IEEE 802.11ac: From Channelization to Multi-User MIMO

Oscar Bejarano and Edward W. Knightly, Rice University
Minyoung Park, Intel Corporation

ABSTRACT

The IEEE 802.11ac amendment has been proposed to enhance the throughput of IEEE 802.11n beyond gigabit-per-second rates. In this article we present an overview of the most important features proposed in the 802.11ac amendment, including channel bonding mechanisms and multi-user MIMO.

INTRODUCTION

Mobile data traffic is projected to experience an 18-fold increase between 2011 and 2016 due to the growth of mobile subscribers and bandwidth demands to support data-hungry applications [1]. Consequently, there is a need for devices and standards capable of coping with the next-generation mobile networks, which will require very high data rates to sustain video, voice, live gaming, and augmented reality applications, among others. To this end, the IEEE 802.11ac Task Group (TGac) is working on an amendment that has the goal of reaching maximum aggregate network throughputs of at least 1 Gb/s on bands below 6 GHz, excluding the 2.4 GHz band. In particular, the standard envisions a maximum medium access control (MAC) throughput of at least 500 Mb/s for a single user, and at least 1 Gb/s in the case of multiple users. This represents an increase of five times the maximum achievable rate (per user) compared to the previous amendment, 802.11n [2]. In contrast to all previous amendments of the 802.11 standard, 802.11ac is aimed at improving total network throughput as well as individual link performance [3]. Due to the significant rate increase achieved by 802.11ac, the term very high throughput (VHT) is also used in reference to this new amendment. In Table 1 we present a summary of the evolution of the 802.11 standard by comparing some of the main characteristics of each generation. The table shows differences between 802.11b, a, g, n, and ac.

Several modifications have been proposed in order to reach gigabit throughput rates. In this article we explore the different features and enhancements that differentiate 802.11ac from previous standards. More specifically, we describe the key modifications to both the physical layer (PHY) and MAC. While many of the specifications of 802.11n have been kept for 802.11ac (e.g., static and dynamic channel bonding and simultaneous data streams), these have been enhanced to allow support for wider channels as well as more data streams, among others.

The two main features that allow 802.11ac to achieve gigabit transmission rates are:
• Static and dynamic channel bonding
• Multi-user multiple-input multiple-output (MU-MIMO)

To enable these two features, substantial modifications are required at the PHY. For the most part, at the MAC level, the proposed changes are needed to guarantee compatibility with the modified PHY. More specifically, key features proposed in the 802.11ac amendment include the following [4].

Mandatory support for 20, 40, and 80 MHz channels, and optional support for 160 MHz as well as 80 + 80 MHz channel widths (contiguous and non-contiguous, respectively). Additionally, it proposes a request to send/clear to send (RTS/CTS) mechanism for both static and dynamic bandwidth reservation [4].

802.11ac introduces MU-MIMO by proposing a unique explicit feedback protocol that enables transmit beamforming. This is in contrast to previous standards, where different single-user beamforming methods were introduced, but none of them were mandated for certification. This led to lack of interoperability among different manufacturers. Moreover, the number of spatial streams is increased from four in 802.11n to eight in 802.11ac.

In terms of modulation and coding schemes, 802.11ac mandates single spatial stream modulation of up to 64-quadrature amplitude modulation (QAM) with 5/6 coding rate, and binary convolutional coding. It also allows higher constellation density (256-QAM with 3/4 and 5/6 coding rate), and the use of space-time block coding (STBC) and low density parity check coding (LDPC) as options. Moreover, the standard specifies the use of different frame aggregation schemes. In particular, it proposes mandatory use of frame aggregation to increase channel utilization and MAC efficiency.

The 802.11ac amendment is being developed to address different types of usage models. The main categories are wireless display, in-home distribution of HDTV and other content, rapid

Feature/IEEE standard	802.11b	802.11g/a	802.11n	802.11ac
Maximum data rate per stream (Mb/s)	11	54	>100	>500 (Assuming 80 MHz channels)
Frequency band	2.4 GHz	2.4 GHz/5 GHz	2.4 GHz and 5 GHz	5 GHz
Channel width (MHz)	20	20/20	20 and 40 (40 is optional)	20,40,80, 160, and 80+80 (last two are optional)
Antenna technology	Single-input single-output (SISO)	SISO	Multiple-Input Multiple-Output (MIMO)	MIMO/MU-MIMO
Transmission technique	Direct sequence spread spectrum (DSSS)	DSSS and orthogonal frequency-division multiplexing (OFDM)	OFDM	OFDM
Maximum number of spatial streams	1	1	4	8
Beamforming-capable	No	No	Yes	Yes
Date ratified by IEEE	1999	2003/1999	2009	Expected 2014

Table 1. *Comparison of IEEE802.11 standards.*

upload and download of large files to/from servers, backhaul traffic, campus and auditorium deployments, and manufacturing floor automation [5]. Notice that 802.11ac stations are compatible with legacy devices. That is, the new amendment defines features in addition to 802.11n, which means that an 802.11ac conforming station can also support all the mandatory features defined in 802.11n.

In this article we assume the reader is familiar with the 802.11 standard; otherwise, we suggest first looking into [6].

CHANNELIZATION

One of the most important enhancements to the 802.11ac amendment is the support for wider channels as well as both static and dynamic channel access. We dedicate this section to describing these 802.11ac features.

CHANNEL WIDTHS SUPPORTED

The amendment mandates that all devices support 20, 40, and 80 MHz channels. In addition, it provides optional support for operation on 160 MHz channels. 80 and 160 MHz can be formed by a combination of two adjacent non-overlapping 40 and 80 MHz channels, respectively. The amendment also specifies that two non-adjacent 80 MHz channels can be used to form a 160 MHz one. More important, a device operating on non-contiguous 80 + 80 MHz should be capable of communicating with devices operating on contiguous 160 MHz if the former segments are placed in frequency to match the tone allocation of the latter case. In Fig. 1 we show the channel allocation for the U.S. region.

PRIMARY AND SECONDARY SUBCHANNELS

Similar to 802.11n, channels consisting of 40 MHz or wider always require a primary 20 MHz wide subchannel. Additionally, 80 MHz channels
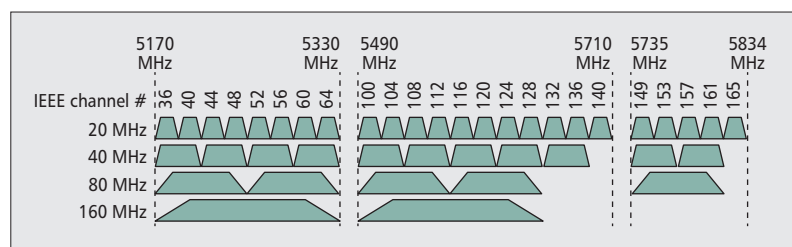
Figure 1. *Channel allocation in the United States [7].*

have a primary 40 MHz (which includes the primary 20 MHz) subchannel and a secondary 40 MHz subchannel. The same applies to 160 MHz and 80 + 80 MHz channels, which consist of primary and secondary 80 MHz subchannels. In Fig. 2 we depict this relationship between the primary and secondary subchannels based on the different bandwidth options.

In all cases, the primary suchannel is used for carrier sensing in order to guarantee that no other device is transmitting. The presence of the 20 MHz primary subchannel is also necessary to guarantee coexistence and backward compatibility with legacy 802.11 devices. Only the primary subchannel performs full clear channel assessment (CCA), which involves packet detection starting with the preamble. In contrast, the secondary subchannel is not required to perform full CCA.

The CCA sensitivity of the primary subchannel is –82 dBm for a valid 802.11 20 MHz signal, –79 dBm for a valid 802.11 40 MHz signal, –76 dBm for a valid 80 MHz signal, and –73 dBm for a valid 160 MHz one. On the other hand, for the secondary subchannel the sensitivity was improved from –62 dBm to –72 dBm for both 20 and 40 MHz channels, compared to 802.11n (and –69 dBm for 80 MHz channels). According to [7], an

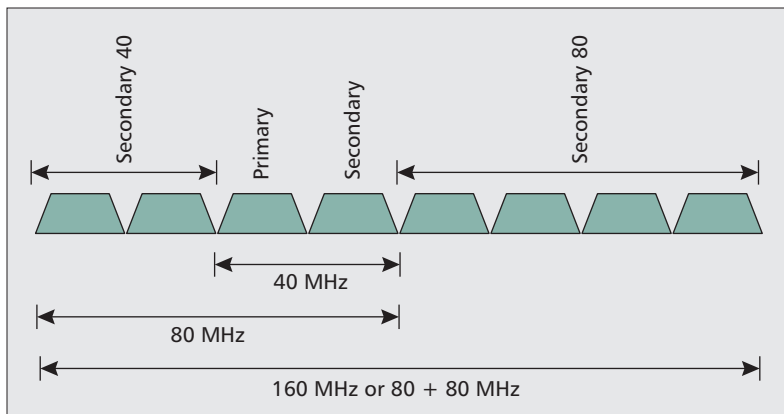**Figure 2.** *Primary and secondary channel selection [4].*

802.11ac device should detect whether the primary subchannel is busy within 4 μs with a probability greater than 90 percent . In contrast, on the secondary subchannel the device has up to 25 μs to detect if it is busy with the same probability.

STATIC AND DYNAMIC CHANNEL ACCESS

IEEE 802.11ac extends the channel access policies proposed in 802.11n to the case of 80 and 160 MHz channels. In order for an 802.11ac station to be able to transmit an 80 MHz PHY convergence procedure protocol data unit (PPDU), two conditions must be true:

• The primary channel follows enhanced distributed channel access (EDCA) rules, so it needs to be idle for distributed coordination function inter-framing spacing (DIFS) plus the backoff counter duration.

• All three secondary subchannels must have been idle for a duration of point coordination function inter-framing spacing (PIFS) immediately preceding the expiration of the backoff counter [4].

Should any of the secondary subchannels be busy, the station can follow either static or dynamic channel access rules as dictated in 802.11n:

• Static channel access: Consider an 802.11ac station trying to transmit on 80 MHz. If the secondary subchannel is busy, the station will choose a random backoff period within the current contention window size to restart the contention process and continue to attempt only until no interferer is present in any of the subchannels. Notice that with a large number of legacy stations, the probability of accessing the medium with such a wide channel will be diminished.

• Dynamic channel access: The 802.11ac station may attempt to transmit over a narrower channel using 20 or 40 MHz instead. This will depend on each subchannel's CCA. This is clearly a more flexible approach, which allows for more efficient resource allocation because the station can still transmit over a fraction of the original bandwidth. All transmissions always have to include the primary channel in order to inform the receiver of which channels the transmitter will use [8].

IMPLICATIONS OF CHANNELIZATION RULES ON THE MAC

Subchannel collisions can occur due to the sensitivity level of the secondary subchannel being higher than that of the primary. Consider the scenario where a legacy 802.11a/n station transmits on a 20 MHz channel, and an 802.11ac station transmits on an 80 MHz one. If the legacy device's transmission takes place on one of the secondary subchannels of the 802.11ac station, and has a sensitivity between –82 and –72 dBm, the 802.11ac station will determine that the secondary subchannel is idle and transmit an 80 MHz signal, which in turn may collide with the transmission of the legacy device.

In 802.11ac, primary and secondary subchannels have different rules for setting the network allocation vector (NAV); specifically, when a station that is not the intended receiver overhears a packet in the primary subchannel, it is decoded and, based on the MAC header, the NAV setting is updated in order for that station to defer its transmission. In contrast, for wider channels, if the reception occurred on the secondary subchannel, the device is not required to set the NAV [4]. This means that unless the physical carrier sensing mechanism is able to detect an ongoing transmission on a secondary channel, a transmission on that same channel may lead to a collision.

RTS/CTS MECHANISM ENHANCED

If an 802.11ac access point (AP) is nearby other legacy APs, it is possible that the 20 MHz primary channel of any of the latter ones is anywhere within 80 or 160 MHz of the former one. This means that the different APs and their clients can transmit at overlapping times on different subchannels, thus leading to collisions or deferrals [9].

To overcome this problem, 802.11ac defines a handshake to properly handle both static and dynamic channel allocation. This handshake consists of a modified RTS/CTS mechanism that provides information about the current amount of available bandwidth. We show how the enhanced RTS/CTS protocol works with the following example (depicted in Fig. 3) [9]: Consider a scenario in which an initiating AP wants to transmit data to an associated client through an 80 MHz channel. The AP first checks if the channel is idle. If it is, it transmits multiple RTS in the 802.11a PPDU format (one RTS for each 20 MHz subchannel). Therefore, it is expected that every nearby device (legacy or 802.11ac) can receive an RTS on its primary channel. Each of these devices then sets its NAV. Before a client replies with a CTS, it checks if any of the subchannels in the 80 MHz band is busy. The client only replies with a CTS on those subchannels that are idle, and reports the total bandwidth of the replicated CTS. As with the RTS, the CTS is sent in an 802.11a PPDU format and is replicated across the different 20 MHz subchannels that are idle.

Notice in Fig. 3 the two different cases. In Fig. 3a no interference is present at either the initiating AP or its client. On the other hand, in Fig. 3b a nearby AP is already transmitting
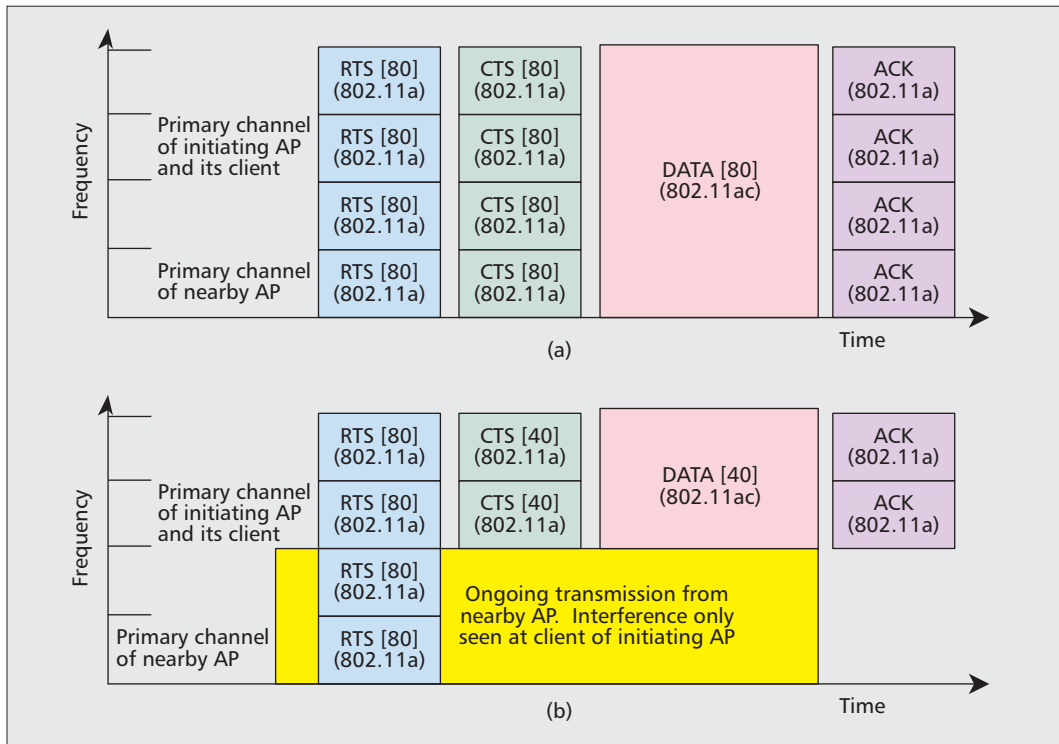
Figure 3. *Enhanced RTS/CTS mechanism [9]: a) no interference case; b) interference case.*

before the initiating AP starts; however, it is only interfering with the client. Therefore, the client has to inform the AP by replying with a CTS only on the idle subchannels.

ENABLING MULTIPLE DATA STREAMS VIA DOWNLINK MIMO

MIMO was first introduced to the IEEE 802.11 standards with the 802.11n amendment [2]. This technique consists of a physical layer scheme where both the transmitter and the receiver employ multiple antennas. The 802.11n standard supports a maximum of four MIMO streams that can serve a single user at a time (a technique known as single-user MIMO, SU-MIMO) with spatial multiplexing of up to four spatial streams. In contrast, 802.11ac is the first 802.11 amendment to introduce multi-user MIMO to serve multiple stations simultaneously. Moreover, it increases the number of streams allowed for SU-MIMO from four to eight. We dedicate this section to providing a brief introduction to these two techniques, which are supported by 802.11ac. It is important to mention that the implementation of MU-MIMO and a technique we explore in a later section, known as transmit beamforming, are necessary to achieve the maximum throughput gains targeted by 802.11ac. Notice that the amendment supports SU-MIMO for both uplink and downlink, but only downlink MU-MIMO is supported.

SINGLE-USER MIMO

SU-MIMO exploits the presence of multiple transmit and receive antennas to improve both the capacity and the reliability of a transmission. By using space-time codes, an SU-MIMO system provides diversity gains, thus increasing reliability. On the other hand, by transmitting different information on different streams, it can provide significant multiplexing gains, thus increasing link capacity [10]. This scheme is depicted in Fig. 4a. Observe that a multi-antenna station communicates with a single user at a time. That is, transmissions to different users are orthogonal in time as in time-division multiple access (TDMA).

MULTI-USER MIMO

MU-MIMO is defined by the standard as a "technique where multiple stations, each with potentially multiple antennas, transmit and/or receiver independent data streams simultaneously [4]." That is, MU-MIMO allows stations having multiple antennas to transmit several data streams to multiple users at the same time over the same frequency channel. For example, if an AP has four antennas, it can serve four single-antenna users at a time, or two users each having two antennas by sending up to one stream per receiving antenna (over the same frequency). In Fig. 4b we illustrate the basic idea behind MU-MIMO. Notice that the AP can serve several single- or multi-antenna users simultaneously. Ideally, the number of simultaneous data streams enabled by MU-MIMO techniques is only constrained by the minimum number of antennas at either the AP or the receiver side (e.g., observe in Fig. 4b that although there are five potential receiving antennas distributed among three different users, the AP has only four antennas for transmission).

In multi-user mode, the 802.11ac amendment supports up to four streams serving four different users simultaneously, or up to four streams per single user. Furthermore, the amendment
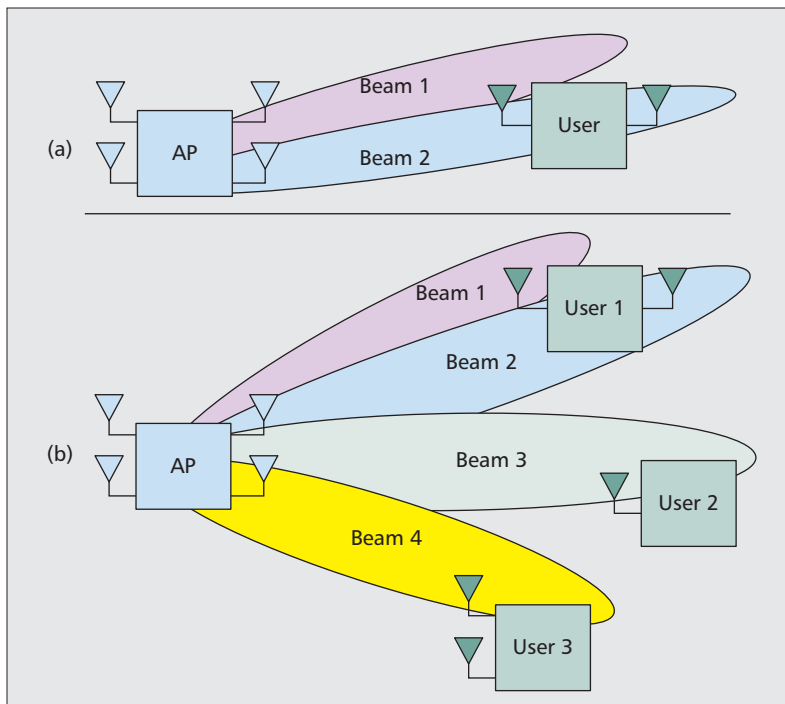
Figure 4. *Different MIMO concepts: a) SU-MIMO beamforming; b) downlink MU-MIMO beamforming.*

specifies support for a different modulation and coding rate for each station being served in a downlink MU-MIMO transmission.

An MU-MIMO transmitting station requires knowledge of the channel state information (CSI) by all users in order to decrease the amount of inter-user interference generated by the multiple simultaneous streams. To achieve this, most existing approaches use a combination of strategies such as feedback, where the transmitting station obtains a measure of the CSI, and data precoding, where that information is used to perform inter-user interference cancellation at the transmitter side. 802.11ac specifies a single compressed beamforming method that relies on the use of explicit feedback to implement MU-MIMO (a technique known as MU-beamforming).

TRANSMIT BEAMFORMING

IEEE 802.11ac specifies a unique transmit beamforming method based on explicit feedback to enable both SU- and MU-MIMO. In particular, beamforming allows a station to transmit multiple simultaneous data streams to a single user or multiple users. Different beamforming methods were introduced in 802.11n, but none of them were mandated for certification, which led to the development of multiple non-interoperable techniques by several chipset manufacturers [11]. As a consequence, the 802.11ac draft now defines a unique feedback protocol to guarantee interoperability between different beamforming implementations by different manufacturers. Although the sounding protocol proposed in 802.11ac closely resembles the *explicit compressed feedback* method proposed in 802.11n, it is not backward compatible with 802.11n devices [9].

In general, SU- and MU-MIMO beamforming techniques are employed by transmitting stations to steer signals based on knowledge of the channel in order to improve reception of PPDUs at a receiver. As the name indicates, in SU-MIMO beamforming, the different spatial streams transmitted are received at only one station. On the other hand, MU-MIMO beamforming implies division of the space-time streams among several receiving stations.

Beamforming is directly enabled by the support of *sounding*. Sounding is the term used to denote the process performed by the transmitter (e.g., the AP in a downlink transmission) to acquire CSI from each user by sending training symbols and waiting for the receivers to provide explicit feedback containing a measure of the channel. This feedback is then used to create a *weight* or *steering matrix* that will be used to precode the data transmission by creating a set of steered beams to optimize reception at one or multiple receivers. The station sounding the channel and transmitting data using a steering matrix is known as the 802.11ac beamformer (VHT beamformer). On the other hand, the served 802.11ac station replying with the feedback is called the VHT *beamformee*. In Fig. 4b we illustrate the concept of MU-beamforming. Observe that a multi-antenna AP serves a set of users by forming various beams, each transmitting a different data stream.

Explicit Beamforming Feedback in 802.11ac [4] — Every time a steering matrix is computed from new channel measurements, the current steering matrix is replaced for the next data transmission. With this method, a beamformee generates and replies to a predefined channel sounding packet with a compressed beamforming feedback matrix in the form of angles quantized according to [4]. The number of bits used for quantization is chosen by the beamformee based on the indication given by the beamformer of whether the feedback is for an SU-MIMO or MU-MIMO transmission. Notice that the method the beamformer uses to calculate the steering matrix is implementation specific and is not defined by the standard.

Protocol Description — The beamforming protocol proposed for 802.11ac works as follows. Before every beamformed transmission, the beamformer transmits a VHT null data packet announcement (NDPA) frame that contains the addresses of the transmitter (AP) and the set of beamformees, as well as a sequence number that identifies the VHT NDP announcement. The main purpose of this announcement is to notify the receive station that it should be ready to prepare a beamforming report frame. After a short inter-frame space (SIFS), the AP transmits a VHT NDP frame in order to sound the channel. Based on the NDP frame, the station will prepare the information that will be carried by the beamforming report. The VHT NDP frame has the same format as the VHT PPDU but does not include a data field. Therefore, the targeted receptors only use a preamble to measure the channel between the AP and themselves.
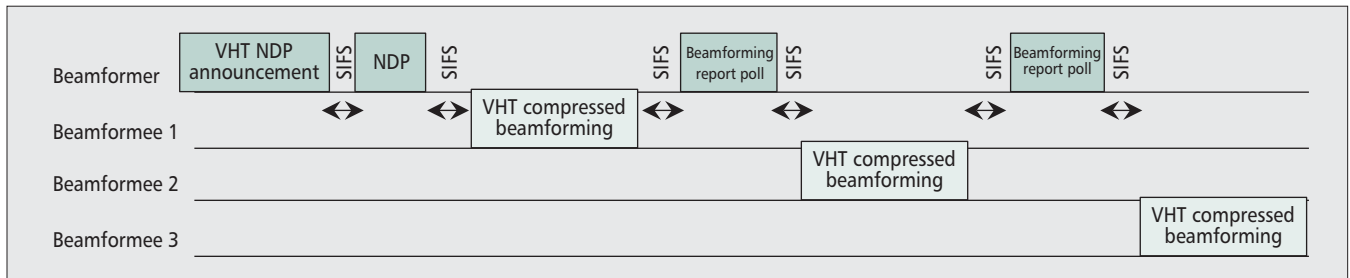
Figure 5. *Sounding protocol with multiple beamformees [4].*

More specifically, upon reception of the VHT NDP frame, each beamformee removes the space-time stream cyclic shift diversity (CSD) applied to the signals transmitted. The CSD consists of a signal shaping technique where different phase shifts are applied to the same signal across different transmit chains. After removing the CSD, the targeted beamformees are required to reply with a VHT compressed beamforming frame. The first intended station replies immediately, whereas the others have to wait to be polled by the AP (using a beamforming report poll). The most relevant information carried by the VHT compressed beamforming frame is the following:

• The *VHT MIMO control field*, which contains the dimensions of the matrix, an indicator of the width of the channel in which the measurements used to create the feedback matrix were taken, and information indicating the size of the codebook entries

• The *VHT compressed beamforming report* containing the compressed beamforming feedback matrix in the form of two angles, as well as the signal-to-noise ratio (SNR) of each space-time stream averaged over all subcarriers used

• The *MU exclusive beamforming report* carrying explicit information used by a multi-user beamformer in order to create the steering matrices

Notice that when transmitting a VHT NDP announcement frame with multiple stations as destinations (contrary to SU-MIMO, where only one station is served at a time), beamforming report poll frames need to be transmitted in order to retrieve a compressed beamforming report from each of the stations. In Fig. 5 we present an example of the VHT sounding protocol used when the AP deals with more than one VHT beamformee. In terms of power, regulation limits the transmit power based on the number of antennas used at the transmitter; therefore, transmit beamforming does not increase the maximum distance range [11].

The implementation of transmit beamforming is optional in the 802.11ac standard [9].

OTHER PHY AND MAC ENHANCEMENTS

Other modifications to the PHY and MAC have been proposed in the 802.11ac amendment. In this section we present three of the most relevant changes.

HIGHER ORDER MODULATION

The modulation and coding scheme (MCS) determines the modulation and coding used to transmit the data portion of the PPDU. At the PHY, the amendment includes mandatory support for binary phase shift keying (BPSK), quaternary PSK (QPSK), 16-QAM, and 64-QAM modulation schemes (thus keeping the same modulation, interleaving, and coding architecture of 802.11n), plus two additional (and optional) 256-QAM with 3/4 and 5/6 coding rates. IEEE 802.11n allows the use of unequal modulations; that is, a single user can receive at one MCS in one stream and at a different one in another stream [7]. In contrast, 802.11ac only supports unequal modulation in the multi-user case, not for a single user.

The use of a particular MCS is dependent on a combination of the channel bandwidth used for the data transmission as well as the number of spatial streams. While most channel bandwidth options can support the highest MCS (i.e., 256-QAM), these are not valid for a certain specific number of spatial streams. Notice that some MCSs are not valid due to the fact that some of the parameters such as the number of data bits per symbol need to be integers, and for certain coding rates this is not feasible. For a complete list of the different supported combinations and the parameters for each MCS, refer to [4].

CODING TECHNIQUES

Forward error correction (FEC) is enabled in 802.11ac. Two different schemes are proposed for this. The first one is binary convolutional coding (BCC), and it is mandatory. On the other hand, LDPC coding is optional. FEC is used by 802.11ac with a coding rate of 1/2, 2/3, 3/4, or 5/6. Furthermore, STBC is also an optional feature of the amendment. Compared to 802.11n, the new standard proposes fewer modes that only include 2 × 1, 4 × 2, and also 3 × 2 and 4 × 3 as extension modes [9]. Each different mode corresponds to a different combination of transmitting and receiving antennas. For example, a 4 × 2 STBC mode indicates that four antennas are used at the transmitter, whereas only two are used on the receiver side.

FRAME AGGREGATION

At the MAC, the standard specifies the use of different frame aggregation schemes, and capability negotiations to indicate channel width, for example, as well as VHT frame formats to enable the operation of the new PHY enhancements. In particular, it proposes mandatory use

of frame aggregation via an aggregate-MAC PDU (A-MPDU), which was introduced in 802.11n. A-MPDUs are enhanced in 802.11ac by increasing their size, thus packing several MPDUs within a single PPDU. This in turn increases channel utilization and MAC efficiency.

CONCLUSION

In this article we present a detailed description of what we consider are the most important enhancements proposed in the 802.11ac amendment. These modifications are key to attaining the performance gains dictated by this new amendment. We identify the changes to the channelization techniques as well as the MU-MIMO capabilities as the paramount strategies for reaching gigabit wireless transmissions. Moreover, we present a comparison of the capabilities of 802.11ac stations with those of legacy devices.

REFERENCES

[1] Cisco Inc., "Cisco Visual Networking iIndex: Global Mobile Data Traffic Forecast Update, 2011–2016," http://cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html.
[2] IEEE 802.11n-2009, "IEEE Standard for Local and Metropolitan Area Networks — Telecommunications and Information Exchange between Systems — Local and Metropolitan Area Networks — Specific Requirements — Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput," 2009.
[3] R. Van Nee, "Breaking the Gigabit-Per-Second Barrier With 802.11ac," *IEEE Wireless Commun.*, vol. 18, Apr. 2011, p. 4.
[4] IEEE 802.11ac/D3.0, "Draft Standard for Information Technology — Telecommunications and Information Exchange Between Systems – Local and Metropolitan Area Networks — Specific Requirements — Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 4: Enhancements for Very High Throughput for Operation in Bands Below 6 GHz."
[5] R. de Vegt, IEEE 802.11-09/0161r2, "802.11ac Usage Models Document," Jan. 22, 2009, https://mentor.ieee.org/802.11/dcn/09/11-09-0161-02-00ac-802-11ac-usage-model-document.ppt.
[6] M. Ergen, "IEEE 802.11 Tutorial," UC Berkeley, 2002, p. 70.
[7] Rohde & Schwarz, "802.11ac Technology Introduction White Paper," Mar. 2012, http://cdn.rohde-schwarz.com/dl_downloads/dl_application/application_notes/1ma192/1MA192_7e.pdf.
[8] M. Park, "IEEE 802.11ac: Dynamic Bandwidth Channel Access," *Proc. IEEE ICC*, 2011.
[9] Cisco, "802.11ac: The Fifth Generation of Wi-Fi Technical White Paper," Aug. 2012, http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps11983/white_paper_c11-713103.html.
[10] D. Gesbert *et al.*, "Shifting the MIMO Paradigm," *IEEE Signal Processing Mag.*, vol.24, no. 5, Sept. 2007, pp. 36–46.
[11] Qualcomm, "IEEE802.11ac: The Next Evolution of Wi-Fi Standards," Aug. 21, 2012, http://www.qualcomm.com/media/documents/ieee80211ac-next-evolution-wi-fi.

BIOGRAPHIES

OSCAR BEJARANO (obejarano@rice.edu) received a B.S. degree from the University of Texas at Austin in 2009 and an M.S. degree from Rice University, Houston, Texas, in 2011. Since then, he has been pursuing a Ph.D. degree in the Rice Networks Group at Rice University; all degrees are in electrical and computer engineering. His current research interests include MAC layer protocol design for MU-MIMO networks.

MINYOUNG PARK (minyoung.park@intel.com) received B.S. and M.S. degrees from Yonsei University, Seoul, Korea in 1993 and 1995, respectively, and a Ph.D. degree from The University of Texas at Austin in 2005, all in electrical engineering. From 1995 to 2001, he worked as a senior research engineer at LG Electronics Inc. Since 2005, he has been a wireless system architect at Intel Corporation. His research interest is in physical and MAC layers of wireless communications systems.

EDWARD KNIGHTLY [F] (knightly@rice.edu) is a professor of electrical and computer engineering at Rice University. He received his Ph.D. and M.S. from the University of California at Berkeley, and his B.S. from Auburn University. He is a Sloan Fellow and a recipient of the National Science Foundation CAREER Award. He received the best paper award from ACM MobiCom 2008 and serves on the IMDEA Networks Scientific Council.