and multicast traffic when the station wakes up for its listen interval. The station and the access point may configure this mode of traffic delivery, e.g., during configuration setup for the power-save mode. This traffic delivery mode may be selectively applied to stations desiring deep sleep and may be readily supported by the access point when the network broadcast and multicast traffic is segregated from the application broadcast and multicast traffic.

[0074] **FIG. 10** shows a process 1000 for receiving broadcast and multicast traffic by a station in the power-save mode. The station receives a slow DTIM sent at a slower rate than a regular DTIM, with the regular DTIM indicating first broadcast and multicast traffic and the slow DTIM indicating second broadcast and multicast traffic. (block 1012). The first broadcast and multicast traffic may comprise network broadcast

and multicast traffic, application broadcast and multicast traffic, and/or other broadcast and multicast traffic for a WLAN. The second broadcast and multicast traffic may comprise network broadcast and multicast traffic and/or other broadcast and multicast traffic that might be interest to stations in the power-save mode. The station receives the second broadcast and multicast traffic (e.g., used for maintaining network connectivity) as indicated by the slow DTIM (block 1014). The slow DTIM may be sent in every N regular DTIM, in every listen interval for the station, etc. The station may sleep through the regular DTIM (block 1016).

[0075] **FIG. 11** shows a design of an apparatus 1100 for receiving broadcast and multicast traffic in the power-save mode. Apparatus 1100 includes means for receiving a slow DTIM sent at a slower rate than a regular DTIM, with the regular DTIM indicating first broadcast and multicast traffic and the slow DTIM indicating second broadcast and multicast traffic (module 1112), means for receiving the second broadcast and multicast traffic (e.g., used for maintaining network connectivity), as indicated by the slow DTIM (module 1114), and means for sleeping through the regular DTIM (module 1116). Modules 1112 to 1116 may comprise processors, electronics devices, hardware devices, electronics components, logical circuits, memories, etc., or any combination thereof.

[0076] **FIG. 12** shows a block diagram of access point 110 and station 120, which may be one of the stations in FIG. 1. On the downlink, at access point 110, a transmit (TX) data processor 1212 receives traffic data from a data source 1210 for the stations scheduled for transmission, control data (e.g., the maximum listen interval, association timeout, response frames, etc.) from a controller/processor 1220, and scheduling information (e.g., TIM, DTIM, slow DTIM, etc.) from a scheduler 1224 via controller/processor 1220. TX data processor 1212 processes (e.g., encodes, interleaves, modulates, and scrambles) the traffic data for each station based on a rate selected for that station, processes the control data and scheduling information, and generates data chips. A transmitter (TMTR) 1214 processes (e.g., converts to analog, amplifies, filters, and upconverts) the data chips and generates a downlink signal, which is transmitted via an antenna 1216 to the stations.

[0077] At station 120, an antenna 1252 receives the downlink signal from access point 110 and provides a received signal. A receiver (RCVR) 1254 processes the received signal and provides samples. A receive (RX) data processor 1256 processes

(e.g., descrambles, demodulates, deinterleaves, and decodes) the samples, provides decoded data for station 120 to a data sink 1258, and provides control data and scheduling information to a controller/processor 1260.

[0078] On the uplink, at station 120, a TX data processor 1272 receives traffic data from a data source 1270 and control data (e.g., listen interval, request frames, etc.) from controller/processor 1260. TX data processor 1272 processes the traffic and control data based on a rate selected for the station and generates data chips. A transmitter 1274 processes the data chips and generates an uplink signal, which is transmitted via antenna 1252 to access point 110.

[0079] At access point 110, antenna 1216 receives the uplink signals from the station 120 and other stations. A receiver 1230 processes a received signal from antenna 1216 and provides samples. An RX data processor 1232 processes the samples and provides decoded data for each station to a data sink 1234 and provides control data to controller/processor 1220.

[0080] Controllers/processors 1220 and 1260 direct operation at access point 110 and station 120, respectively. Scheduler 1224 may perform scheduling for the stations and may also perform scheduling for broadcast and multicast traffic sent using the DTIM and slow DTIM. Scheduler 1224 may reside at access point 110, as shown in FIG. 12, or at another network entity.

[0081] The power saving techniques described herein may be implemented by various means. For example, these techniques may be implemented in hardware, firmware, software, or a combination thereof. For a hardware implementation, the processing units used to perform the techniques at a station may be implemented within one or more application specific integrated circuits (ASICs), digital signal processors (DSPs), digital signal processing devices (DSPDs), programmable logic devices (PLDs), field programmable gate arrays (FPGAs), processors, controllers, micro-controllers, microprocessors, electronic devices, other electronic units designed to perform the functions described herein, or a combination thereof. The processing units used to perform the techniques at an access point may be implemented within one or more ASICs, DSPs, processors, etc.

[0082] For a firmware and/or software implementation, the power saving techniques may be implemented with modules (e.g., procedures, functions, etc.) that perform the functions described herein. The firmware and/or software codes may be stored in a

memory (e.g., memory 1222 or 1262 in FIG. 12) and executed by a processor (e.g., processor 1220 or 1260). The memory may be implemented within the processor or external to the processor.

[0083] The previous description of the disclosure is provided to enable any person skilled in the art to make or use the disclosure. Various modifications to the disclosure will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other variations without departing from the spirit or scope of the disclosure. Thus, the disclosure is not intended to be limited to the examples described herein but is to be accorded the widest scope consistent with the principles and novel features disclosed herein.

WHAT IS CLAIMED IS:

CLAIMS

1. An apparatus comprising:

a processor configured to determine a maximum listen interval supported by an access point, to select a listen interval based on the maximum listen interval, and to send the selected listen interval to the access point; and

a memory coupled to the processor.

2. The apparatus of claim 1, wherein the processor is configured to receive a beacon sent by the access point and to obtain the maximum listen interval from the beacon.

3. The apparatus of claim 1, wherein the processor is configured to send a probe request to the access point and to receive a probe response with the maximum listen interval from the access point.

4. The apparatus of claim 1, wherein the processor is configured to select an initial listen interval, to send an association request with the initial listen interval to the access point, and to receive an association response with the maximum listen interval from the access point.

5. The apparatus of claim 1, wherein the processor is configured to receive data from the access point at time instances determined based on the selected listen interval and to sleep between the time instances.

6. The apparatus of claim 1, wherein the selected listen interval is equal to the maximum listen interval.

7. A method comprising:

determining a maximum listen interval supported by an access point;

selecting a listen interval based on the maximum listen interval; and

sending the selected listen interval to the access point.

8. The method of claim 7, wherein the determining the maximum listen interval comprises

receiving a beacon sent by the access point, and

obtaining the maximum listen interval from the beacon.

9. The method of claim 7, further comprising:

receiving data from the access point at time instances determined based on the selected listen interval and

sleeping between the time instances.

10. An apparatus comprising:

means for determining a maximum listen interval supported by an access point;

means for selecting a listen interval based on the maximum listen interval; and

means for sending the selected listen interval to the access point.

11. A processor-readable medium including instructions stored thereon, comprising:

a first instruction set for determining a maximum listen interval supported by an access point;

a second instruction set for selecting a listen interval based on the maximum listen interval; and

a third instruction set for sending the selected listen interval to the access point.

12. An apparatus comprising:

a processor configured to send a plurality of requests for a plurality of listen interval values, to receive a response accepting one of the plurality of requests, to receive another response denying another one of the plurality of requests, and to determine a listen interval for use based on the received responses; and

a memory coupled to the processor.

13. The apparatus of claim 12, wherein the processor is configured to send one request for one listen interval value at a time until the response accepting one of the plurality of requests and the response denying another one of the plurality of requests are both received.

14. The apparatus of claim 12, wherein the processor is configured to send one request for one listen interval value at a time, starting with a request for a largest listen interval value and concluding with a request for a smallest listen interval value, until the response accepting one of the plurality of requests is received.

15. The apparatus of claim 12, wherein the processor is configured to send one request for one listen interval value at a time, starting with a request for a smallest listen interval value and concluding with a request for a largest listen interval value, until the response denying another one of the plurality of requests is received.

16. An apparatus comprising:

a processor configured to send a maximum listen interval supported by an access point, to receive from a station a listen interval selected based on the maximum listen interval, and to send data to the station based on the listen interval; and

a memory coupled to the processor.

17. The apparatus of claim 16, wherein the processor is configured to send the maximum listen interval in a beacon or a frame to the station.

18. An apparatus comprising:

a processor configured to establish an association with an access point, to receive data based on a first listen interval negotiated with the access point during establishment of the association, to renegotiate with the access point for a second listen interval without disassociating with the access point, and to receive data based on the second listen interval after the renegotiation; and

a memory coupled to the processor.

19. The apparatus of claim 18, wherein the processor is configured to establish states during the establishment of the association and to retain the states during the renegotiation for the second listen interval.

20. The apparatus of claim 18, wherein the processor is configured to send a control frame with the second listen interval to the access point and to receive an indication of acceptance of the second listen interval from the access point.

21. The apparatus of claim 18, wherein the processor is configured to wake up for beacon frames determined by the first or second listen interval and to sleep between the beacon frames when no data is sent to the apparatus.

22. A method comprising:

establishing an association with an access point;

receiving data based on a first listen interval negotiated with the access point during establishment of the association;

renegotiating with the access point for a second listen interval without disassociating with the access point; and

receiving data based on the second listen interval after the renegotiation.

23. The method of claim 22, further comprising:

establishing states during establishment of the association; and

retaining the states during the renegotiation for the second listen interval.

24. An apparatus comprising:

means for establishing an association with an access point;

means for receiving data based on a first listen interval negotiated with the access point during establishment of the association;

means for renegotiating with the access point for a second listen interval without disassociating with the access point; and

means for receiving data based on the second listen interval after the renegotiation.

25. A processor-readable medium including instructions stored thereon, comprising:

a first instruction set for establishing an association with an access point;

a second instruction set for receiving data based on a first listen interval negotiated with the access point during establishment of the association;

a third instruction set for renegotiating with the access point for a second listen interval without disassociating with the access point; and

a fourth instruction set for receiving data based on the second listen interval after the renegotiation.

26. An apparatus comprising:

a processor configured to establish an association for a station, to send data to the station based on a first listen interval negotiated with the station during establishment of the association, to renegotiate with the station for a second listen interval without disassociating the station, and to send data to the station based on the second listen interval after the renegotiation; and

a memory coupled to the processor.

27. The apparatus of claim 26, wherein the processor is configured to establish states for the station during the establishment of the association and to retain the states for the station during the renegotiation for the second listen interval.

28. An apparatus comprising:

a processor configured to establish an association with an access point, to determine an association timeout for the access point, and to become active at least once in every association timeout to keep the association with the access point alive; and

a memory coupled to the processor.

29. The apparatus of claim 28, wherein the processor is configured to receive a beacon sent by the access point and to obtain the association timeout from the beacon.

30. The apparatus of claim 28, wherein the processor is configured to receive a frame from the access point and to obtain the association timeout from the received frame.

31. The apparatus of claim 28, wherein the processor is configured to negotiate with the access point for a listen interval during the establishment of the association and to sleep for a period longer than the listen interval.

32. The apparatus of claim 28, wherein the processor is configured to negotiate with the access point for a listen interval during the establishment of the association and to sleep for a period longer than the listen interval without informing the access point.

33. A method comprising:

establishing an association with an access point;

determining an association timeout for the access point; and

becoming active at least once in every association timeout to keep the association with the access point alive.

34. The method of claim 33, wherein the determining the association timeout comprises

receiving a beacon sent by the access point, and

obtaining the association timeout from the beacon.

35. The method of claim 33, further comprising:

negotiating with the access point for a listen interval during the establishment of the association; and

sleeping for a period longer than the listen interval.

36. An apparatus comprising:

means for establishing an association with an access point;

means for determining an association timeout for the access point; and

means for becoming active at least once in every association timeout to keep the association with the access point alive.

37. A processor-readable medium including instructions stored thereon, comprising:

a first instruction set for establishing an association with an access point;

a second instruction set for determining an association timeout for the access point; and

a third instruction set for becoming active at least once in every association timeout to keep the association with the access point alive.

38. An apparatus comprising:

a processor configured to establish an association for a station, to send an association timeout supported by an access point, and to receive from the station at least one transmission in every association timeout to keep the association alive; and

a memory coupled to the processor.

39. The apparatus of claim 38, wherein the processor is configured to send the association timeout in a beacon or a frame to the station.

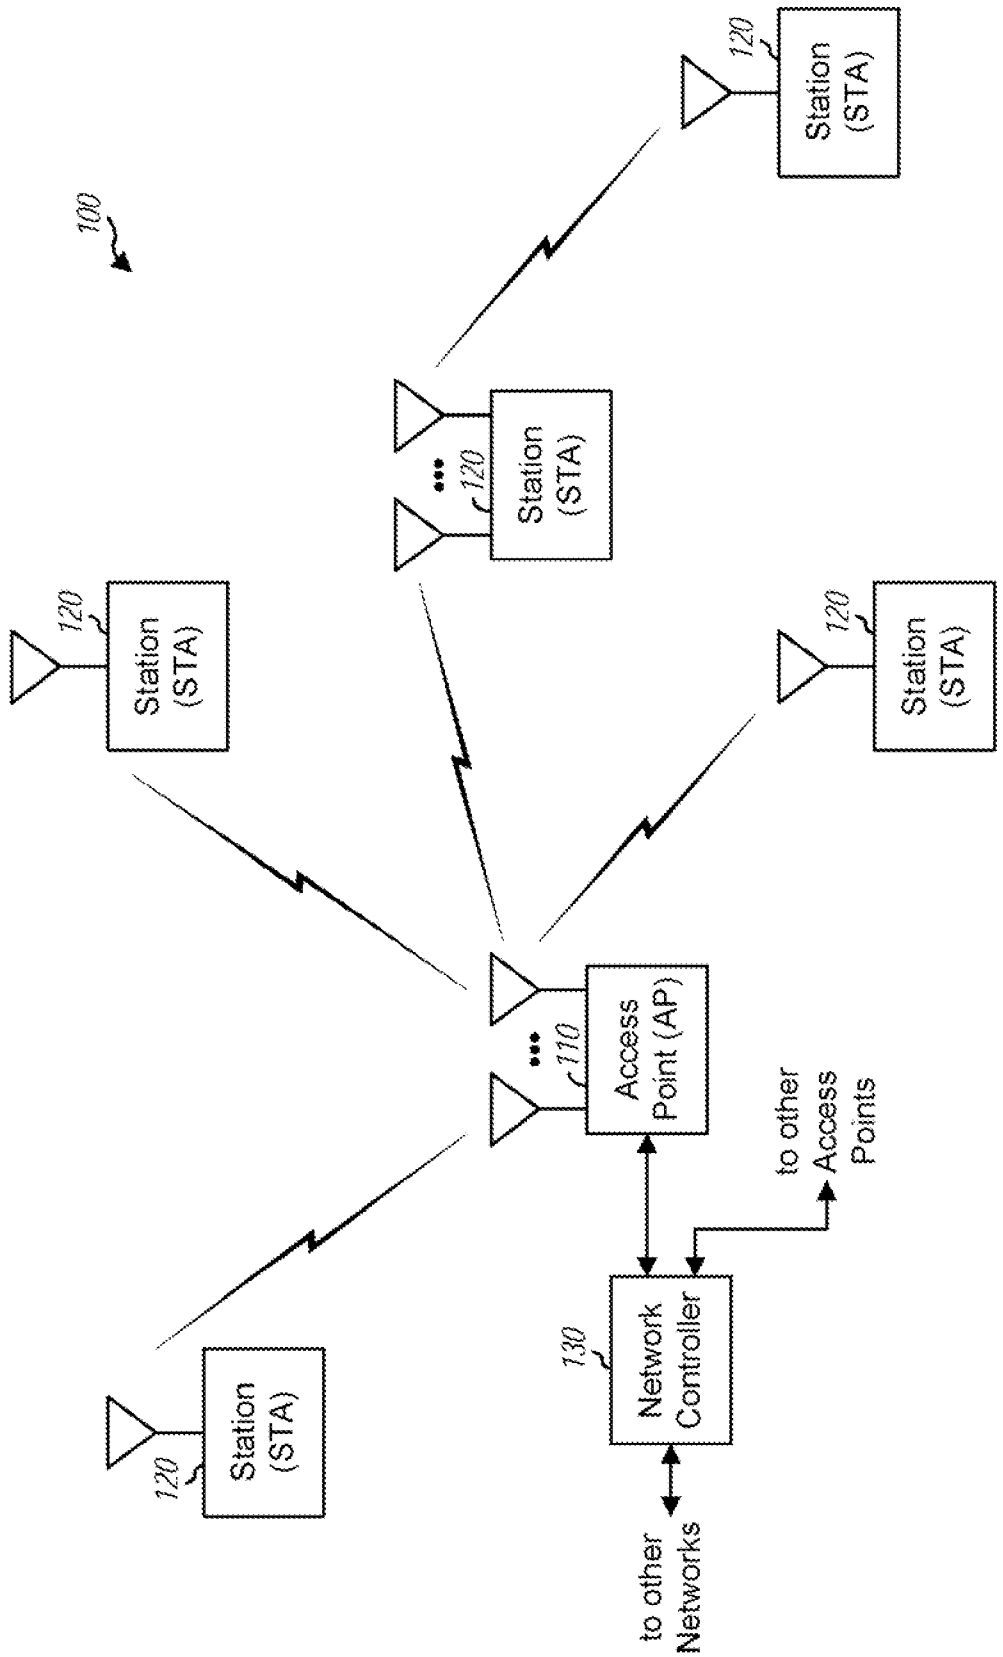
100
Station (STA)
120
Access Point (AP)
110
Network Controller
130
to other Networks
to other Access Points

FIG. 1

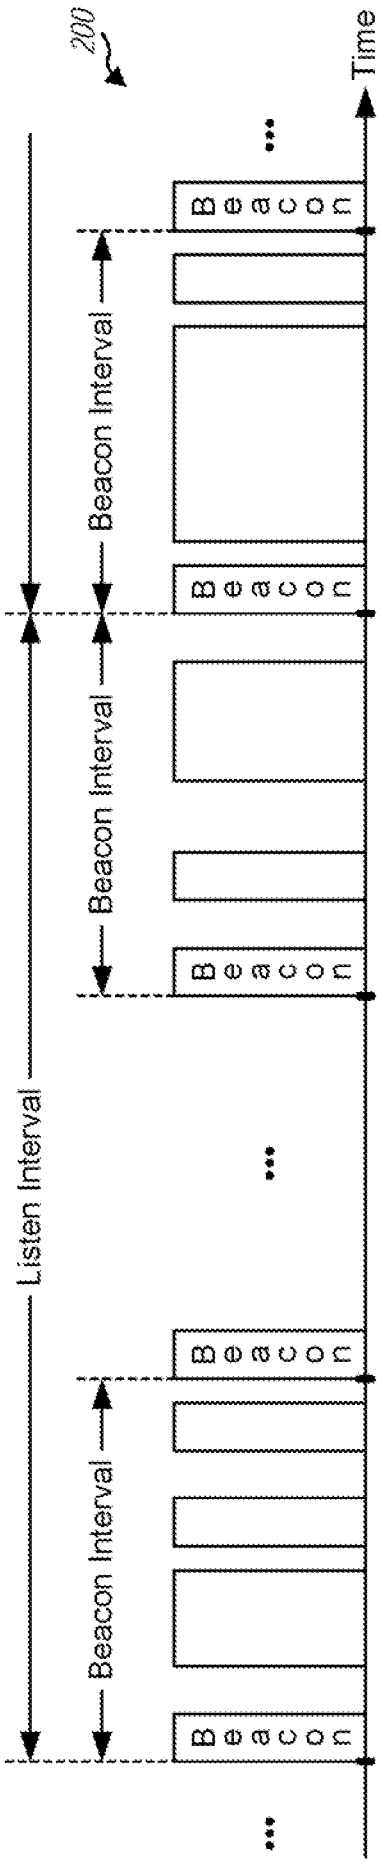
200
Time
Beacon
Beacon Interval
Listen Interval

FIG. 2

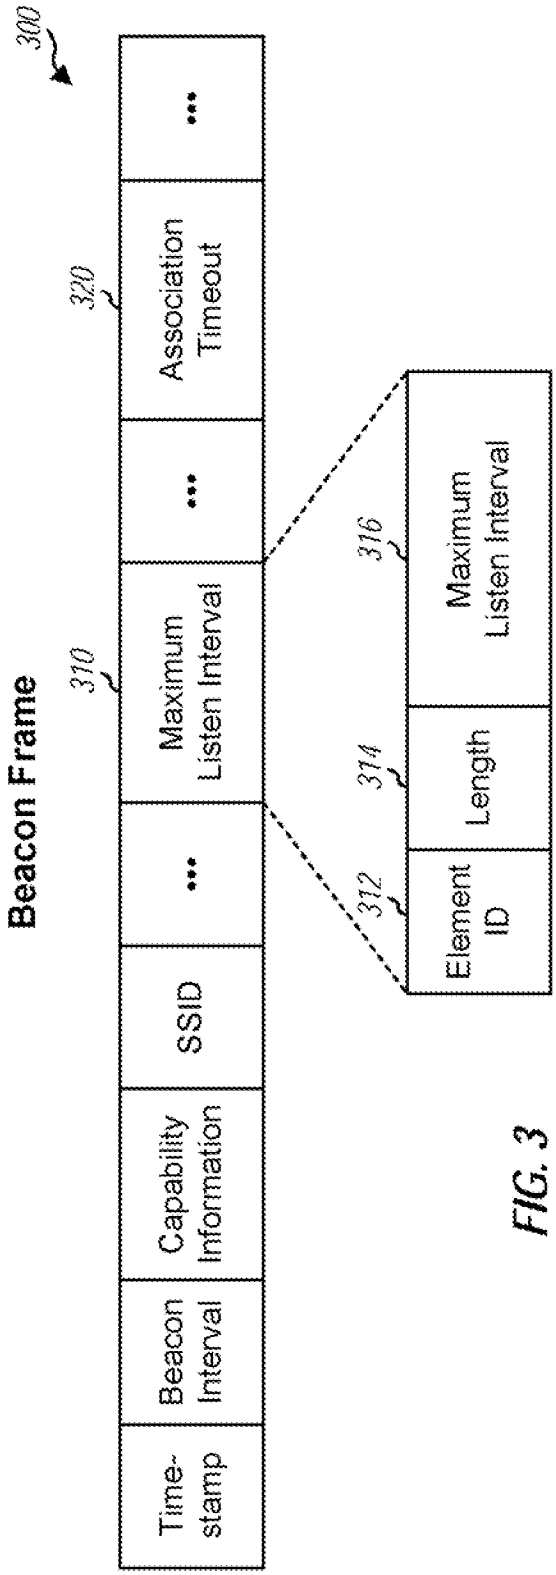
300
Beacon Frame
Time-stamp
Beacon Interval
Capability Information
SSID
Maximum Listen Interval
Association Timeout
310
320
312
314
316
Element ID
Length
Maximum Listen Interval

FIG. 3

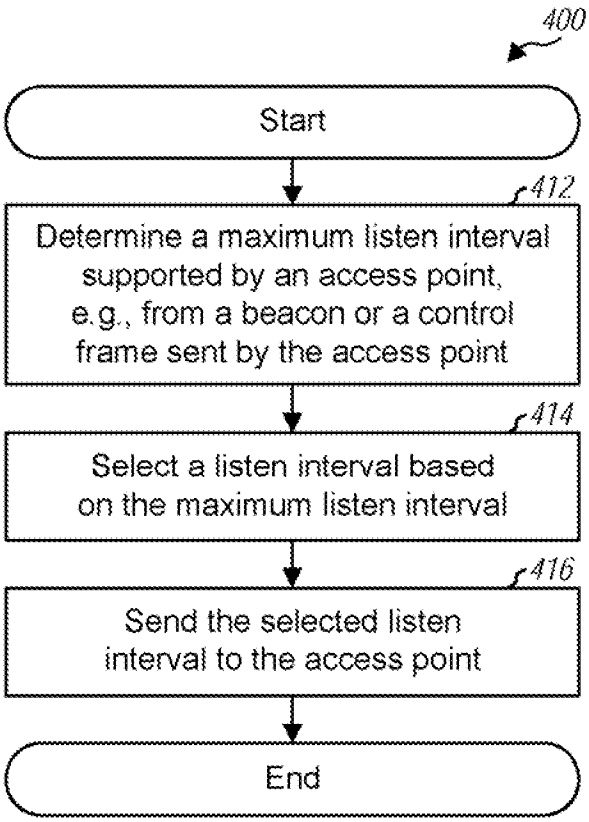
400
Start
412
Determine a maximum listen interval supported by an access point, e.g., from a beacon or a control frame sent by the access point
414
Select a listen interval based on the maximum listen interval
416
Send the selected listen interval to the access point
End

FIG. 4

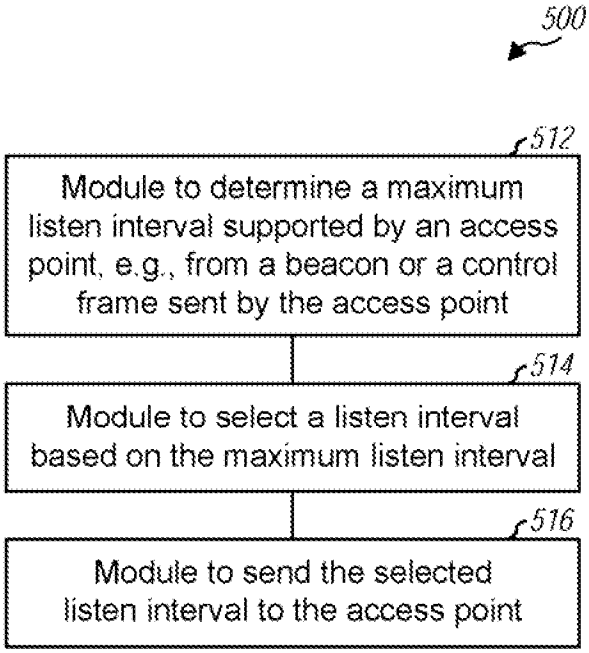
500
512
Module to determine a maximum listen interval supported by an access point, e.g., from a beacon or a control frame sent by the access point
514
Module to select a listen interval based on the maximum listen interval
516
Module to send the selected listen interval to the access point

FIG. 5

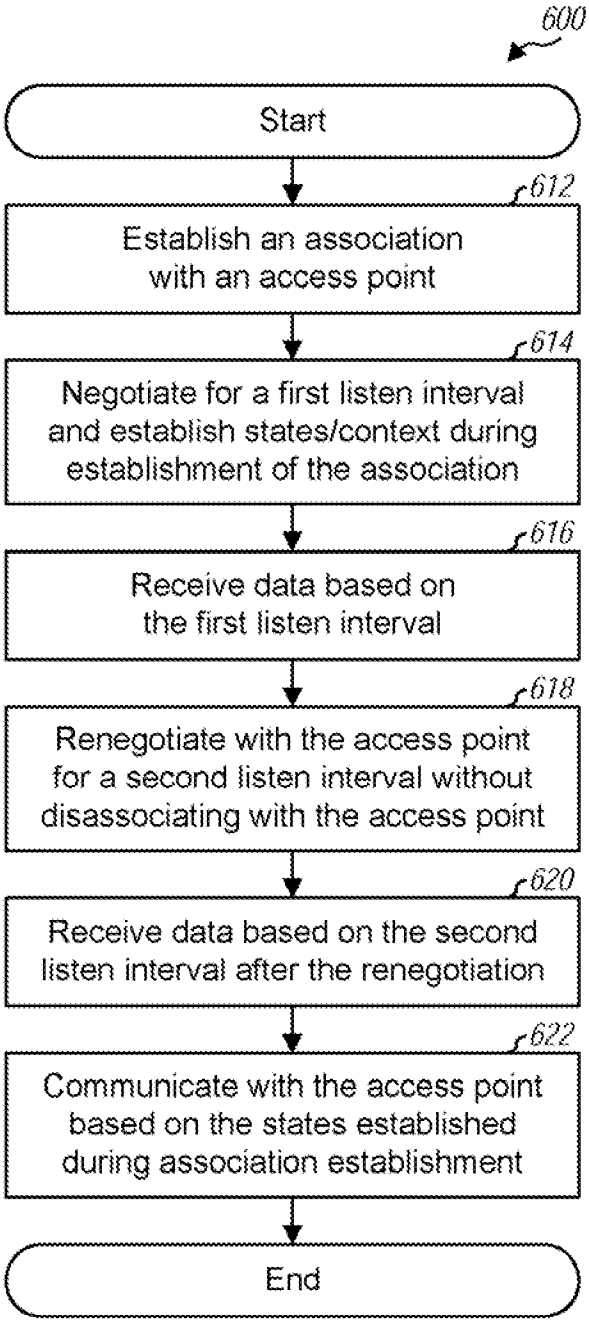
600
Start
612
Establish an association with an access point
614
Negotiate for a first listen interval and establish states/context during establishment of the association
616
Receive data based on the first listen interval
618
Renegotiate with the access point for a second listen interval without disassociating with the access point
620
Receive data based on the second listen interval after the renegotiation
622
Communicate with the access point based on the states established during association establishment
End

FIG. 6

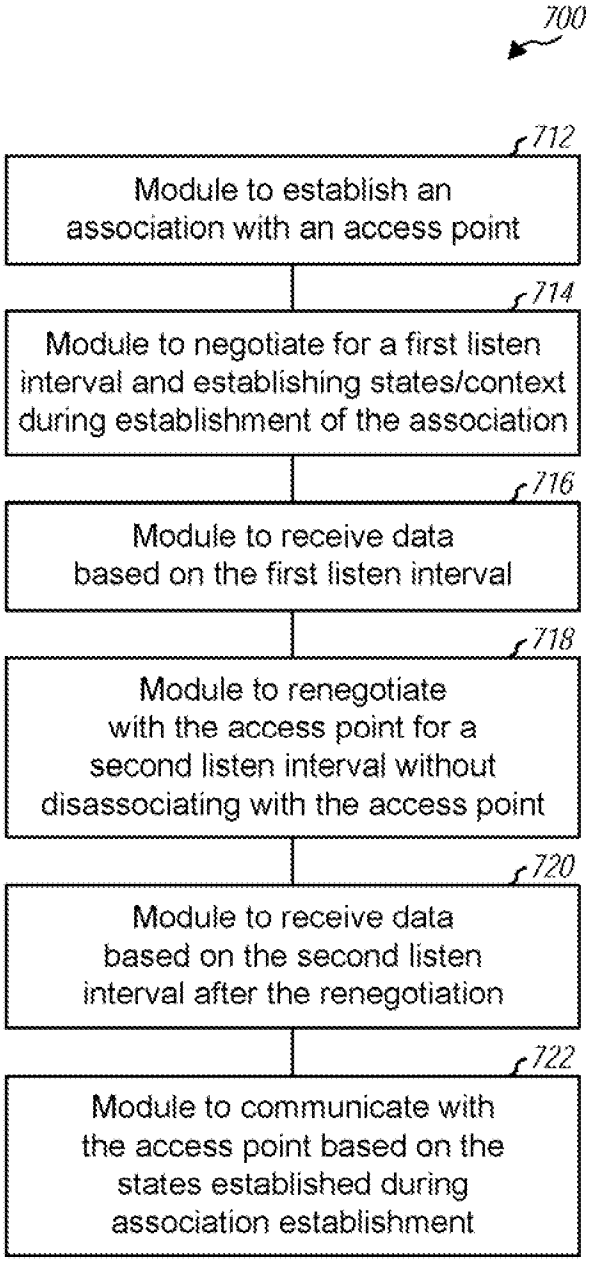
700
712
Module to establish an association with an access point
714
Module to negotiate for a first listen interval and establishing states/context during establishment of the association
716
Module to receive data based on the first listen interval
718
Module to renegotiate with the access point for a second listen interval without disassociating with the access point
720
Module to receive data based on the second listen interval after the renegotiation
722
Module to communicate with the access point based on the states established during association establishment

FIG. 7

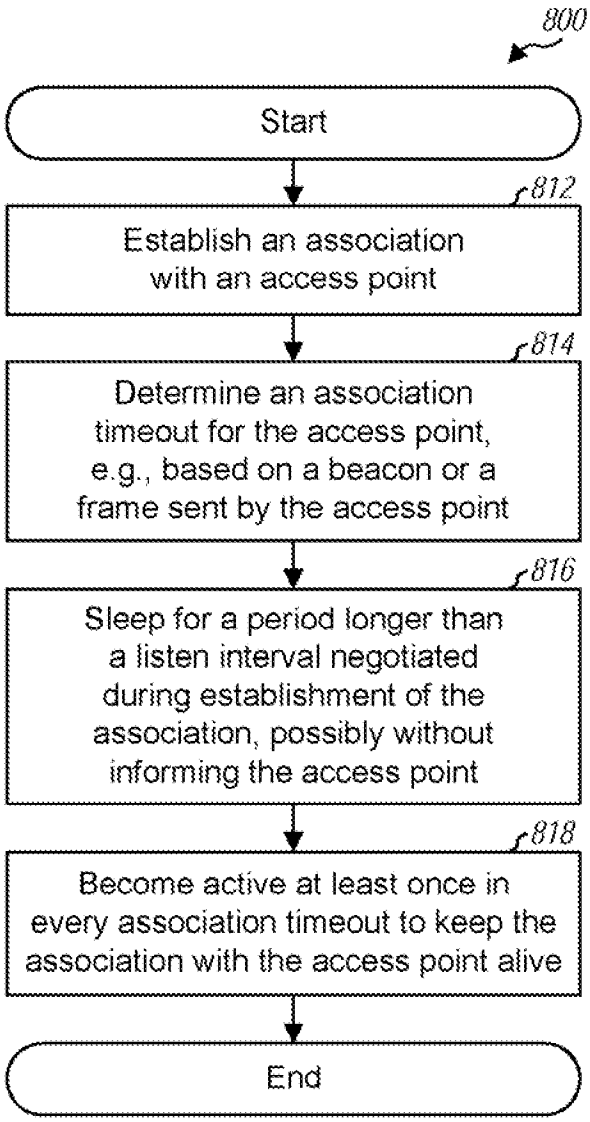
800
Start
812
Establish an association with an access point
814
Determine an association timeout for the access point, e.g., based on a beacon or a frame sent by the access point
816
Sleep for a period longer than a listen interval negotiated during establishment of the association, possibly without informing the access point
818
Become active at least once in every association timeout to keep the association with the access point alive
End

FIG. 8

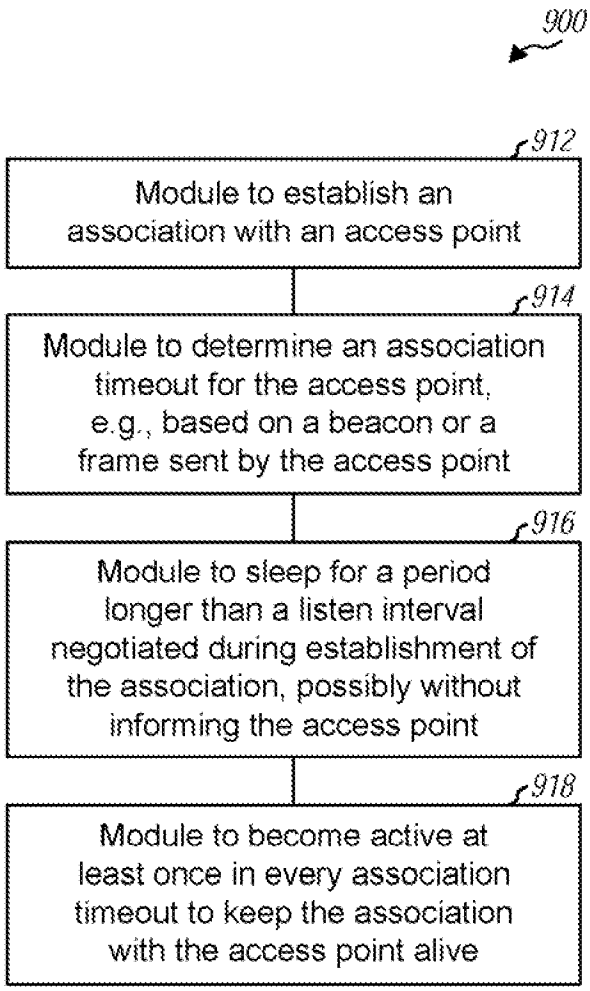
900
912
Module to establish an association with an access point
914
Module to determine an association timeout for the access point, e.g., based on a beacon or a frame sent by the access point
916
Module to sleep for a period longer than a listen interval negotiated during establishment of the association, possibly without informing the access point
918
Module to become active at least once in every association timeout to keep the association with the access point alive

FIG. 9

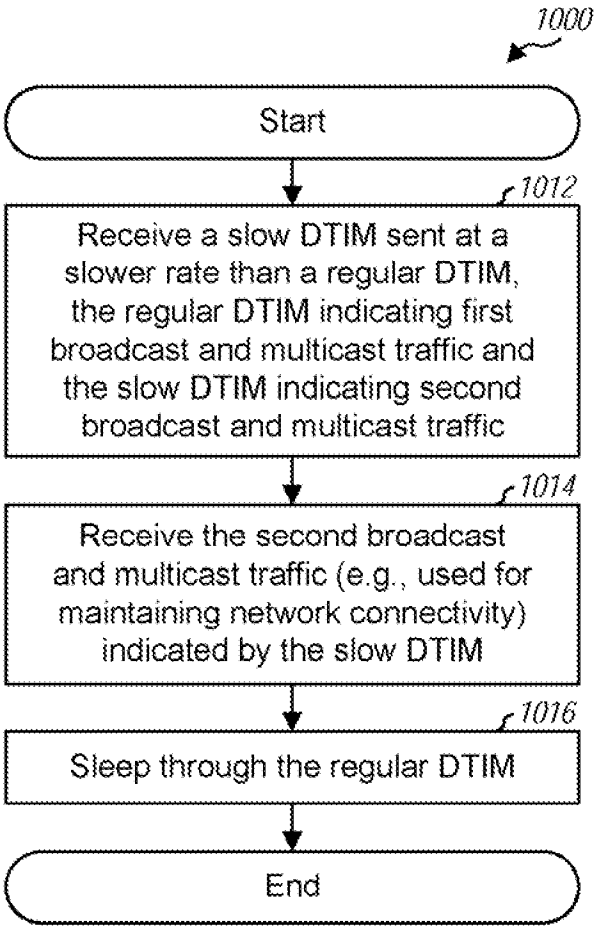
1000
Start
1012
Receive a slow DTIM sent at a slower rate than a regular DTIM, the regular DTIM indicating first broadcast and multicast traffic and the slow DTIM indicating second broadcast and multicast traffic
1014
Receive the second broadcast and multicast traffic (e.g., used for maintaining network connectivity) indicated by the slow DTIM
1016
Sleep through the regular DTIM
End

FIG. 10

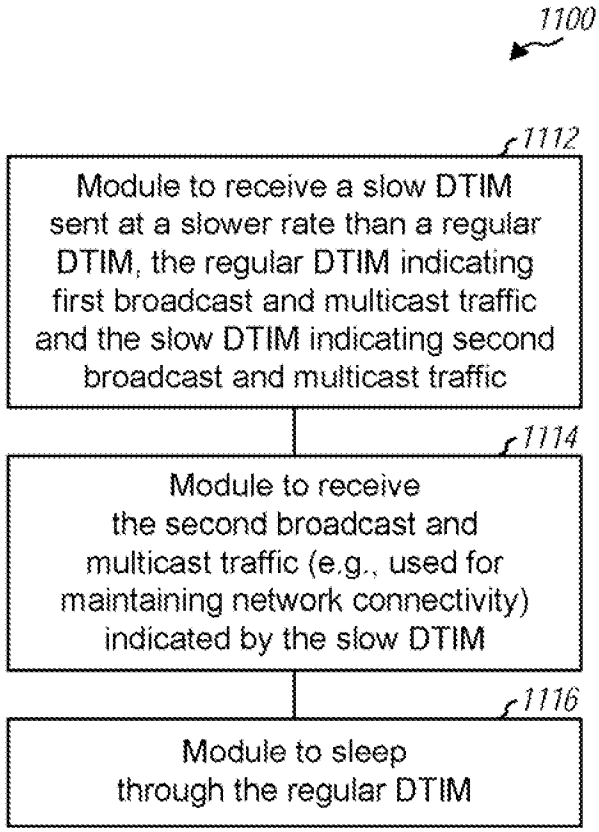
1100
1112
Module to receive a slow DTIM sent at a slower rate than a regular DTIM, the regular DTIM indicating first broadcast and multicast traffic and the slow DTIM indicating second broadcast and multicast traffic
1114
Module to receive the second broadcast and multicast traffic (e.g., used for maintaining network connectivity) indicated by the slow DTIM
1116
Module to sleep through the regular DTIM

FIG. 11

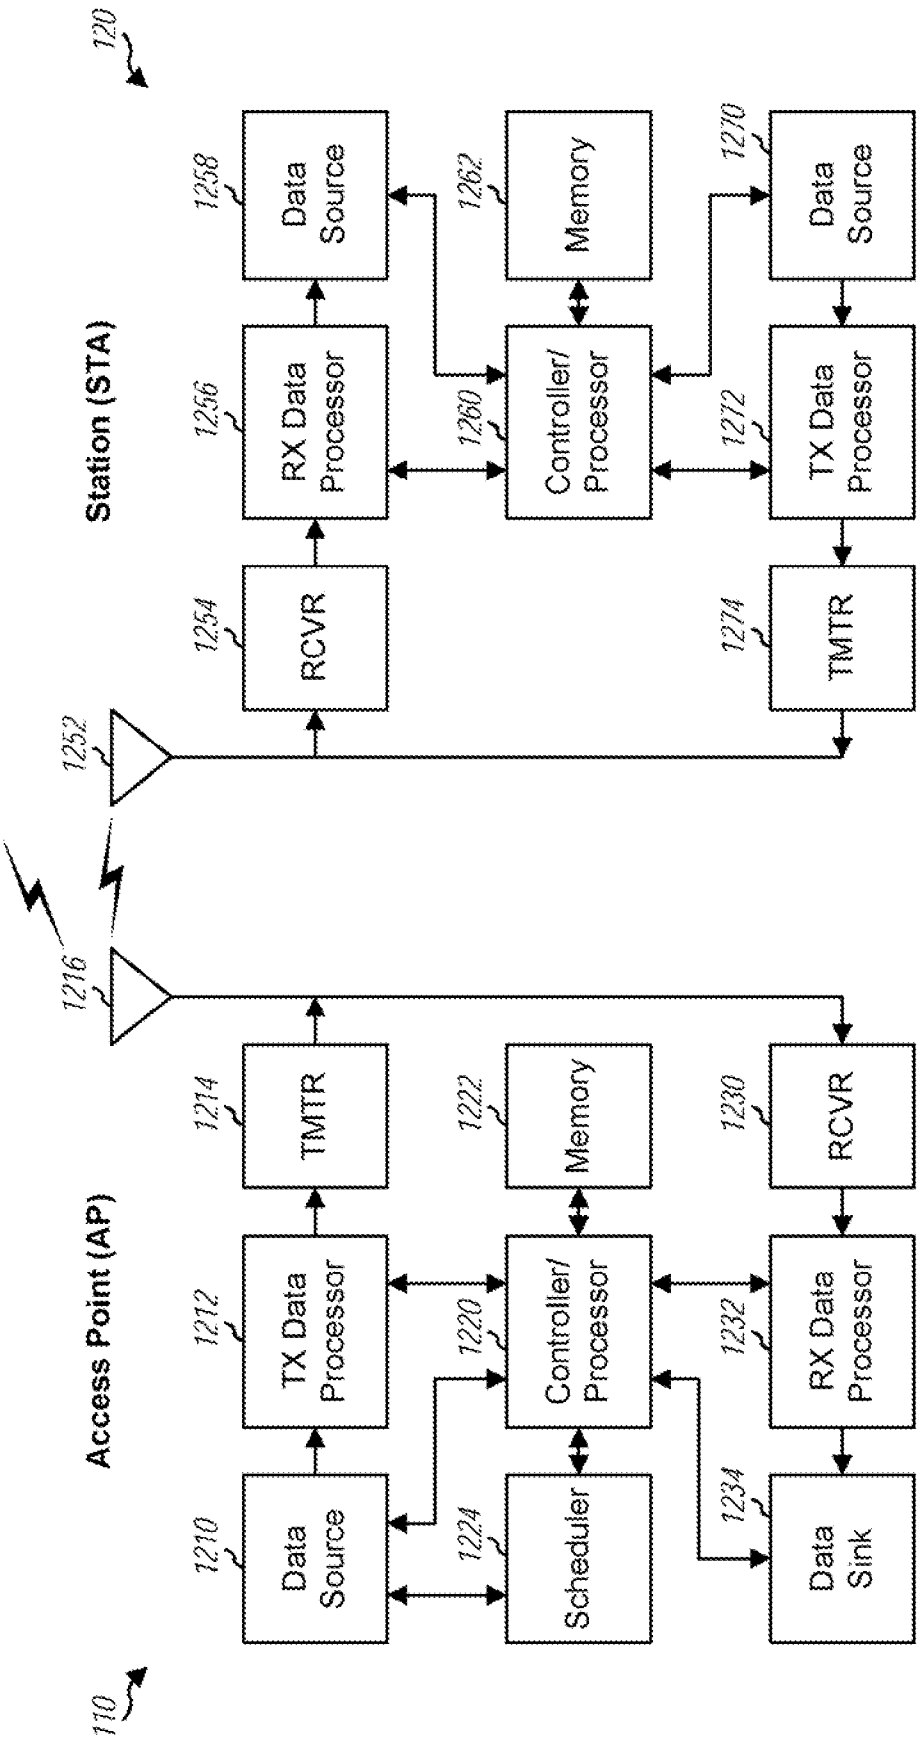
120
Station (STA)
1258
Data Source
1262
Memory
1270
Data Source
1256
RX Data Processor
1260
Controller/ Processor
1272
TX Data Processor
1254
RCVR
1274
TMTR
1252
1216
Access Point (AP)
1214
TMTR
1222
Memory
1230
RCVR
1212
TX Data Processor
1220
Controller/ Processor
1232
RX Data Processor
1210
Data Source
1224
Scheduler
1234
Data Sink
110

FIG. 12