

ENERGY MANAGEMENT IN WIRELESS NETWORKS

CROSS-REFERENCE TO RELATED APPLICATION(S)

[0001] This application claims priority to U.S. Patent Application No. 18/477,043, filed September 28, 2023, which is assigned to the assignee hereof and hereby expressly incorporated by reference in its entirety as if fully set forth below and for all applicable purposes.

TECHNICAL FIELD

[0002] This disclosure relates generally to wireless communication, and more specifically, to techniques for power savings mechanisms in wireless networks.

DESCRIPTION OF THE RELATED TECHNOLOGY

[0003] A wireless local area network (WLAN) may be formed by one or more wireless access points (APs) that provide a shared wireless communication medium for use by multiple client devices also referred to as wireless stations (STAs). The basic building block of a WLAN conforming to the Institute of Electrical and Electronics Engineers (IEEE) 802.11 family of standards is a Basic Service Set (BSS), which is managed by an AP. Each BSS is identified by a Basic Service Set Identifier (BSSID) that is advertised by the AP. An AP periodically broadcasts beacon frames to enable any STAs within wireless range of the AP to establish or maintain a communication link with the WLAN.

SUMMARY

[0004] The systems, methods and devices of this disclosure each have several innovative aspects, no single one of which is solely responsible for the desirable attributes disclosed herein.

[0005] One innovative aspect of the subject matter described in this disclosure may be implemented as a method for wireless communications at a first wireless node. The method includes obtaining information regarding a topology of a mesh network; and performing, based on the information, one or more actions to adjust power consumption at one or more wireless nodes of the mesh network.

[0006] Another innovative aspect of the subject matter described in this disclosure may be implemented at an apparatus for wireless communications. The apparatus includes at least one processor; memory coupled with the processor; and instructions stored in the memory and executable by the processor to cause the apparatus to: obtain

information regarding a topology of a mesh network; and perform, based on the information, one or more actions to adjust power consumption at one or more wireless nodes of the mesh network.

[0007] In some examples, the one or more actions comprise steering one or more clients away from a first wireless node of the mesh network to a second wireless node of the mesh network.

[0008] Other aspects provide: an apparatus operable, configured, or otherwise adapted to perform any one or more of the aforementioned methods and/or those described elsewhere herein; a non-transitory, computer-readable media comprising instructions that, when executed by a processor of an apparatus, cause the apparatus to perform the aforementioned methods as well as those described elsewhere herein; a computer program product embodied on a computer-readable storage medium comprising code for performing the aforementioned methods as well as those described elsewhere herein; and/or an apparatus comprising means for performing the aforementioned methods as well as those described elsewhere herein. By way of example, an apparatus may comprise a processing system, a device with a processing system, or processing systems cooperating over one or more networks.

[0009] Details of one or more implementations of the subject matter described in this disclosure are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages will become apparent from the description, the drawings and the claims. Note that the relative dimensions of the following figures may not be drawn to scale.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] Figure 1 shows a pictorial diagram of an example wireless communication network.

[0011] Figure 2 shows an example protocol data unit (PDU) usable for communications between a wireless access point (AP) and one or more wireless stations (STAs).

[0012] Figure 3 shows a hierarchical format of an example physical layer PDU (PPDU) usable for communications between a wireless AP and one or more wireless STAs.

[0013] Figure 4 shows a pictorial diagram of another example wireless communication network.

[0014] Figure 5 shows an example scenario in which power saving mechanisms proposed herein may be utilized.

[0015] Figure 6 shows an example mesh network in which power saving mechanisms proposed herein may be utilized.

[0016] Figure 7 shows an example architecture that may implement power saving mechanisms proposed herein.

[0017] Figure 8 shows an example scenario with centralized energy management, in accordance with aspects of the present disclosure.

[0018] Figures 9-14 show example power savings mechanisms using centralized energy management, in accordance with aspects of the present disclosure.

[0019] Figures 15 and 16 show diagrams illustrating examples of power saving mechanisms, in accordance with certain aspects of the present disclosure.

[0020] Figure 17 shows a flowchart illustrating an example process performable by a first wireless node configured as an AP.

[0021] Figure 18 shows a block diagram of an example wireless communication device.

[0022] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0023] The following description is directed to some particular examples for the purposes of describing innovative aspects of this disclosure. However, a person having ordinary skill in the art will readily recognize that the teachings herein can be applied in a multitude of different ways. Some or all of the described examples may be implemented in any device, system or network that is capable of transmitting and receiving radio frequency (RF) signals according to one or more of the Institute of Electrical and Electronics Engineers (IEEE) 802.11 standards, the IEEE 802.15 standards, the Bluetooth® standards as defined by the Bluetooth Special Interest Group (SIG), or the Long Term Evolution (LTE), 3G, 4G or 5G (New Radio (NR)) standards promulgated by the 3rd Generation Partnership Project (3GPP), among others. The described examples can be implemented in any device, system or network that is capable of transmitting and receiving RF signals according to one or more of the following technologies or techniques: code division multiple access (CDMA), time division multiple access (TDMA), frequency division multiple access (FDMA), orthogonal FDMA (OFDMA), single-carrier FDMA (SC-FDMA), spatial division multiple access (SDMA), rate-

splitting multiple access (RSMA), multi-user shared access (MUSA), single-user (SU) multiple-input multiple-output (MIMO) and multi-user (MU)-MIMO. The described examples also can be implemented using other wireless communication protocols or RF signals suitable for use in one or more of a wireless personal area network (WPAN), a wireless local area network (WLAN), a wireless wide area network (WWAN), a wireless metropolitan area network (WMAN), or an internet of things (IOT) network.

[0024] Various aspects relate generally to wireless communication. Some aspects more specifically relate to power savings mechanisms in wireless networks. In some cases, the power savings mechanisms proposed herein may be used to activate/de-activate links between wireless nodes in wireless networks. As used herein, “wireless node” generally refers to any type of device capable of communicating wirelessly, such as an access point (AP), a wireless station (STA) serving as an AP (an AP-STA), a STA that is not serving as an AP (a non-AP STA), or a user equipment (UE).

[0025] Advances in wireless technology often come with some form of cost. For example, features added in each new Wi-Fi generation tend to add a greater energy toll on the access point (AP). For example, WiFi-6e (an enhanced version of 802.11ax) added an additional link (6GHz), WiFi-7 (802.11be) added multiple concurrent links to support multilink operation (MLO), and WiFi-8 plans to add even higher bandwidth (e.g., 640MHz or more).

[0026] This additional power, in terms of operating watts, increases cost both to build the product and operate the product. For example, potentially expensive heat sinks and power-adapters used to accommodate the increased power may end up dominating the bill of material (BOM) cost, resulting in increased product cost and/or reduced profit margins.

[0027] The higher energy cost may be felt by the end user, which may impact purchasing decisions. As illustrated in the example 500 of Figure 5, a typical home may include a root AP 502 and multiple repeaters 512, forming a mesh network. Each such node may draw significant power when in an active state, increasing consumer energy bills. In some cases, the network energy cost may approach or exceed that of an appliance. Further, increased energy adds to the global carbon footprint in scenarios, for example, with always-on high output power multi-radio nodes, multi-nodes mesh networks, and the like.

[0028] One approach to energy management is de-centralized, with individual nodes (e.g., APs serving as repeaters) performing actions to reduce power consumption. For

example, each individual AP may calculate metrics per-link and perform actions to reduce power consumption based on the metrics. As an example, based on metrics indicative of link loading, an AP could activate/de-activate certain links or ports and/or adjust parameters such as operating bandwidth (BW) and modulation and coding schemes (MCS).

[0029] In some cases, applying energy management in individual nodes may be less than optimal for a mesh network. For example, mesh networks in a home environment may involve coordination between nodes. As such, mesh networks may benefit from topology awareness to make effective energy management/saving decisions that consider the mesh network as a whole (e.g., to achieve optimal whole home coverage).

[0030] Aspects of the present disclosure, however, provide a centralized energy management that may consider a topology of a mesh network when making power savings decisions. Such a centralized energy management approach may utilize power savings mechanisms that may help mitigate the increase in power consumption associated with advanced network features.

[0031] According to certain aspects, a first wireless node (e.g., an AP serving as a root node or a controller entity) may obtain information regarding a topology of a mesh network that includes a number of repeaters (APs). The information may include, for example, capability/status of repeaters, information regarding clients connected to the repeaters, position/location of the repeaters, and/or position/location of the clients. Based on the information, the first wireless node may perform one or more actions to adjust power consumption at one or more wireless nodes of the mesh network.

[0032] As a result, aspects of the present disclosure may help reduce overall power consumption in a network. The power savings mechanisms proposed herein may help reduce energy costs associated with operating the network. The power savings mechanisms proposed herein may also allow help avoid costly parts (e.g., expensive heat sinks and power-adapters), which may help reduce the build cost.

EXAMPLE WIRELESS COMMUNICATION NETWORK

[0033] Figure 1 shows a block diagram of an example wireless communication network 100. According to some aspects, the wireless communication network 100 can be an example of a wireless local area network (WLAN) such as a Wi-Fi network (and will hereinafter be referred to as WLAN 100). For example, the WLAN 100 can be a network implementing at least one of the IEEE 802.11 family of wireless communication

protocol standards (such as that defined by the IEEE 802.11-2020 specification or amendments thereof including, but not limited to, 802.11ay, 802.11ax, 802.11az, 802.11ba, 802.11bd, 802.11be, 802.11bf, and the 802.11 amendment associated with Wi-Fi 8). The WLAN 100 may include numerous wireless communication devices such as a wireless access point (AP) 102 and multiple wireless stations (STAs) 104. While only one AP 102 is shown in Figure 1, the WLAN network 100 also can include multiple APs 102. AP 102 shown in Figure 1 can represent various different types of APs including but not limited to enterprise-level APs, single-frequency APs, dual-band APs, standalone APs, software-enabled APs (soft APs), and multi-link APs. The coverage area and capacity of a cellular network (such as LTE, 5G NR, etc.) can be further improved by a small cell which is supported by an AP serving as a miniature base station. Furthermore, private cellular networks also can be set up through a wireless area network using small cells.

[0034] Each of the STAs 104 also may be referred to as a mobile station (MS), a mobile device, a mobile handset, a wireless handset, an access terminal (AT), a user equipment (UE), a subscriber station (SS), or a subscriber unit, among other examples. The STAs 104 may represent various devices such as mobile phones, personal digital assistant (PDAs), other handheld devices, netbooks, notebook computers, tablet computers, laptops, chromebooks, extended reality (XR) headsets, wearable devices, display devices (for example, TVs (including smart TVs), computer monitors, navigation systems, among others), music or other audio or stereo devices, remote control devices (“remotes”), printers, kitchen appliances (including smart refrigerators) or other household appliances, key fobs (for example, for passive keyless entry and start (PKES) systems), Internet of Things (IoT) devices, and vehicles, among other examples. The various STAs 104 in the network are able to communicate with one another via the AP 102.

[0035] A single AP 102 and an associated set of STAs 104 may be referred to as a basic service set (BSS), which is managed by the respective AP 102. Figure 1 additionally shows an example coverage area 108 of the AP 102, which may represent a basic service area (BSA) of the WLAN 100. The BSS may be identified or indicated to users by a service set identifier (SSID), as well as to other devices by a basic service set identifier (BSSID), which may be a medium access control (MAC) address of the AP 102. The AP 102 may periodically broadcast beacon frames (“beacons”) including the BSSID to enable any STAs 104 within wireless range of the AP 102 to “associate” or re-associate with the

AP 102 to establish a respective communication link 106 (hereinafter also referred to as a “Wi-Fi link”), or to maintain a communication link 106, with the AP 102. For example, the beacons can include an identification or indication of a primary channel used by the respective AP 102 as well as a timing synchronization function for establishing or maintaining timing synchronization with the AP 102. The AP 102 may provide access to external networks to various STAs 104 in the WLAN via respective communication links 106.

[0036] To establish a communication link 106 with an AP 102, each of the STAs 104 is configured to perform passive or active scanning operations (“scans”) on frequency channels in one or more frequency bands (for example, the 2.4 GHz, 5 GHz, 6 GHz or 60 GHz bands). To perform passive scanning, a STA 104 listens for beacons, which are transmitted by respective APs 102 at a periodic time interval referred to as the target beacon transmission time (TBTT) (measured in time units (TUs) where one TU may be equal to 1024 microseconds (μ s)). To perform active scanning, a STA 104 generates and sequentially transmits probe requests on each channel to be scanned and listens for probe responses from APs 102. Each STA 104 may identify, determine, ascertain, or select an AP 102 with which to associate in accordance with the scanning information obtained through the passive or active scans, and to perform authentication and association operations to establish a communication link 106 with the selected AP 102. The AP 102 assigns an association identifier (AID) to the STA 104 at the culmination of the association operations, which the AP 102 uses to track the STA 104.

[0037] As a result of the increasing ubiquity of wireless networks, a STA 104 may have the opportunity to select one of many BSSs within range of the STA or to select among multiple APs 102 that together form an extended service set (ESS) including multiple connected BSSs. An extended network station associated with the WLAN 100 may be connected to a wired or wireless distribution system that may allow multiple APs 102 to be connected in such an ESS. As such, a STA 104 can be covered by more than one AP 102 and can associate with different APs 102 at different times for different transmissions. Additionally, after association with an AP 102, a STA 104 also may periodically scan its surroundings to find a more suitable AP 102 with which to associate. For example, a STA 104 that is moving relative to its associated AP 102 may perform a “roaming” scan to find another AP 102 having more desirable network characteristics such as a greater received signal strength indicator (RSSI) or a reduced traffic load.

[0038] In some cases, STAs 104 may form networks without APs 102 or other equipment other than the STAs 104 themselves. One example of such a network is an ad hoc network (or wireless ad hoc network). Ad hoc networks may alternatively be referred to as mesh networks or peer-to-peer (P2P) networks. In some cases, ad hoc networks may be implemented within a larger wireless network such as the WLAN 100. In such examples, while the STAs 104 may be capable of communicating with each other through the AP 102 using communication links 106, STAs 104 also can communicate directly with each other via direct wireless communication links 110. Additionally, two STAs 104 may communicate via a direct communication link 110 regardless of whether both STAs 104 are associated with and served by the same AP 102. In such an ad hoc system, one or more of the STAs 104 may assume the role filled by the AP 102 in a BSS. Such a STA 104 may be referred to as a group owner (GO) and may coordinate transmissions within the ad hoc network. Examples of direct wireless communication links 110 include Wi-Fi Direct connections, connections established by using a Wi-Fi Tunneled Direct Link Setup (TDLS) link, and other P2P group connections.

[0039] The APs 102 and STAs 104 may function and communicate (via the respective communication links 106) according to one or more of the IEEE 802.11 family of wireless communication protocol standards. These standards define the WLAN radio and baseband protocols for the PHY and MAC layers. The APs 102 and STAs 104 transmit and receive wireless communications (hereinafter also referred to as “Wi-Fi communications” or “wireless packets”) to and from one another in the form of PHY protocol data units (PPDUs). The APs 102 and STAs 104 in the WLAN 100 may transmit PPDUs over an unlicensed spectrum, which may be a portion of spectrum that includes frequency bands traditionally used by Wi-Fi technology, such as the 2.4 GHz band, the 5 GHz band, the 60 GHz band, the 3.6 GHz band, and the 900 MHz band. Some examples of the APs 102 and STAs 104 described herein also may communicate in other frequency bands, such as the 5.9 GHz and the 6 GHz bands, which may support both licensed and unlicensed communications. The APs 102 and STAs 104 also can communicate over other frequency bands such as shared licensed frequency bands, where multiple operators may have a license to operate in the same or overlapping frequency band or bands.

[0040] Each of the frequency bands may include multiple sub-bands or frequency channels. For example, PPDUs conforming to the IEEE 802.11n, 802.11ac, 802.11ax and 802.11be standard amendments may be transmitted over the 2.4 GHz, 5 GHz or 6 GHz bands, each of which is divided into multiple 20 MHz channels. As such, these PPDUs

are transmitted over a physical channel having a minimum bandwidth of 20 MHz, but larger channels can be formed through channel bonding. For example, PPDU may be transmitted over physical channels having bandwidths of 40 MHz, 80 MHz, 160 or 320 MHz by bonding together multiple 20 MHz channels.

[0041] Each PPDU is a composite structure that includes a PHY preamble and a payload in the form of a PHY service data unit (PSDU). The information provided in the preamble may be used by a receiving device to decode the subsequent data in the PSDU. In instances in which PPDU are transmitted over a bonded channel, the preamble fields may be duplicated and transmitted in each of the multiple component channels. The PHY preamble may include both a legacy portion (or “legacy preamble”) and a non-legacy portion (or “non-legacy preamble”). The legacy preamble may be used for packet detection, automatic gain control and channel estimation, among other uses. The legacy preamble also may generally be used to maintain compatibility with legacy devices. The format of, coding of, and information provided in the non-legacy portion of the preamble is associated with the particular IEEE 802.11 protocol to be used to transmit the payload

[0042] Figure 2 shows an example protocol data unit (PDU) 200 usable for wireless communication between a wireless AP 102 and one or more wireless STAs 104. For example, the PDU 200 can be configured as a PPDU. As shown, the PDU 200 includes a PHY preamble 202 and a PHY payload 204. For example, the preamble 202 may include a legacy portion that itself includes a legacy short training field (L-STF) 206, which may consist of two symbols, a legacy long training field (L-LTF) 208, which may consist of two symbols, and a legacy signal field (L-SIG) 210, which may consist of two symbols. The legacy portion of the preamble 202 may be configured according to the IEEE 802.11a wireless communication protocol standard. The preamble 202 also may include a non-legacy portion including one or more non-legacy fields 212, for example, conforming to one or more of the IEEE 802.11 family of wireless communication protocol standards.

[0043] The L-STF 206 generally enables a receiving device to perform coarse timing and frequency tracking and automatic gain control (AGC). The L-LTF 208 generally enables a receiving device to perform fine timing and frequency tracking and also to perform an initial estimate of the wireless channel. The L-SIG 210 generally enables a receiving device to determine (for example, obtain, select, identify, detect, ascertain, calculate, or compute) a duration of the PDU and to use the determined duration to avoid transmitting on top of the PDU. The legacy portion of the preamble, including the L-STF 206, the L-LTF 208 and the L-SIG 210, may be modulated according to a binary phase

shift keying (BPSK) modulation scheme. The payload 204 may be modulated according to a BPSK modulation scheme, a quadrature BPSK (Q-BPSK) modulation scheme, a quadrature amplitude modulation (QAM) modulation scheme, or another appropriate modulation scheme. The payload 204 may include a PSDU including a data field (DATA) 214 that, in turn, may carry higher layer data, for example, in the form of MAC protocol data units (MPDUs) or an aggregated MPDU (A-MPDU).

[0044] Figure 3 shows a hierarchical format of an example PPDU usable for communications between a wireless AP 102 and one or more wireless STAs 104. As described, each PPDU 300 includes a PHY preamble 302 and a PSDU 304. Each PSDU 304 may represent (or “carry”) one or more MAC protocol data units (MPDUs) 316. For example, each PSDU 304 may carry an aggregated MPDU (A-MPDU) 306 that includes an aggregation of multiple A-MPDU subframes 308. Each A-MPDU subframe 306 may include an MPDU frame 310 that includes a MAC delimiter 312 and a MAC header 314 prior to the accompanying MPDU 316, which includes the data portion (“payload” or “frame body”) of the MPDU frame 310. Each MPDU frame 310 also may include a frame check sequence (FCS) field 318 for error detection (for example, the FCS field may include a cyclic redundancy check (CRC)) and padding bits 320. The MPDU 316 may carry one or more MAC service data units (MSDUs) 316. For example, the MPDU 316 may carry an aggregated MSDU (A-MSDU) 322 including multiple A-MSDU subframes 324. Each A-MSDU subframe 324 contains a corresponding MSDU 330 preceded by a subframe header 328 and in some cases followed by padding bits 332.

[0045] Referring back to the MPDU frame 310, the MAC delimiter 312 may serve as a marker of the start of the associated MPDU 316 and indicate the length of the associated MPDU 316. The MAC header 314 may include multiple fields containing information that defines or indicates characteristics or attributes of data encapsulated within the frame body 316. The MAC header 314 includes a duration field indicating a duration extending from the end of the PPDU until at least the end of an acknowledgment (ACK) or Block ACK (BA) of the PPDU that is to be transmitted by the receiving wireless communication device. The use of the duration field serves to reserve the wireless medium for the indicated duration, and enables the receiving device to establish its network allocation vector (NAV). The MAC header 314 also includes one or more fields indicating addresses for the data encapsulated within the frame body 316. For example, the MAC header 314 may include a combination of a source address, a transmitter address, a receiver address or a destination address. The MAC header 314 may further include a frame control field

containing control information. The frame control field may specify a frame type, for example, a data frame, a control frame, or a management frame.

[0046] Some APs and STAs may implement techniques for spatial reuse that involve participation in a coordinated communication scheme. According to such techniques, an AP may contend for access to a wireless medium to obtain control of the medium for a TXOP. The AP that wins the contention (hereinafter also referred to as a “sharing AP”) may select one or more other APs (hereinafter also referred to as “shared APs”) to share resources of the TXOP. The sharing and shared APs may be located in proximity to one another such that at least some of their wireless coverage areas at least partially overlap. Some examples may specifically involve coordinated AP TDMA or OFDMA techniques for sharing the time or frequency resources of a TXOP. To share its time or frequency resources, the sharing AP may partition the TXOP into multiple time segments or frequency segments each including respective time or frequency resources representing a portion of the TXOP. The sharing AP may allocate the time or frequency segments to itself or to one or more of the shared APs. For example, each shared AP may utilize a partial TXOP assigned by the sharing AP for its uplink or downlink communications with its associated STAs.

[0047] In some examples of such TDMA techniques, each portion of a plurality of portions of the TXOP includes a set of time resources that do not overlap with any time resources of any other portion of the plurality of portions. In such examples, the scheduling information may include an indication of time resources, of multiple time resources of the TXOP, associated with each portion of the TXOP. For example, the scheduling information may include an indication of a time segment of the TXOP such as an indication of one or more slots or sets of symbol periods associated with each portion of the TXOP such as for multi-user TDMA.

[0048] In some other examples of OFDMA techniques, each portion of the plurality of portions of the TXOP includes a set of frequency resources that do not overlap with any frequency resources of any other portion of the plurality of portions. In such implementations, the scheduling information may include an indication of frequency resources, of multiple frequency resources of the TXOP, associated with each portion of the TXOP. For example, the scheduling information may include an indication of a bandwidth portion of the wireless channel such as an indication of one or more subchannels or resource units (RUs) associated with each portion of the TXOP such as for multi-user OFDMA.

[0049] In this manner, the sharing AP's acquisition of the TXOP enables communication between one or more additional shared APs and their respective BSSs, subject to appropriate power control and link adaptation. For example, the sharing AP may limit the transmit powers of the selected shared APs such that interference from the selected APs does not prevent STAs associated with the TXOP owner from successfully decoding packets transmitted by the sharing AP. Such techniques may be used to reduce latency because the other APs may not need to wait to win contention for a TXOP to be able to transmit and receive data according to conventional CSMA/CA or EDCA techniques. Additionally, by enabling a group of APs associated with different BSSs to participate in a coordinated AP transmission session, during which the group of APs may share at least a portion of a single TXOP obtained by any one of the participating APs, such techniques may increase throughput across the BSSs associated with the participating APs and may also achieve improvements in throughput fairness. Furthermore, with appropriate selection of the shared APs and the scheduling of their respective time or frequency resources, medium utilization may be maximized or otherwise increased while packet loss resulting from OBSS interference is minimized or otherwise reduced. Various implementations may achieve these and other advantages without requiring that the sharing AP or the shared APs be aware of the STAs associated with other BSSs, without requiring a preassigned or dedicated master AP or preassigned groups of APs, and without requiring backhaul coordination between the APs participating in the TXOP.

[0050] In some examples in which the signal strengths or levels of interference associated with the selected APs are relatively low (such as less than a given value), or when the decoding error rates of the selected APs are relatively low (such as less than a threshold), the start times of the communications among the different BSSs may be synchronous. Conversely, when the signal strengths or levels of interference associated with the selected APs are relatively high (such as greater than the given value), or when the decoding error rates of the selected APs are relatively high (such as greater than the threshold), the start times may be offset from one another by a time period associated with decoding the preamble of a wireless packet and determining, from the decoded preamble, whether the wireless packet is an intra-BSS packet or is an OBSS packet. For example, the time period between the transmission of an intra-BSS packet and the transmission of an OBSS packet may allow a respective AP (or its associated STAs) to decode the preamble of the wireless packet and obtain the BSS color value carried in the wireless

packet to determine whether the wireless packet is an intra-BSS packet or an OBSS packet. In this manner, each of the participating APs and their associated STAs may be able to receive and decode intra-BSS packets in the presence of OBSS interference.

[0051] In some examples, the sharing AP may perform polling of a set of un-managed or non-co-managed APs that support coordinated reuse to identify candidates for future spatial reuse opportunities. For example, the sharing AP may transmit one or more spatial reuse poll frames as part of determining one or more spatial reuse criteria and selecting one or more other APs to be shared APs. According to the polling, the sharing AP may receive responses from one or more of the polled APs. In some specific examples, the sharing AP may transmit a coordinated AP TXOP indication (CTI) frame to other APs that indicates time and frequency of resources of the TXOP that can be shared. The sharing AP may select one or more candidate APs upon receiving a coordinated AP TXOP request (CTR) frame from a respective candidate AP that indicates a desire by the respective AP to participate in the TXOP. The poll responses or CTR frames may include a power indication, for example, an RX power or RSSI measured by the respective AP. In some other examples, the sharing AP may directly measure potential interference of a service supported (such as UL transmission) at one or more APs, and select the shared APs based on the measured potential interference. The sharing AP generally selects the APs to participate in coordinated spatial reuse such that it still protects its own transmissions (which may be referred to as primary transmissions) to and from the STAs in its BSS. The selected APs may then be allocated resources during the TXOP as described above.

[0052] Retransmission protocols, such as hybrid automatic repeat request (HARQ), also may offer performance gains. A HARQ protocol may support various HARQ signaling between transmitting and receiving wireless communication devices as well as signaling between the PHY and MAC layers to improve the retransmission operations in a WLAN. HARQ uses a combination of error detection and error correction. For example, a HARQ transmission may include error checking bits that are added to data to be transmitted using an error-detecting (ED) code, such as a cyclic redundancy check (CRC). The error checking bits may be used by the receiving device to determine if it has properly decoded the received HARQ transmission. In some examples, the original data (information bits) to be transmitted may be encoded with a forward error correction (FEC) code, such as using a low-density parity check (LDPC) coding scheme that systematically encodes the information bits to produce parity bits. The transmitting device may transmit

both the original information bits as well as the parity bits in the HARQ transmission to the receiving device. The receiving device may be able to use the parity bits to correct errors in the information bits, thus avoiding a retransmission.

[0053] Implementing a HARQ protocol in a WLAN may improve reliability of data communicated from a transmitting device to a receiving device. The HARQ protocol may support the establishment of a HARQ session between the two devices. Once a HARQ session is established, If a receiving device cannot properly decode (and cannot correct the errors) a first HARQ transmission received from the transmitting device, the receiving device may transmit a HARQ feedback message to the transmitting device (for example, a negative acknowledgement (NACK)) that indicates at least part of the first HARQ transmission was not properly decoded. Such a HARQ feedback message may be different than the traditional Block ACK feedback message type associated with conventional ARQ. In response to receiving the HARQ feedback message, the transmitting device may transmit a second HARQ transmission to the receiving device to communicate at least part of further assist the receiving device in decoding the first HARQ transmission. For example, the transmitting device may include some or all of the original information bits, some or all of the original parity bits, as well as other, different parity bits in the second HARQ transmission. The combined HARQ transmissions may be processed for decoding and error correction such that the complete signal associated with the HARQ transmissions can be obtained.

[0054] In some examples, the receiving device may be enabled to control whether to continue the HARQ process or revert to a non-HARQ retransmission scheme (such as an ARQ protocol). Such switching may reduce feedback overhead and increase the flexibility for retransmissions by allowing devices to dynamically switch between ARQ and HARQ protocols during frame exchanges. Some implementations also may allow multiplexing of communications that employ ARQ with those that employ HARQ.

[0055] Some wireless communication devices (including both APs and STAs) are capable of multi-link operation (MLO). In some examples, MLO supports establishing multiple different communication links (such as a first link on the 2.4 GHz band, a second link on the 5 GHz band, and the third link on the 6 GHz band) between the STA and the AP. Each communication link may support one or more sets of channels or logical entities. In some cases, each communication link associated with a given wireless communication device may be associated with a respective radio of the wireless communication device, which may include one or more transmit/receive (Tx/Rx) chains,

include or be coupled with one or more physical antennas, or include signal processing components, among other components. An MLO-capable device may be referred to as a multi-link device (MLD). For example, an AP MLD may include multiple APs each configured to communicate on a respective communication link with a respective one of multiple STAs of a non-AP MLD (also referred to as a “STA MLD”). The STA MLD may communicate with the AP MLD over one or more of the multiple communication links at a given time.

[0056] One type of MLO is multi-link aggregation (MLA), where traffic associated with a single STA is simultaneously transmitted across multiple communication links in parallel to maximize the utilization of available resources to achieve higher throughput. That is, during at least some duration of time, transmissions or portions of transmissions may occur over two or more links in parallel at the same time. In some examples, the parallel wireless communication links may support synchronized transmissions. In some other examples, or during some other durations of time, transmissions over the links may be parallel, but not be synchronized or concurrent. In some examples or durations of time, two or more of the links may be used for communications between the wireless communication devices in the same direction (such as all uplink or all downlink). In some other examples or durations of time, two or more of the links may be used for communications in different directions. For example, one or more links may support uplink communications and one or more links may support downlink communications. In such examples, at least one of the wireless communication devices operates in a full duplex mode. Generally, full duplex operation enables bi-directional communications where at least one of the wireless communication devices may transmit and receive at the same time.

[0057] MLA may be implemented in a number of ways. In some examples, MLA may be packet-based. For packet-based aggregation, frames of a single traffic flow (such as all traffic associated with a given traffic identifier (TID)) may be sent concurrently across multiple communication links. In some other examples, MLA may be flow-based. For flow-based aggregation, each traffic flow (such as all traffic associated with a given TID) may be sent using a single one of multiple available communication links. As an example, a single STA MLD may access a web browser while streaming a video in parallel. The traffic associated with the web browser access may be communicated over a first communication link while the traffic associated with the video stream may be communicated over a second communication link in parallel (such that at least some of

the data may be transmitted on the first channel concurrently with data transmitted on the second channel).

[0058] In some other examples, MLA may be implemented as a hybrid of flow-based and packet-based aggregation. For example, an MLD may employ flow-based aggregation in situations in which multiple traffic flows are created and may employ packet-based aggregation in other situations. The determination to switch among the MLA techniques or modes may additionally or alternatively be associated with other metrics (such as a time of day, traffic load within the network, or battery power for a wireless communication device, among other factors or considerations).

[0059] To support MLO techniques, an AP MLD and a STA MLD may exchange supported MLO capability information (such as supported aggregation type or supported frequency bands, among other information). In some examples, the exchange of information may occur via a beacon signal, a probe request or probe response, an association request or an association response frame, a dedicated action frame, or an operating mode indicator (OMI), among other examples. In some examples, an AP MLD may designate a given channel in a given band as an anchor channel (such as the channel on which it transmits beacons and other management frames). In such examples, the AP MLD also may transmit beacons (such as ones which may contain less information) on other channels for discovery purposes.

[0060] MLO techniques may provide multiple benefits to a WLAN. For example, MLO may improve user perceived throughput (UPT) (such as by quickly flushing per-user transmit queues). Similarly, MLO may improve throughput by improving utilization of available channels and may increase spectral utilization (such as increasing the bandwidth-time product). Further, MLO may enable smooth transitions between multi-band radios (such as where each radio may be associated with a given RF band) or enable a framework to set up separation of control channels and data channels. Other benefits of MLO include reducing the ON time of a modem, which may benefit a wireless communication device in terms of power consumption. Another benefit of MLO is the increased multiplexing opportunities in the case of a single BSS. For example, multi-link aggregation may increase the number of users per multiplexed transmission served by the multi-link AP MLD.

[0061] Figure 4 shows a pictorial diagram of another example wireless communication network 400. According to some aspects, the wireless communication network 400 can be an example of a mesh network, an IoT network or a sensor network

in accordance with one or more of the IEEE 802.11 family of wireless communication protocol standards (including the 802.11ah amendment). The wireless network 400 may include multiple wireless communication devices 414. The wireless communication devices 414 may represent various devices such as display devices (for example, TVs, computer monitors, navigation systems, among others), music or other audio or stereo devices, remote control devices (“remotes”), printers, kitchen or other household appliances, among other examples.

[0062] In some examples, the wireless communication devices 414 sense, measure, collect or otherwise obtain and process data and then transmit such raw or processed data to an intermediate device 412 for subsequent processing or distribution. Additionally or alternatively, the intermediate device 412 may transmit control information, digital content (for example, audio or video data), configuration information or other instructions to the wireless communication devices 414. The intermediate device 412 and the wireless communication devices 414 can communicate with one another via wireless communication links 416. In some examples, the wireless communication links 416 include Bluetooth links or other PAN or short-range communication links.

[0063] In some examples, the intermediate device 412 also may be configured for wireless communication with other networks such as with a Wi-Fi WLAN or a wireless (for example, cellular) wide area network (WWAN), which may, in turn, provide access to external networks including the Internet. For example, the intermediate device 412 may associate and communicate, over a Wi-Fi link 418, with an AP 402 of a WLAN network, which also may serve various STAs 404. In some examples, the intermediate device 412 is an example of a network gateway, for example, an IoT gateway. In such a manner, the intermediate device 412 may serve as an edge network bridge providing a Wi-Fi core backhaul for the IoT network including the wireless communication devices 414. In some examples, the intermediate device 412 can analyze, preprocess and aggregate data received from the wireless communication devices 414 locally at the edge before transmitting it to other devices or external networks via the Wi-Fi link 418. The intermediate device 412 also can provide additional security for the IoT network and the data it transports.

[0064] Some wireless communication devices (including both APs and STAs) are capable of multi-link operation (MLO). In some examples, MLO supports establishing multiple different communication links (such as a first link on the 2.4 GHz band, a second link on the 5 GHz band, and the third link on the 6 GHz band) between the STA and the

AP. Each communication link may support one or more sets of channels or logical entities. In some cases, each communication link associated with a given wireless communication device may be associated with a respective radio of the wireless communication device, which may include one or more transmit/receive (Tx/Rx) chains, include or be coupled with one or more physical antennas, or include signal processing components, among other components. An MLO-capable device may be referred to as a multi-link device (MLD). For example, an AP MLD may include multiple APs each configured to communicate on a respective communication link with a respective one of multiple STAs of a non-AP MLD (also referred to as a “STA MLD”). The STA MLD may communicate with the AP MLD over one or more of the multiple communication links at a given time.

[0065] One type of MLO is multi-link aggregation (MLA), where traffic associated with a single STA is simultaneously transmitted across multiple communication links in parallel to maximize the utilization of available resources to achieve higher throughput. That is, during at least some duration of time, transmissions or portions of transmissions may occur over two or more links in parallel at the same time. In some examples, the parallel wireless communication links may support synchronized transmissions. In some other examples, or during some other durations of time, transmissions over the links may be parallel, but not be synchronized or concurrent. In some examples or durations of time, two or more of the links may be used for communications between the wireless communication devices in the same direction (such as all uplink or all downlink). In some other examples or durations of time, two or more of the links may be used for communications in different directions. For example, one or more links may support uplink communications and one or more links may support downlink communications. In such examples, at least one of the wireless communication devices operates in a full duplex mode. Generally, full duplex operation enables bi-directional communications where at least one of the wireless communication devices may transmit and receive at the same time.

[0066] MLA may be implemented in a number of ways. In some examples, MLA may be packet-based. For packet-based aggregation, frames of a single traffic flow (such as all traffic associated with a given traffic identifier (TID)) may be sent concurrently across multiple communication links. In some other examples, MLA may be flow-based. For flow-based aggregation, each traffic flow (such as all traffic associated with a given TID) may be sent using a single one of multiple available communication links. As an

example, a single STA MLD may access a web browser while streaming a video in parallel. The traffic associated with the web browser access may be communicated over a first communication link while the traffic associated with the video stream may be communicated over a second communication link in parallel (such that at least some of the data may be transmitted on the first channel concurrently with data transmitted on the second channel).

[0067] In some other examples, MLA may be implemented as a hybrid of flow-based and packet-based aggregation. For example, an MLD may employ flow-based aggregation in situations in which multiple traffic flows are created and may employ packet-based aggregation in other situations. The determination to switch among the MLA techniques or modes may additionally or alternatively be associated with other metrics (such as a time of day, traffic load within the network, or battery power for a wireless communication device, among other factors or considerations).

[0068] To support MLO techniques, an AP MLD and a STA MLD may exchange supported MLO capability information (such as supported aggregation type or supported frequency bands, among other information). In some examples, the exchange of information may occur via a beacon signal, a probe request or probe response, an association request or an association response frame, a dedicated action frame, or an operating mode indicator (OMI), among other examples. In some examples, an AP MLD may designate a given channel in a given band as an anchor channel (such as the channel on which it transmits beacons and other management frames). In such examples, the AP MLD also may transmit beacons (such as ones which may contain less information) on other channels for discovery purposes.

[0069] MLO techniques may provide multiple benefits to a WLAN. For example, MLO may improve user perceived throughput (UPT) (such as by quickly flushing per-user transmit queues). Similarly, MLO may improve throughput by improving utilization of available channels and may increase spectral utilization (such as increasing the bandwidth-time product). Further, MLO may enable smooth transitions between multi-band radios (such as where each radio may be associated with a given RF band) or enable a framework to set up separation of control channels and data channels. Other benefits of MLO include reducing the ON time of a modem, which may benefit a wireless communication device in terms of power consumption. Another benefit of MLO is the increased multiplexing opportunities in the case of a single BSS. For example, multi-link

aggregation may increase the number of users per multiplexed transmission served by the multi-link AP MLD.

CENTRALIZED ENERGY MANAGEMENT IN MESH NETWORKS

[0070] Advances in wireless technology often come with some form of cost. Aspects of the present disclosure, however, provide power savings mechanisms that may help mitigate the increase in power consumption associated with advanced network features.

[0071] Various networks may have certain energy consumption guidelines for APs. For example, referring to mesh network 600 of Figure 6, various APs, such as root AP 512 and repeaters 514, may have certain defined energy key performance indicators (KPIs). At the AP level, individual nodes may perform actions to reduce power consumption. For example, each individual AP may calculate metrics per-link and perform actions to reduce power consumption based on the metrics. As an example, an AP could activate/de-activate certain links (e.g., used for a backhaul 516 or to support clients 518) or ports and/or adjust parameters, such as operating BW and MCS.

[0072] As noted above, applying energy management in individual nodes may be less than optimal for a mesh network. For example, mesh networks in a home environment may involve coordination between nodes. As such, mesh networks may benefit from topology awareness to make effective energy management decisions that consider holistic knowledge about the mesh network topology and capabilities.

[0073] Aspects of the present disclosure, however, provide a centralized energy management that may consider a topology of a mesh network when making power savings decisions. Such a centralized energy management approach may utilize power savings mechanisms that may help mitigate the increase in power consumption associated with advanced network features.

[0074] In some cases, the power savings mechanisms proposed herein may be implemented as part of an overall AP architecture, such as the architecture 700 shown in Figure 7. As illustrated, the architecture 700 may include an energy service 712 implemented in user space 710. The energy service 712 may support a service aware (SA) framework designed to help ensure certain quality of service (QoS) objectives, such as those related to service level agreements (SLAs), are met, within reasonable energy constraints.

[0075] The energy service generally strives to achieve the access point's power-save goals aligning with the SA-framework's SLA-goals. Such an energy service may be

achieved via various algorithms described herein. In some cases, at a driver level 720, the power saving may involve various power saving mechanisms in a WLAN driver, Ethernet driver, or other driver platform. The WLAN driver power saving mechanisms may include band coalescing 722, which generally refers to inactivating one link, effectively merging that link with one or more other links. Such coalescing may be achieved using mechanisms such as TID to link mapping (T2LM), BSS transmission management (BTM), and/or link utilization (LU) mechanisms. Additional power savings mechanisms 724 may include RF chain reduction and bandwidth reduction. In some cases, a driver (e.g., 802.3az) may be used for certain signaling and control. At the firmware level 730, power saving may be achieved via various PS mechanisms 732, such as chain reduction, rate control, and dynamic bandwidth reduction.

[0076] As illustrated in the example mesh network 800 of Figure 8, the topology aware centralized energy management proposed herein may be combined with AP-level energy management, since Energy Policy for a single node may not suffice for a mesh network.

[0077] As shown at 802, AP-level energy management may be performed at an individual repeater 514, and may include link/port management, adjusting processing (e.g., CPU core/clock gating), and/or adjusting parameters, such as operating BW, MCS, and the like. In some cases, AP-level energy management may be performed in addition to, or as an alternative to, centralized energy management if one or more conditions are met.

[0078] As shown at 804, centralized energy management may be performed at a root AP 512 (or controller entity) in a manner that considers a topology of the mesh network 800. For example, AP 512 may perform topology-aware energy management by steering clients 518 to certain APs, transitioning certain APs to low power states, and/or reducing the amount of control messages exchanged between APs via a backhaul 516. The exchange of such control messages may be referred to as chatter, and reducing the amount of chatter may be referred to as chatter control.

[0079] As noted above, energy aware repeater transition may be achieved by steering one or more clients away from a specific mesh node to provide an opportunity for that node to enter a power savings or standby mode (such modes may be defined, for example, by placing one or more components in a low power state, deactivating links, or performing some other action, as will be described in greater detail below). In some cases, clients may be steered from nodes that are farther away (e.g., end nodes) to nodes that are closer

to the Root (in terms of hops), which gives an opportunity for the end nodes to exercise energy saving policies.

[0080] Examples of energy aware repeater transition and AP/client steering may be understood with reference to example scenarios 900 and 1000 depicted in Figure 9 and Figure 10.

[0081] Referring first to Figure 9, as shown at 902, client 1 and client 2 are steered away from Repeater 2. Client 1 is steered from Repeater 2 to the Root node, while Client 2 is steered away from Repeater 2 to Repeater 1. As a result, as illustrated at 904, Repeater 2 is no longer serving any clients and can transition to standby.

[0082] In some cases, the root node (or controller entity) may select the particular clients and APs for steering based on one or more suitability criteria (based on topology information).

[0083] For example, in the example shown in Figure 9, client 1 and client 2 are within the coverage area of the Root and Repeater 1, making steering an option. Further, the Root may realize that these are the only two clients served by Repeater 2, meaning Repeater 2 may transition to standby after the steering, as indicated at 904.

[0084] In some cases, distance from the root node may be considered when performing AP/client steering.

[0085] For example, in the example shown in Figure 10 in which Repeater 1 and Repeater 2 are arranged in a daisy-chained topology of multiple hops, Repeater 2 is farther away from the Root than Repeater 1. Thus, as indicated at 1002, client 1 may be steered to Repeater 1 rather than Repeater 2, allowing Repeater 2 to transition to standby. While an alternative would have been to steer client 2 to Repeater 2, because Repeater 1 is used in the hop between the Root and Repeater 2, it may not be possible to transition Repeater 1 to standby.

[0086] Another example of topology aware decision making for energy management may involve topology aware backhaul management. As an example, if a mesh is a daisy chain topology of multiple hops, as in the Example of Figure 10, energy saving rules may not be enforced on intermediate backhaul links, if end nodes have service level agreement (SLA) sensitive clients.

[0087] For example, in the scenario 1100 illustrated in Figure 11, instead of keeping a multi-link operation (MLO) backhaul, a backhaul could be reduced to SLO if a given repeater has transitioned to standby mode. In the example illustrated in Figure 11, a 3-link MLO backhaul is maintained between the Root and Repeater 1. As illustrated at 1102,

however, the backhaul between Repeater 1 and Repeater 2 may be reduced to SLO when Repeater 2 has transitioned to a standby mode.

[0088] On the other hand, in the scenario 1200 illustrated in Figure 12, a backhaul may be maintained as MLO even if a repeater is serving no clients. In the illustrated example, Repeater 1 is not serving any clients directly, but is serving as a virtual AP (VAP) so Repeater 1 may not be transitioned to standby. As a result, as indicated at 1202, the links for the backhaul between Repeater 1 and the root may not be reduced.

[0089] According to certain aspects, a controller (e.g., root node or other centralized controller entity) may be configured to select various operating parameters to optimize energy savings opportunities. For example, the controller may attempt to select a best channel, bandwidth, and/or number of spatial streams (NSS) for specific nodes in an effort to facilitate energy saving opportunities in the mesh network. For example, the controller may increase BW or NSS on a specific repeater, so that a repeater can serve additional nearby clients, to allow another repeater to transition to standby.

[0090] Another example of topology aware decision making for energy management may involve adaptive power control. As an example, if multiple mesh nodes are in close proximity, a controller may reduce a maximum transmission (TX) power on certain nodes to conserve energy, without impacting a desired coverage area (e.g., a whole home coverage).

[0091] In the scenario 1300A illustrated in Figure 13A, the root, repeater 1, and repeater 2 are in close proximity and provide coverage of a home 1300. As illustrated at 1302, Tx power of Repeater 2 may be reduced without impacting coverage. In other words, as illustrated in scenario 1300B illustrated in Figure 13B, even with a reduced coverage area of Repeater 2 (due to reduced Tx power), the home 1300 is still covered by the mesh network.

[0092] As noted above, another example of topology aware decision making involves (smart) chatter control, based on topology information. For example, in the scenario 1400 illustrated in Figure 14, Repeater 2 is being lightly used as indicated at 1402. As a result, as indicated at 1404, the root node may reduce control message chatter between the root and Repeater 2.

[0093] In some cases, a controller may determine a frequency of Request-Response of periodic control messages (chatter) for mesh node communication in order to reduce chatter based on a topology aware power save policy. The controller may control a Request timer in which repeaters respond to controller requests. Example of control

messages that may be considered chatter include, but are not limited to: topology request/response messages, link statistics request/response messages, client statistics request/response messages, and other types of request/response messages. In some cases, mesh topology aware energy management may be implemented using various types of communication and messaging. For example, mesh topology aware energy management may be implemented using multi-AP protocol messages (e.g., according to IEEE 1905.1) including Type-length-values (TLVs).

[0094] In some cases, topology aware decision making may consider location information associated with one or more of the clients. For example, a controller may consider location information of one or more clients relative to one or more nodes (repeaters), when making steering decisions. In some cases, the controller may also consider a capabilities match between a client (that is a candidate for steering) and an AP (that is a potential to steer the client to).

[0095] Topology aware energy management proposed herein may be understood with reference to the flow diagram 1500 of Figure 15, which illustrates example operations for AP steering. Operations in Figure 15 may be performed, for example, by a root node.

[0096] As shown at 1502, a root node of a mesh network may be powered on and, at 1504, the root node may associate with other nodes in the mesh network. At 1506, the root node builds a detailed mesh layout, gathering topology information for use in making topology-aware energy management decisions.

[0097] For example, the root node may obtain topology information via topology discover requests and responses. The topology information obtained may include, for example, a number of nodes in the mesh network and various information per node. The per-node information may include a number of interfaces and types, link statistics, channel supervision/selection/channel availability check (CAC). The information may also include AP statistics and telemetry information, such as a number of clients, STA link metrics, operating BW, classified flows and SLAs. The information may also include per-client information, such as client capabilities, organizationally unique identifier (OUI) IDs (used for classification of devices), and/or a neighborhood map.

[0098] As illustrated at 1508, the root node/controller may calculate a weight or metric per node (periodically or based on certain triggers). If periodically, the calculation may be based on a timer. Triggers may include certain topology notification events.

[0099] The weights may be used to determine if one or more prerequisites for AP steering have been met. If the prerequisites are met, as determined at 1510, the controller

may initiate AP steering, at 1512. In some cases, AP steering prerequisites may be based on a STAs airtime utilization and SLAs. For example, lower throughput, higher latency clients may be steered to repeaters that can potentially operate on fewer links to reduce energy consumption.

[0100] In some cases, it may be beneficial for a root to be able to trigger a repeater to exit a standby mode and transition to an active state. Considerations for transitioning a repeater from a standby mode to an active state may be understood with reference to the flow diagram 1600 of Figure 16.

[0101] As shown at 1602, a root node of a root network may periodically monitor SLAs across the mesh network. The example assumes, at 1604, that an SLA breach warning is triggered based on the monitoring.

[0102] The root node determines, at 1606, if transitioning a repeater (or repeaters) to an active state would help address the SLA breach. If so, at 1608, the root node may send an activation to the repeater(s) in standby. After bringing repeaters to active, the root node may re-assess the topology and re-assign STAs to repeaters to optimize energy savings. For example, at 1610, the root node may steer one or more candidate stations to an activated repeater.

[0103] In some cases, repeaters in standby mode may transition to an active state upon receiving certain defined packets, such as wake-on-LAN or wake-on-WLAN packets. In such cases, the repeaters may communicate this change in state, for example, via a topology notification message to the root (and/or other repeaters).

[0104] Figure 17 shows a flowchart illustrating a process 1700 performable at a first wireless node, according to some aspects of the present disclosure. The operations of the process 1700 may be implemented by the first wireless node or its components as described herein. For example, the process 1700 may be performed by a wireless communication device, such as a wireless communication device 1800 described with reference to Figure 18, operating as or within the first wireless node. In some examples, the process 1700 may be performed by the first wireless node, such as one of the wireless APs 102 described with reference to Figure 1. The operations of the process 1700 and subject matter of the present disclosure may be applicable for different systems.

[0105] At 1710, the process 1700 includes obtaining information regarding a topology of a mesh network.

[0106] At 1720, the process 1700 includes performing, based on the information, one or more actions to adjust power consumption at one or more wireless nodes of the mesh network.

[0107] In some aspects, the at least one action comprises steering one or more clients away from a first wireless node of the mesh network to a second wireless node of the mesh network.

[0108] In some aspects, the method 1700 further includes selecting the second wireless node based on one or more suitability criteria.

[0109] Figure 18 shows a block diagram of a wireless communication device 1800 (such as an AP or non-AP STA), according to some aspects of the present disclosure. In one example, the wireless communication device 1800 is configured or operable to perform a process 1700 described with reference to Figure 17. In various examples, the wireless communication device 1700 can be a chip, SoC, chipset, package or device that may include: one or more modems (such as a Wi-Fi (IEEE 802.11) modem or a cellular modem such as 3GPP 4G LTE or 5G compliant modem); one or more processors, processing blocks or processing elements (collectively “the processor”); one or more radios (collectively “the radio”); and one or more memories or memory blocks (collectively “the memory”).

[0110] In some examples, the wireless communication device 1800 can be a device for use in an AP, such as AP 102 described with reference to Figure 1. In some other examples, the wireless communication device 1800 can be an AP that includes a chip, SoC, chipset, package or device as well as multiple antennas. The wireless communication device 1800 is capable of transmitting and receiving wireless communications in the form of, for example, wireless packets. For example, the wireless communication device can be configured or operable to transmit and receive packets in the form of physical layer PPDU and MPDU conforming to one or more of the IEEE 802.11 family of wireless communication protocol standards. In some examples, the wireless communication device 1800 also includes or can be coupled with an application processor which may be further coupled with another memory. In some examples, the wireless communication device 1800 further includes at least one external network interface that enables communication with a core network or backhaul network to gain access to external networks including the Internet.

[0111] The wireless communication device 1800 includes obtaining component 1802, outputting component 1804, performing component 1806, steering component 1808, selecting component 1810, and adjusting component 1812.

[0112] Portions of one or more of the components 1802, 1804, 1806, 1808, 1810, and/or 1812 may be implemented at least in part in hardware or firmware. In some examples, at least some of the components 1802, 1804, 1806, 1808, 1810, and/or 1812 are implemented at least in part by a processor and as software stored in a memory. For example, portions of one or more of the components 1802, 1804, 1806, 1808, 1810, and/or 1812 can be implemented as non-transitory instructions (or “code”) executable by the processor to perform the functions or operations of the respective module.

[0113] In some implementations, the processor may be a component of a processing system. A processing system may generally refer to a system or series of machines or components that receives inputs and processes the inputs to produce a set of outputs (which may be passed to other systems or components of, for example, the wireless communication device 1800). For example, a processing system of the wireless communication device 1800 may refer to a system including the various other components or subcomponents of the wireless communication device 1800, such as the processor, or a transceiver, or a communications manager, or other components or combinations of components of the wireless communication device 1800. The processing system of the wireless communication device 1800 may interface with other components of the wireless communication device 1800, and may process information received from other components (such as inputs or signals) or output information to other components. For example, a chip or modem of the wireless communication device 1800 may include a processing system, a first interface to output information and a second interface to obtain information. In some implementations, the first interface may refer to an interface between the processing system of the chip or modem and a transmitter, such that the wireless communication device 1800 may transmit information output from the chip or modem. In some implementations, the second interface may refer to an interface between the processing system of the chip or modem and a receiver, such that the wireless communication device 1800 may obtain information or signal inputs, and the information may be passed to the processing system. A person having ordinary skill in the art will readily recognize that the first interface also may obtain information or signal inputs, and the second interface also may output information or signal outputs.

[0114] The obtaining component 1802 may be capable of, configured to, or operable to at least obtain a first request that requests the first wireless node to modify the wireless operating parameters.

[0115] The outputting component 1804 may be capable of, configured to, or operable to at least output, for transmission, information that indicates at least one of (i) one or more criteria that, when met by a second wireless node, indicate the second wireless node is allowed to request that the first wireless node modify one or more wireless operating parameters, or (ii) one or more options for modifying the wireless operating parameters of the first wireless node.

[0116] Various components of the wireless communication device 1800 may provide means for performing the process 1700 described with reference to Figure 17, or any aspect related to it. Means for receiving or obtaining may include transceivers and/or antenna(s) of the AP 102 described with reference to Figure 1.

[0117] Means for obtaining, means for performing, means for outputting, means for selecting, means for adjusting, means for reducing, and/or means for steering may include one or more processors (such as a receive processor, a controller, and/or a transmit processor) of the AP 102 described with reference to Figure 1.

[0118] In some cases, rather than actually transmitting, for example, signals and/or data, the wireless communication device 1800 may have an interface to output signals and/or data for transmission (means for outputting). For example, a processor may output signals and/or data, via a bus interface, to a radio frequency (RF) front end of the wireless communication device 1800 for transmission. In various aspects, the RF front end may include various components, including transmit and receive processors, transmit and receive MIMO processors, modulators, demodulators, and the like.

[0119] In some cases, rather than actually receiving signals and/or data, the wireless communication device 1800 may have an interface to obtain the signals and/or data received from another device (means for obtaining). For example, a processor may obtain (or receive) the signals and/or data, via a bus interface, from an RF front end of the wireless communication device 1800 for reception. In various aspects, the RF front end may include various components, including transmit and receive processors, transmit and receive MIMO processors, modulators, demodulators, and the like.

Example Clauses

[0120] Implementation examples are described in the following numbered clauses:

[0121] Clause 1: A method for wireless communications at a wireless node, comprising: obtaining information regarding a topology of a mesh network; and performing, based on the information, one or more actions to adjust power consumption at one or more wireless nodes of the mesh network.

[0122] Clause 2: The method of Clause 1, wherein the one or more actions comprise: steering one or more clients away from a first wireless node of the mesh network to a second wireless node of the mesh network.

[0123] Clause 3: The method of Clause 2, further comprising selecting the second wireless node based on one or more suitability criteria.

[0124] Clause 4: The method of Clause 3, wherein the one or more suitability criteria comprises the second wireless node being in closer proximity to a root node than one or more other wireless nodes of the mesh network.

[0125] Clause 5: The method of Clause 3, wherein the one or more actions further comprise adjusting one or more parameters to place the first wireless node in an energy saving state after steering the one or more clients away from the first wireless node.

[0126] Clause 6: The method of Clause 3, wherein: the information comprises location information associated with the one or more clients relative to the one or more wireless nodes; or the one or more suitability criteria involves the location information.

[0127] Clause 7: The method of Clause 6, wherein the suitability criteria also involves a capabilities match between the one or more clients and the second wireless node.

[0128] Clause 8: The method of any one of Clauses 1-7, wherein: the information indicates the topology of the mesh network comprises a daisy chain of multiple hops; and the one or more actions comprise adjusting energy saving rules related to at least one backhaul link for a wireless node that is an end node of the daisy chain, if at least one condition is met.

[0129] Clause 9: The method of Clause 8, wherein the at least one condition relates to a service level agreement (SLA) of a client served by the end node.

[0130] Clause 10: The method of Clause 8, wherein at least one of: the at least one condition is met if the end node has entered an energy saving state; or the one or more actions comprise reducing a number of links in the at least one backhaul link.

[0131] Clause 11: The method of any one of Clauses 1-10, wherein the one or more actions comprise selecting one or more transmission parameters for one or more of the one or more wireless nodes of the mesh network, based on the information.

[0132] Clause 12: The method of Clause 11, wherein the one or more transmission parameters comprise at least one of a channel, a bandwidth, a modulation and coding scheme (MCS), or a number of spatial streams.

[0133] Clause 13: The method of Clause 12, wherein the channel, bandwidth, or number of spatial streams are selected to allow the first wireless node of the mesh network to serve a number of clients so a second wireless node of the mesh network can enter an energy saving state.

[0134] Clause 14: The method of any one of Clauses 1-13, wherein: the one or more actions comprise reducing a transmit power on the one or more wireless nodes.

[0135] Clause 15: The method of Clause 14, wherein: the information indicates at least relative location of wireless nodes in the mesh network; and the transmit power is reduced, based on the at least the relative location of the wireless nodes, in a manner that maintains a coverage area of the mesh network.

[0136] Clause 16: The method of any one of Clauses 1-15, wherein: the one or more actions comprise adjusting at least one of an amount of control message signaling or frequency of control message signaling between wireless nodes of the mesh network, the adjustment being based on an energy savings policy.

[0137] Clause 17: The method of Clause 16, wherein the control message signaling comprises at least one of: a topology request, a topology response, a link status request, a link status response, a client status request, a status client response, an auto channel selection (ACS) request, an ACS response, a channel availability check (CAC) request, or a CAC response.

[0138] Clause 18: An apparatus, comprising: at least one memory comprising executable instructions; and one or more processors configured to execute the executable instructions and cause the apparatus to perform a method in accordance with any one of Clauses 1-17.

[0139] Clause 19: An apparatus, comprising means for performing a method in accordance with any one of Clauses 1-17.

[0140] Clause 20: A non-transitory computer-readable medium comprising executable instructions that, when executed by one or more processors of an apparatus, cause the apparatus to perform a method in accordance with any one of Clauses 1-17.

[0141] Clause 21: A computer program product embodied on a computer-readable storage medium comprising code for performing a method in accordance with any one of Clauses 1-17.

[0142] Clause 22: An access point (AP) comprising: at least one transceiver, at least one memory comprising executable instructions; and one or more processors configured to execute the executable instructions and cause the AP to perform a method in accordance with any one of Clauses 1-17, wherein the at least one transceiver is configured to receive the information.

ADDITIONAL CONSIDERATIONS

[0143] As used herein, the term “determine” or “determining” encompasses a wide variety of actions and, therefore, “determining” can include calculating, computing, processing, deriving, investigating, looking up (such as via looking up in a table, a database or another data structure), inferring, ascertaining, measuring, and the like. Also, “determining” can include receiving (such as receiving information), accessing (such as accessing data stored in memory), transmitting (such as transmitting information) and the like. Also, “determining” can include resolving, selecting, obtaining, choosing, establishing and other such similar actions.

[0144] As used herein, a phrase referring to “at least one of” a list of items refers to any combination of those items, including single members. As an example, “at least one of: a, b, or c” is intended to cover: a, b, c, a-b, a-c, b-c, and a-b-c. As used herein, “or” is intended to be interpreted in the inclusive sense, unless otherwise explicitly indicated. For example, “a or b” may include a only, b only, or a combination of a and b.

[0145] As used herein, “based on” is intended to be interpreted in the inclusive sense, unless otherwise explicitly indicated. For example, “based on” may be used interchangeably with “based at least in part on,” “associated with”, or “in accordance with” unless otherwise explicitly indicated. Specifically, unless a phrase refers to “based on only ‘a,’” or the equivalent in context, whatever it is that is “based on ‘a,’” or “based at least in part on ‘a,’” may be based on “a” alone or based on a combination of “a” and one or more other factors, conditions or information.

[0146] As used herein, “a processor,” “at least one processor” or “one or more processors” generally refers to a single processor configured to perform one or multiple operations or multiple processors configured to collectively perform one or more operations. In the case of multiple processors, performance the one or more operations could be divided amongst different processors, though one processor may perform multiple operations, and multiple processors could collectively perform a single operation. Similarly, “a memory,” “at least one memory” or “one or more memories”

generally refers to a single memory configured to store data and/or instructions, multiple memories configured to collectively store data and/or instructions.

[0147] The various illustrative components, logic, logical blocks, modules, circuits, operations and algorithm processes described in connection with the examples disclosed herein may be implemented as electronic hardware, firmware, software, or combinations of hardware, firmware or software, including the structures disclosed in this specification and the structural equivalents thereof. The interchangeability of hardware, firmware and software has been described generally, in terms of functionality, and illustrated in the various illustrative components, blocks, modules, circuits and processes described above. Whether such functionality is implemented in hardware, firmware or software depends upon the particular application and design constraints imposed on the overall system.

[0148] Various modifications to the examples described in this disclosure may be readily apparent to persons having ordinary skill in the art, and the generic principles defined herein may be applied to other examples without departing from the spirit or scope of this disclosure. Thus, the claims are not intended to be limited to the examples shown herein, but are to be accorded the widest scope consistent with this disclosure, the principles and the novel features disclosed herein.

[0149] Additionally, various features that are described in this specification in the context of separate examples also can be implemented in combination in a single implementation. Conversely, various features that are described in the context of a single implementation also can be implemented in multiple examples separately or in any suitable subcombination. As such, although features may be described above as acting in particular combinations, and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[0150] Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. Further, the drawings may schematically depict one or more example processes in the form of a flowchart or flow diagram. However, other operations that are not depicted can be incorporated in the example processes that are schematically illustrated. For example, one or more additional operations can be performed before, after, simultaneously, or between any of the illustrated operations. In some circumstances, multitasking and parallel processing may be advantageous.

Moreover, the separation of various system components in the examples described above should not be understood as requiring such separation in all examples, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

CLAIMS

1. An apparatus for wireless communications, comprising:
at least one memory comprising instructions; and one or more processors configured to execute the instructions and cause the apparatus to:
obtain information regarding a topology of a mesh network; and
perform, based on the information, one or more actions to adjust power consumption at one or more wireless nodes of the mesh network.
2. The apparatus of claim 1, wherein the one or more actions comprise: steering one or more clients away from a first wireless node of the mesh network to a second wireless node of the mesh network.
3. The apparatus of claim 2, wherein the one or more processors are further configured to execute the instructions and cause the apparatus to:
select the second wireless node based on one or more suitability criteria.
4. The apparatus of claim 3, wherein the one or more suitability criteria comprises the second wireless node being in closer proximity to a root node than one or more other wireless nodes of the mesh network.
5. The apparatus of claim 3, wherein the one or more actions further comprise adjusting one or more parameters to place the first wireless node in an energy saving state after steering the one or more clients away from the first wireless node.
6. The apparatus of claim 3, wherein the information comprises location information associated with the one or more clients relative to the one or more wireless nodes; or the one or more suitability criteria involves the location information.
7. The apparatus of claim 6, wherein the suitability criteria also involves a capabilities match between the one or more clients and the second wireless node.

8. The apparatus of claim 1, wherein the information indicates the topology of the mesh network comprises a daisy chain of multiple hops; and the one or more actions comprise adjusting energy saving rules related to at least one backhaul link for a wireless node that is an end node of the daisy chain, if at least one condition is met.
9. The apparatus of claim 8, wherein the at least one condition relates to a service level agreement (SLA) of a client served by the end node.
10. The apparatus of claim 8, wherein at least one of: the at least one condition is met if the end node has entered an energy saving state; or the one or more actions comprise reducing a number of links in the at least one backhaul link.
11. The apparatus of claim 1, wherein the one or more actions comprise selecting one or more transmission parameters for one or more of the one or more wireless nodes of the mesh network, based on the information.
12. The apparatus of claim 11, wherein the one or more transmission parameters comprise at least one of a channel, a bandwidth, a modulation and coding scheme (MCS), or a number of spatial streams.
13. The apparatus of claim 12, wherein the channel, bandwidth, or number of spatial streams are selected to allow the first wireless node of the mesh network to serve a number of clients so a second wireless node of the mesh network can enter an energy saving state.
14. The apparatus of claim 1, wherein the one or more actions comprise reducing a transmit power on the one or more wireless nodes.
15. The apparatus of claim 14, wherein the information indicates at least relative location of wireless nodes in the mesh network; and the transmit power is reduced, based on the at least the relative location of the wireless nodes, in a manner that maintains a coverage area of the mesh network.

16. The apparatus of claim 1, wherein the one or more actions comprise adjusting at least one of an amount of control message signaling or frequency of control message signaling between wireless nodes of the mesh network, the adjustment being based on an energy savings policy.

17. The apparatus of claim 16, wherein the control message signaling comprises at least one of: a topology request, a topology response, a link status request, a link status response, a client status request, a status client response, an auto channel selection (ACS) request, an ACS response, a channel availability check (CAC) request, or a CAC response.

18. A method for wireless communications at a wireless node, comprising:
obtaining information regarding a topology of a mesh network; and
performing, based on the information, one or more actions to adjust power consumption at one or more wireless nodes of the mesh network.

19. The method of claim 18, wherein the one or more actions comprise: steering one or more clients away from a first wireless node of the mesh network to a second wireless node of the mesh network.

20. An access point (AP), comprising:
at least one transceiver; at least one memory comprising executable instructions;
and one or more processors configured to execute the executable instructions and cause the AP to:

receive, via the at least one transceiver, information regarding a topology of a mesh network; and

perform, based on the information, one or more actions to adjust power consumption at one or more wireless nodes of the mesh network.