

LINK AGGREGATION WITH FLOATING PRIMARY LINK

TECHNICAL FIELD

[0001] The present implementations relate generally to wireless communications, and more specifically, to link aggregation in wireless networks.

DESCRIPTION OF THE RELATED TECHNOLOGY

[0002] A wireless local area network (WLAN) may be formed by one or more access points (APs) that provide a shared wireless communication medium for use by a number of client devices also referred to as stations (STAs). The shared wireless medium may encompass multiple frequency bands (such as the 2.4 GHz, 5 GHz, or 6 GHz frequency bands) or multiple channels of one or more frequency bands. Link aggregation is a technique that may allow the AP to communicate with a particular STA over multiple concurrent communication “links.” For example, to improve data throughput, the shared wireless medium may be divided into a primary link and one or more secondary links. The primary and secondary links may be of various bandwidths, for example, by bonding a number of 20 MHz-wide channels together to form 40 MHz-wide channels, 80 MHz-wide channels, or 160 MHz-wide channels.

[0003] The basic building block of a WLAN conforming to the Institute of 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 service set identifier (SSID) that is advertised by the AP. Each AP periodically broadcasts beacon frames on the primary link to enable any STAs within wireless range of the AP to establish and maintain a communication link with the WLAN. In a typical WLAN, only one wireless device (AP or STA) may use the wireless medium at any given time. To prevent multiple devices from accessing the shared wireless medium at the same time (which may result in collisions), wireless devices may contend with each other for access to the wireless medium. For example, each wireless device may have a Media Access Controller (MAC) that may be used to contend for access on a particular link.

[0004] A wireless device with only one MAC may contend for access on only one of the aggregated links. For example, the wireless device may contend for access on the primary link, and may use the secondary links only if it gains access to the shared wireless medium on the primary link. If the primary link is busy (because another device in the vicinity is transmitting on the same wireless channel), no other wireless devices may access the shared wireless medium even if one or more

secondary links are idle. Thus, even with link aggregation, the performance of the wireless network may be limited by the load on the primary link.

SUMMARY

[0005] 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.

[0006] One innovative aspect of the subject matter of this disclosure can be implemented in a method for wireless communication by a wireless communication device. In some implementations, the method may include steps of establishing a primary communication link with a second wireless device, the primary communication link being associated with one or more first wireless channels, access to the primary communication link being based on contention; selectively accessing a secondary communication link with the second wireless device based at least in part on having access to the primary communication link, the secondary communication link being associated with one or more second wireless channels; and selectively reassigning the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on one or more channel conditions of the first wireless channels or the second wireless channels.

[0007] Another innovative aspect of the subject matter described in this disclosure can be implemented in a wireless device. The wireless device includes one or more processors and a memory. The memory stores instructions that, when executed by the one or more processors, cause the wireless device to establish a primary communication link with a second wireless device, the primary communication link being associated with one or more first wireless channels, access to the primary communication link being based on contention; selectively access a secondary communication link with the second wireless device based at least in part on having access to the primary communication link, the secondary communication link being associated with one or more second wireless channels; and selectively reassign the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on one or more channel conditions of the first wireless channels or the second wireless channels.

[0008] The one or more channel conditions may include an indication of an amount of activity on the first or second wireless channels. For example, the amount of activity may be based on a frequency of busy channel conditions on the first or second wireless channels or a number of failed attempts to access the first or second wireless channels. In some implementations, the primary communication link may be reassigned to the one or more second wireless channels when the indication of the amount of activity on the first wireless channels is above a first threshold level, the indication of the amount of activity on the second wireless channels is below a second threshold level, the indication

of the amount of activity on the first wireless channels exceeds the amount of activity on the second wireless channels, or a combination thereof.

[0009] In some implementations of the methods and wireless communication devices, selectively reassigning the primary communication link includes transmitting primary link information to the second wireless device. For example, the primary link information may indicate a change in the primary communication link. In some aspects, the primary link information may be provided in a first beacon frame broadcast on the one or more first wireless channels. In some other aspects, the primary link information may be further provided in a second beacon frame broadcast on the one or more second wireless channels.

[0010] In some implementations of the methods and wireless communication devices, transmitting the primary link information includes transmitting a data frame, a management frame, or a control frame, where the data frame, management frame or control frame includes an Operating Mode Indication (OMI) control field that includes the primary link information. In some other implementations of the methods and wireless communication devices, transmitting the primary link information includes receiving an uplink transmission from the second wireless device and transmitting a trigger message to the second wireless device in response to the uplink transmission, where the trigger message includes the primary link information. Still further, in some other implementations of the methods and wireless communication devices, transmitting the primary link information includes transmitting a trigger message to the second wireless device based at least in part on a target wake time (TWT) service period for the second wireless device, where the trigger message includes the primary link information.

[0011] In some implementations of the methods and wireless communication devices, the primary link information may further indicate a primary-switching time at which the change is to occur. Accordingly, the primary communication link may be reassigned to the one or more second wireless channels at the primary-switching time. In some aspects, the primary-switching time may coincide with at least one of a target beacon transmission time (TBTT) or a TWT of the wireless device. In some other aspects, the primary link information may further include a countdown timer indicating a number of TBTTs or TWTs remaining until the primary-switching time.

[0012] Another innovative aspect of the subject matter described in this disclosure can be implemented in a method performed by a wireless station (STA). The method may include steps of establishing a primary communication link with an access point (AP), the primary communication link being associated with one or more first wireless channels, access to the primary communication link being based on contention; selectively accessing a secondary communication link with the AP based at least in part on having access to the primary communication link, the secondary communication link being associated with one or more second wireless channels; receiving primary link information from

the AP indicating a change in the primary communication link; and dynamically changing the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on the primary link information.

[0013] Another innovative aspect of the subject matter described in this disclosure can be implemented in a wireless device. The wireless device includes one or more processors and a memory. The memory stores instructions that, when executed by the one or more processors, cause the wireless device to establish a primary communication link with an access point (AP), the primary communication link being associated with one or more first wireless channels, access to the primary communication link being based on contention; selectively access a secondary communication link with the AP based at least in part on having access to the primary communication link, the secondary communication link being associated with one or more second wireless channels; receive primary link information from the AP indicating a change in the primary communication link; and dynamically change the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on the primary link information.

[0014] In some implementations of the methods and wireless communication devices, receiving the primary link information includes receiving a beacon frame broadcast by the AP, where the beacon frame includes the primary link information. In some other implementations of the methods and wireless communication devices, receiving the primary link information includes receiving a data frame, a management frame, or a control frame transmitted by the AP, where the data frame, management frame, or control frame includes an Operating Mode Indication (OMI) control field that includes the primary link information.

[0015] In some implementations of the methods and wireless communication devices, receiving the primary link information from the AP includes waking up from a power save state and receiving a trigger message from the AP after waking up from the power save state, where the trigger message includes the primary link information. Still further, in some other implementations of the methods and wireless communication devices, receiving the trigger message includes a step of initiating an uplink transmission to the AP upon waking up from the power save state, where the trigger message is received from the AP in response to the uplink transmission.

[0016] In some implementations of the methods and wireless communication devices, the primary link information may further indicate a primary-switching time at which the change is to occur. Accordingly, the primary communication link may be changed to the one or more second wireless channels at or after the primary-switching time. In some aspects, the primary-switching time may coincide with at least one of a TBTT or a TWT of the STA. In some other aspects, the primary link information may further include a countdown timer indicating a number of TBTTs or TWTs remaining until the primary-switching time.

[0017] Still further, in some implementations of the methods and wireless communication devices, dynamically changing the primary communication link includes entering a power save state after receiving the primary link information; waking up from the power save state at or after the primary-switching time; and using the one or more second wireless channels as the primary communication link upon waking up from the power save state.

[0018] 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

[0019] Figure 1 shows a block diagram of a wireless system.

[0020] Figure 2 shows a timing diagram depicting an example multi-link communication.

[0021] Figures 3A and 3B show timing diagrams depicting example operations for changing the primary link in a multi-link environment.

[0022] Figures 4A–4C show timing diagrams depicting example operations for notifying a wireless station (STA) of a primary link change.

[0023] Figure 5 shows a timing diagram depicting an example operation for changing the primary link in a multi-link communication.

[0024] Figure 6 shows a block diagram of an example wireless station for use in wireless communication according to some implementations.

[0025] Figure 7 shows a block diagram of an example access point for use in wireless communication according to some implementations.

[0026] Figure 8 shows a flowchart illustrating an example process for configuring a multi-link environment according to some implementations.

[0027] Figure 9 shows a flowchart illustrating an example process for selectively reassigning the primary link in a multi-link environment according to some implementations.

[0028] Figure 10 shows a flowchart illustrating an example process for dynamically changing the primary link of a STA according to some implementations.

[0029] Figure 11 shows a flowchart illustrating an example process for dynamically changing the primary link of a STA according to some other implementations.

[0030] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0031] The following description is directed to certain implementations for the purposes of describing the 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. The described implementations can 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 standards, among others. The described implementations 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), frequency division multiple access (FDMA), time division multiple access (TDMA), orthogonal frequency division multiple access (OFDMA), single-user (SU) multi-input multiple-output (MIMO) and multi-user (MU) MIMO. The described implementations 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), or an internet of things (IOT) network.

[0032] Link aggregation is a technique that may allow wireless devices to communicate with one another via multiple concurrent communication links. More specifically, link aggregation allows a wireless device to aggregate data transmissions over a primary communication link and one or more secondary communication links, thereby increasing throughput by spreading the transmissions over a wider bandwidth. For example, a wireless device with multiple Media Access Controllers (MACs) may contend for access to the shared wireless medium on multiple communication links, concurrently and independently of one another. However, a wireless device with a single MAC typically contends for access on the primary link, and may use the secondary links only if it gains access to the shared wireless medium on the primary link. Because some wireless devices have only one MAC, the increase in throughput achievable through link aggregation may be limited by the load on the primary link. Thus, the implementations described herein may enable wireless devices in a multi-link environment to dynamically change their primary links depending on the channel conditions in the shared wireless medium.

[0033] In some implementations, an access point (AP) may monitor the channel conditions on a primary link and one or more secondary links of a multi-link environment, and may signal a change in

the primary link (such as to one of the secondary links) when one or more conditions are met. For example, the change in primary link may be triggered when the amount of traffic or interference on the current primary link exceeds a threshold level, a secondary link has less traffic or interference than the current primary link, or any combination thereof. The AP may signal the change in primary link to one or more wireless stations (STAs) in its Basic Service Set (BSS) via broadcast or unicast communication frames. Aspects of the disclosure recognize that STAs typically listen for beacon frames broadcast by the AP at Target Beacon Transmission Times (TBTTs). Thus, in some aspects, the AP may signal the change in primary link via beacon frames broadcast periodically at Target Beacon Transmission Times (TBTTs). In some other aspects, the AP may signal the change in primary link via unicast management or control frames targeted at STAs that may have missed one or more beacons broadcast by the AP.

[0034] Particular implementations of the subject matter described in this disclosure can be implemented to realize one or more of the following potential advantages. The throughput of communications in multi-link environments (such as provided by aggregating a primary link with one or more secondary links) may be improved. For example, by dynamically changing the primary link based at least in part on the channel conditions of the various links, an AP may enable its associated STAs to contend for medium access on wireless channels that would otherwise remain idle (or underutilized) as secondary links. This may allow for more optimized usage of the available bandwidth in the shared wireless medium. Furthermore, by signaling the change in primary link via beacons broadcast at one or more TBTTs, the AP may ensure that the change occurs at a time when most (if not all) of its associated STAs are likely to be listening for updates regarding the BSS. This may reduce the amount of time and resources needed to notify each STA in the BSS of the change in primary link.

[0035] In the following description, numerous specific details are set forth such as examples of specific components, circuits, and processes to provide a thorough understanding of the present disclosure. The term “HE” may refer to a high efficiency frame format or protocol defined, for example, by the IEEE 802.11ax specification. Thus, the term “HE AP” may refer to an AP that operates according to the IEEE 802.11ax specification, and the term “HE STA” may refer to STAs that operate according to the IEEE 802.11ax specification.

[0036] The term “multi-link” or “ML” may refer to various link aggregation techniques that may be used to aggregate transmissions across multiple communication links. Thus, the term “ML AP” may refer to an AP that is capable of link aggregation, and the term “ML STA” may refer to STAs that are capable of link aggregation. The term “legacy STA” may refer to any STA that does not support link aggregation. The term “primary link” or “primary communication link” may refer to a set or subset of wireless channels on which peer devices (such as a STA and an AP) contend for access to a shared wireless medium. The term “secondary link” or “secondary communication link” may refer to any set or subset of wireless channels that are accessible by the peer devices for purposes of link aggregation.

[0037] In addition, although described herein in terms of exchanging data frames between wireless devices, the implementations may be applied to the exchange of any data unit, packet, or frame between wireless devices. Thus, the term “frame” may include any frame, packet, or data unit such as, for example, protocol data units (PDUs), MAC protocol data units (MPDUs), aggregated MPDUs (A-MPDUs), and physical layer convergence procedure protocol data units (PPDUs).

[0038] Figure 1 shows a block diagram of a wireless system 100. The wireless system 100 is shown to include an access point (AP) 110 and a number of wireless stations STA1–STA3. Although only three wireless stations STA1–STA3 are shown in the example of Figure 1 for simplicity, it is to be understood that the wireless system 100 may include any number of STAs.

[0039] The wireless stations STA1–STA3 may include any suitable Wi-Fi enabled wireless device including, for example, a cell phone, personal digital assistant (PDA), tablet device, laptop computer, other user equipment (UE), access terminal (AT) or the like. The AP 110 may be any suitable device that allows one or more wireless devices to connect to a network (such as a local area network (LAN), wide area network (WAN), metropolitan area network (MAN), or the Internet) using Wi-Fi, Bluetooth, or any other suitable wireless communication standards. More specifically, the AP 110 may correspond to, or provide, a basic service set (BSS). A BSS represents a basic building block of a wireless network, and may thus include a single AP (such as AP 110) and one or more associated STAs (such as STA1–STA3). In some implementations, the AP 110 may be any suitable wireless device (such as a wireless STA) acting as a software-enabled access point (“SoftAP”). The AP 110 and stations STA1–STA3 may each include one or more transceivers, one or more processing resources (such as processors or ASICs), one or more memory resources, and a power source.

[0040] In some implementations, the AP 110 may be capable of establishing multiple communication links with each of the wireless stations STA1–STA3. For example, the AP 110 may communicate with each of the wireless stations STA1–STA3 via a primary link and one or more secondary links. In the example of Figure 1, each of the wireless stations STA1–STA3 is shown to have only one secondary link. However, in actual implementations the AP 110 may establish any number of secondary links with any of the wireless stations STA1–STA3. In some implementations, the communication links may encompass different frequency bands. For example, the primary link may comprise one or more channels of the 5 GHz frequency band and the secondary link may comprise one or more channels of the 6 GHz frequency band. In some other implementations, the communication links may encompass different channels of the same frequency band. For example, the primary link may comprise a first subset of channels in the 6 GHz frequency band and the secondary link may comprise a second subset of channels in the 6 GHz frequency band. Still further, in some implementations, the communication links may encompass different channels of different frequency bands. For example, the primary link may comprise a first subset of channels in the 5 GHz frequency

band and a first subset of channels in the 6 GHz frequency band, and the secondary link may comprise a second subset of channels in the 5 GHz frequency band and a second subset of channels in the 6 GHz frequency band.

[0041] Each of the wireless devices (such as AP 110 and STA1–STA3) may have one or more media access controllers (MACs) that may be used to contend for access to the shared wireless medium. For example, the MAC may implement carrier sense multiple access collision avoidance (CSMA/CA) techniques to listen to the wireless medium to determine when the wireless medium is idle. When the wireless medium has been idle for a given duration, the MAC may contend for medium access (such as by waiting a random “back-off” period before attempting to transmit on the wireless medium). The winning device may be granted exclusive use of the shared wireless medium for a period of time (commonly referred to as a transmit opportunity or TXOP), during which only the winning device may transmit (or receive) data over the shared wireless medium.

[0042] A wireless device with multiple MACs (or multiple wireless radios) may contend for medium access on each of the communication links independently. For example, a different MAC may be assigned to each of the different communication links. As used herein, the term “multi-MAC device” may refer to any device that uses multiple MACs to communicate over multiple communication links. In some implementations, the AP 110 may be a multi-MAC device. For example, the AP 110 may include a first MAC configured to access the wireless medium via the primary link and a second MAC configured to access the wireless medium via the secondary link. More specifically, the first MAC may contend for medium access (on the primary link) when the primary link is idle, independent of the traffic conditions on the secondary link. Similarly, the second MAC may contend for medium access (on the secondary link) when the secondary link is idle, independent of the traffic conditions on the primary link. Thus, as long as one of the communications links is idle (such as the primary link or the secondary link), a multi-MAC device may gain access to the shared wireless medium. However, many wireless devices currently have only one MAC.

[0043] A wireless device with a single MAC (or a single wireless radio) may contend for medium access on only one of the communication links at any given time. For example, the wireless device may include a wideband radio that can be configured to communicate on each of the primary and secondary links. As described above, each of the primary and secondary links may be of various bandwidths, for example, by bonding a number of 20 MHz-wide channels together to form 40 MHz-wide channels, 80 MHz-wide channels, or 160 MHz-wide channels. It is noted that the next generation of Wi-Fi may allow a 320 MHz-wide channel (by accessing the 6 GHz frequency band), which may allow for link aggregation across multiple secondary channels. Thus, the MAC associated with the wideband radio may be assigned to cover a wide range of frequencies encompassed by the primary link

and one or more secondary links. As used herein, the term “single-MAC device” may refer to any device that uses a single MAC to communicate over multiple communication links.

[0044] In some implementations, one or more of the wireless stations STA1–STA3 may be a single-MAC device. For example, the single MAC may be configured to contend for medium access on the primary link (such as by waiting a random back-off period before attempting to transmit on an idle channel). However, a single-MAC device may attempt to access the secondary link only after it has gained access to the wireless medium on the primary link. For example, while operating on the primary link, the single-MAC device may check for a signal on the secondary link to determine whether the secondary link is idle. In some aspects, the single-MAC device may use energy detection (ED) or power detection (PD) techniques to sense the level of activity on the secondary link to determine if, and when, link aggregation is possible. If the activity on the secondary link is sufficiently low (such as below a threshold), the single-MAC device may access the secondary link, concurrently with the primary link, to transmit or receive wireless communications using link aggregation techniques. In contrast with the primary link, access to the secondary link is not contention-based. For example, the single-MAC device need not wait a random back-off period (after determining that the secondary link is idle) before attempting to transmit on the secondary link. However, the single-MAC device may use link aggregation only when the primary link and the secondary link are concurrently idle.

[0045] In some implementations, each of the wireless devices in the WLAN 120 (including the AP 110 and each of the wireless stations STA1–STA3) may be a single-MAC device. Accordingly, only the primary link may be used to contend for access to the shared wireless medium. In some other implementations, the WLAN 120 may include a combination of multi-MAC devices and single-MAC devices. For example, the AP 110 may be a multi-MAC device (having two or more wireless radios) and the wireless stations STA1–STA3 may be single-MAC devices (having a single wideband radio). As noted above, a multi-MAC device may contend on each of the communication links, independently, and may thus have a competitive advantage over single-MAC devices in obtaining access to the shared wireless medium (via at least one of the communication links). To ensure each of the wireless devices in the WLAN 120 has a fair chance of accessing the shared wireless medium, multi-MAC devices may also be configured to contend for medium access on only the primary link. In some implementations, where the AP 110 is a multi-MAC device, the AP 110 may contend for medium access on only the primary link when attempting to transmit downlink (DL) data. However, when attempting to transmit management and control frames (such as beacons, probe responses, and the like) or responding to uplink (UL) transmissions by a particular STA, the AP 110 may contend for medium access on each of the primary and secondary links separately and concurrently.

[0046] Aspects of this disclosure recognize that, because single-MAC devices must contend for medium access on only the primary link, the overall throughput of the WLAN 120 may be limited by

the load on the primary link. For example, as long as the primary link is busy (due to ongoing transmissions by a wireless device in the current BSS or an overlapping BSS), no other wireless device may utilize the shared wireless medium even if the secondary link is idle. Thus, to improve the throughput and channel utilization in a multi-link environment, the implementations described herein may dynamically change (or reallocate) the primary link based at least in part on channel conditions on each of the communication links. For example, when the current primary link is heavily loaded (due to heavy traffic or interference on the corresponding channel), and the load on the secondary link is relatively light, the primary and secondary links may switch roles. More specifically, single-MAC devices (and some multi-MAC devices) may contend for medium access on the wireless channels associated with the existing secondary link (now the new primary link), and may utilize the wireless channels associated with the existing primary link (now the new secondary link) for purposes of link aggregation only.

[0047] Figure 2 shows a timing diagram 200 depicting an example multi-link (ML) communication. For purposes of discussion, the example ML communication may be performed by a STA in a wireless network. With reference for example to the wireless system 100 of Figure 1, the STA may be one of the wireless stations STA1–STA3. However, in actual implementations, the example ML communication depicted in Figure 2 may be performed by any of the wireless devices (including the AP 110) in the WLAN 120. In some implementations, the STA may be a single-MAC device that is configured to communicate with an AP via multiple communication links. For example, a first link (Link1) may encompass a first set of channels spanning one or more frequency bands and a second link (Link2) may encompass a second set of channels spanning the same or different frequency bands. Although only two communication links are shown in Figure 2, it is to be understood that the STA may establish any number of communication links with the AP.

[0048] At time t_0 , Link1 may be configured as the primary link and Link2 may be configured as the secondary link. At time t_1 , the STA attempts to access the wireless medium. For example, the STA may have UL data to transmit to the AP or another STA in the wireless network. Since Link1 is the primary link, the STA may attempt to access the shared wireless medium via Link1. In the example of Figure 2, Link1 may be busy when the STA attempts to access the wireless medium at time t_1 . For example, using energy or power detection techniques, the STA may sense that the energy or power in the wireless channel associated with Link1 exceeds a threshold level. The activity on Link1 may be caused by transmissions from other wireless devices in the vicinity of the STA (such as APs or STAs in the same BSS or an overlapping BSS) or various sources of wireless interference. Since Link1 is busy, the STA may refrain from even contending for access to the shared wireless medium at time t_1 .

[0049] In some implementations, the STA may attempt to access the shared wireless medium for at least a threshold duration (such as from times t_1 to t_2). For example, the STA may continue

monitoring the activity (energy or power) level of Link1, for the threshold duration, to determine whether the communications link is available for medium access. However, in the example of Figure 2, Link1 remains busy for at least the entirety of the threshold duration. In some implementations, the threshold duration may correspond with the duration of a beacon interval. It is noted that Link2 is idle for the duration in which the STA attempts to access the wireless medium on Link1 (from times t_1 to t_2). However, because the STA is a single-MAC device and Link2 is a secondary link, the STA may be unable to contend for medium access on Link2. In some implementations, the STA may change or switch its primary link (from Link1 to Link2) if one or more conditions are satisfied when the threshold duration expires. For example, the STA may switch its current primary link with a secondary link if the load (activity or interference) on the primary link is above a threshold level, the load on the secondary link is below a threshold level, the load on the secondary link is less than the load on the primary link, or any combination thereof.

[0050] Thus, rather than wait for the current primary link (Link1) to become idle, the STA may change its primary link from Link1 to Link2 at time t_2 . Upon switching the primary link to Link2, the STA may detect that Link2 is idle and may thus contend for access to the shared wireless medium on Link 2. For example, the STA may wait a random back-off time before initiating a UL transmission (TX) on Link2. In the example of Figure 2, the STA wins access to the shared wireless medium, at time t_3 , and is thus able to successfully transmit a frame (or PPDU) of UL data on Link2. While communicating on Link2, the STA may monitor the activity level on Link1 to determine whether the communication link is available for link aggregation. Since Link1 is busy at time t_3 , the STA may transmit the UL frame only on Link2. At time t_4 , the STA may transmit another frame of UL data on the shared wireless medium. Since Link1 is still busy at time t_4 , the STA may again transmit the UL frame only on Link2.

[0051] At time t_5 , the STA may transmit another frame of UL data on the shared wireless medium. At this time, Link1 is no longer busy. Thus, the STA may transmit a frame of UL data on Link2 and another frame of UL data, concurrently, on Link1. For example, the STA may use link aggregation techniques to aggregate the remaining UL data for transmission across the primary link (Link2) and the secondary link (Link1). Upon completion of the multi-link (ML) transmission at time t_5 , the STA's access to the shared wireless medium may be terminated. For example, the STA may have no more UL data in its queue or the TXOP assigned to the STA may have expired. Accordingly, the STA may relinquish control of Link2 to allow other devices to access the shared wireless medium.

[0052] At time t_6 , the STA may once again attempt to access the shared wireless medium. Since Link2 is the current primary link, the STA may attempt to access the shared wireless medium via Link2. However, in the example of Figure 2, Link2 may be busy when the STA attempts to access the wireless medium at time t_6 . For example, using energy or power detection techniques, the STA may sense that

the energy or power in the wireless channel associated with Link2 exceeds a threshold level. The activity on Link2 may be caused by transmissions from other wireless devices in the vicinity of the STA (such as APs or STAs in the same BSS or an overlapping BSS) or various sources of wireless interference. Since Link2 is busy, the STA may refrain from contending for access to the shared wireless medium at time t_6 .

[0053] In some implementations, the STA may attempt to access the shared wireless medium for at least a threshold duration (such as from times t_6 to t_7). However, in the example of Figure 2, Link2 remains busy for at least the entirety of the threshold duration. It is noted that Link1 is idle for the duration in which the STA attempts to access the wireless medium on Link2 (from times t_6 to t_7). Thus, in some implementations, the STA may change or switch its primary link (from Link2 to Link1) if one or more conditions are satisfied when the threshold expires. For example, the STA may switch its current primary link with a secondary link if the load on the primary link is above a threshold level, the load on the secondary link is below a threshold level, the load on the secondary link is less than the load on the primary link, or any combination thereof.

[0054] Thus, rather than wait for the current primary link (Link2) to become idle, the STA may once again switch its primary link from Link2 back to Link1 at time t_7 . Upon switching the primary link to Link1, the STA may detect that Link1 is idle and may thus contend for access to the shared wireless medium on Link1. In the example of Figure 2, the STA wins access to the shared wireless medium, at time t_8 , and is thus able to successfully transmit a frame (or PPDU) of UL data on Link1. While communicating on Link1, the STA may monitor the activity level on Link2 to determine whether the communication link is available for link aggregation. Since Link2 is busy at time t_8 , the STA may transmit the UL frame only on Link1. At time t_9 , the STA may transmit another frame of UL data on the shared wireless medium. At this time, Link2 is no longer busy. Thus, the STA may transmit a frame of UL data on Link1 and another frame of UL data, concurrently, on Link2 (using link aggregation).

[0055] In the example ML communication described above, the role of “primary link” is not fixed or permanently assigned to any particular one of the communication links Link1 or Link2. Rather, the primary link may be dynamically changing (switching between Link1 and Link2) based, at least in part, on the channel conditions of the various communication links. By using a “floating” primary link, the STA (and other wireless devices in the network) may adapt to changing channel conditions in the shared wireless medium, for example, to ensure efficient utilization of the available communication links or wireless channels. For example, when the current primary link becomes overloaded (compared to a secondary link), the STA may change or reconfigure its primary link so that a less-loaded communication link becomes the new primary link (such as at times t_2 and t_7 of Figure 2). This allows the STA to take advantage of the available resources on the less-loaded communication link

to gain access to the shared wireless medium. In contrast, a STA that uses a fixed primary link may have to wait a relatively long time for the primary link to become idle before it can even contend for access to the shared wireless medium.

[0056] The decision to switch the primary link with a secondary link may depend on the channel conditions associated with each of the communication links. For example, it may be desirable to change the primary link if a secondary link is less busy or underutilized compared to the current primary link. On the other hand, it may not be desirable to change the primary link if the secondary link(s) is even more heavily loaded than the current primary link. In some wireless networks, the AP may have knowledge of the channel conditions on each of the communication links. For example, the AP may monitor the traffic patterns or activity of its associated STAs on each of the primary link and the secondary link. Alternatively, or in addition, the AP may learn the channel conditions on each of the communication links by correlating with past channel conditions on the primary link and the secondary link.

[0057] In some implementations, the AP may manage or control the assignment of the primary link in a multi-link environment. For example, the AP may monitor the channel conditions on each of the communication links, and may selectively assign (and re-assign) the role of primary link to one of the communication links based at least in part on the channel conditions associated with each of the communication links. In some aspects, the AP may signal changes in the primary link via beacon frames broadcast at target beacon transmission times (TBTTs). For example, beacon frames are typically broadcast at regular intervals (referred to as “beacon intervals”) to enable any STAs in the vicinity of the AP to establish or maintain a connection to the wireless network. Accordingly, most (if not all) STAs in the wireless network may be expected to listen for such beacon frames from the AP at the start of each beacon interval. In some aspects, the AP may dynamically change or reassign the primary link to optimize throughput across the multiple communication links on a per-TBTT basis.

[0058] Figure 3A shows a timing diagram 300A depicting an example operation for changing the primary link in a multi-link environment. The example operation may be performed by an AP in a wireless network such as, for example, the AP 110 of Figure 1. In some implementations, the AP may be a single-MAC device that uses a single MAC to communicate with one or more associated STAs via multiple communication links. In some other implementations, the AP may be a multi-MAC device that uses multiple MACs to communicate with one or more associated STAs via respective communication links. For example, a first link (Link1) may encompass a first set of channels spanning one or more frequency bands and a second link (Link2) may encompass a second set of channels spanning the same or different frequency bands. Although only two communication links are shown in Figure 3A, it is to be understood that the AP may establish any number of communication links with its associated STAs.

[0059] At time t_0 , the AP may broadcast respective beacon frames (Beacon1 and Beacon2) on each of the communication links Link1 and Link2. For example, time t_0 may coincide with a TBTT or the start of a beacon interval (from times t_0 to t_1). In some aspects, although not shown for simplicity, the AP may contend for medium access at the start of each beacon interval. During this time, Link1 may be configured as the primary link and Link2 may be configured as the secondary link. Thus, if the AP is a single-MAC device, the AP may contend for medium access on Link1 to broadcast Beacon1 and, upon gaining access to shared wireless medium via Link1, the AP may use link aggregation to broadcast Beacon2 on Link2. On the other hand, if the AP is a multi-MAC device, the AP may contend for medium access on each of the communication links Link1 and Link2, concurrently, and may independently broadcast the beacon frames Beacon1 and Beacon2 upon gaining access to the respective communication links.

[0060] It is noted that ML STAs in the wireless network may listen to the primary link (Link1) for beacon frames (Beacon1) broadcast by the AP. However, the AP also may support legacy STAs that may not be capable of ML functionality. More specifically, legacy STAs may not be capable of concurrently communicating over multiple communication links. Thus, legacy STAs may not distinguish between a primary link or secondary link. In some instances, such legacy STAs may be associated with the AP on Link2, and may thus only listen to Link2 for beacon frames (Beacon2) broadcast by the AP. Accordingly, it may be desirable for the AP to broadcast beacon frames on each of the available communication links (including the primary link and any secondary links).

[0061] In some implementations, the AP may monitor the channel conditions on each of the communication links Link1 and Link2, for the duration of the beacon interval (from times t_0 to t_1), to determine whether the quality or throughput of ML communications can be improved by dynamically changing the primary link. More specifically, in determining whether to change the primary link, the AP may consider the load (activity or interference) on Link1 and Link2. For example, it may be desirable to change the primary link if the load on the primary link is above a threshold level, the load on the secondary link is below a threshold level, the load on the secondary link is less than the load on the primary link, or any combination thereof.

[0062] At time t_1 , the AP may switch the primary link from Link1 to Link2. For example, time t_1 may coincide with a TBTT or the start of a subsequent beacon interval (from times t_1 to t_2). Thus, the AP may signal the change in primary link via one or more beacon frames broadcast at the start of the beacon interval. Since Link1 is configured as the primary link at this time (and thus ML STAs are currently listening to Link1 for beacon information), the AP may provide primary link (PL) information in Beacon1, broadcast on Link1, at time t_1 . For example, the PL information may be provided in a new element or field in the beacon frame that can be interpreted or decoded by ML STAs. In some aspects, the PL information may indicate that the primary link has been changed. In some other aspects (where

there are multiple secondary links), the PL information may specify which secondary link is to become the new primary link.

[0063] Upon receiving the PL information on Link1, any ML STAs in the wireless network may switch their primary links from Link1 to Link2. Thus, the ML STAs may subsequently listen for beacon frames and contend for medium access on Link2. In some implementations, the AP may include the PL information in other management or control frames (such as probe responses, association responses, reassociation responses, and the like) transmitted to any associated STAs after the change in primary link occurs at time t_1 . For example, the AP may signal the current primary link (Link2) to any STAs that may have recently joined the wireless network or otherwise missed Beacon1 broadcast at time t_1 . In some implementations, the AP may continue to monitor the channel conditions on each of the communication links Link1 and Link2, for the duration of the beacon interval (from times t_1 to t_2), to determine whether the quality or throughput of ML communications can be improved by dynamically changing the primary link.

[0064] At time t_2 , coinciding with the next TBTT, the AP may decide to maintain the primary link on Link2. For example, the AP may determine that the quality or throughput of communications would not be improved by switching the primary link back to Link1. Thus, the AP may broadcast respective beacon frames Beacon1 and Beacon2 on each of the communication links Link1 and Link2, at time t_2 , without signaling a change in primary link. It is noted that Link2 is now configured as the primary link and Link1 is now configured as the secondary link. Thus, if the AP is a single-MAC device, the AP may contend for medium access on Link2 to broadcast Beacon2 and, upon gaining access to the shared wireless medium via Link2, the AP may use link aggregation to broadcast Beacon1 on Link1. On the other hand, if the AP is a multi-MAC device, the AP may contend for medium access on each of the communication links Link1 and Link2, concurrently, and may independently broadcast the beacon frames Beacon1 and Beacon2 upon gaining access to the respective communication links.

[0065] At time t_3 , coinciding with the next TBTT, the AP may switch the primary link from Link2 back to Link1. For example, the AP may signal the change in primary link via one or more beacon frames broadcast at the start of the beacon interval. Since Link2 is configured as the primary link at this time (and thus ML STAs are currently listening to Link2 for beacon information), the AP may provide PL information in Beacon2, broadcast on Link2, at time t_3 . As described above, the PL information may indicate that the primary link has changed or may specify which secondary link is to become the new primary link. Upon receiving the PL information on Link2, any ML STAs in the wireless network may switch their primary links from Link2 back to Link1. Thus, the ML STAs may subsequently listen for beacon frames and contend for medium access on Link1. In some implementations, the AP also may change the PL information in other management or control frames transmitted to any associated STAs after the change in primary link occurs at time t_1 .

[0066] At time t_4 , coinciding with the next TBTT, the AP may once again switch the primary link from Link1 to Link2. Since Link1 is configured as the primary link at this time, the AP may provide PL information in Beacon1, broadcast on Link1, at time t_4 . As described above, the PL information may indicate that the primary link has changed or may specify which secondary link is to become the new primary link. Upon receiving the PL information on Link1, any ML STAs in the wireless network may switch their primary links from Link1 to Link2. Thus, the ML STAs may subsequently listen for beacon frames and contend for medium access on Link2. In some implementations, the AP also may change the PL information in other management or control frames transmitted to any associated STAs after the change in primary link occurs at time t_4 .

[0067] As shown in Figure 3A, the example implementations may enable the AP to change the primary link used in ML communications as often as every beacon interval, depending on the current channel conditions. It is noted, however, that some associated ML STAs may be in a power save (PS) state or may otherwise fail to receive the PL information signaled in a beacon frame (due to interference or service interruptions). For example, an ML STA with its primary link set to Link1 may be in a power save state when Beacon1 (carrying the PL information) is broadcast at time t_1 . As a result, if the ML STA wakes up thereafter (between times t_1 and t_3), the ML STA may continue attempting to access the shared wireless medium via Link1 (and thus experience relatively poor performance) even though the primary link has been switched to Link2. This may result in inefficient use of the shared wireless medium. Thus, in some implementations, the AP may specify the current primary link via beacon frames broadcast on each of the communication links at the start of each beacon interval.

[0068] Figure 3B shows a timing diagram 300B depicting another example operation for changing the primary link in a multi-link environment. The example operation may be performed by an AP in a wireless network such as, for example, the AP 110 of Figure 1. In some implementations, the AP may be a single-MAC device that uses a single MAC to communicate with one or more associated STAs via multiple communication links. In some other implementations, the AP may be a multi-MAC device that uses multiple MACs to communicate with one or more associated STAs via respective communication links. For example, a first link (Link1) may encompass a first set of channels spanning one or more frequency bands and a second link (Link2) may encompass a second set of channels spanning the same or different frequency bands. Although only two communication links are shown in Figure 3B, it is to be understood that the AP may establish any number of communication link with its associated STAs.

[0069] At time t_0 , the AP may broadcast respective beacon frames (Beacon1 and Beacon2) on each of the communication links Link1 and Link2. For example, time t_0 may coincide with a TBTT or the start of a beacon interval (from times t_0 to t_1). In some implementations, the AP may provide primary link (PL) information indicating the current primary link (Link1) in one or more beacon frames

broadcast at the start of the beacon interval. For example, the PL information may be included as a new element or field in each beacon frame that can be interpreted or decoded by ML STAs. In some aspects, the AP may include the PL information in the beacon frames broadcast on each of the communication links Link1 and Link2. More specifically, Beacon1 and Beacon2 may each include PL information specifying Link1 as the current primary link.

[0070] Any ML STAs receiving the PL information via Beacon1 may continue using Link1 as their primary link, and any ML STAs receiving the PL information via Beacon2 may switch their current primary link from Link2 to Link1. In some implementations, the AP may include the PL information in other management or control frames (such as probe responses, association response, reassociation responses, and the like) transmitted to any associated STAs, for example, to signal the current primary link (Link1) to any STAs that may have recently joined the wireless network or otherwise missed the beacons broadcast at time t_1 . In some implementations, the AP may monitor the channel conditions on each of the communication links Link1 and Link2, for the duration of the beacon interval (from times t_0 to t_1), to determine whether the quality or throughput of ML communications can be improved by dynamically changing the primary link.

[0071] At time t_1 , the AP may switch the primary link from Link1 to Link2. For example, the AP may determine that the quality or throughput of communications would be improved by switching the primary link to Link2. In some implementations, the AP may signal the change in primary link using the PL information included in each of the beacon frames Beacon1 and Beacon2 broadcast at time t_1 . For example, the PL information may specify a new communication link (Link2) as the current primary link. In some aspects, the AP may include the PL information in the beacon frames broadcast on each of the communication links Link1 and Link2. More specifically, Beacon1 and Beacon2 may each include PL information specifying Link2 as the current primary link. Any ML STAs receiving the PL information via Beacon1 may switch their current primary link from Link1 to Link2, and any ML STAs receiving the PL information via Beacon2 may continue using Link2 as their current primary link.

[0072] At time t_2 , the AP may decide to maintain the primary link on Link2. For example, the AP may determine that the quality or throughput of communications would not be improved by switching the primary link back to Link1. In some implementations, although the primary link does not change, the AP may continue to provide PL information indicating the current primary link (Link2) in each of the beacon frames Beacon1 and Beacon2 broadcast at time t_2 . This may allow any ML STAs that may have missed the beacons broadcast at time t_1 to be notified of the change in primary link. For example, an ML STA with its primary link set to Link1 may have missed Beacon1 (carrying the PL information) at time t_1 . As a result, the ML STA may subsequently attempt to access the shared wireless medium via Link1 for the duration of the beacon interval (from times t_1 to t_2). However, by indicating the current primary link in each of the beacon frames Beacon1 and Beacon2 broadcast at the

start of each beacon interval, the ML STA may be notified of the change in primary link upon receiving Beacon1, at time t_2 , and may thus switch its primary link to Link2 for at least the duration of the next beacon interval (from times t_2 to t_3).

[0073] At time t_3 , the AP may switch the primary link from Link2 back to Link1. For example, the AP may signal the change in primary link using the PL information included in each of the beacon frames Beacon1 and Beacon2 broadcast at time t_3 . More specifically, the PL information may specify a new communication link (Link1) as the current primary link. In some aspects, Beacon1 and Beacon2 may each include PL information specifying Link1 as the current primary link. Any ML STAs receiving the PL information via Beacon1 may continue using Link1 as their primary link, and any ML STAs receiving the PL information via Beacon2 may switch their current primary link from Link2 to Link1.

[0074] At time t_4 , the AP may once again switch the primary link from Link1 to Link2. For example, the AP may signal the change in primary link using the PL information included in each of the beacon frames Beacon1 and Beacon2 broadcast at time t_4 . More specifically, the PL information may specify a new communication link (Link2) as the current primary link. In some aspects, Beacon1 and Beacon2 may each include PL information specifying Link2 as the current primary link. Any ML STAs receiving the PL information via Beacon1 may switch their current primary link from Link1 to Link2, and any ML STAs receiving the PL information via Beacon2 may continue using Link2 as their primary link.

[0075] Figure 4A shows a timing diagram 400A depicting an example operation for notifying a STA of a primary link change. The example operation may be performed during ML communications between an AP and a STA. With reference for example to Figure 1, the AP may correspond to AP 110 and the STA may correspond to any one of the wireless stations STA1–STA3. In some implementations, the STA may be a single-MAC device that is configured to communicate with the AP via multiple communication links. Although only two communication links (Link1 and Link2) are shown in Figure 4A, it is to be understood that the STA may establish any number of communication links with the AP. For simplicity, only the AP's activity on Link1 is shown in example of Figure 4A.

[0076] At time t_0 , Link1 may be configured as the primary link and Link2 may be configured as the secondary link. Furthermore, time t_0 may coincide with a TBTT or the start of a beacon interval (from times t_0 to t_2). Thus, the AP may broadcast a beacon frame (Beacon1) on Link1 at time t_0 . In some implementations, Beacon1 may include PL information specifying Link1 as the current primary link. The STA may already be configured to utilize Link1 as its current primary link. Thus, the STA may receive Beacon1 on Link1 and continue utilizing Link1 as its primary link. In the example of Figure 4A, the STA may enter a power save (PS) state, at time t_1 , after receiving the PL information specifying Link1 as the primary link.

[0077] At time t_2 , the AP may switch the primary link from Link1 to Link2. For example, the AP may determine that the quality or throughput of communications would be improved by switching the primary link to Link2. In some implementations, the AP may signal the change in primary link using the PL information included in Beacon1 broadcast at time t_2 . For example, the PL information may specify a new communication link (Link2) as the current primary link. In the example of Figure 4A, the STA is still in a power save state when the AP broadcasts Beacon1, at time t_2 , and may thus be unaware of the change in primary link.

[0078] At time t_3 , the STA wakes up from the power save state and attempts to access the shared wireless medium. Since the STA is unaware that the primary link has been changed to Link2, STA may sense that Link1 is idle and contend for medium access on Link1. Upon gaining access to the shared wireless medium, at time t_4 , the STA may transmit a frame of UL data on Link1. In the example of Figure 4A, the STA may not have any additional UL data to transmit, and may thus relinquish control of the shared wireless medium after transmitting the single frame of UL data.

[0079] At time t_5 , the STA may attempt to access the shared wireless medium once again on Link1. However, at this time Link1 is busy. For example, using energy or power detection techniques, the STA may sense that the energy or power in the wireless channel associated with Link1 exceeds a threshold level. The activity on Link1 may be caused by transmission from other wireless devices in the vicinity of the STA (such as APs or STAs in the same BSS or an overlapping BSS) or various sources of wireless interference. Since Link1 is busy, the STA may refrain from even contending for access to the shared wireless medium at time t_5 .

[0080] In the example of Figure 4A, the STA may continue monitoring the activity level of Link1, from times t_5 to t_6 , to determine whether the communications link is available for medium access. Then, at time t_6 , the STA may detect that Link1 is idle and once again contend for medium access on Link1. Upon gaining access to the shared wireless medium, at time t_7 , the STA may transmit a frame of UL data on Link1. In the example of Figure 4A, the STA may not have enough time to transmit any additional frames of UL data before the start of the next beacon interval (when the AP is expected to broadcast beacon frames), and may thus relinquish control of the shared wireless medium after transmitting the single frame of UL data.

[0081] At time t_8 , the AP may broadcast another beacon frame (Beacon1) on Link1. In the example of Figure 4A, the AP may decide to maintain the primary link on Link2. For example, the AP may determine that the quality or throughput of communications would not be improved by switching the primary link back to Link1. In some implementations, even though the primary link does not change, the AP may continue to provide PL information indicating the current primary link (Link2) in Beacon1 broadcast at time t_8 . Since the STA is awake at this time, the STA may receive Beacon1 from the AP and subsequently switch its primary link from Link1 to Link2.

[0082] In the example of Figure 4A, the STA must wait the majority of a beacon interval (from times t_3 to t_8) until it is able to switch its primary link from Link1 to Link2. Moreover, a significant portion of that time (from times t_3 to t_8) is spent monitoring the busy channel conditions on Link1. Thus, it may be desirable to implement a more efficient mechanism for signaling a change in primary link to any associated STAs that may have missed one or more beacon frames.

[0083] Figure 4B shows a timing diagram 400B depicting another example operation for notifying a STA of a primary link change. The example operation may be performed during ML communications between an AP and a STA. With reference for example to Figure 1, the AP may correspond to AP 110 and the STA may correspond to any one of the wireless stations STA1–STA3. In some implementations, the STA may be a single-MAC device that is configured to communicate with the AP via multiple communication links. Although only two communication links (Link1 and Link2) are shown in Figure 4B, it is to be understood that the STA may establish any number of communication links with the AP. For simplicity, only the AP's activity on Link1 is shown in the example of Figure 4B.

[0084] At time t_0 , Link1 may be configured as the primary link and Link2 may be configured as the secondary link. Furthermore, time t_0 may coincide with a TBTT or the start of a beacon interval (from times t_0 to t_2). Thus, the AP may broadcast a beacon frame (Beacon1) on Link1 at time t_0 . In some implementations, Beacon1 may include PL information specifying Link1 as the current primary link. The STA may already be configured to utilize Link1 as its current primary link. Thus, the STA may receive Beacon1 on Link1 and continue utilizing Link1 as its primary link. In the example of Figure 4B, the STA may enter a power save (PS) state, at time t_1 , after receiving the PL information specifying Link1 as the primary link.

[0085] At time t_2 , the AP may switch the primary link from Link1 to Link2. For example, the AP may determine that the quality or throughput of communications would be improved by switching the primary link to Link2. In some implementations, the AP may signal the change in primary link using the PL information included in Beacon1 broadcast at time t_2 . For example, the PL information may specify a new communication link (Link2) as the current primary link. In the example of Figure 4B, the STA is still in a power save state when the AP broadcasts Beacon1, at time t_2 , and may thus be unaware of the change in primary link.

[0086] At time t_3 , the STA wakes up from the power save state and attempts to access the shared wireless medium. Since the STA is unaware that the primary link has been changed to Link2, the STA may sense that Link1 is idle and contend for medium access on Link1. Upon gaining access to the shared wireless medium, at time t_4 , the STA may transmit a frame of UL data via the shared wireless medium. In the example of Figure 4B, the STA may not have any additional UL data to transmit, and

may thus relinquish control of the shared wireless medium after transmitting the single frame of UL data.

[0087] In some implementations, the AP may respond to the UL transmission with a change-in-primary-link (CPL) trigger message. The CPL trigger message may include PL information specifying the current primary link (Link2). In some aspects, the CPL trigger message may be signaled using Operating Mode Indication (OMI) procedures defined by the IEEE 802.11ax specification. For example, data frames encoded in accordance with the HE frame format may include an OMI control field (provided in an HE Control element) that may indicate changes to one or more operating parameters for wireless communications between the transmitting device and the receiving device. When the AP receives an HE data frame from the STA, the AP may send a response frame (such as a QoS Null frame) having an OMI control field indicating the change in primary link. Thus, the PL information (specifying Link2 as the current primary link) may be included in the OMI control field of the response frame transmitted by AP. In some implementations, the STA may be configured to initiate a UL transmission upon waking up from a power save state (even if it has no UL data to send) to elicit an OMI response from the AP. For example, the STA may send a PS Poll request to the AP upon waking up from a power save state, and the AP may respond with a QoS Null frame with an OMI control field indicating the change in primary link.

[0088] Upon receiving the CPL trigger message on Link1, the STA may subsequently switch its primary link from Link1 to Link2. Then, at time t_5 , the STA may attempt to access the shared wireless medium on Link2. The STA may sense that Link2 is idle and contend for medium access on Link2. Upon gaining access to the shared wireless medium, at time t_6 , the STA may transmit another frame of UL data via the shared wireless medium. In the example of Figure 4B, Link1 is busy at time t_6 . Thus, the STA may transmit the UL frame only on Link2. At time t_7 , the STA may transmit another frame of UL data via the shared wireless medium. Since Link1 is still busy at time t_7 , the STA may again transmit the UL frame only on Link2.

[0089] At time t_8 , the STA may transmit another frame of UL data on the shared wireless medium. At this time, Link1 is no longer busy. Thus, the STA may transmit a frame of UL data on Link2 and another frame of UL data, concurrently, on Link1. For example, the STA may use link aggregation techniques to aggregate multiple frames of UL data for transmission across the primary link (Link2) and the secondary link (Link1). In the example of Figure 4B, the STA may not have enough time to transmit any additional frames of UL data before the start of the next beacon interval (when the AP is expected to broadcast beacon frames). Thus, the STA may relinquish control of the shared wireless medium upon completion of the multi-link (ML) transmission.

[0090] At time t_9 , the AP may broadcast another beacon frame (Beacon1) on Link1. In the example of Figure 4B, the AP may decide to maintain the primary link on Link2. For example, the AP

may determine that the quality or throughput of communications would not be improved by switching the primary link back to Link1. In some implementations, even though the primary link does not change, the AP may continue to provide PL information indicating the current primary link (Link2) in Beacon1 broadcast at time t_9 . Since the STA has already switched its primary link to Link2, the STA may not receive Beacon 1 on Link1. Rather, the STA may continue utilizing Link2 as its primary link.

[0091] Figure 4C shows a timing diagram 400C depicting another example operation for notifying a STA of a primary link change. The example operation may be performed during ML communications between an AP and a STA. With reference for example to Figure 1, the AP may correspond to AP 110 and the STA may correspond to any one of the wireless stations STA1–STA3. In some implementations, the STA may be a single-MAC device that is configured to communicate with the AP via multiple communication links. Although only two communication links (Link1 and Link2) are shown in Figure 4C, it is to be understood that the STA may establish any number of communication links with the AP. For simplicity, only the AP's activity on Link1 is shown in the example of Figure 4C.

[0092] At time t_0 , Link1 may be configured as the primary link and Link2 may be configured as the secondary link. Furthermore, time t_0 may coincide with a TBTT or the start of a beacon interval (from times t_0 to t_2). Thus, the AP may broadcast a beacon frame (Beacon1) on Link1 at time t_0 . In some implementations, Beacon1 may include PL information specifying Link1 as the current primary link. The STA may already be configured to utilize Link1 as its current primary link. Thus, the STA may receive Beacon1 on Link1 and continue utilizing Link1 as its primary link. In some implementations, Beacon1 may further include a Target Wake Time (TWT) element specifying a TWT service period for the STA. For example, a STA operating in a TWT mode may remain in a power save state for relatively long durations (such as one or more beacon intervals), waking up only during its assigned TWT service periods. Thus, the STA may enter a power save state, at time t_1 , after receiving the PL information specifying Link1 as the primary link.

[0093] At time t_2 , the AP may switch the primary link from Link1 to Link2. For example, the AP may determine that the quality or throughput of communications would be improved by switching the primary link to Link2. In some implementations, the AP may signal the change in primary link using the PL information included in Beacon1 broadcast at time t_2 . For example, the PL information may specify a new communication link (Link2) as the current primary link. In the example of Figure 4C, the STA is still in a power save state when the AP broadcasts Beacon1, at time t_2 , and may thus be unaware of the change in primary link.

[0094] At time t_3 , the STA wakes up at its scheduled TWT. In some implementations, the AP may be aware of the STA's TWT schedule and may thus transmit a CPL trigger message to the STA, at time t_4 , once the STA wakes up from its power save state. The CPL trigger message may include PL

information specifying the current primary link (Link2). In some implementations, the CPL trigger message may be provided in a probe response or action frame (or other management or control frame) transmitted at the start of the TWT service period to notify the STA of the current primary link. In some aspects, the CPL trigger message may be signaled on each of the communication links (Link1 and Link2) to ensure that the STA receives the PL information regardless of its current primary link.

[0095] Upon receiving the CPL trigger message on Link1, the STA may subsequently switch its primary link from Link1 to Link2. Then, at time t_5 , the STA may attempt to access the shared wireless medium on Link2. The STA may sense that Link2 is idle and contend for medium access on Link2. Upon gaining access to the shared wireless medium, at time t_6 , the STA may transmit a frame of UL data via the shared wireless medium. In the example of Figure 4C, Link1 is busy at time t_6 . Thus, the STA may transmit the UL frame only on Link2. At times t_7 and t_8 , the STA may transmit additional frames of UL data via the shared wireless medium. Since Link1 remains busy throughout this period (times t_7 and t_8) the STA may again transmit the UL frames only on Link2.

[0096] At time t_9 , the STA may transmit another frame of UL data on the shared wireless medium. At this time, Link1 is no longer busy. Thus, the STA may transmit a frame of UL data on Link2 and another frame of UL data, concurrently, on Link1. For example, the STA may use link aggregation techniques to aggregate multiple frames of UL data for transmission across the primary link (Link2) and the secondary link (Link1). In the example of Figure 4C, the STA may not have enough time to transmit any additional frames of UL data before the start of the next beacon interval (when the AP is expected to broadcast beacon frames). Thus, the STA may relinquish control of the shared wireless medium upon completion of the multi-link (ML) transmission.

[0097] At time t_{10} , the AP may broadcast another beacon frame (Beacon1) on Link1. In the example of Figure 4C, the AP may decide to maintain the primary link on Link2. For example, the AP may determine that the quality or throughput of communications would not be improved by switching the primary link back to Link1. In some implementations, even though the primary link does not change, the AP may continue to provide PL information indicating the current primary link (Link2) in Beacon1 broadcast at time t_{10} . Since the STA has already switched its primary link to Link2, the STA may not receive Beacon 1 on Link1. Rather, the STA may continue utilizing Link2 as its primary link.

[0098] Figure 5 shows a timing diagram 500 depicting an example operation for changing the primary link in a ML communication. The operation may be performed during ML communications between an AP and a number of associated wireless stations STA1 and STA2. With reference for example to Figure 1, the AP may correspond to AP 110 and each of the wireless stations STA1 or STA2 may correspond to any of the wireless stations STA1–STA3. In some implementations, each of the wireless stations STA1 and STA2 may be a single-MAC device that is configured to communicate with the AP via multiple communication links. In the example of Figure 5, the ML communications between

the AP and each of the wireless stations STA1 and STA2 may occur on a first communication link (Link1).

[0099] At time t_0 , the AP may broadcast a beacon frame (Beacon1) on Link1. At this time, Link1 may be set as the primary link. In some implementations, Beacon1 may include PL information indicating a change in primary link that is scheduled to take effect at a future primary-switching time (TBTT_{PS}), as well as a countdown towards the primary-switching time. In some aspects, the PL information also may specify the communication link (Link2) that will become the new primary link at the primary-switching time. The countdown may indicate the number of TBTT periods or beacon intervals remaining before the change in primary link is scheduled to occur. Thus, in some aspects, Beacon1 broadcast at time t_0 may include a countdown timer indicating that the change in primary link is to occur in three successive beacon intervals (Count=3). In the example of Figure 5, STA1 is awake at time t_0 , whereas STA2 is in a power saving state. As a result, STA1 may receive the PL information, at time t_0 , whereas STA2 may not. Upon receiving the PL information, STA1 may be configured to switch its primary link from Link1 to Link2 at the indicated primary-switching time (after three successive TBTTs).

[00100] At time t_1 , STA2 wakes up from its power save state to access the shared wireless medium. In some aspects, time t_1 may coincide with a DTIM interval for which most (if not all) of the STAs in the wireless network are scheduled to listen for beacon frames from the AP. In some other aspects, time t_1 may coincide with the a TWT (or the start of a TWT service period) for which STA2 is assigned. Then, at time t_2 , the AP may broadcast another Beacon1 on Link1. In some implementations, Beacon1 broadcast at time t_2 may include updated PL information. As described above, the PL information may include the new primary link (Link2) that is scheduled to take effect at the primary-switching time (TBTT_{PS}), as well as an updated countdown timer (Count=2). The wireless stations STA1 and STA2 may each receive the PL information at time t_2 . Upon receiving the PL information, STA2 may be configured to switch its primary link from Link1 to Link2 at the indicated primary-switching time (after two successive TBTTs).

[00101] At time t_3 , STA2 returns to a power saving state. Then, at time t_4 , the AP broadcasts another Beacon1 on Link1. In some implementations, Beacon1 broadcast at time t_4 may include updated PL information. The PL information may indicate the new primary link (Link2) that is scheduled to take effect at the primary-switching time (TBTT_{PS}), as well as an updated countdown timer (Count=1). STA1 may receive this PL information, whereas STA2 may not (since it is in a power saving state). Nonetheless, STA2 may continue counting down the number of beacon intervals until the change in primary link is to occur (at TBTT_{PS}), for example, based on the countdown timer received, at time t_2 , via PL information provided in Beacon1.

[00102] Finally, at time t_5 , the AP broadcasts another Beacon1 on Link1. In some implementations, Beacon1 broadcast at time t_5 may include updated PL information. More specifically, in the example of Figure 5, time t_5 coincides with the primary-switching time (TBTT_{PS}). Thus, the PL information may indicate the new primary link (Link2), as well as an updated countdown timer (Count=0) indicating that the new primary link is to take effect at this time. STA1 may receive the updated PL information via Beacon1 and may switch its primary link from Link1 to Link2, accordingly, at time t_5 . However, STA2 may remain in a power saving state and therefore may not receive the updated PL information transmitted at time t_5 . Nonetheless, STA2 also may switch its primary link from Link1 to Link2, at time t_5 (or the next time it wakes up from the power save state), based on its own internal countdown towards the primary-switching time (TBTT_{PS}).

[00103] By scheduling changes in primary link to occur at particular epochs, the wireless stations STA1 and STA2 are not required to be awake at the primary-switching time (TBTT_{PS}) in order to carry out the change in primary link. Accordingly, such implementations may provide a greater degree of flexibility for the wireless stations STA1 and STA2 to operate in accordance with their pre-configured power save schedules. For example, the AP may accommodate individual power save schedules of the respective STAs by providing all of the information necessary to execute the change in primary link (at the primary-switching time) when the STAs are awake and listening to the AP. In some implementations, after the primary-switching time has passed, subsequent beacon transmissions by the AP may include information indicating that a change in primary link occurred at a prior time (such as described with respect to Figures 4A–4C). This may further ensure that each STA in the wireless network is properly notified of the change in primary link, including any STAs that may not have received the PL information prior to the primary-switching time (TBTT_{PS}).

[00104] Figure 6 shows a block diagram of an example wireless station (STA) 600 for use in wireless communication according to some implementations. In some implementations, the STA 600 may be a wireless device (such as a ML STA) that is capable of communicating over multiple communication links, concurrently, using link aggregation techniques. For example, the STA 600 may be an example implementation of any of the STAs of Figure 1, Figure 2, Figures 4A–4C, or Figure 5. The STA 600 may include a PHY 610, a MAC 620, a processor 630, a memory 640, and a number of antennas 650(1)–650(n).

[00105] The PHY 610 may include a number of transceivers 612 and a baseband processor 614. The transceivers 612 may be coupled to the antennas 650(1)–650(n), either directly or through an antenna selection circuit (not shown for simplicity). The transceivers 612 may be used to communicate wirelessly with one or more APs, with one or more STAs, or with other suitable devices. The baseband processor 614 may be used to process outgoing signals received from the processor 630 or the memory 640 and to forward the processed signals to the transceivers 612 for transmission via one or more of the

antennas 650(1)–650(n), and may be used to process incoming received from one or more of the antennas 650(1)–650(n) via the transceivers 612 and to forward the processed signals to the processor 630 or the memory 640.

[00106] Although not shown in Figure 6, for simplicity, the transceivers 612 may include any number of transmit chains to process and transmit outgoing signals to other wireless devices via the antennas 650(1)–650(n), and may include any number of receive chains to process incoming signals received from the antennas 650(1)–650(n). Thus, in some implementations, the STA 600 may be configured for MIMO operations including, for example, single-user MIMO (SU-MIMO) operations and multi-user (MU-MIMO) operations. In addition, the STA 600 may be configured for OFDMA communications or other suitable multiple access mechanisms, for example, as may be specified by any of the IEEE 802.11 standards.

[00107] The MAC 620 may include at least a number of contention engines 622 and frame formatting circuitry 624. The contention engines 622 may contend for access to the shared wireless medium on a primary communication link, and may store packets for transmission over the shared wireless medium. In some implementations, the contention engines 622 may be separate from the MAC 620. Still further, in some implementations, the contention engines 622 may be implemented as one or more software modules (stored in the memory 640 or in memory provided within the MAC 620). The frame formatting circuitry 624 may be used to create or format frames received from the processor 630 or the memory 640 (such as by adding MAC headers to PDUs provided by the processor 630), and may be used to re-format frames received from the PHY 610 (such as by stripping the MAC headers from frames received from the PHY 610).

[00108] The memory 640 may include an AP profile data store 641 that stores profile information for a plurality of BSSs. The profile information for a particular BSS may include, for example, the BSSID, MAC address, channel information, received signal strength indicator (RSSI) values, goodput values, channel state information (CSI), supported data rates, connection history with the BSS, a trustworthiness value of the BSS (indicating a level of confidence about the BSS's location or other properties associated with the BSS), and any other suitable information pertaining to or describing the operation of the BSS. In some implementations, the profile information for a particular BSS may include one or more wireless channels to be used as a primary communication link (such as for contention). In some other implementations, the profile information for a particular BSS may include wireless channels to be used as one or more secondary communication links (such as for link aggregation).

[00109] The memory 640 also may include a non-transitory computer-readable medium (one or more nonvolatile memory elements, such as EPROM, EEPROM, Flash memory, a hard drive, and the like) that may store at least the following software (SW) modules:

- a channel monitoring SW module 642 to monitor one or more channel conditions of a plurality of wireless channels available for wireless communications;
- a link configuration SW module 643 to establish a primary communication link on a first subset of the wireless channels and to selectively access one or more secondary communication links on one or more subsets of the remaining wireless channels, the link configuration SW module 643 including:
 - a primary link information (PLI) retrieval submodule 644 to acquire primary link information from an AP (such as upon waking up from a power save state);
 - a primary link (PL) reassignment submodule 645 to dynamically change the primary communication link from the first subset of wireless channels to another subset of the wireless channels;
 - a primary switching (PS) timer submodule 646 to countdown towards a time at which the change in the primary communication link is scheduled to occur; and
- a link aggregation SW module 647 to access one or more of the secondary communication links, concurrently with the primary communication link, using link aggregation techniques, the link aggregation SW module 647 including:
 - a secondary link (SL) selection submodule 648 to select one or more of the secondary communication links for link aggregation based, at least in part, on the channel conditions of the secondary communication links.

Each software module includes instructions that, when executed by the processor 630, cause the STA 600 to perform the corresponding functions.

[00110] For example, the processor 630 may execute the channel monitoring SW module 642 to monitor one or more channel conditions of a plurality of wireless channels available for wireless communications. The processor 630 may execute the link configuration SW module 643 to establish a primary communication link on a first subset of the wireless channels and to selectively access one or more secondary communication links on one or more subsets of the remaining wireless channels. In executing the link configuration SW module 643, the processor 630 may further execute the PLI retrieval submodule 644, the PL switching submodule 645, or the PS timer submodule 646.

[00111] The processor 630 may execute the PLI retrieval submodule 644 to acquire primary link information from an AP. For example, execution of the PLI retrieval submodule 644 may cause the STA 600 to initiate an uplink transmission (such as a PS Poll request) to the AP upon waking up from a power save state and receive a trigger message (such as a QoS Null frame), including the primary link information, from the AP in response to the uplink transmission (which may be transmitted using OMI signaling procedures). In some aspects, the primary link information may indicate a change in the

primary communication link. In some other aspects, the primary link information may indicate the subset of wireless channels to be used for the (new) primary communication link. For example, the new primary communication link may be implemented on one or more wireless channels previously associated with a secondary communication link.

[00112] The processor 630 may execute the PL reassignment submodule 645 to dynamically change the primary communication link from the first subset of wireless channels to another subset of the wireless channels. In some implementations, execution of the PL reassignment submodule 645 may cause the STA 600 to change the primary communication link based, at least in part, on the one or more channel conditions measured by the channel monitoring SW module 642. In some other implementations, execution of the PL reassignment submodule 645 may cause the STA 600 to change the primary communication link based, at least in part, on primary link information received from the AP. In some aspects, the primary link information may be provided in beacon frames broadcast by the AP. In some other aspects, the primary link information may be provided in trigger messages sent by the AP to the STA 600 (where the trigger messages may be transmitted using OMI signaling procedures).

[00113] The processor 630 may execute the PS timer submodule 646 to countdown towards a time at which the change in the primary communication link is scheduled to occur. For example, in some implementations, the primary link information received from the AP may indicate a future time (the primary switching time) at which the change in primary link is to be implemented. The primary switching time may coincide with at least one of a TBTT or a TWT of the STA 600. Thus, in some aspects, the primary link information may further include a countdown timer indicating a number of TBTTs or TWTs remaining until the primary-switching time. In some aspects, the STA 600 may update the countdown timer based on a number of TBTT or TWT intervals that have elapsed. In some other aspects, the STA 600 may update the countdown timer based on updated primary link information received from the AP.

[00114] The processor 630 also may execute the link aggregation SW module 647 to access one or more of the secondary communication links, concurrently with the primary communication link, using link aggregation techniques. In some aspects, the STA 600 may access the secondary communication links only after the STA 600 has already gained access to the shared wireless medium via the primary communication link (such as through contention). In executing the link aggregation SW module 647 the processor 630 may further execute the SL selection submodule 648 to select one or more of the secondary communication links for link aggregation based, at least in part, on the channel conditions of the secondary communication links. For example, execution of the SL selection submodule 648 may cause the STA 600 to determine an availability of one or more of the secondary links based at least in part on the one or more channel conditions measured by the channel monitoring SW module 642.

[00115] Figure 7 shows a block diagram of an example access point (AP) 700 for use in wireless communication according to some implementations. In some implementations, the AP 700 may be a wireless device (such as a ML AP) that is capable of communicating over multiple communication links, concurrently, using link aggregation techniques. For example, the AP 700 may be an example implementation of any of the APs of Figure 1 or Figures 3A–5. The AP 700 may include a PHY 710, a MAC 720, a processor 730, a memory 740, and a number of antennas 750(1)–750(n).

[00116] The PHY 710 may include a number of transceivers 712 and a baseband processor 714. The transceivers 712 may be coupled to the antennas 750(1)–750(n), either directly or through an antenna selection circuit (not shown for simplicity). The transceivers 712 may be used to communicate wirelessly with one or more STAs, with one or more APs, or with other suitable devices. The baseband processor 714 may be used to process outgoing signals received from the processor 730 or the memory 740 and to forward the processed signals to the transceivers 712 for transmission via one or more of the antennas 750(1)–750(n), and may be used to process incoming signals received from one or more of the antennas 750(1)–750(n) via the transceivers 712 and to forward the processed signals to the processor 730 or the memory 740.

[00117] Although not shown in Figure 7, for simplicity, the transceivers 712 may include any number of transmit chains to process and transmit outgoing signals to other wireless devices via the antennas 750(1)–750(n), and may include any number of receive chains to process incoming signals received from the antennas 750(1)–750(n). Thus, in some implementations, the AP 700 may be configured for MIMO operations including, for example, single-user MIMO (SU-MIMO) operations and multi-user (MU-MIMO) operations. In addition, the AP 700 may be configured for OFDMA communications or other suitable multiple access mechanisms, for example, as may be specified by any of the IEEE 802.11 standards.

[00118] The MAC 720 may include at least a number of contention engines 722 and frame formatting circuitry 724. The contention engines 722 may contend for access to the shared wireless medium on a primary communication link, and may store packets for transmission over the shared wireless medium. In some implementations, the contention engines 722 may be separate from the MAC 720. Still further, in some implementations, the contention engines 722 may be implemented as one or more software modules (stored in the memory 740 or in memory provided within the MAC 720). The frame formatting circuitry 724 may be used to create or format frames received from the processor 730 or the memory 740 (such as by adding MAC headers to PDUs provided by the processor 730), and may be used to re-format frames received from the PHY 710 (such as by stripping the MAC headers from frames received from the PHY 710). In the example of Figure 7, the AP 700 is shown to include a single MAC 720. However, in actual implementations, the AP 700 may be a multi-MAC device,

including a plurality of MACs that may be used to contend on multiple communication links, concurrently, of a multi-link environment.

[00119] The memory 740 may include a STA profile data store 741 that stores profile information for a plurality of STAs. The profile information for a particular STA may include, for example, its MAC address, supported data rates, connection history with the AP 700, one or more resource units (RUs) allocated to the STA, and any other suitable information pertaining to or describing the operation of the STA. In some implementations, the profile information for a particular STA may include one or more wireless channels to be used as a primary communication link (such as for contention). In some other implementations, the profile information for a particular STA may include wireless channels to be used as one or more secondary communication links (such as for link aggregation).

[00120] The memory 740 also may include a non-transitory computer-readable medium (one or more nonvolatile memory elements, such as EPROM, EEPROM, Flash memory, a hard drive, and the like) that may store at least the following software (SW) modules:

- a channel monitoring SW module 742 to monitor one or more channel conditions of a plurality of wireless channels available for wireless communications;
- a link configuration SW module 743 to establish a primary communication link on a first subset of the wireless channels and to selectively access one or more secondary communication links on one or more subsets of the remaining wireless channels, the link configuration SW module 743 including:
 - a primary link (PL) reassignment submodule 744 to dynamically change the primary communication link from the first subset of wireless channels to another subset of the wireless channels;
 - a primary link information (PLI) advertisement submodule 745 to communicate primary link information to other wireless devices (such as STAs) in the vicinity of the AP 700; and
 - a primary switching (PS) timer submodule 746 to countdown towards a time at which the change in the primary communication link is scheduled to occur; and
- a link aggregation SW module 747 to access one or more of the secondary communication links, concurrently with the primary communication link, using link aggregation techniques, the link aggregation SW module 747 including:

- a secondary link (SL) selection submodule 748 to select one or more of the secondary communication links for link aggregation based, at least in part, on the channel conditions of the secondary communication links.

Each software module includes instructions that, when executed by the processor 730, cause the AP 700 to perform the corresponding functions.

[00121] For example, the processor 730 may execute the channel monitoring SW module 742 to monitor one or more channel conditions of a plurality of wireless channels available for use in the BSS. The processor 730 may execute the link configuration SW module 743 to establish a primary communication link on a first subset of the wireless channels and to selectively access one or more secondary communication links on one or more subsets of the remaining wireless channels. In executing the link configuration SW module 743, the processor 730 may further execute the PL switching submodule 744, the PLI advertisement submodule 745, or the PS timer submodule 746.

[00122] The processor 730 may execute the PL reassignment submodule 744 to dynamically change the primary communication link from the first subset of wireless channels to another subset of the wireless channels. In some implementations, execution of the PL reassignment submodule 744 may cause the AP 700 to change the primary communication link based, at least in part, on the one or more channel conditions measured by the channel monitoring SW module 742. For example, the AP 700 may change the primary communication link from the first subset of wireless channels to a second subset of wireless channels (which may be associated with a secondary communication link) if the amount of activity on the first subset of wireless channels is above a threshold amount, the amount of activity on the second subset of wireless channels is below a second threshold amount, the amount of activity on the first subset of wireless channels exceeds the amount of activity on the second wireless channels, or any combination thereof. The amount of activity may be based on a frequency of busy channel conditions on one or more of the wireless channels or a number of failed attempts to access one or more of the wireless channels

[00123] The processor 730 may execute the PLI advertisement submodule 745 to communicate primary link information to other wireless devices (such as STAs) in the vicinity of the AP 700. In some aspects, the primary link information may indicate a change in the primary communication link. In some other aspects, the primary link information may indicate the subset of wireless channels to be used for the (new) primary communication link. For example, the new primary communication link may be implemented on one or more wireless channels previously associated with a secondary communication link. In some implementations, execution of the PLI advertisement submodule 745 may cause the AP 700 to broadcast the primary link information, via beacon frames, on the primary communication link. In some other implementations, execution of the PLI advertisement submodule 745 may cause the AP 700 to broadcast the primary link information, via beacon frames, on one or

more of the secondary communication links. Still further, in some implementations, execution of the PLI advertisement submodule 745 may cause the AP 700 to send trigger messages containing the primary link information to one or more STAs (where the trigger messages may be transmitted using OMI signaling procedures).

[00124] The processor 730 may execute the PS timer submodule 746 to countdown towards a time at which the change in the primary communication link is scheduled to occur. For example, in some implementations, the primary link information may indicate a future time (the primary switching time) at which the change in primary link is to be implemented. The primary switching time may coincide with at least one of a TBTT or a TWT of one or more STAs associated with the AP 700. Thus, in some aspects, the primary link information may further include a countdown timer indicating a number of TBTTs or TWTs remaining until the primary-switching time. In some aspects, the AP 700 may update the countdown timer provided with the primary link information sent to the STAs based on a number of TBTT or TWT intervals that have elapsed.

[00125] The processor 730 also may execute the link aggregation SW module 747 to access one or more of the secondary communication links, concurrently with the primary communication link, using link aggregation techniques. In some aspects, the AP 700 may access the secondary communication links only after the AP 600 has already gained access to the shared wireless medium via the primary communication link (such as through contention). In executing the link aggregation SW module 747 the processor 730 may further execute the SL selection submodule 748 to select one or more of the secondary communication links for link aggregation based, at least in part, on the channel conditions of the secondary communication links. For example, execution of the SL selection submodule 748 may cause the AP 700 to determine an availability of one or more of the secondary links based at least in part on the one or more channel conditions measured by the channel monitoring SW module 742.

[00126] Figure 8 shows a flowchart illustrating an example process 800 for configuring a multi-link environment according to some implementations. The process 800 may be performed by any wireless device capable of link aggregation or otherwise communicating in a multi-link environment. With reference for example to Figure 2, the process 800 may be performed by a ML STA to dynamically change its primary link based at least in part on one or more channel conditions of the wireless channels associated with the primary communication link or one or more secondary communication links. With reference for example to Figures 3A–5, the process 800 also may be performed by a ML AP to dynamically change the primary link that is used by other wireless devices (such as ML STAs) in its vicinity.

[00127] The wireless device may establish a primary communication link with a second wireless device (810). The primary communication link may be associated with one or more first wireless

channels. More specifically, access to the primary communication link may be based on contention. In some implementations, the wireless device may be a single-MAC device and may therefore contend for access to a shared wireless medium (with the second wireless device) only on the primary communication link.

[00128] The wireless device may selectively access a secondary communication link with the second wireless device based at least in part on having access to the primary communication link (820). The secondary communication link may be associated with one or more second wireless channels. More specifically, the wireless device may send or receive communications over the primary communication link and the secondary communication link (when available), concurrently, using link aggregation techniques. However, in some implementations, the wireless device may not access the secondary communication link unless the wireless device has also gained access to the primary communication link (such as through contention).

[00129] The wireless device may selectively reassign the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on one or more channel conditions of the first wireless channels or the second wireless channels (830). For example, when the load or activity on the first wireless channels is relatively high (due to heavy data traffic or interference) while the load or activity on the second wireless channels is relatively low, the first and second wireless channels may switch roles. As a result, the wireless device may subsequently contend for medium access on the one or more second wireless channels (the new primary communication link) and may attempt to access the one or more first wireless channels (the new secondary communication link) only after it has gained access to the primary communication link.

[00130] Figure 9 shows a flowchart illustrating an example process 900 for selectively reassigning the primary link in a multi-link environment according to some implementations. The process 900 may be performed by a ML AP capable of link aggregation. With reference for example to Figures 3A–5, the process 900 may be performed by the AP to dynamically change the primary link that is used by other wireless devices (such as STAs) in its vicinity.

[00131] The AP may monitor channel conditions of a primary communication link and one or more secondary communication links (910). The channel conditions may correspond to an amount of activity (due to traffic or interference) on the primary communication link or the one or more secondary communication links. For example, the amount of activity may be based on a frequency of busy channel conditions on one or more of the communication links or a number of failed attempts to access one or more of the communication links. In some aspects, the AP may monitor the traffic patterns or activity of its associated STAs on each of the primary link and the secondary link. In some other aspects, the AP may learn the channel conditions on each of the communication links by correlating with past channel conditions on the primary link and the secondary link.

[00132] The AP may determine, based on the channel conditions, whether a change-in-primary-link (CPL) condition is satisfied (920). In some aspects, the CPL condition may be satisfied when the amount of activity on the primary communication link is above a threshold amount. In some other aspects, the CPL condition may be satisfied when the amount of activity on a secondary communication link is below a threshold amount. Still further, in some aspects, the CPL condition may be satisfied when the amount of activity on the primary communication link exceeds the amount of activity on one or more secondary communication links. As long as the CPL condition is not satisfied (as tested at 920), the AP may continue monitoring the channel conditions of the communication links (910).

[00133] If the CPL condition is satisfied (as tested at 920), the AP may select a new subset of wireless channels for the primary communication link (930). In some implementations, the AP may select the new subset of wireless channels based, at least in part, on the monitored channel conditions of the primary communications link and the one or more secondary communication links. For example, when the load or activity on a first subset of wireless channels associated with the primary communication link is relatively high (due to heavy data traffic or interference) while the load or activity on a second subset of wireless channels associated with the secondary communication link is relatively low, the AP may reassign the second subset of wireless channels to the primary communication link (and may reassign the first subset of wireless channels to the secondary communication link).

[00134] The AP may further transmit primary link information to other wireless devices indicating a change to the primary communication link (940). For example, the primary link information may indicate a change to the primary communication link. Alternatively, or in addition, the primary link information may indicate the new wireless channels (or secondary communication link) to be subsequently used as the primary communication link. In some implementations, the AP may provide the primary link information in beacon frames broadcast at TBTT intervals. In some aspects, the AP may broadcast the primary link information via beacon frames on the primary communication link (such as described with respect to Figure 3A). In some other aspects, the AP may broadcast the primary link information via beacon frames on the secondary communication links (such as described with respect to Figure 3B).

[00135] In some other implementations, the AP may provide the primary link information in trigger messages sent to one or more individual STAs. For example, the trigger messages may be transmitted using OMI signaling procedures. In some aspects, the trigger messages may be sent in response to uplink transmission received from the one or more STAs (such as described with respect to Figure 4B). In some other aspects, the trigger messages may be sent based on a TWT of the one or more STAs (such as described with respect to Figure 4C).

[00136] The AP may change the primary communication link to the new wireless channels at the appropriate time (950). In some implementations, the change in the primary communication link may be effective immediately (such as at the time of transmission or reception of the primary link information). In some other implementations, the change in the primary communication link may take effect at a future primary switching time. For example, the primary switching time may be scheduled to coincide with a subsequent TBTT or TWT of the one or more STAs. Thus, in some aspects, the primary link information transmitted by the AP may include an indication (such as a countdown timer) as to when the primary switching time is to occur (such as described with respect to Figure 5).

[00137] Figure 10 shows a flowchart illustrating an example process 1000 for dynamically changing the primary link of a STA according to some implementations. The process 1000 may be performed by a ML STA capable of link aggregation. With reference for example to Figure 2 and Figures 4A–5, the process 1000 may be performed by a STA to dynamically change its primary link based at least in part on primary link information received from an AP.

[00138] The STA may establish a primary communication link with the AP (1010). The primary communication link may be associated with one or more first wireless channels. More specifically, access to the primary communication link may be based on contention. In some implementations, the STA may be a single-MAC device and may therefore contend for access to a shared wireless medium (with the AP) only on the primary communication link.

[00139] The STA may selectively access a secondary communication link with the AP based at least in part on having access to the primary communication link (1020). The secondary communication link may be associated with one or more second wireless channels. More specifically, the STA may send or receive communications over the primary communication link and the secondary communication link (when available), concurrently, using link aggregation techniques. However, in some implementations, the STA may not access the secondary communication link unless the STA already has already gained access to the primary communication link (such as through contention).

[00140] The STA may receive primary link information from the AP indicating a change in the primary communication link (1030). For example, the primary link information may indicate a change to the primary communication link. Alternatively, or in addition, the primary link information may indicate a new subset of wireless channels (or a secondary communication link) to be subsequently used as the primary communication link. In some implementations, the primary link information may be provided in beacon frames broadcast at TBTT intervals (such as described with respect to Figures 3A and 3B). In some other implementations, the primary link information may be provided in trigger messages sent to the STA (such as described with respect to Figures 4B and 4C).

[00141] The STA may dynamically change the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on the

primary link information (1040). As a result of the change, the STA may subsequently contend for medium access on the one or more second wireless channels (the new primary communication link) and may attempt to access the one or more first wireless channels (the new secondary communication link) only after it has gained access to the primary communication link.

[00142] Figure 11 shows a flowchart illustrating an example process 1100 for dynamically changing the primary link of a STA according to some other implementations. The process 1100 may be performed by a ML STA capable of link aggregation. With reference for example to Figure 2 and Figures 4A–5, the process 1100 may be performed by a STA to dynamically change its primary link based at least in part on primary link information received from an AP.

[00143] The STA wakes up from a power save state (1110) and initiates an uplink transmission to the AP (1120). For example, the STA may wake up from the power save state when it has uplink data to send to the AP. The STA also may wake up from the power save state at scheduled times or epochs (such as TBTTs or TWTs). In some implementations, the STA may initiate an uplink transmission to the AP, upon waking up from the power save state, even if the STA has no uplink data to send. For example, in some aspects, the STA may send a PS Poll request to the AP upon waking up from the power save state (such as described with respect to Figure 4B).

[00144] The STA may receive a trigger message from the AP in response to the uplink transmission (1130) and may determine primary link information from the received trigger message (1140). In some implementations, the trigger message may include primary link information indicating a change in the primary communication link or the (new) wireless channels associated with the primary communication link. In some implementations, the trigger message may be transmitted by the AP using OMI signaling procedures. For example, the AP may respond to the uplink transmission by sending a QoS Null frame back to the STA. In some aspects, the primary link information may be included in an OMI control field of the QoS Null frame.

[00145] The STA may change the primary communication link to the new wireless channels at the appropriate time (1150). In some implementations, the change in the primary communication link may be effective immediately (such as at the time of transmission or reception of the primary link information). In some other implementations, the change in the primary communication link take effect at a future primary switching time. For example, the primary switching time may be scheduled to coincide with a subsequent TBTT or TWT of the STA. Thus, in some aspects, the primary link information received from the AP may include an indication (such as a countdown timer) as to when the primary switching time is to occur (such as described with respect to Figure 5).

[00146] 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. For example, “at least one of: a, b, or c” is

intended to cover the possibilities of: a only, b only, c only, a combination of a and b, a combination of a and c, a combination of b and c, and a combination of a and b and c.

[00147] The term “wireless station” or “STA,” as used herein, also may refer to as a user equipment (UE), a subscriber station, a mobile unit, a subscriber unit, a wireless unit, a remote unit, a mobile device, a wireless device, a wireless communications device, a remote device, a mobile subscriber station, an access terminal, a mobile terminal, a wireless terminal, a remote terminal, a handset, a user agent, a mobile client, a client, or some other suitable terminology.

[00148] The various illustrative logics, logical blocks, modules, circuits and algorithm processes described in connection with the implementations disclosed herein may be implemented as electronic hardware, computer software, or combinations of both. The interchangeability of hardware 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 or software depends upon the particular application and design constraints imposed on the overall system.

[00149] The hardware and data processing apparatus used to implement the various illustrative logics, logical blocks, modules and circuits described in connection with the aspects disclosed herein may be implemented or performed with a general purpose single- or multi-chip processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein. A general purpose processor may be a microprocessor, or, any conventional processor, controller, microcontroller, or state machine. A processor also may be implemented as a combination of computing devices such as, for example, a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration. In some implementations, particular processes and methods may be performed by circuitry that is specific to a given function.

[00150] In one or more aspects, the functions described may be implemented in hardware, digital electronic circuitry, computer software, firmware, including the structures disclosed in this specification and their structural equivalents thereof, or in any combination thereof. Implementations of the subject matter described in this specification also can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions, encoded on a computer storage media for execution by, or to control the operation of, data processing apparatus.

[00151] If implemented in software, the functions may be stored on or transmitted over as one or more instructions or code on a computer-readable medium. The processes of a method or algorithm disclosed herein may be implemented in a processor-executable software module which may reside on a

computer-readable medium. Computer-readable media includes both computer storage media and communication media including any medium that can be enabled to transfer a computer program from one place to another. A storage media may be any available media that may be accessed by a computer. By way of example, and not limitation, such computer-readable media may include RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium that may be used to store desired program code in the form of instructions or data structures and that may be accessed by a computer. Also, any connection can be properly termed a computer-readable medium. Disk and disc, as used herein, includes compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk, and blu-ray disc where disks usually reproduce data magnetically, while discs reproduce data optically with lasers. Combinations of the above should also be included within the scope of computer-readable media. Additionally, the operations of a method or algorithm may reside as one or any combination or set of codes and instructions on a machine readable medium and computer-readable medium, which may be incorporated into a computer program product.

[00152] Various modifications to the implementations described in this disclosure may be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other implementations without departing from the spirit or scope of this disclosure. Thus, the claims are not intended to be limited to the implementations shown herein, but are to be accorded the widest scope consistent with this disclosure, the principles and the novel features disclosed herein.

[00153] Certain features that are described in this specification in the context of separate implementations 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 implementations separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain 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.

[00154] 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 more example processes in the form of a 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 certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system components in the implementations described above should not be understood as requiring such

separation in all implementations, 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. Additionally, other implementations are within the scope of the following claims. In some cases, the actions recited in the claims can be performed in a different order and still achieve desirable results.

CLAIMS

1. A method for wireless communication by a wireless communication device, comprising:

establishing a primary communication link between the wireless communication device and a second wireless device, the primary communication link being associated with one or more first wireless channels, access to the primary communication link being based on contention among the wireless communication device and the second wireless device;

selectively accessing a secondary communication link between the wireless communication device and the second wireless device based at least in part on having access to the primary communication link, the secondary communication link being associated with one or more second wireless channels; and

selectively reassigning the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on one or more channel conditions of the first wireless channels or the second wireless channels.

2. The method of claim 1, wherein the one or more channel conditions includes an indication of an amount of activity on the first or second wireless channels, and wherein the selectively reassigning comprises:

reassigning the primary communication link to the one or more second wireless channels when the indication of the amount of activity on the first wireless channels is above a first threshold level, the indication of the amount of activity on the second wireless channels is below a second threshold level, the indication of the amount of activity on the first wireless channels exceeds the amount of activity on the second wireless channels, or a combination thereof.

3. The method of claim 2, wherein the amount of activity is based on a frequency of busy channel conditions on the first or second wireless channels or a number of failed attempts to access the first or second wireless channels.

4. The method of claim 1, wherein the selectively reassigning comprises:
transmitting primary link information to the second wireless device, the primary link

information indicating a change in the primary communication link.

5. The method of claim 4, wherein transmitting the primary link information includes broadcasting the primary link information in a first beacon frame on the one or more first wireless channels.

6. The method of claim 5, wherein the selectively reassigning further comprises:
broadcasting the primary link information, in a second beacon frame, on the one or more second wireless channels.

7. The method of claim 4, wherein transmitting the primary link information includes transmitting a data frame, a management frame, or a control frame, the data frame, management frame or control frame including an Operating Mode Indication (OMI) control field that includes the primary link information.

8. The method of claim 4, further comprising receiving an uplink transmission from the second wireless device, wherein transmitting the primary link information includes transmitting a trigger message to the second wireless device in response to the uplink transmission, the trigger message including the primary link information.

9. The method of claim 4, wherein transmitting the primary link information includes transmitting a trigger message to the second wireless device based at least in part on a target wake time (TWT) service period for the second wireless device, the trigger message including the primary link information.

10. The method of claim 4, wherein the primary link information further indicates a primary-switching time at which the change is to occur, and wherein the selectively reassigning comprises reassigning the primary communication link to the one or more second wireless channels at the primary-switching time.

11. The method of claim 10, wherein the primary-switching time coincides with at least

one of a target beacon transmission time (TBTT) or a target wake time (TWT) service period of the second wireless device, and wherein the primary link information further includes a countdown timer indicating a number of TBTTs or TWTs remaining until the primary-switching time.

12. A wireless communication device, comprising:
one or more processors; and
a memory storing instructions that, when executed by the one or more processors, cause the wireless communication device to:

establish a primary communication link between the wireless communication device and a second wireless device, the primary communication link being associated with one or more first wireless channels, access to the primary communication link being based on contention among the wireless communication device and the second wireless device;

selectively access a secondary communication link between the wireless communication device and the second wireless device based at least in part on having access to the primary communication link, the secondary communication link being associated with one or more second wireless channels; and

selectively reassign the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on one or more channel conditions of the first wireless channels or the second wireless channels.

13. The wireless communication device of claim 12, wherein the one or more channel conditions includes an indication of an amount of activity on the first or second wireless channels, and wherein execution of the instructions for selectively reassigning the primary communication link causes the wireless communication device to reassign the primary communication link to the one or more second wireless channels when the indication of the amount of activity on the first wireless channels is above a first threshold level, the indication of the amount of activity on the second wireless channels is below a second threshold level, the indication of the amount of activity on the first wireless channels exceeds the amount of activity on the second wireless channels, or a combination thereof.

14. The wireless communication device of claim 12, wherein the execution of the

instructions for selectively reassigning the primary communication link causes the wireless communication device to:

transmit primary link information to the second wireless device, the primary link information indicating a change in the primary communication link.

15. The wireless communication device of claim 14, wherein execution of the instructions for transmitting the primary link information causes the wireless communication device to broadcast the primary link information in a beacon frame on the one or more first wireless channels or the one or more second wireless channels.

16. The wireless communication device of claim 14, wherein execution of the instructions for transmitting the primary link information causes the wireless communication device to transmit a trigger message to the second wireless device in response to an uplink transmission or based at least in part on a target wake time (TWT) service period for the second wireless device, the trigger message including the primary link information.

17. The wireless communication device of claim 14, wherein the primary link information further indicates a primary-switching time at which the change is to occur, and wherein execution of the instructions for selectively reassigning the primary communication link further causes the wireless communication device to reassign the primary communication link to the one or more second wireless channels at the primary-switching time, wherein the primary-switching time coincides with at least one of a target beacon transmission time (TBTT) or a TWT of the second wireless device, and wherein the primary link information further includes a countdown timer indicating a number of TBTTs or TWTs remaining until the primary-switching time.

18. A method performed by a wireless station (STA), comprising:
establishing a primary communication link between the STA and an access point (AP), the primary communication link being associated with one or more first wireless channels, access to the primary communication link being based on contention among the STA and the AP;
selectively accessing a secondary communication link between the STA and the AP based

at least in part on having access to the primary communication link, the secondary communication link being associated with one or more second wireless channels;

receiving primary link information from the AP indicating a change in the primary communication link; and

dynamically changing the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on the primary link information.

19. The method of claim 18, wherein the receiving comprises receiving a beacon frame broadcast by the AP, the beacon frame including the primary link information.

20. The method of claim 18, wherein the receiving comprises receiving a data frame, a management frame, or a control frame transmitted by the AP, the data frame, management frame, or control frame including an Operating Mode Indication (OMI) control field that includes the primary link information.

21. The method of claim 18, wherein the receiving comprises:
waking up from a power save state; and
receiving a trigger message from the AP after waking up from the power save state, the trigger message including the primary link information.

22. The method of claim 21, wherein receiving the trigger message comprises:
initiating an uplink transmission to the AP upon waking up from the power save state, the trigger message being received from the AP in response to the uplink transmission.

23. The method of claim 18, wherein the primary link information further indicates a primary-switching time at which the change is to occur, and wherein the dynamically changing comprises changing the primary communication link to the one or more second wireless channels at or after the primary-switching time.

24. The method of claim 23, wherein the primary-switching time coincides with at least

one of a target beacon transmission time (TBTT) or a target wake time (TWT) of the STA, and wherein the primary link information further includes a countdown timer indicating a number of TBTTs or TWTs remaining until the primary-switching time.

25. The method of claim 23, wherein the dynamically changing further comprises:
entering a power save state after receiving the primary link information;
waking up from the power save state at or after the primary-switching time; and
using the one or more second wireless channels for the primary communication link upon waking up from the power save state.

26. A wireless station (STA), comprising:
one or more processors; and
a memory storing instructions that, when executed by the one or more processors, cause the STA to:

establish a primary communication link between the STA and an access point (AP), the primary communication link being associated with one or more first wireless channels, access to the primary communication link being based on contention among the STA and the AP;

selectively access a secondary communication link between the STA and the AP based at least in part on having access to the primary communication link, the secondary communication link being associated with one or more second wireless channels;

receive primary link information from the AP indicating a change in the primary communication link; and

dynamically change the primary communication link from the one or more first wireless channels to the one or more second wireless channels based at least in part on the primary link information.

27. The wireless station of claim 26, wherein execution of the instructions for receiving the primary link information from the AP causes the STA to:

wake up from a power save state;

initiate an uplink transmission to the AP upon waking up from the power save state; and

receive a trigger message from the AP in response to the uplink transmission, the trigger message including the primary link information.

28. The wireless station of claim 26, wherein the primary link information further indicates a primary-switching time at which the change is to occur, and wherein execution of the instructions for dynamically changing the primary communication link causes the STA to change the primary communication link to the one or more second wireless channels at or after the primary-switching time.

29. The wireless station of claim 28, wherein the primary-switching time coincides with at least one of a target beacon transmission time (TBTT) or a target wake time (TWT) of the STA, and wherein the primary link information further includes a countdown timer indicating a number of TBTTs or TWTs remaining until the primary-switching time.

30. The wireless station of claim 28, wherein execution of the instructions for dynamically changing the primary communication link further causes the STA to:

- enter a power save state after receiving the primary link information;
- wake up from the power save state at or after the primary-switching time; and
- use the one or more second wireless channels for the primary communication link upon waking up from the power save state.