Emerging Technologies for WLAN

VK Jones and Hemanth Sampath

The authors are with Qualcomm.

ABSTRACT

New technologies continue to be introduced for WLAN applications at a robust pace. We review the value proposition for some of the key features of 802.11ac such as larger bandwidth, higher order modulation, and MIMO and MU-MIMO transmission modes, explaining how each feature translates to improved user experience. We present channel measurements and prototype performance data to demonstrate the gains of MIMO and MU-MIMO in an indoor environment. Next, we discuss the value proposition of some key features of 802.11ah. Measurement data of 802.11ah performance is provided, showing how it is a compelling technology for the growing Internet of Things market. We conclude with a preview of emerging technologies that promise to improve user experience — 802.11ai for fast network acquisition and 802.11ax for high-efficiency networking in dense indoor and outdoor networks.

INTRODUCTION

The WLAN market is experiencing unprecedented growth. Almost every home and office in advanced technology markets utilizes a WLAN, and deployments are rapidly proliferating in public areas of congregation like cafés, hotels, and airports. Wireless operators are embracing WLAN for cellular offload as attach rates in smartphones have reached 100 percent. In addition, there is rapid proliferation of WLAN enabled devices across many consumer electronic categories, as consumers demand that their entertainment devices be "connected."

New applications for WLAN continue to emerge. Wireless audio and video streaming are becoming increasingly popular with consumers. Service providers desire WLAN to support multiple high definition (HD) streams in a home, moving content to and from devices like tablets, set-top boxes, digital media adapters, media centers, or televisions. Phone and tablet manufacturers want devices to interact wirelessly to local services or peripherals such as video cameras. In addition, there are emerging use cases like Miracast [1] that provide means for "mirroring" a small screen to a larger video screen.

In order to pave the way for new device categories and new application use cases, there are exciting new technologies emerging for WLAN that will address the need for increased network capacity, longer range, lower power consumption, and ease of use.

KEY TECHNICAL FEATURES OF 802.11AC

MANDATED USE OF THE 5GHZ BAND

The WLAN market is now transitioning from IEEE 802.11n to 802.11ac [2], due to the promise of higher throughput and more reliable performance available in the 5 GHz unlicensed band. 802.11ac mandates the use of the 5 GHz band, a band with more significant spectrum availability compared to the commonly used 2.4 GHz band. While 802.11a mandated use of the 5 GHz band, the attractiveness of 802.11a was limited due to the fact that 802.11g offered similar data rates. 5 GHz was optional for 802.11n; thus, lower-cost 2.4-GHz-only solutions have been more popular than dual-band devices. In contrast, the higher bandwidth modes of 802.11ac cannot be used in the 2.4 GHz band, so the transition from single-band 802.11n to dual-band 802.11ac is being driven by a desire for higher data rates.

In the United States today, there is approximately 500 MHz of available unlicensed spectrum between 5.15 and 5.725 GHz. In Europe and Japan, there is approximately 400 MHz of spectrum available. In India and China, there is 260 and 100 MHz of spectrum available. In contrast, the 2.4 GHz band can accommodate only three non-overlapping 20-MHz-wide channels, which has led to many competing devices per channel and heavy levels of interference. The larger spectrum availability in the 5 GHz band provides for more network capacity, and leads to fewer competing devices per channel and thus reduced interference compared to traditional 802.11g and single-band 802.11n networks.

MAXIMUM DATA RATE

One key feature of 802.11ac is the increase in data rate compared to 802.11n. This is achieved through the use of expanded channel bandwidth and higher-order modulation. Figure 1a shows the peak data rate and the per-spatial stream (SS) data rate for various WLAN standards that have evolved over the years. 802.11b and 802.11g support peak data rates of 11 and 54 Mb/s, respectively. 802.11n increased the peak data rate to 600 Mb/s. 802.11ac has further increased the peak data rate to 6.9 Gb/s, over ten times that of 802.11n.

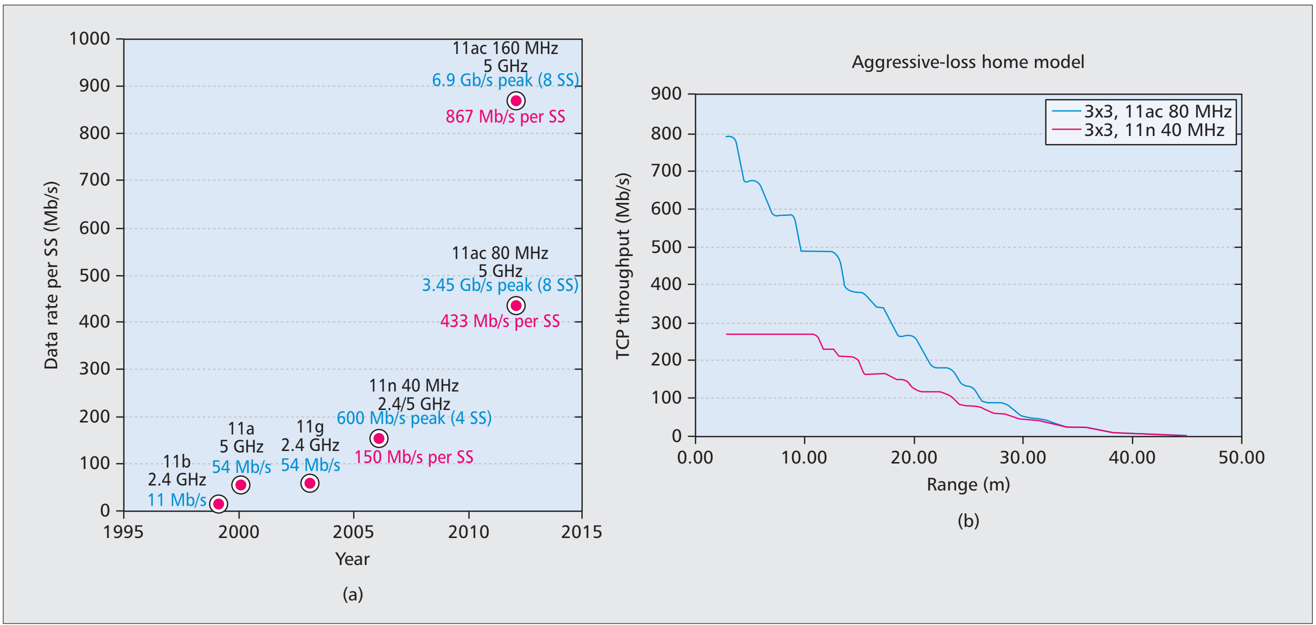

Figure 1. a) Peak data rate and the per-spatial stream (SS) data rate of WLAN technologies; b) rate-over-range of three spatial stream 802.11ac and 802.11n devices.

Wider Bandwidth and Higher Order Modulation

802.11ac introduces three new expanded channel bandwidth modes: 80, 160, and 80+80 MHz, while also defining support for 20 and 40 MHz bandwidth modes to match channel bandwidth modes of 802.11n and 802.11g. The 80 and 160 MHz transmissions use contiguous spectrum, while the 80+80 MHz mode allows the construction of the transmitted signal to occupy separate 80 MHz segments. The 80 MHz transmissions use 234 data tones per orthogonal frequency-division multiplexing (OFDM) symbol, which more than doubles the data rate compared to an 802.11n 40 MHz transmission that only uses 108 data tones. The 160 and 80+80 MHz transmissions use exactly twice the number of data tones as the 80 MHz transmission, thereby doubling the data rate further.

802.11ac introduces the use of 256-quadrature amplitude modulation (QAM) [2]. Two forward error correction (FEC) coding rates are defined: 3/4-rate and 5/6-rate. As a comparison, 802.11n supports up to 64-QAM with these same coding rates. Thus, 802.11ac achieves a 33 percent increase in peak data rate over 802.11n.

The combination of higher-order modulation and increased channel bandwidth enables an 802.11ac device to support approximately three to six times higher data rate compared to an 802.11n device for the same number of antennas or spatial streams (SSs). 802.11n achieves a maximum of 150 Mb/s per SS (108 data tones in 40 MHz of bandwidth with a maximum of 5.0 bits per tone). This results in a maximum data rate of 600 Mb/s, assuming the maximum supported 4 SS multiple-input multiple-output (MIMO) transmission of 802.11n. 802.11ac reaches 433 Mb/s per SS using 80 MHz channel bandwidth (234 data tones with 6.67 bits per tone), and 867 Mb/s per spatial-stream using 160 or 80+80 MHz bandwidth. This results in a maximum data rate of 6.9 Gb/s, assuming the maximum supported 8 SS MIMO transmission of 802.11ac.

While 6.9 Gb/s is an eye-catching maximum data rate, a more commonly cited advantage of 802.11ac is the ability to cross the 1 Gb/s barrier with small form-factor devices [6]. A two-antenna 802.11ac device (using a maximum of two SSs) can support a maximum data rate of 1.73 Gb/s. Furthermore, an 802.11ac device can surpass the data rate of an 802.11n device, but with much lower complexity and cost. For example, an 802.11n device requires three antennas (three spatial streams) to achieve a similar data rate (450 Mb/s) as a single-antenna 802.11ac device.

Rate-Range

In addition to increasing the maximum data rate, these enhancements also lead to improved rate-over-range performance of 802.11ac compared to 802.11n. Figure 1b shows simulated performance of both technologies, using a path loss model validated with measurements from a large home. Both technologies support three spatial streams, and have three transmit and receive antennas (3×3). The figure shows the TCP/IP throughput vs. the distance between the wireless devices.

From Fig. 1b, it can be seen that the 802.11ac devices can connect at twice the range of the 802.11n devices, at the maximum TCP/IP throughput of the 802.11n device (approximately 280 Mb/s). For an end user, this translates to 802.11ac devices experiencing higher throughputs across most locations in a home/office environment. Another observation from Fig. 1b is that the peak rate of the 802.11ac device is three times that of the 802.11n device. For the end user, this translates to 802.11ac-enabled devices experiencing much higher throughput for in-room and peer-to-peer scenarios.

Improved 802.11ac data rates translate to improved application layer throughputs for the same latency. For example, a single-antenna 802.11ac 80 MHz device (433 Mb/s data rate) can support approximately 200 Mb/s application throughput with a 20 ms latency limit. In comparison, a single antenna 802.11n 40 MHz device (150 Mb/s data rate) can support approximately 70 Mb/s throughput with the same latency constraint. Improved 802.11ac data rates also translate to a reduction in application layer latency relative to 802.11n. A single-antenna 802.11ac device (433 Mb/s data rate) can support 100 Mb/s throughput with less than 5 ms latency. In comparison, a single-antenna 802.11n 40 MHz device (150 Mb/s data rate) requires more than 100 ms latency.

Another important benefit of improved data rate is improved battery life due to lower joules-per-bit consumption in a variety of applications like video streaming, Miracast [1] and file transfers. In a typical commercial handset with an integrated CPU running a file transfer application, the measured time for a file transfer was 0.3 s using 802.11ac with 80 MHz bandwidth. That same file size consumed 0.84 s when using 802.11n with 40 MHz bandwidth. This reduction in device on time to complete a task resulted in 2.3 times lower energy consumption when using 802.11ac relative to 802.11n, despite the fact that the instantaneous power consumption for 802.11ac was slightly higher than when using 802.11n.

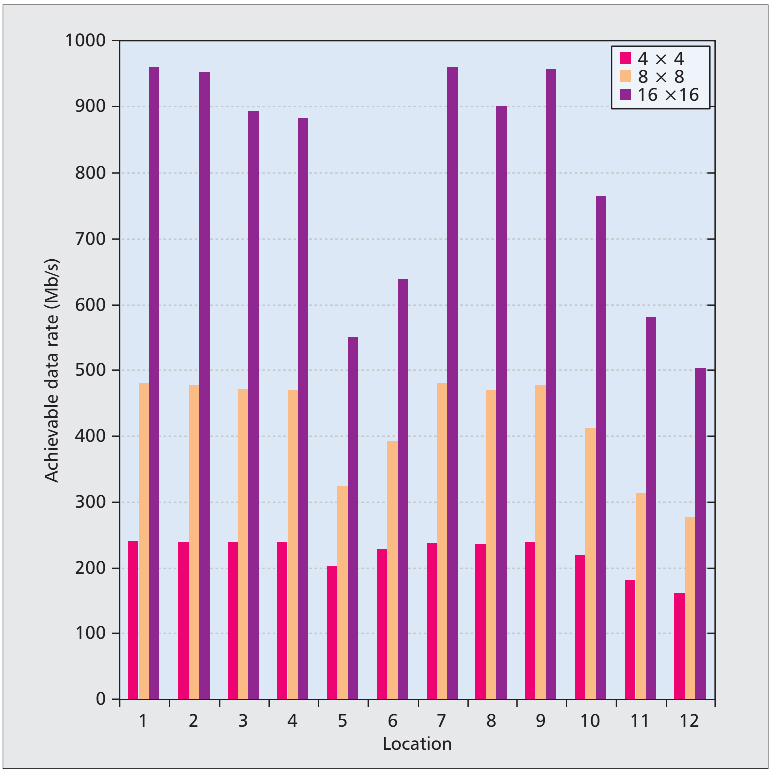

Figure 2. MIMO channel measurements in indoor environments.

Eight Spatial Streams

802.11ac introduces support for up to eight SSs, compared to 802.11n, which defines up to four SSs. In 802.11ac, equal modulation is applied to all SSs for a particular user. Specifically, the transmitter bits are encoded, interleaved and modulated according to one of 10 prescribed modulation and coding scheme (MCSs) and then spatially mapped to physical antennas. The spatial mapping between the SSs and antennas is implementation-specific, and may be frequency-dependent and include transmit steering or precoding matrices [2].

At Qualcomm Incorporated, we conducted extensive indoor channel measurements [3, 4] in an enterprise setting to assist in the construction of the IEEE 802.11ac channel model. Figure 2 illustrates some results from a measurement campaign in one of our office buildings. More than 3500 MIMO channel measurements were made in 12 locations, spanning both line-of-sight (LOS) and non-line-of-sight (NLOS) scenarios, with up to 98 dB of path loss. Measurements were made with up to 16 antennas at both ends of the link, so a comparison can be made of achievable throughput for various numbers of antenna configurations. Figure 2 shows results for 4 × 4, 8 × 8, and 16 × 16 achievable data rates. The achievable data rate was computed by recording the MIMO channel realization across 64 tones in a 20 MHz bandwidth, computing the receiver post-processing signal-to-interference-plus-noise ratio (SINR) assuming a minimum mean square error (MMSE) receiver, and mapping it to the 802.11n physical layer (PHY) rate table using 52 tones for data.

It can be seen from Fig. 2 that data rate scales with MIMO channel dimension beyond 4 × 4 to a large number of antennas. This is attributed to the rich indoor scattering environment, leading to a large number of supported spatial modes even for a 16 × 16 antenna configuration. In most locations, antennas spaced 1/2 wavelength apart observe channel realizations with low correlations that provide almost linear growth of achievable data rate vs antenna dimension. While eight-SS support may seem particularly challenging to accommodate in a small form-factor device, it should be noted that using co-located cross-polarized antennas can save significant space, and that many retail WLAN access points (APs) have more than 10 antennas today. In fact, with additional channel measurements, we observed that using co-located cross-polarized antennas lead to minimal loss in channel capacity in an indoor environment. It was our desire to exploit this spatial capacity, while recognizing the feasibility issues of using more than two antennas at a station device, which led us to promote multi-user MIMO (MU-MIMO) for 802.11ac.

Multi-User MIMO

In 802.11n, a transmission to a device happens using single-user (SU) MIMO modes, where the data rate to a device scales with the minimum number of antennas of each devices. An 802.11n access point (AP) must transmit data using time-division multiplexing (TDM) to different devices, attempting to divide up the network throughput between stations. Unfortunately, 802.11n network capacity is then limited by lower-cost devices that have a smaller number of antennas.

802.11ac MU-MIMO transmission modes

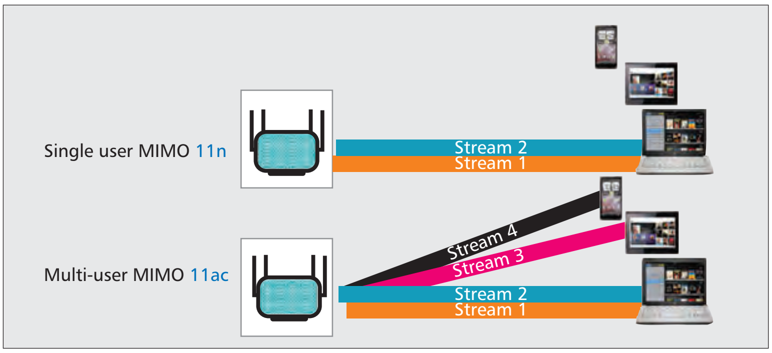

Figure 3. Illustration of MU-MIMO 802.11ac and SU-MIMO 802.11n.

allow simultaneous transmissions to multiple devices using space-division multiplexing (SDM), as depicted in Fig. 3. This significantly improves the spectral efficiency of a WLAN when there are stations with limited numbers of antennas. Essentially, MU-MIMO captures the maximum spatial capacity without requiring all individual stations to have a large number of antennas. The 802.11ac standard allows a MU-MIMO transmission to be sent to up to four simultaneous stations. Each station may receive up to four SSs, but the total number of SSs in a MU-MIMO transmission does not exceed eight total SSs summed across all stations.

As an example, consider a WLAN network where the AP, utilizing four antennas, serves downlink traffic to a laptop with two antennas and a handset with one antenna. Even though an AP may support four SS MIMO (i.e., 1.73 Gb/s data rate in 80 MHz), only two SS transmissions (i.e. 867 Mb/s) can be supported to the laptop, and only a single SS can be transmitted to the handset (i.e., 433 Mb/s) due to the antenna limitations at the stations. Without MU-MIMO, the AP must TDM data to both devices and incur some additional medium access contention overhead. Hence, the network capacity with both devices drops to less than the average of the two rates (e.g., less than 650 Mb/s). The single-antenna handset device creates a capacity bottleneck for the two-antenna device.

Using MU-MIMO, the AP can transmit data to both devices simultaneously, avoiding additional contention overhead and increasing the network data rate to 1300 Mb/s. MU-MIMO delivers MIMO network capacity, but without the requirement of many antennas at the stations.

An important benefit of an 802.11ac network is the reduction of station-side complexity. Thus, 802.11ac can be viewed as a technology that lowers network deployment costs. This is because in an 802.11ac MU-MIMO network, performance is not sacrificed with the addition of stations with fewer antennas. Fewer antennas means lower-cost station devices with fewer analog chains and with digital circuitry built for fewer SSs. As a result, an 802.11ac device can have a smaller number of antennas than an 802.11n device, but support much higher network throughput compared to an 802.11n device. A Qualcomm over-the-air MU-MIMO prototype has demonstrated that a lower-complexity 802.11ac network with a 4 × 4 AP and three single-antenna stations has similar network throughput as a more complex 802.11n network with a 4 × 4 AP and three 3 × 3 station devices, even when both types of networks use the same channel bandwidth.

Transmit beamforming and MU-MIMO require knowledge of the channel state to compute a precoding or steering matrix that is applied to the transmitted signal to optimize reception at one or more stations. The 802.11ac standard employs an explicit channel sounding and feedback protocol that works as follows. Upon gaining access to the medium, the AP transmits a Null Data Packet (NDP) Announcement Frame, identifying the stations for which it wants to collect channel state information for a subsequent MU-MIMO transmission, followed by an NDP frame. The first identified station in the NDP Announcement Frame then estimates the downlink channel from the NDP and feeds back the channel state information (CSI) to the AP. The remaining identified stations send CSI feedback upon subsequently being polled. The AP can then calculate the MU-MIMO precoding weights and use them in subsequent MU-MIMO data transmissions to the relevant stations. One positive consequence of the 802.11ac MU-MIMO specification is that it led to a unification of an underlying transmit beamforming feedback mechanism. This fact accelerated the industry adoption of transmit beamforming, and cleared up the confusion surrounding the multiple defined methods of 802.11n.

MU-MIMO Performance

During the development of the 802.11ac specification, we validated the performance of MU-MIMO over the air (OTA), using custom software and utilizing a four-antenna 802.11n mini-PCIe card containing the Qualcomm WCN 1320 chip (Fig. 4a).

Similar to how 802.11ac is defined, the AP transmitted one NDP to multiple stations. Each station in the network then estimated the downlink channel from the NDP and sequentially transmitted back the CSI data to the AP using an 802.11n packet. The AP used the CSI to calculate the MU-MIMO MMSE precoder and then transmitted a single precoded SS to each station simultaneously. The OTA prototype was equipped with a tool to plot a raw packet-by-packet data rate over time. Physical layer data rates per packet were calculated by determining the receiver post-processing SINR and mapping SINR to an appropriate 802.11n or 802.11ac MCS. Much of the testing used the four-antenna AP device in Fig. 4a, but to test the full capabilities of the 802.11ac MU-MIMO specification, we also built a custom eight-antenna AP that could enable transmissions of up to eight SSs distributed across multiple stations.

Figure 4b shows a representative snapshot of 802.11n SU-MIMO, 802.11ac SU-MIMO, and 802.11ac MU-MIMO PHY data rates. The network configuration uses a four-antenna AP, three one-antenna stations, 40 MHz bandwidth (BW), and long guard interval. Figure 4b shows data collected at in-room distances between the AP and stations, so there was a high signal-to -

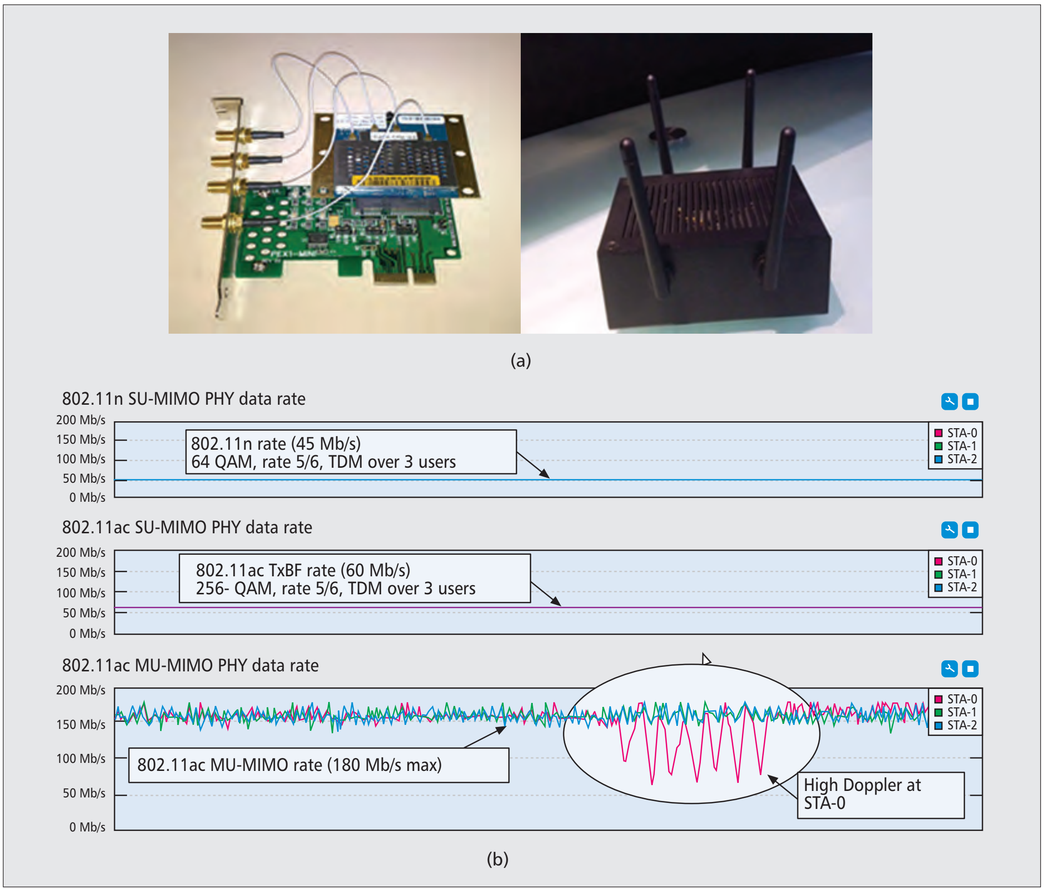

Figure 4. a) 4 × 4 802.11n mini–PCIe card and MU–MIMO access point; b) time snapshot from OTA MU–MIMO prototype in LOS, high SNR scenario.

noise ratio (SNR) for each link. It should be noted that no MAC layer adaptive rate control or rate averaging across packets was employed. It can be seen from Fig. 4b that the 802.11n SU-MIMO transmissions are saturated at the maximum 802.11n MCS of 64-QAM rate 5/6 (135 Mb/s), leading to an effective data rate per station (STA) of 45 Mb/s, including the TDM factor of three. The 802.11ac SU-MIMO transmissions are saturated at the maximum MCS of 256-QAM rate 5/6 (180 Mb/s), leading to an effective data rate per STA of 60 Mb/s, including the TDM factor of three.

MU-MIMO provides an improvement of approximately three times the data rate over baseline 802.11n transmissions in this scenario. These gains were realized even when the AP and stations were in LOS channel conditions, and even with stations placed as close together as possible. This is in alignment with channel measurement data presented earlier, where it was explained that even a half-wavelength antenna spacing leads to a very high number of available spatial dimensions. In practice, despite close proximity or LOS channel conditions, there remains enough local scattering from walls, floors, ceilings, and the device itself to create a rich multipath environment.

Another important finding, illustrated in Fig. 4b, is that MU-MIMO transmissions are robust to a noisy or rogue station that may send outdated or otherwise inaccurate CSI feedback (due to either Doppler or poor implementation). In the data-rate-versus-time plot, STA-0 was subjected to a high Doppler event by rapidly waving the handheld device. As a result, the CSI feedback from STA-0 was typically outdated once the MU-MIMO transmission occurred. This reduction in CSI fidelity for STA-0 directly impacts the STA-0 data rate and not that of other STAs in the MU-MIMO group. Intuitively, this can be explained as follows. In high SNR regimes, MMSE-based transmit precoding schemes approach the zero-forcing precoder, where transmissions to a given station occur in the null-space of other stations. If a station's CSI (and hence its null-space information) is incorrect at the AP, only that station suffers from interference as transmissions to other stations will no longer be in the null-space of the station.

Figure 5a shows a representative plot of the median system throughput over a two-hour period of time vs. the total number of MU-MIMO SSs (Nss) used across all stations. We see that maximum system throughput is obtained when the total number of MU-MIMO SSs is approximately 75 percent of the total number of AP antennas. For an AP with four antennas, throughput is maximized with three total SSs. An eight antenna AP can maximize throughput using six total SSs. For less than 75 percent loading, system throughput is proportional to the

MU-MIMO provides an improvement of approximately three times the data rate over baseline 802.11n transmissions in this scenario. These gains were realized even when the AP and stations were in LOS channel conditions, and even with stations placed as close together as possible.

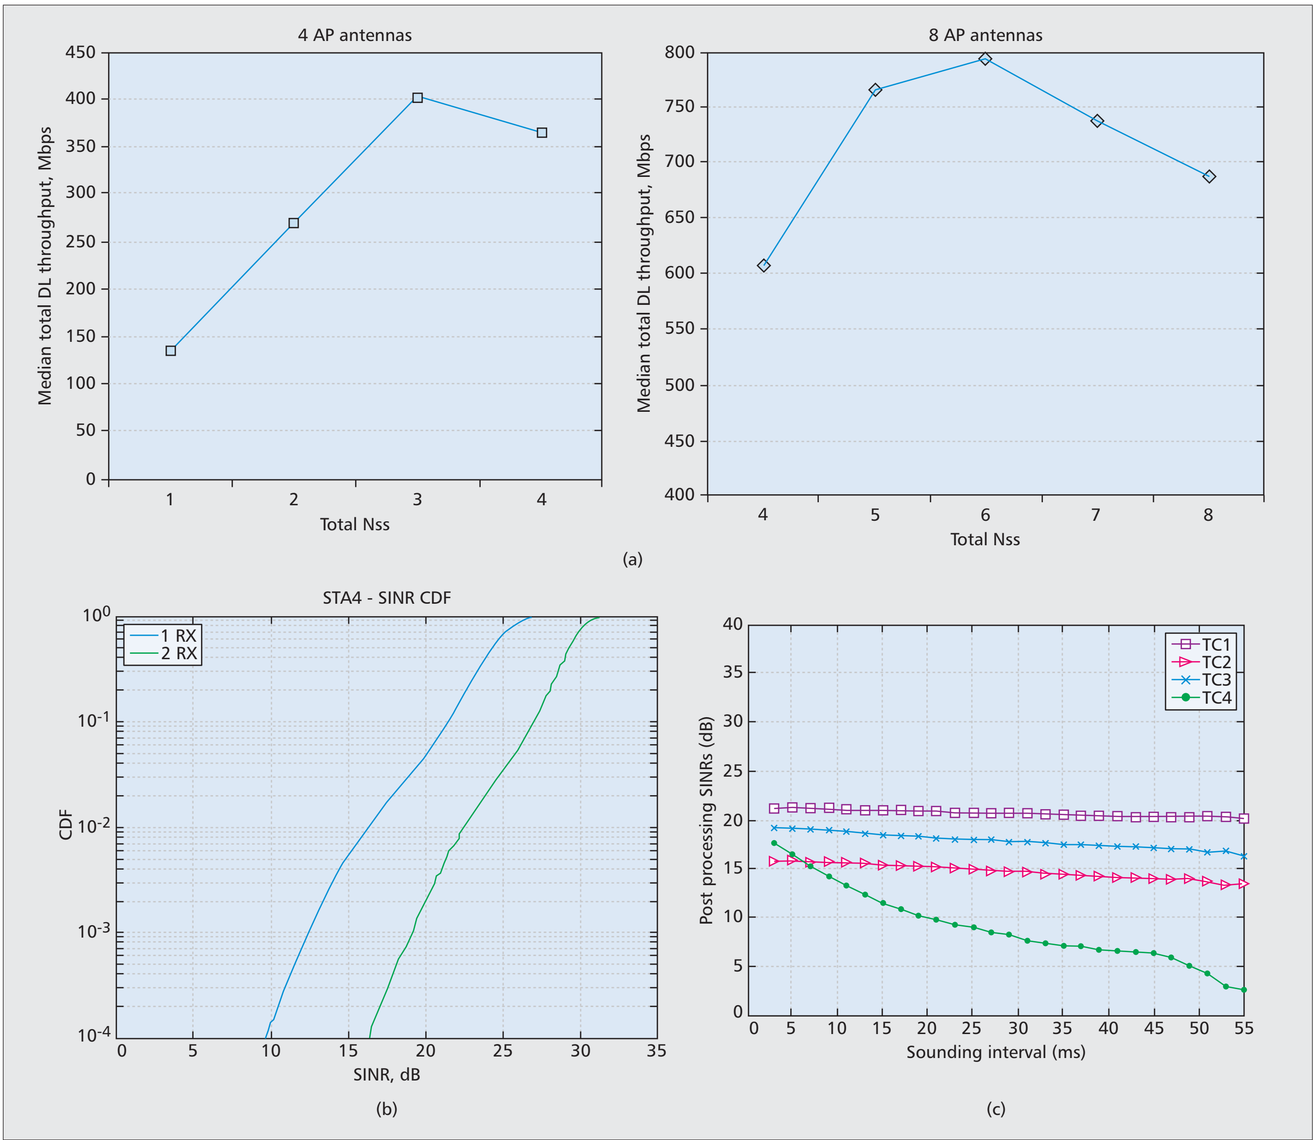

Figure 5. a) Median throughput vs. total number of MU-MIMO spatial streams (Nss) across all STAs; b) CDF of MU-MIMO post-processing SINRs with 1 Rx antenna and 2 Rx antennas (interference nulling enabled); c) MU-MIMO post-processing receiver SINRs, over time, for different user test cases (TC): TC1 — watching a movie; TC2 — typing on the keypad; TC3 — holding the device to the ear; TC4 — walking.

number of SSs, and inter-stream interference tends to be negligible. However, beyond 75 percent loading, the throughput loss due to inter-stream interference exceeds any additional throughput gains due to increased number of spatial streams. That level of inter-stream interference is dependent on the noise floor of the station devices, along with the transmit error floors due to common analog impairments in an AP. This demonstration of MU-MIMO employed devices that have 3 dB margin to the error vector magnitude requirements in the 802.11ac specification.

If a station has additional antennas beyond the number of spatial streams assigned to it in a MU-MIMO transmission, the surplus antenna(s) can be used to cancel interference created by MU-MIMO streams intended for other stations [5]. This is enabled by the design of the long training field(s) (LTFs) in 802.11ac, allowing the receiver to estimate the channel corresponding to all MU-MIMO spatial streams. Figure 5b demonstrates the improvement of post-processing receiver SINR due to interference nulling. The figure shows the cumulative distribution function (CDF) of post-processing SINR collected from a network configuration where three stations each receive one spatial stream from a four-antenna AP. One of the stations has two receive antennas and uses the extra antenna to cancel any interference from other streams. The SINR is logged from the stations during a two-hour period in a busy lab environment. In this scenario, interference nulling provides a 5 dB median SINR gain for the station with two receiver antennas.

Figure 5c demonstrates another observation about the performance of MU-MIMO with handheld devices, which will exhibit repeated small movements associated with stationary use cases such as watching a movie, typing the keypad, and holding the device to the ear. For reference, performance with a station experiencing walking speeds was also recorded. For each test case, multiple tests were run across a variety of LOS and NLOS conditions, with varying path loss. The stations are actual handset mockups in order to encourage users to interact as if they were using a real phone. In addition, different test subjects were used to gather diverse data. For each test, post processing SINRs were recorded for a minute of time. The main observation is that for typical stationary use cases, degradation in the post-processing SINR is at most 3 dB for a CSI feedback delay of 50 ms. Therefore, MU-MIMO works well with handheld devices and typical use cases.

For the higher Doppler rate walking test, where the test subject is walking with the mockup antenna placed to the user's ear, the degradation in SINR is ~10 dB for a CSI feedback delay of 50 ms. With immediate CSI feedback, the performance degradation was less than 3 dB. Thus, for mobility use cases, immediate feedback is required for good performance.

Future technology enhancements like uplink MU-MIMO can be used to reduce CSI feedback overhead in order to improve downlink MU-MIMO performance even in mobility use cases.

802.11ah for Internet of Things and Extended-Range WLAN

Introduction

The wireless industry is poised for high growth in the wireless sensor market, with applications across home and industrial automation, healthcare, energy management, and wearable devices. The home automation category includes devices such as temperature, moisture, and security sensors. They are often small form-factor devices that are battery-powered and headless, but require whole-home coverage (including attics, backyards, basements, and garages). Unfortunately, the market is fragmented with multiple non-interoperable technologies, some with coverage issues and some with non-user-friendly network configuration and deployment issues.

The IEEE 802.11ah Task Group is developing a specification for the license-exempt bands below 1 GHz, targeting such lower-data-rate and longer-range applications for devices commonly referred to as the Internet of Things. 802.11ah enables a "lower band of Wi-Fi," augmenting the traditional 2.4 and 5 GHz bands used today.

802.11ah Features and Supporting Measurements

The 802.11ah Draft Specification [7] defines mandatory 1 and 2 MHz BW modes that are globally interoperable. These modes are ideally suited for devices that require low power consumption and long-range connection to an AP. In addition, there are optional 4, 8, and 16 MHz modes for the more traditional WLAN use cases, which can be used in regulatory domains that allow larger BW. The 802.11ah modes enable a wide range of data rates between 150 kb/s to 78 Mb/s per spatial stream. The lowest data rate of 150 kb/s is enabled by the use of 1 MHz bandwidth, 2× repetition coding, 1/2-rate FEC coding, and binary phase shift keying (BPSK). For four SSs, the maximum data rate is 312 Mb/s. 802.11ah inherits the advanced FEC and spatial diversity schemes from 802.11ac. This rich data set and robustness enable a wide variety of use cases pertaining to the Internet of Things such as sensors, wireless audio, security cameras, wireless video, and Internet connectivity.

Qualcomm Incorporated has conducted measurement studies across several single-family (2400–5000 ft^2) U.S. homes [8] to demonstrate the improved range characteristics of 802.11ah modes in the 900 MHz unlicensed band. Our measurements revealed more than a 10 dB range advantage of 802.11ah over 802.11n and 802.11b. In the measurement campaign, the APs were placed in typical locations, and measurements were made at multiple locations around the home, including backyards, garages, and basements. At each location, the median path loss and SNR were measured in both the 900 MHz and 2.4 GHz bands. These measurements were combined with 802.11n (20 MHz) and 802.11ah (1, 2, 4, and 8 MHz) receiver sensitivity performance targets to estimate the supported physical layer data rates at each location. The data rate computation assumed an AP with two antennas and stations with a single antenna. In addition, an 8 dB obstruction loss was applied to account for possible realistic obstructions like human bodies and furniture. The data rates reflect uplink performance assuming only 4 dBm of transmit power from stations. Figure 6 shows the 802.11ah rate vs. range improvement over 802.11n and 802.11b from three representative homes with varying square footage.

It can be seen that 802.11ah has improved coverage and higher data rates in locations with larger path loss due to the lower-BW physical layer modes [8] and better propagation characteristics at sub-1 GHz frequencies. 802.11n and 802.11b devices experience significant coverage holes for locations with larger range, and would require higher transmit power for whole home coverage. That higher output power comes at a significant disadvantage of higher power consumption for the 802.11n and 802.11b stations. By exploiting the better range capabilities of 802.11ah to reduce output power, 802.11ah stations can achieve multiple-year battery operation for applications that require low duty cycle.

In addition to the link budget benefits described above, 802.11ah accommodates larger delay spread and Doppler spread, making it a favorable technology for outdoor use. Due to longer symbol times, the delay spread tolerance is 10 times that of 802.11ac. Furthermore, 802.11ah has an improved pilot design that enables robust channel tracking throughout a packet.

802.11ah introduces several innovations to enable low-power applications. Included in this list are smaller frame formats that save power [9], new traffic priority rules, and scheduled access for battery-operated devices that improves

The 802.11ah Draft Specification defines mandatory 1MHz and 2MHz BW modes that are globally interoperable. These modes are ideally suited for devices that require low-power consumption and a long-range connection to an AP. In addition, there are optional 4, 8 and 16 MHz modes for the more traditional WLAN use-cases.

latency and reduces collisions. Furthermore, there are more efficient sleep modes that allow a device to remain asleep longer and wake-up more efficiently [10]. In addition, an 802.11ah network can support thousands of devices due to new efficient paging and scheduling mechanisms. 802.11ah also has mechanisms to enable a relay operation that can be used to further extend coverage.

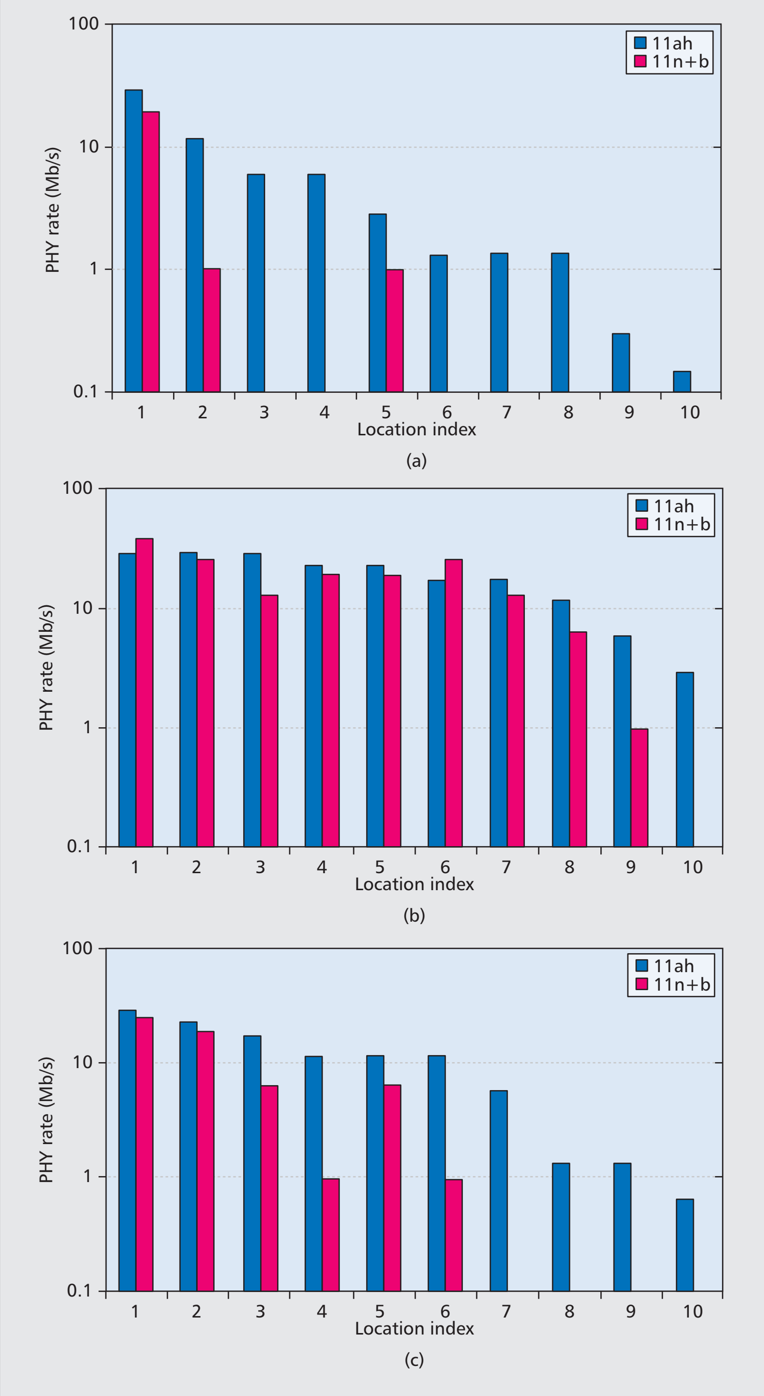

Figure 6. Rate and coverage advantage of 802.11ah over 802.11n and 802.11b in single-family homes: a) 2230 ft^2 home with basement and yard space (2.4 GHz path losses: 76–104 dB); b) 3500 ft^2 single-family home with two floors (2.4 GHz path losses: 72–94 dB); c) 4249 ft^2 home with basement, backyard, and front yard (2.4 GHz path losses: 76–98 dB).

The above mentioned technology enhancements, coupled with an ecosystem of interoperable WLAN devices (including smartphones), make 11ah a compelling technology for the Internet-of-Things.

Emerging Technologies

Due to the overwhelming success of WLAN, improved connection times and throughput performance in crowded networks is an area of concern in the future. Thus, industry efforts have begun to address performance in these scenarios.

802.11ai for Optimized Connectivity Experience

The IEEE 802.11ai Task Group is developing a specification [11] for fast initial link setup (FILS), which promises to address operator concerns about poor offload experience to WLAN. These concerns include significant connection signaling overhead, slow handoff mechanisms between hotspot APs, and "probe storms" that lead to service outage.

802.11ai enhances network connection experience through faster initial link setup time, enabled by defining steps to complete association and authentication with a four-way handshake and IP address assignment in just two round-trip messages [12]. A FILS discovery frame is defined that leads to more efficient and more frequent beacon transmissions, leading to more rapid network detection. AP-to-AP handoff is improved in dense networks by means of a Neighbor AP advertisement and FILS parameter advertisement in beacons and probe responses. Neighbor AP advertisement contains elements such as basic service set identification (BSSID), service set identification (SSID), channel information, and target beacon transmit time (TBTT) offset of neighboring APs. Knowledge of these parameters significantly reduces device network scanning time. FILS parameter advertisement consists of elements such as subnet identifier (ID), authentication domain, and IP address domain of the AP. Knowledge of these parameters helps the device identify the correct APs and enable faster link setup time.

Qualcomm Incorporated has prototyped elements of 802.11ai on commercially available handsets and APs. Assuming a 200 ms authentication server delay, it was observed that 802.11ai significantly reduced handset connection time to a new AP with a different SSID and subnet ID, from approximately 8 s to less than 0.5 s.

802.11ai also reduces probe storms by introducing alternative methods of connection to inefficient Probe Request and Response messages. New efficient messaging is introduced including Broadcast Probe Response, Selective Probe Response based on Probe-Request content, and shorter Probe Response containing only the changed parameters since the last association. The FILS discovery frame, which is sent much more frequently than the legacy beacon, also reduces Probe Request and Response traffic.

802.11AX FOR HIGH-EFFICIENCY WLAN

To address inefficiencies of WLAN in dense indoor and outdoor networks, and improve robustness to interference in the traditional 2.4 and 5 GHz bands, a new IEEE 802.11 Task Group called 802.11ax has been formed. The Task Group is in the early stage of specification development with a projected completion date around the 2019 timeframe. Evaluation metrics include average and 5 percent per station throughput, area throughput, and packet delay and error rate requirements of applications. The Task Group will take into consideration new application trends with bidirectional, uplink-intensive, and peer-to-peer traffic such as video-conferencing, user-generated uploads from wearable devices, smartphones and other consumer electronic devices, display, docking, and scalable peer-to-peer networks.

A key technical direction for the Task Group is to increase parallelization of traffic in the spatial and frequency domains, to achieve at least four times average medium access control (MAC) throughput increase per station over 802.11ac networks. Technologies under consideration include uplink MU-MIMO, orthogonal frequency-division multiple access (OFDMA), an improved OFDM numerology with longer symbol duration and cyclic prefix for outdoor channel support, and enhancements to the legacy clear channel assessment (CCA) schemes.

ACKNOWLEDGMENTS

The authors would like to thank Qualcomm Incorporated's WLAN technology teams for their valuable contributions to the design and prototype of various standards. The authors wish to thank Rolf De Vegt, Richard van Nee, Allert van Zelst, Menzo Wentink, Simone Merlin, Sameer Vermani, Santosh Abraham, George Cherian, Youhan Kim, Bin Tian, Eugene Baik, Alfred Asterjadhi, and Amin Jafarian for making valuable contributions to the development of various IEEE specifications. The authors would like to thank Rahul Tandra, Max Krasnyanskiy, Vinayak Chowdiah, Geert Awater, Anurag Vardhan, Luiza Timariu, Danny Jen, Wenjia Lai, and Greg Breit for their valuable prototype development and performance validation efforts.

The Task Group will take into consideration new application trends with bi-directional, up-link intensive and peer-to-peer traffic such as video-conferencing, user-generated uploads from wearable devices, smartphone and other consumer electronic devices, display, docking and scalable peer-to-peer networks.

REFERENCES

[1] Miracast. http://www.wi-fi.org/discover-wi-fi/wi-fi-certified-miracast
[2] IEEE Std 802.11ac-2013, amendment to IEEE Std 802.11™-2012, Dec. 2013.
[3] G. Breit *et al.*, "802.11ac Channel Modeling," IEEE 802.11-09/0088r1, Jan. 2009.
[4] G. Breit *et al.*, "Multi-User MIMO Channel Measurements," IEEE 802.11-09/0574r0, May 2009.
[5] S. Vermani and A. Van Zelst, "Interference Cancellation for DL MU-MIMO," IEEE 802.11-10/1234r1, Mar. 2010.
[6] VK Jones, "802.11ac Practical Gigabit Wi-Fi," Webinar, May 2, 2012.
[7] IEEE Std 802.11ah/D3.1, "Sub 1 GHz License Exempt Operation," Nov. 2014.
[8] S. Vermani *et al.*, "Preamble Formats for 1 MHz," IEEE 802.11-11/1482r4, Jan. 2012.
[9] S. Merlin *et al.*, "Short MAC Header," IEEE 802.11-12/0857r0, July 2012.
[10] M. Wentink *et al.*, "Low Power Medium Access," IEEE 802.11-12/0114r0, Jan. 2012.
[11] IEEE Std 802.11ai/D3.0, "Fast Initial Link Setup," Sept. 2014.
[12] G. Cherian *et al.*, "Text for FILS Authentication," IEEE 802.11-12/0123r0, Jan. 2012.

BIOGRAPHIES

VK JONES (vkjones@qca.qualcomm.com) is a vice president, technology, at Qualcomm Incorporated where he leads the Wireless LAN System Engineering organization. He came to Qualcomm in 2006 through the acquisition of Airgo Networks Incorporated, a company he co-founded in 2001. Prior to Airgo Networks, he was chief scientist and co-founder of Clarity Wireless, which in 1998 was acquired by Cisco Systems, where he helped establish the first wireless business unit. Prior to founding Clarity Wireless, he served as acting assistant professor, mechanical engineering, Stanford University. He received his Ph.D. and M.S.E.E. from Stanford University in 1995 and 1991, respectively. He received his B.S.E.E. from Purdue University in 1990, where he won numerous basketball championships for Delta Tau Delta.

HEMANTH SAMPATH (hsampath@qti.qualcomm.com) is a principal engineering manager and WLAN systems engineering lead at Qualcomm Research Center. His team focuses on design, standardization, prototype, and product support for next generation WLAN technologies. Prior to his current role, he was the systems lead for a pre-LTE (UMB) chip development project within Qualcomm. From 2001 to 2004, he held technical positions at Iospan Wireless Inc. and Marvell Semiconductor. He received his Ph.D. and M.S.E.E. from Stanford University in 2001 and 1998, respectively. He received his B.S.E.E. from the University of Maryland at College Park in 1996.