# Use MU-MIMO at your own risk — Why we don't get Gb/s Wi-Fi

Hyunwoo Choi, Taesik Gong, Jaehun Kim, Jaemin Shin, Sung-Ju Lee*

*School of Computing, KAIST, Daejeon, Republic of Korea*

A B S T R A C T

With the ever growing popularity of mobile devices, the demand for wireless bandwidth has also increased, with the mobile users now expecting wireless network quality similar to what they experience with wired networks. Wireless LANs have evolved over the last twenty years, with major breakthrough technologies such as OFDM (Orthogonal Frequency Division Multiplexing), MIMO (Multiple Input Multiple Output), and MU (Multi-User)-MIMO. The latest IEEE 802.11ac standard supports up to 6.9 Gb/s theoretical capacity, but it could only be achieved with 8-streams in a "perfect" environment. Commercial 802.11ac *wave 2* APs that include MU-MIMO capability, have only recently been made available in the market. We deployed a few APs from different vendors (that uses chipsets from different vendors) in various office environments and measured user throughput on smartphone mobile devices. We observe an enormous gap between theory and practice, with MU-MIMO often providing less throughput than SU (Single User)-MIMO in various network environments. We analyze the root cause of performance issues and suggest future research directions to achieve Gb/s Wi-Fi in practical deployments.

© 2018 Elsevier B.V. All rights reserved.

## 1. Introduction

IEEE 802.11 Wireless LAN, Wi-Fi, is a widely used wireless communication technology. As the popularity of smartphones and mobile devices has risen, the availability and performance expectation of Wi-Fi connection has also increased, thanks to its ease of deployment and cost effectiveness. The technology itself has evolved over the years, with the link speed significantly being improved. IEEE 802.11ac [1], one of the latest standard specifications, also focuses on physical capacity enhancements similar to its predecessors. From the second wave product of 802.11ac, the maximum downlink capacity almost doubled from the first wave products through the implementation of Multi-User MIMO (MU-MIMO) technology [2].

In MU-MIMO, a wireless Access Point (AP) can transmit multiple data streams *simultaneously* to multiple clients, thus increasing the throughput in proportion to the number of streams. When transmitting eight streams concurrently at 160 MHz channel width, a single AP can achieve up to 6.9 Gb/s capacity as defined in the standard specification. With typical consumer or enterprise APs currently in the market, theoretical maximum capacity with four antennas (four streams and 80 MHz bandwidth) is 1.7 Gb/s.

Since MU-MIMO significantly improves wireless network capacity, major AP and mobile device manufacturers have been adopting the latest WLAN specification to their high-end products. However, there is a large wireless link speed gap between what the manufacturers advertise and the throughput experienced by the users. An industry test report [3] showed that the aggregate throughput of MU-MIMO enabled devices was only about 100 Mb/s with 20 MHz bandwidth. A research study [4] reported that the aggregate throughput of MU-MIMO was less than that of Single-User MIMO (SU-MIMO) at 25% of their experimental settings.

There have been numerous studies that systematically measure and analyze the performance of various wireless products and technologies. The performance of 802.11ac "wave 1" APs have been reported [5,6], but these APs do not implement MU-MIMO. There have been research results that demonstrate the feasibility and effectiveness of MU-MIMO [7–12], but they used research platforms such as WARP (Wireless Open-Access Research Platform) [13], Universal Software Radio Peripheral (USRP) [14], or Microsoft Research Software Radio (Sora) [15] that might not reflect the performance a mobile user experiences in commercial deployments. Some manufacturer's test reports [3,16] measured MU-MIMO throughput, but their aim is to demonstrate the benefit of their specific product. A recent report [17] measured throughput of various MU-MIMO commodity APs, but their experimental setup is limited to a specific lab environment where a wired connection between the AP's antenna and VeriWave [18] equipment is used.

We investigate the user throughput that commercial devices in the current market can achieve through MU-MIMO in various indoor office environment, using various 802.11ac *wave 2* APs (five
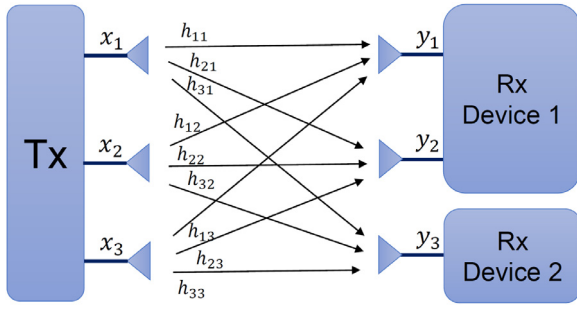
**Fig. 1.** A MU-MIMO system.

different models, three vendors, four chipsets) and smartphones (two different models). The term "user throughput", which we use throughout this paper, represents the actual data rate delivered to the receiver. We measure the throughput in both MU-MIMO and SU-MIMO settings to discover whether and when MU-MIMO provides performance benefits to the users, and analyze the main culprits that cause the throughput gap between theory and practise. We also suggest potential research directions to improve user throughput in WLANs.

Our main findings are as follows:

- We found that none except one Wi-Fi chipsets in our experiment performs MU-MIMO transmission for two antenna-enabled clients group. Moreover, one AP manufacturer, by default, does not activate MU-MIMO.
- We reveal that MU-MIMO throughput has a positive gain over SU-MIMO only in a limited, specific condition, and the maximum MU-MIMO throughput was less than SU-MIMO throughput.
- We discover that commercial APs perform non-optimal user client grouping as channel probing latency exceeds channel coherence time.
- We show that APs fail to adapt to environment changes due to the lack of complete channel state information, resulting in low throughput.
- We identify other issues impacting user throughput, such as APs not adjusting the number of streams, APs selecting wrong transmission rates, and the application not adjusting the socket buffer size.

The rest of the paper is organized as follows. We provide background knowledge on MIMO and MU-MIMO in Section 2 and describe our experiment setup in Section 3. In Section 4, we provide results and analysis of our experiments in various settings. We discuss in Section 5 how to overcome the performance issues and suggest potential research directions. We then conclude in Section 6.

## 2. Background

IEEE 802.11ac extends the existing 802.11n specification and includes wider bandwidth (160 MHz), higher modulation scheme (256 QAM), more spatial streams (up to 8), and downlink MU-MIMO for the Very High Throughput (VHT) communication. We provide the basic background of MIMO, as our study mainly focuses on MIMO and MU-MIMO.

### 2.1. MIMO & MU-MIMO

Multiple Input Multiple Output (MIMO) is a system consisting of a transmitter and a receiver with multiple antennas as shown in Fig. 1. With multiple antennas, an AP can provide transmit diversity by using Cyclic Shift Diversity (CSD), extend coverage by using

transmit beamforming (TxBF), improve the transmission reliability by using Space Time Block Coding (STBC), or increase the data rate by using Spatial Division Multiplexing (SDM).

A key technology of the 802.11ac standard is MU-MIMO. While SU-MIMO supports only one user at a time, MU-MIMO can accommodate multiple users simultaneously so that the throughput, in theory, could be multiple times better than SU-MIMO. MU-MIMO utilizes SDM that sends multiple independent streams concurrently through multiple antennas [19,20]. In other words, when there exist independent propagation paths between each transmitter's antenna and receiver's antenna pair, the network can utilize multiple independent bit streams in that channel. We refer to this individual stream as a "spatial stream." Therefore, when there are independent paths, the maximum throughput can increase as the number of the spatial streams increases.

The possible number of spatial streams is determined by the number of antennas of a transmitter and receivers; it cannot exceed the maximum number of transmitter's antennas and receivers' antennas. For example, if a transmitter has three antennas and two receivers have two antennas respectively, the maximum number of spatial streams is three. Commodity APs that support MU-MIMO usually has four antennas and thus could transmit up to four streams simultaneously.

Wireless signal experiences reflection, refraction, diffraction, and attenuation while travelling over the air, and the received signal gets distorted. This phenomenon is called "fading," and the extent of fading varies with each transmit and receive antenna pair (each spatial stream). A wireless communication system is represented by a mathematical model as shown in Eq. 1, and the MIMO system illustrated in Fig. 1 can be represented as Eq. (2), where $x$ is the transmitted data, $y$ is the received data, and $h$ represents channel fading coefficient [21]. The channel state information, $H$, can be obtained by the channel sounding procedure or by analyzing the uplink channel by the beamformer itself.

$$y = Hx, \tag{1}$$

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}. \tag{2}$$

In MIMO, an antenna receives not only a desired stream for itself, but also other streams sent for other antennas. These other streams act as interference to the antenna, and it is essential to remove interference as well as compensate for distortions. While there are several precoding techniques [22–25], ZFBF (Zero Force Beamforming) [24] has been widely used for its simplicity. In ZFBF, the transmitter calculates the pseudo-inverse matrix, $W_{ZF}$, of $H$ (Eq. 3), and premultiply $x$ by $W_{ZF}$ before the transmission (Eq. (4).

$$W_{ZF} = H^H (HH^H)^{-1}, \tag{3}$$

$$y = HW_{ZF}x. \tag{4}$$

### 2.2. Multi-user beamforming in IEEE 802.11ac

#### 2.2.1. Channel sounding

In order to support 802.11ac MU-MIMO beamforming, a beamformer (i.e., AP) must be informed of the channel state from the beamformees (i.e., clients) as shown in Fig. 2. This channel state information includes the attenuation and the phase difference from the transmitted signal to the received one. With this feedback, a beamformer transmits the desired signal to the beamformees by multiplying a precoding matrix. The beamformer broadcasts a Null Data Packet Announcement (NDPA) packet to the beamformees to notify the beginning of the channel sounding procedure and
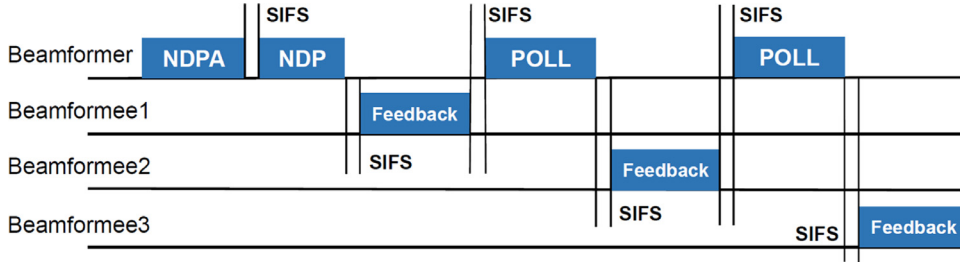
**Fig. 2.** Channel sounding procedure.

transmits a Null Data Packet (NDP) after the Short Inter-Frame Space (SIFS) duration. NDP consists of short and long training symbols.

Each beamformee creates a feedback frame by measuring the channel state from NDP and sends it to the AP (i.e., beamformer). To avoid collision between the feedback frames, the beamformer specifies in NDPA the ID of the beamformees and the sending order of the feedback. The beamformee in the first order sends the feedback first, and the rest send their feedback after receiving the beamforming report poll from the beamformer.

### 2.2.2. Beamforming feedback

A beamformee provides the AP with a beamforming feedback during the channel sounding procedure. The feedback frame consists of VHT MIMO control field, compressed beamforming report, and MU-Exclusive beamforming report [1]. The VHT MIMO control field contains MIMO configurations such as the bandwidth and the resolution of the feedback information. The compressed beamforming report contains average SNR and the compressed form of the channel matrix called the *V-matrix*. The MU-Exclusive beamforming report is a newly added field for MU-MIMO. It contains delta SNR per sub-carrier, which is the difference from the average SNR.

The compressed beamforming feedback scheme differs from the full CSI feedback scheme in that each field of the compressed beamforming is a value for each spatial stream, not the change from a transmit antenna to a receive antenna. The V-matrix is obtained by applying Singular Value Decomposition (SVD) to the full CSI matrix $H$, and it is compressed to several angles by using the Givens rotation method to reduce the amount of required bits to express real values [26].

The delta SNR field of the MU-Exclusive beamforming report is computed by using Eq. (5):

$$\Delta SNR_{k,i} = 10 \ log_{10}\left(\frac{\|H_k V_{k,i}\|^2}{N}\right) - \overline{SNR_i} \qquad (5)$$

where $V_{k,i}$ is the $i$th column of the V-matrix at sub-carrier $k$, $N$ is the noise-plus-interference power, and $\overline{SNR_i}$ is the average SNR of the $i$th spatial stream. The specification document limits the range of the calculated SNR from −8 dB to 7 dB. This field represents the spatial characteristics for each sub-carrier caused by the environment. SNR per sub-carrier can thus be utilized to increase throughput (e.g., use only the sub-carriers with high SNR). The IEEE 802.11ac specification does not detail how this information is exploited in the decision of the beamformer and thus its implementation is chip vendor dependent.

### 2.2.3. User grouping

In MU-MIMO, user grouping involves an AP selecting a proper set of users for transmissions. When a client is associated to an AP, the AP sends a group ID management action frame to notify the client of the group membership along with a symbolic, relative geographical position within the group. If the AP detects the

change in the client's channel state, it can change the information for the client and send this action frame again with the revised information. Based on the group information and channel state, the AP selects clients to perform beamforming and specifies the group information in the PHY header of each data frame. The IEEE 802.11ac specification does not define the user grouping algorithm and leaves it to each vendor's implementation.

### 2.3. Channel capacity

If channel capacity is known in advance, an AP can adapt to channel changes efficiently for user grouping and determining the number of streams [27]. Shannon–Hartley theorem is well known for estimating the theoretical maximum capacity based on the signal and noise power levels, and many studies have utilized the theorem as a basis for channel capacity estimation [4,7,28,29]. In MIMO, channel capacity is calculated as Eq. (6) [30]. Here, the determinant represents the MIMO channel gain:

$$C \ (bps/Hz) = log_2 \ det\left(I_r + \frac{SNR}{N_t}HH^H\right). \qquad (6)$$

MIMO requires the inverse matrix of $H$ (Eq. (2)) for precoding. Since not all $H$ matrices are invertible, a pseudo-inverse matrix is used to retrieve a precoding matrix for beamforming. SVD is a widely used method to find the pseudo-inverse matrix. The channel matrix $H$ is decomposed into three matrices with SVD (Eq. (7)). $U$ and $V$ are unitary matrices, and $S$ is a diagonal matrix that consists of singular values. The $V$ matrix is delivered to the beamformer and used for precoding. The Hermitian transpose of the $U$ matrix is multiplied by the received signal in the beamformee. By the property of the unitary matrix, $U^H U$ and $V^H V$ become identical matrices. As a result, the diagonal matrix remains, so that the input streams are delivered to each beamformee without interfering with each other in the ideal beamforming (Eq. (8)).

$$H = USV^H, \qquad (7)$$

$$y \ = \ U^H(USV^H)Vx \ = \ Sx. \qquad (8)$$

There are two important terms, a rank indicator and a condition number, which describe the channel capacity. The rank indicator is the number of non-zero singular values of a matrix, the $S$ matrix in Eq. (7). It indicates the number of independent communication channels in a wireless system. For example, the rank indicator one in $2 \times 2$ MIMO channel means merely one stream is available in this channel. The other term, a condition number, is the ratio of the largest singular value to the smallest singular value. It represents the quality of spatial multiplexing. The maximum performance is achieved when the condition number is one.

## 3. Experiment settings

We describe how we set up our MU-MIMO and SU-MIMO testbed with commercial 802.11ac devices.

**Table 1**
802.11ac wave2 APs in our testbed where $N \times N$: $M$ in the MIMO column is $N$ transmit and receive chains and $M$ spatial streams supported by each device.

| Model ID | AP vendor | Wi-Fi chipset | Bandwidth | Modulation | Support. rate (5 GHz only) | MIMO | User/group |
|---|---|---|---|---|---|---|---|
| Model A | Vendor1 | Chipset1-A | 80 MHz | 256 QAM | 2.53 (1.73) Gb/s | $4 \times 4$:4 | 3 |
| Model B | Vendor2 | Chipset2 | 80 MHz | 256 QAM | 2.33 (1.73) Gb/s | $4 \times 4$:4 | 3 |
| Model C | Vendor2 | Chipset3 | 80 MHz | 1024 QAM | 3.16 (2.16) Gb/s | $4 \times 4$:4 | 3 |
| Model D | Vendor1 | Chipset4 | 160 MHz | 1024 QAM | 3.20 (2.6) Gb/s | $4 \times 4$:3 | 3 |
| Model E | Vendor3 | Chipset1-B | 80 MHz | 256 QAM | 2.53 (1.73) Gb/s | $4 \times 4$:4 | 3 |

**Table 2**
Smartphones supporting 802.11ac wave2.

| Device name | Wi-Fi chipset | Support. rate | MIMO |
|---|---|---|---|
| Mobile1 | Chipset1-C | 866 Mb/s | $2 \times 2$:2 |
| Mobile2 | Chipset1-C | 866 Mb/s | $2 \times 2$:2 |

### 3.1. Access points & clients

Tables 1 and 2 show the 802.11ac wave2 capable devices and their specification that are used in our experiments. Most major AP manufacturers used Wi-Fi chipsets manufactured by four vendors that support 802.11ac wave2 specification. Since the throughput depends on the performance of the Wi-Fi chipset, we focused on diversifying the chipset manufacturers than AP makers in selecting the APs. All APs are equipped with four antennas, which support up to three clients per group for MU-MIMO transmissions. Unlike APs, there are not many chipset types that support MU-MIMO for clients. By early 2017, only one chipset (Chipset1-C) supported the 802.11ac wave2 specification. All clients we used are equipped with two antennas. We used these devices as single and multiple antenna clients by changing the Wi-Fi configuration on the device.

The data rates shown in Tables 1 and 2 are theoretical physical layer capacities that can be reached by using dual-band (2.4 GHz and 5 GHz) networks, in which AP manufacturers advertise. However, these APs cannot achieve the advertised throughput due to the following hardware limitations: (i) Model C and Model D must use 160 MHz bandwidth and 1024-QAM modulation for the maximum speed. However, since there is no client that supports such features, the actual capacity of these APs is 1.7 Gb/s at 5 GHz band. (ii) Except for the APs using Chipset1-A and Chipset1-B, no APs support the grouping of $2 \times 2$:2 clients. (iii) The wired network port (WAN/LAN) supports up to 1 Gb/s, which causes the speed of the wireless downlink to be limited to 1 Gb/s. Vendor2 APs support the link aggregation (IEEE 802.3ad) for providing 2 Gb/s of WAN/LAN link speed, but Vendor2 by default did not activate the MU-MIMO option due to performance issues. Therefore, the actual maximum capacity of these APs in 5 GHz band is 1 Gb/s. Since the chipsets not manufactured by Chipset1 vendor do not perform MU-MIMO transmission with the $2 \times 2$:2 clients, we used Model E for our experiments.

### 3.2. Network configuration & environment

We constructed a network testbed as shown in Fig. 3. One wired client (a laptop) and multiple wireless clients (smartphones) are connected to a single AP, and the wired client generates network traffic towards all wireless clients. The AP then transmits the received Ethernet traffic to smartphones through the wireless interface. We used iperf to generate traffic and measure throughput. We set the total sending rate not to exceed 700 Mb/s and set the maximum sending rate for each client not to exceed 350 Mb/s.

In order to measure accurate Wi-Fi performance, experiments should be done in an environment where interference from other infrastructure is removed, such as in a shield room. However,

measuring the achievable performance of the device itself is not our goal, and a shield room might not be a suitable environment for measuring user throughput. We thus conducted all our experiments in a real campus environment. We mainly measured Wi-Fi throughput in indoor environments as MIMO has been developed for the indoor environment that is typical of WLAN deployments [21,31]. We used multiple classrooms, with each room size of 8 m × 10 m (Fig. 4). We measured MU-MIMO and SU-MIMO throughput with varying the position of clients, the number of clients, the number of client antennas, and obstacles. To minimize interference from other Wi-Fi devices in different infrastructure, all experiments were conducted at night-time, where there was no student (hence only the beacon frames from other networks exist). We confirmed that other Wi-Fi network operated by the university switched to channels in the UNII-1 band due to auto channel selection when we intensively used the UNII-3 band. Few APs remaining on the same frequency band did not affect our results as they only sent beacons. We also experimented in an outdoor open space (Fig. 5), which did not have multipath signals.

### 3.3. Log analysis

Due to the lack of detailed documentation and access to the driver and firmware source code of the Wi-Fi chipset on commercial APs, we could not determine the exact algorithm of the AP operations. To analyze the experiment results without the exact knowledge of the AP's operation, we collected three types of external logs from the client devices and wireless packets.

#### 3.3.1. Throughput measurement
For throughput measurement, we logged the data rate on the clients and extracted the output files after the experiments. We added the timestamp to iperf's throughput logs and used them for post processed time synchronization.

#### 3.3.2. Sniffing wireless packets
By sniffing the wireless packets exchanged between an AP and clients, we can infer the operation of the AP. Tamosoft's Commview for Wi-Fi [32] is a software solution that captures Wi-Fi packets using an embedded Wi-Fi network interface card. This software not only captures IEEE 802.11 packets but also measures PHY information such as RSSI and transmission rate. Although MU-MIMO data packets cannot be captured due to the hardware limitation, we can infer most of the operations with control and management frames. By analyzing these frames, we extract following information.

(i) Precise transmission rate and error rate: 802.11ac uses Ack and BlockAck mechanisms that contain the sequence number of data packets. We can measure the exact transmission rate at the frame level by counting the changes of the sequence number.

(ii) MIMO transmission mode: as discussed in Section 2, an AP groups multiple clients and measures wireless channel to each client for MU-MIMO transmissions. By extracting the sequence of channel sounding frames, we discover which transmission mode (SU-MIMO or MU-MIMO) the AP used.
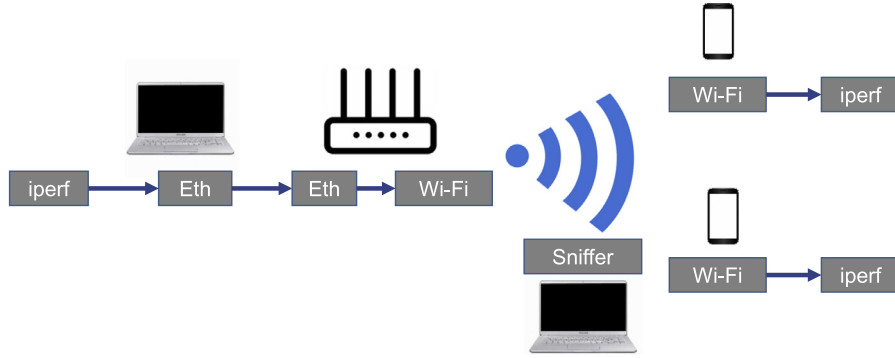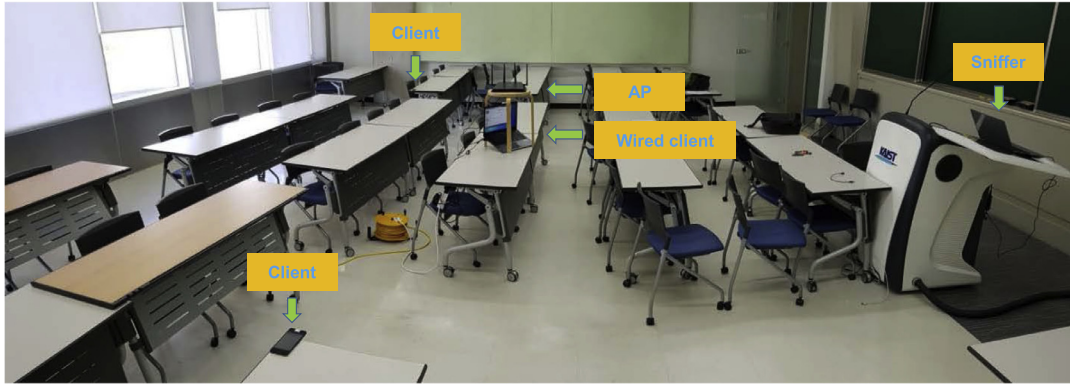
**Fig. 3.** Network configuration.



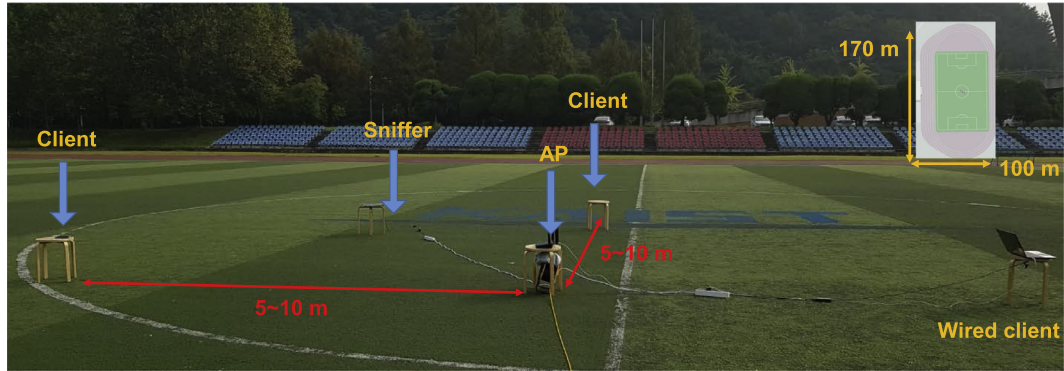**Fig. 4.** Indoor classroom environment.



**Fig. 5.** Outdoor open environment.

(iii) User grouping operation: AP sets the group membership and user position for each client using the group ID management frame, and uses this group information for subsequent MU-MIMO transmissions. By extracting these frames and channel sounding sequence, we can infer the user grouping policy of the AP.

(iv) Correlation between the clients: clients feed wireless channel information, including SNR and V-matrix per sub-carrier to the AP. The similarity of the V-matrix represents the correlation between the clients [4].

(v) Channel sounding overhead: the overhead can be calculated through measuring the size and time occupancy of the channel sounding packets.

### 3.3.3. Wi-Fi driver log on client devices

Once user grouping is performed, the AP informs each client the number of spatial streams to send and the MCS index to modulate the data. Knowledge of these parameters is significant for an-

alyzing WLAN performance. We obtain this information from the client device by printing the Wi-Fi driver log to the Android kernel message (kmsg). From the kernel driver source code, we discovered that the Wi-Fi firmware on Mobile1 indicates its PHY information to the Wi-Fi driver every three seconds. Although the reporting interval is large, this data is still useful for result analysis.

## 4. Experimental results & analysis

We present the experimental results and analysis. We measured downlink user application throughput in various network environments to evaluate how much throughput is increased by applying MU-MIMO to commercial devices compared to SU-MIMO. We then analyze the result and highlight the limitations of current commercial WLAN deployments.
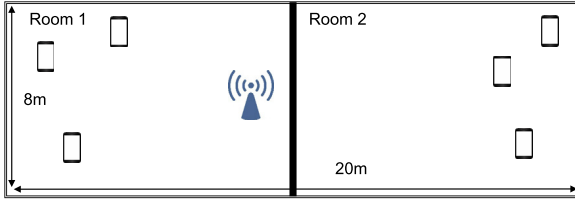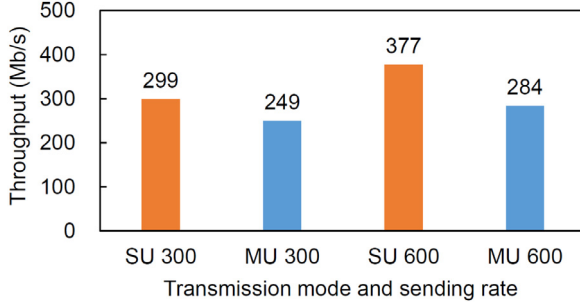
Fig. 6. Experimental setup for the SU vs. MU.



Fig. 7. Throughput comparison between SU- and MU-MIMO.

## 4.1. Throughput

We placed an AP and six clients as shown in Fig. 6 and generated UDP traffic with the rate of 50 Mb/s and 100 Mb/s per client. We measured aggregate user throughput of SU-MIMO and MU-MIMO, respectively. Through sniffer log analysis, we observed that the AP sent data packets to two or three clients simultaneously during MU-MIMO transmissions. However, overall MU-MIMO throughput was surprisingly lower than SU-MIMO (Fig. 7). We repeated the same experiments with changing the position of the clients and the size of the experimental space (1–3 room), but MU-MIMO throughput was consistently lower than SU-MIMO. In addition, when we analyzed the real-time throughput of each client (Fig. 8), MU-MIMO transmission was very unstable with high fluctuations compared to SU-MIMO; four clients even experienced zero throughput for several seconds during MU-MIMO transmission (Fig. 8(b)).

This poor MU-MIMO performance result is contrary to the results from previous studies [3,4]. The biggest difference of our experiments from previous studies is that we used $2 \times 2$ antenna-enabled clients, as opposed to $1 \times 1$. With $2 \times 2$ clients, the number of data streams simultaneously transmitted by the AP increases by one or two compared to when only $1 \times 1$ clients exist. It is also twice as large as the SU-MIMO transmission mode. Therefore, MU-MIMO throughput gain should be positive in our environment as well.

We thus conducted additional experiments to further analyze this phenomenon. We placed an AP and clients as shown in Fig. 9 and generated UDP traffic with the rate of 200 Mb/s per client. In this experiment, we used only three clients to reduce the overhead caused by channel sounding and user grouping. We first placed three $2 \times 2$ clients and changed their antenna setting from $2 \times 2$ to $1 \times 1$ one by one so that the number of spatial streams the AP simultaneously transmits is gradually reduced. That is, the total number of spatial streams is four in the first experiment, and then two or three in the fourth experiment.
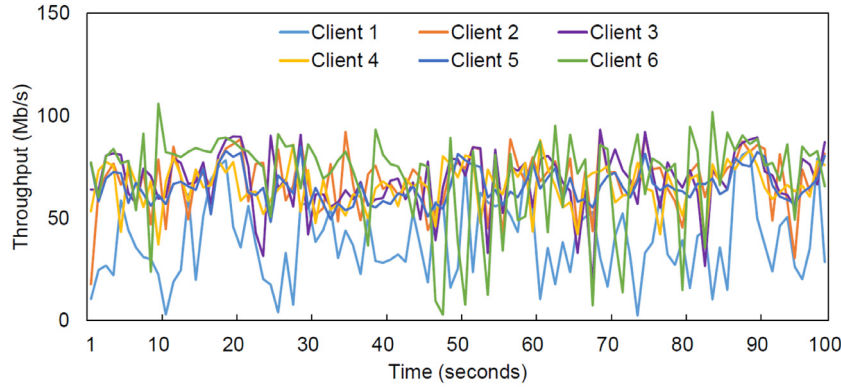


(a) Real time throughput in SU-MIMO.
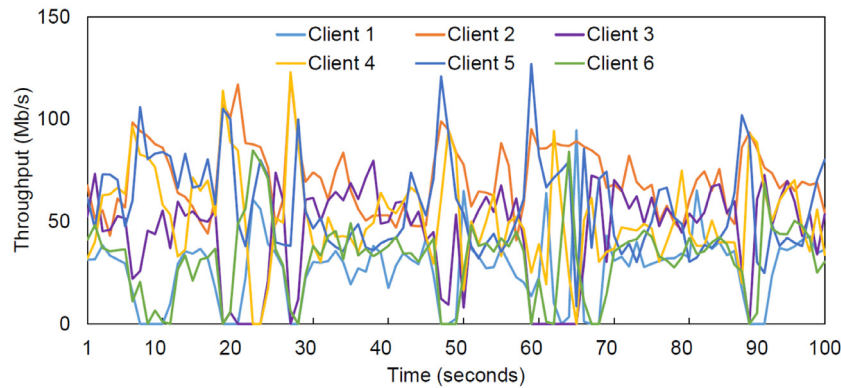


(b) Real time throughput in MU-MIMO.

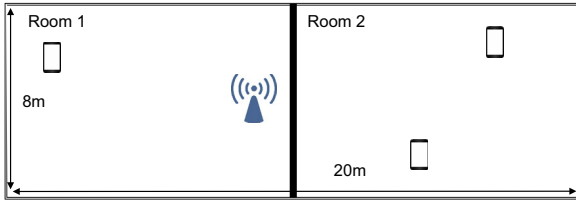Fig. 8. Real time throughput comparison.

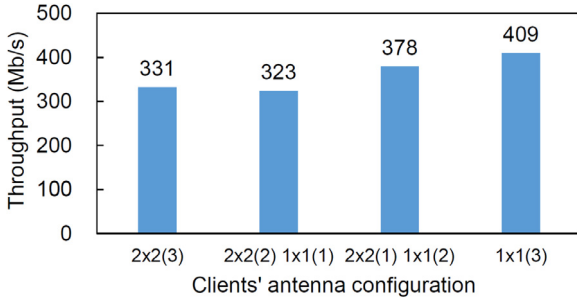Fig. 9. Experiment setup with $1 \times 1$ clients.



Fig. 10. MU-MIMO throughput with different antenna configuration, where $N \times N$ is the transmit and receive chain supported in clients and the numbers in parentheses indicate the number of clients corresponding to the configuration.
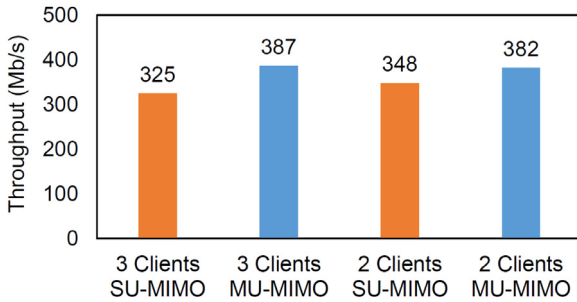


Fig. 11. Throughput comparison between SU- and MU-MIMO with $1 \times 1$ clients.

We observe in Fig. 10 that the aggregate throughput increases as the number of receive antennas decreases with the same number of clients. We then placed two or three $1 \times 1$ clients to the same position in Fig. 9 and measured MU-MIMO and SU-MIMO throughput. As shown in Fig. 11, MU-MIMO throughput was better than SU-MIMO, with the gain rate of about 1.15. SU-MIMO throughput shown in Fig. 11 was the maximum throughput that can be achieved in this experiment setup as the theoretical rate of a single stream is 433 Mb/s, and the maximum throughput of the Mobile1 device was about 350 Mb/s with the single antenna configuration. MU-MIMO throughput exceeds the maximum achievable throughput of SU-MIMO, although the gain was small. Through additional experiments, we observed that MU-MIMO throughput increased up to 478 Mb/s in our experiment setup, and the MU-MIMO gain was about 1.38.

Although our experiment results could not represent the Wi-Fi performance in all environments, it is clear that MU-MIMO performance is significantly influenced by the number of spatial streams of the clients and MU-MIMO has positive throughput gain over SU-MIMO only in limited conditions. The result of Figs. 7 and 10 in particular contradict common expectations. This is a very critical problem as most current smartphones and laptops supporting MU-MIMO in the market are equipped with two antennas.

In the following sections, we analyze the factors that affect MU-MIMO performance and identify what are the major issues causing below expectation performance.

**Table 3**
Channel occupation ratio (%) of the channel sounding frames with different number of clients.

| Number of clients | 2 | 3 | 4 | 10 |
|---|---|---|---|---|
| Channel occupation ratio (%) | 1.6 | 5.4 | 9.2 | 16.3 |

**Table 4**
Channel occupation ratio (%) of the channel sounding frames with different sending rates.

| Sending rate (Mb/s) | 100 | 200 | 500 | 700 |
|---|---|---|---|---|
| Channel occupation ratio(%) | 1.1 | 16.3 | 13.9 | 11 |

### 4.2. Channel sounding overhead

For MIMO transmissions, an AP has to measure the wireless channel state for each client. During channel measurement, data communication cannot occur concurrently, and thus large measurement overhead directly results in poor throughput. In addition, as the control and management frames are transmitted using low and reliable rates, frequent measurements would consume precious wireless channel.

In MU-MIMO, managing channel sounding overhead is a major challenge; the size of a feedback frame is up to $1.7 \times$ larger than that of SU-MIMO, and an AP must collect the feedback from multiple clients at once. Moreover, the AP measures the channel state more frequently due to its high sensitivity to the channel state variations. We thus empirically measured how large the channel sounding overhead is.

We extracted the channel sounding frames from the sniffer log described in Section 3.3.2 and summed up the difference of the captured time between NDPA and the compressed beamforming feedback frame of the last beamformee (Fig. 12). We could not measure the duration of NDPA frame due to the absence of a reference point to calculate the time difference, but it is negligible as its frame size is small compared to the feedback frames.

We measured the channel sounding overhead as the number of clients and the traffic rate increased. Table 3 shows the maximum channel sounding overhead when two, three, four, and ten clients are present. Table 4 shows the sounding overhead of different sending rates in the ten clients case. As the number of clients increases, the channel sounding overhead (air time occupation) also increases. It increased up to 16.3% with ten client devices. However, channel sounding overhead was not proportional to the sending rate. When the overall sending rate was lower than 100 Mb/s, the overhead was very low (1.1%) as AP performed SU-MIMO. When the sending rate was 200 Mb/s or higher, the AP performed MU-MIMO transmissions and the overhead exceeded 10%. Interestingly, the maximum occupation occurred when the overall sending rate was 200 Mb/s. This could depend on the algorithm used by the chipset manufacturer, and we infer that the AP measured the channel more frequently because there exists more channel idle time.

One might argue that the channel sounding overhead was not a major culprit of low throughput in our experiments. Note however that the overhead was measured at a *static* environment. In typical indoor office/campus environments where people (and mobile devices) move, an AP has to perform channel sounding more often to adapt to the change of the channel condition. In our experiment, one channel sounding took about 1.7 ms with 117 Mb/s transmission rate and 2.4 ms with 29.5 Mb/s rate, which is similar to the duration of a burst transmission of an A-MPDU. It suggests that the overhead could be increased up to about 50% of the wireless channel occupation at the worst case if the AP does not control channel
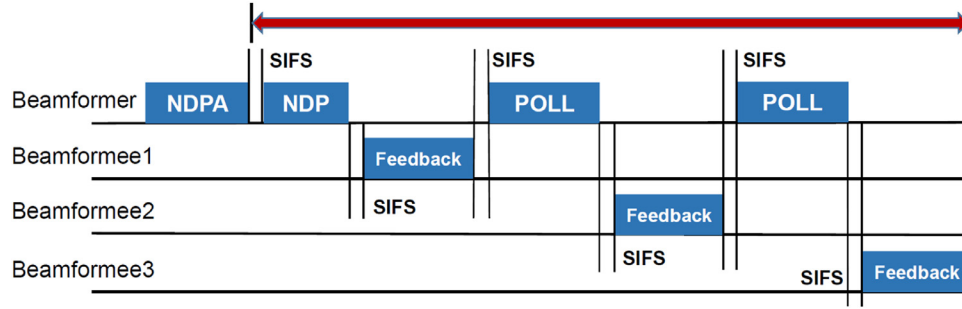
**Fig. 12.** Measurement of the air time occupation of channel sounding frames.

**Table 5**
Predefined group information in Model E AP, where P$x$ is the user position assigned by the AP and C$x$ indicates clients.

| Group ID | P0 | P1 | P2 | P3 |
|---|---|---|---|---|
| Group 1 | C1, C5 | C2, C6 | C3 | C4 |
| Group 2 | C1, C3 | C2, C4 | C5 | C6 |
| Group 3 | C1, C2 | C3, C4 | C5, C6 | |
| Group 4 | C1, C3, C5 | C2, C4, C6 | | |
| Group 5 | C1, C2, C5, C6 | C3, C4 | | |
| Group 6 | C1, C2, C3, C4 | C5, C6 | | |
| Group 7 | C1, C3, C5 | C2, C4, C6 | | |
| Group 8 | C1, C2, C5, C6 | C3, C4 | | |
| Group 9 | C1, C2, C3, C4, C5, C6 | | | |
| Group 10 | C1, C2 ,C3, C4 | C5, C6 | | |
| Group 11 | C1, C3, C5 | C2, C4, C6 | | |
| Group 12 | C1, C2, C5, C6 | C3, C4 | | |
| Group 13 | C1, C2, C3, C4, C5, C6 | | | |
| Group 14 | C1, C2, C3, C4 | C5, C6 | | |
| Group 15 | C1, C2, C3, C4, C5, C6 | | | |



(a) Original placement of clients.



(b) After changing clients' placement.

**Fig. 13.** Experiment setup for user grouping.

sounding overhead. This long channel sounding duration not only limits throughput but also prevents correct user grouping.
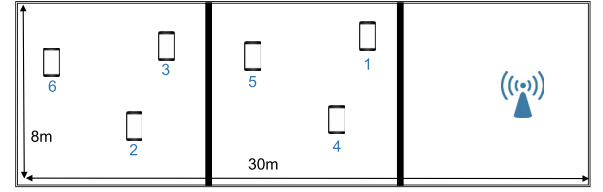
### 4.3. User grouping

Most current commercial APs can simultaneously transmit data to up to three clients. When a network has more than three clients, an AP must select and group the clients. Grouping clients is one of the most important features that the manufacturers must consider to maximize MU–MIMO performance. When the AP does not find the best combination of clients, interference between spatial streams increases and network performance suffers. The 802.11ac standard does not define the grouping method, and thus each chipset manufacturer develops its own algorithm.
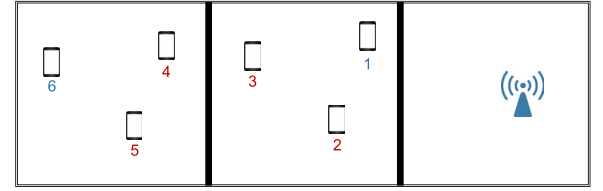
We analyze the user grouping behavior by extracting the group information from the control and management frames exchanged between the AP and the clients. That is, we capture the group ID management action frame to infer the AP's group management policy (see Section 2.2.3). We also investigate the channel sounding sequence and the BlockAck frames to identify which group the AP sends data to.

We found Vendor3's Model E AP uses a predefined group (Table 5). The AP activates 15 groups when there is at least one MU-MIMO client and sets the user position with a predefined order. For example, if one client (C1) associates to the AP, it is assigned to position 0 of all 15 groups. If another client (C2) associates to the AP, it is assigned to position 0 or 1 of all 15 groups. That is, the AP assigns all clients to all groups, as shown in Table. 5. The AP retransmitted the group ID management frame to a client when it moved, but it did not change the membership status and the position array.

This allocation scheme covers all the combinations according to the number of stream. Specifically, as there are one to four available streams and six users in this example, the total combinations can be represented as $^6C_1 + {}^6C_2 + {}^6C_3 + {}^6C_4$. The predefined table covers all the combinations. For example, when the AP decides to send three streams to C1, C3, and C4, the AP uses Group 1 and matches P0 for C1, P2 for C3, and P3 for C4 (does not use P1). That is, this AP can select a proper group among multiple candidates. Note that this operation is Chipset1-B specific. Chipset3 and Chipset4 use one group while Chipset2 uses nine groups by default.

When we generated UDP traffic, the AP transmitted data after selecting two or three clients in each group. To check how the AP selects the clients, we first placed three clients in each of the two classrooms (Fig. 13(a)) and counted the number of two client pairs at each MU-MIMO transmission. We then changed the location of four clients (Fig. 13(b)) and measured the number of two client pairs. Table 6 shows the result of this experiment. In Fig. 13(a) environment, the number of transmissions to C1–C2, C2–C3, and C1–C3 pairs was the greatest. On the other hand, in Fig. 13(b) environment, the number of transmissions to C1–C4, C4–C6, and C1–C6 pairs was the greatest. This indicates that the AP changed the client combination based on channel conditions. However, this behavior resulted in poor fairness. Channel capacity was concentrated on top three devices. As a result, the aggregate throughput of other clients was only about 10 Mb/s despite the 300 Mb/s traffic rate. In addition, there are many non-dominant pairs despite being in a static environment. For example, in Fig. 13(a) environment, C2 was grouped 1244 times with clients other than C1 and C3, which is similar to the number of dominant (C1–C2 or C2–C3) pairs, while C6 was evenly grouped with all other clients.

From our experiment result, we infer that the criterion the commercial APs use for user grouping is inaccurate. The AP must measure the channel state for all clients at the same time to dis-

**Table 6**
The number of two client pairs during MU-MIMO transmissions.

| With the original placement. | | | | | | |
|---|---|---|---|---|---|---|
| | C1 | C2 | C3 | C4 | C5 | C6 |
| C1 | – | 1274 | 1983 | 1235 | 306 | 303 |
| C2 | – | – | 1519 | 441 | 327 | 476 |
| C3 | – | – | – | 1215 | 381 | 413 |
| C4 | – | – | – | – | 203 | 299 |
| C5 | – | – | – | – | – | 248 |
| C6 | – | – | – | – | – | – |
| With the changed placement. | | | | | | |
| | C1 | C2 | C3 | C4 | C5 | C6 |
| C1 | – | 472 | 118 | 1898 | 248 | 3299 |
| C2 | – | – | 27 | 297 | 114 | 662 |
| C3 | – | – | – | 73 | 46 | 104 |
| C4 | – | – | – | – | 161 | 2121 |
| C5 | – | – | – | – | – | 356 |
| C6 | – | – | – | – | – | – |

cover the optimal combination of clients. However in practice, receiving feedback from all clients at every channel measurement not only causes large sounding overhead but also degrades performance due to the delayed channel response. We found that all commercial APs in our experiments involved up to three clients for a single channel measurement, and therefore many measurement cycles are required to measure all clients. For example, 120 measurements are needed for 10 clients, which takes about 288 ms. It is clearly not feasible to measure the entire channel within a channel coherence time using this channel sounding method. If the AP measures all clients for every channel sounding, this time could be reduced. However, as it is possible for APs to transmit concurrently to up to only three clients, collecting feedback from every client for every data transfer could be excessive. Besides, the overhead increases linearly with the number of clients, which results in overdue channel state information for the duration of collecting the feedback. Therefore, the AP needs to rely on partial channel information when grouping users.

### 4.4. Environmental impact

Wireless communication is highly influenced by the environment due to the physical characteristics of the RF signals [33]. In particular, when a radio signal propagates indoors, replicated signals are generated on many multiple paths. WLAN has used this property to increase SNR at the receiver since 802.11n and also in 802.11ac. MU-MIMO exploits the difference of the propagation paths between an AP and the clients, as the path to each client exhibits different channel states based on position [21]. Therefore, the surroundings of where the AP and the clients are located, and their placement have a significant impact on performance. Theoretically, the performance is better when the wireless channel to each client is sufficiently different from other clients. We thus investigate the impact of the environment on Wi-Fi throughput and how the commercial AP copes with the environmental effect.

We measured MU-MIMO and SU-MIMO throughput in indoor (Fig. 14) and outdoor (Fig. 5) environments. In this experiment, we used only two clients to eliminate the effects of channel sounding and user grouping. The first four experiments (Env.1 ∼ 4) were conducted in a single room with changing the location of the obstacles. We located the obstacles considering SNR and presence of a line-of-sight (LoS) path. For example, Env.1 has high SNR and low channel difference. Env.2 is an environment where the SNR of a client is lower than another, and the channel difference is higher than that of Env.1. Env.4 has the highest SNR and the lowest channel difference among the four environments. Env.5 and 6

were conducted in two classrooms to create a large channel difference. We placed one client device in each classroom. In addition, we measured Wi-Fi throughput in an outdoor open space without any objects. In this outdoor open environment, only the angle-of-arrival (AoA) and angle-of-departure (AoD) diversity caused by the AP's antenna arrangement and the clients' location would cause channel quality differences to different clients. Furthermore, the average SNR in open space was about 12 dB lower than in indoor environment due to the lack of the multipath signals. Average SNR reported by the clients through channel sounding feedback was about 43 dB in Env.1, 37 dB in Env.5, and 31 dB in open space.

The aggregate throughput of MU-MIMO and SU-MIMO measured at these different environments are shown in Fig. 15. The total UDP sending rate in these experiments was 700 Mb/s. Similar to previous experimental results, MU-MIMO was always outperformed by SU-MIMO. SU-MIMO throughput was proportional to SNR throughout our experiments and the throughput was high when an LoS path exists. When there are more obstacles blocking the LoS path, the aggregate throughput suffered (Env.1−3). In Env.1 and 4 where SNR was the highest, the aggregate throughput reached the maximum achievable throughput.

On the other hand, with MU-MIMO, only having high SNR did not result in higher throughput; the aggregate throughput when blocking only one LoS path (Env.2) was higher than when there were all LoS paths (Env.1) or only non-LoS paths (Env.3) for two clients. That is, MU-MIMO throughput was influenced more by the channel state difference between the clients. However, as we see from Env5's throughput, only having large channel difference did not result in higher throughput. MU-MIMO throughput was the best in Env.6, which had large channel difference *and* high SNR due to the guard walls. In Env.4, the AP actually transmitted data in SU-MIMO mode. That is, *both large channel difference and high SNR are required for high MU-MIMO performance*. This result shows that in MU-MIMO, we could not expect high throughput when we locate our mobile device close to the AP. In an open space, SU-MIMO and MU-MIMO throughput were the lowest as SNR was low and the channel difference between the clients was small.

We also observed that MU-MIMO throughput improved up to 160 Mb/s by simply adding three guards (Env.5 and 6). This difference is much higher than with SU-MIMO. This result indicates that MU-MIMO has a larger variation than SU-MIMO due to the spatial changes and users. We therefore cannot expect stable throughput in environments where channel state is dynamically changed by the movement of users and mobile devices.

In typical WLAN deployments, APs are installed on a ceiling or top of a wall to provide a better signal to clients. The AP thus has a LoS path with most clients within the area, which could benefit SU-MIMO, but not MU-MIMO. Although the MU-MIMO performance does not depend solely on the LoS path, this consideration is necessary when deploying MU-MIMO. In our experiment for example, simply blocking the LoS path increased MU-MIMO throughput by 80 Mb/s.

### 4.5. Transmission mode selection

MU-MIMO capable APs must determine the optimal transmission mode (SU-MIMO or MU-MIMO) based on channel conditions. However, we observed in Section 4.4 that Model E AP maintained MU-MIMO transmissions in most environments even when it would clearly be better to use SU-MIMO. In our experiments, Model E AP selected the SU-MIMO mode only in Env.4, and used MU-MIMO in all other experiments even when the throughput was up to 180 Mb/s lower than SU-MIMO. For further analysis, we investigate the transmission mode selection behavior of Model E AP and provide a theoretical analysis of the challenges in mode selection for the APs.
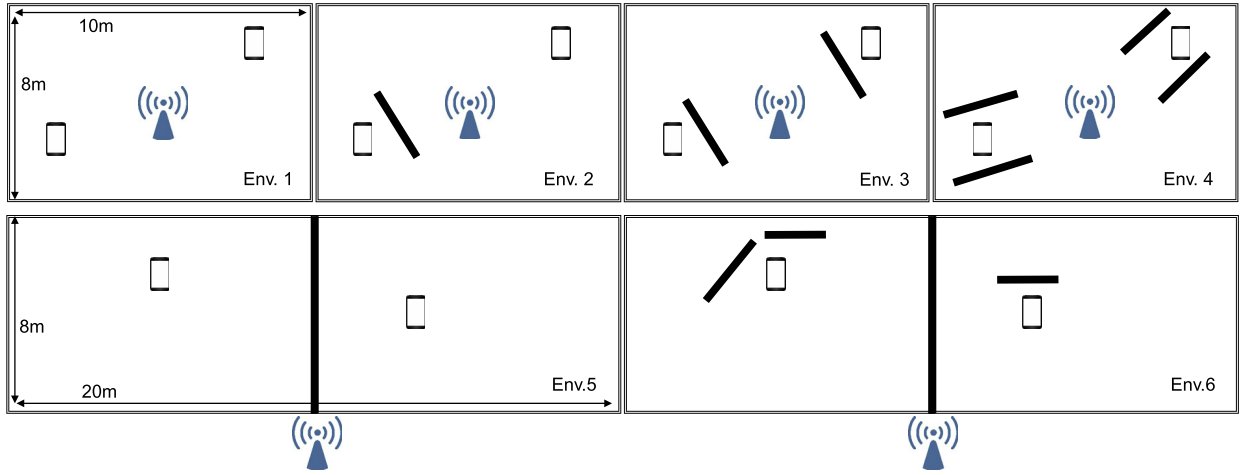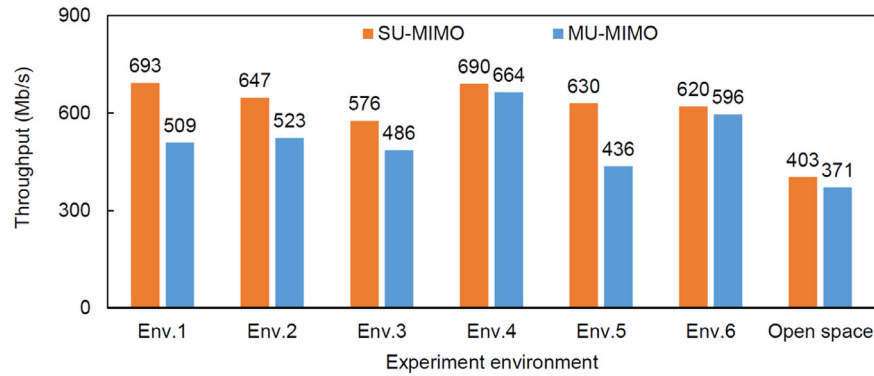
**Fig. 14.** Experiment setup for a multipath effect.



**Fig. 15.** Aggregate throughput at each experimental setting in indoor and outdoor environments.
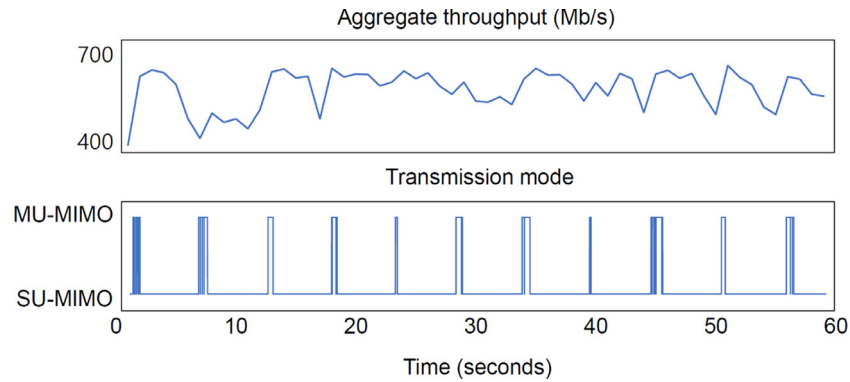


**Fig. 16.** Aggregate throughput and transmission mode in the Env.4 experiment.

Fig. 16 shows the aggregate throughput and transmission mode in Env.4. The AP changed the transmission mode several times at first, and then kept the SU-MIMO mode. The AP switched to MU-MIMO but soon switched back to SU-MIMO. The AP repeated this behavior every five seconds. We investigated the number of transmitted frames and the frame error rate by analyzing the Block-Ack frames (Fig. 17). Before the AP switched to MU-MIMO, the data rate reached the maximum achievable throughput and SNR was constant over time. The AP might have switched to MU-MIMO for efficient channel utilization as the channel state was stable, or simply for a trial. However, after the transition to MU-MIMO, the frame error rate sharply increased. The AP tried several transmissions with short-period channel sounding, and finally switched back to SU-MIMO within 0.3 s. The error rate decreased

immediately after switching to SU-MIMO. Due to the repeated transition overhead, the throughput was 26 Mb/s lower than transmitting only in SU-MIMO mode.

At other experiments, the error rate did not increase as sharply as in Env.4, and the AP kept the MU-MIMO transmission mode despite the low throughput. From these behaviors, we infer that Model E AP determines the transmission mode using a trial-and-error method based on packet error rate, without estimating the capacity of the channel. That is, as the AP does not know which transmission mode is better, it could make a wrong decision that results in lower throughput.

Accurate channel capacity estimation is required for transmission mode selection. It is known that the channel capacity can be estimated using the Shannon–Hartley theorem, singular values,
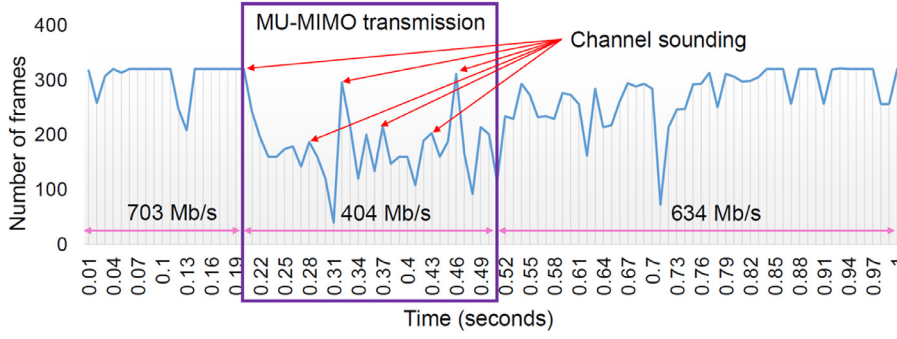
**Fig. 17.** Frame rate during the transmission mode changes in the Env.4 experiment.
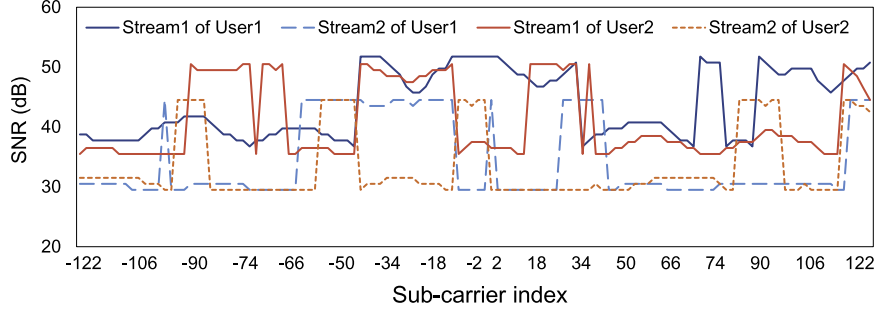


**Fig. 18.** SNR distribution in the Env.5 experiment.

and interference as described in Section 2.3. We first check the Shannon–Hartley theorem based on our experiment data. In Eq. 6, the determinant of $HH^H$ is expressed as the determinant of the square of the singular value matrix Eq. 9). Meanwhile, the singular values are delivered to the AP through the beamforming feedback in the form of SNR (Eqs. 5 and (10). Therefore, the capacity can be estimated by using SNR feedback.

$$HH^H = USV^HVSU^H = US^2U^H$$
$$\therefore \ det(HH^H) = det(US^2U^H) = det(S^2), \tag{9}$$

$$\|HV\|^2 = \|USV^HV\|^2 = \|US\|^2 = \|S\|^2. \tag{10}$$

However, the AP cannot directly estimate the channel capacity using SNR from the clients, as they are singular values of each client's channel matrix, not the entire channel matrix $H$. Furthermore, as the Shannon–Hartley capacity predicts the theoretical maximum value simply based on the power level, it is inadequate to accurately predict the capacity.

The rank indicator and the condition number (see Section 2.3) are not suitable for capacity estimation for the same reason. For example in Fig. 18, the SNR of the four streams is high at the 120th sub-carrier, and the ratio of the largest SNR to the smallest SNR is close to one, but this does not mean that it is in good channel state for MU-MIMO; the SNRs are derived by the SVD of the individual channel matrix $H$ in each client, not the entire channel matrix $H$.

Knowing the level of interference between the clients would be more helpful in estimating the channel capacity in MU-MIMO. However, there is no way for clients to measure the interference by other streams toward other clients, and thus clients' feedback does not contain interference information. Therefore, the APs must estimate interference from the correlation between individual clients' feedback. However, simply using signal-to-interference-plus-noise ratio (SINR) instead of SNR in the Shannon–Hartley theorem does not result in accurate capacity estimation as it also uses only the power level to estimate the capacity.

Therefore, due to its fundamental limitations, existing channel capacity estimation methods are not suitable for transmission mode selection. Commercial APs thus have no choice but to use trial-and-error based on throughput or PER, which results in poor mode selection and switching. Moreover, a finer-grained control of mode switching threshold is needed to limit wrong transmission mode selections.

### 4.6. Adjusting the number of streams

The number of permitted spatial streams (NSS) in a wireless channel is based on the channel state. An AP selects the clients and notify them of how many spatial streams the AP sends by using the PHY header of each data frame. The APs might adjust the number of streams based on the channel state, but as we observed in Section 4.1, the throughput with $1 \times 1$ clients was higher than with $2 \times 2$ clients (Fig. 10). It suggests that the AP did not adjust NSS; otherwise, the aggregate throughput for $2 \times 2$ clients would be similar to the throughput for $1 \times 1$ clients by sending one stream to each client.

In order to evaluate whether the APs adjust the number of streams, we modified the Wi-Fi driver of Mobile1 smartphone to extract the number of received streams. We found that Model E AP always transmitted four streams to $2 \times 2$ clients during MU-MIMO transmission even when there were severe interference. It means Chipset1-B, the Wi-Fi chipset of Model E AP, did not adjust the number of streams depending on the channel state. However, this fact does not mean that the AP is incapable of adjusting the NSS, as the AP can send data to a group containing three $2 \times 2$ clients. It might rather be due to improper threshold value for adjusting the number of streams.

We also discovered that except for Chipset1-A and Chipset1-B, none of the chipsets we experimented with activated MU-MIMO when there were only two antennas-enabled clients, even though some are capable of transmitting four spatial streams. If we placed only $1 \times 1$ clients or mixed $1 \times 1$ and $2 \times 2$ clients in the network, the APs performed MU-MIMO transmissions. Therefore, these APs might not have the functionality of adjusting the NSS. This

functionality is quite critical for MU-MIMO performance, and we hope these APs will have it implemented in their next release.

### 4.7. Other factors

In addition to the issues discussed above, we observed two more factors that degrade MU-MIMO user throughput. One is the socket buffer size in application. The application using wireless network needs to allocate proper socket buffer size based on the link status. When we used the default socket buffer size of iperf (244 KB), the throughput for single client device did not reach 300 Mb/s. When changed to 3 MB, the throughput reached 600 Mb/s. A large buffer is required for high throughput in wireless networks with unstable channel quality, but it might be costly for the server to allocate large buffer to all clients.

Another is rate adaptation. From the Wi-Fi driver log in the clients, we observed the AP transmitted data using high modulation schemes with MCS 7 − 9, despite the high error rate. The SNR values the clients reported to the AP was high enough to use high MCS, but interference was also high. Transmitting data using a high modulation scheme in a noisy environment leads to re-transmissions, which results in lower throughput. Commercial chipset and AP vendors should reconsider their rate selection algorithm, as well as estimating noise and interference.

## 5. Discussion

MU-MIMO that simultaneously transmits to multiple wireless devices is a breakthrough technology. It enhances channel efficiency and increases the network capacity. However, our experiment results show that current commercial MU-MIMO devices do not fulfill the theoretical potential, and in fact, often yielded poor throughput compared with SU-MIMO. We discuss practical issues, challenges and recent research effort in achieving very high throughput WLAN.

As shown in Section 4.4, MU-MIMO is very sensitive to the environment, including where the wireless devices are placed. In addition to the dynamic changes of the channel condition, variations due to spatial structures also pose a challenge. As WLANs cannot completely eliminate the effects of physical constraints caused by the spatial characteristics such as path loss and fading, they must accurately measure the characteristics of a given environment and maximize the performance with optimal configuration settings and algorithms.

In order to account for the environmental effect, APs rely on channel sounding measurements. However, as elaborated in Sections 4.2 and 4.3, channel sounding overhead could degrade throughput and become the bottleneck of efficient user grouping. The issue is more severe in MU-MIMO as the required channel sounding overhead increases with the number of clients. There have been efforts in minimizing channel sounding overhead without compromising the feedback quality [8,34]. A recent approach pre-sounds the environment to remove the channel sounding overhead [7], but it is vulnerable to environmental changes. Another approach is to estimate the entire channel by extrapolating a small number of channel measurements [35]. Although it requires the receivers to have the precoding matrix, it significantly reduces the channel sounding overhead. An alternative is to use implicit feedback instead of explicit feedback. Implicit feedback leverages channel reciprocity between the uplink and the downlink, to estimate the downlink quality without receiving a large amount of feedback from clients. With the recent trend of increasing number of available antennas, such as 5G massive MIMO [36] and 802.11ax [37], we believe implicit feedback could be a viable solution moving forward for sounding overhead reduction and user grouping.

In addition to the channel sounding overhead, computational overhead for user grouping also remains as an issue. It is a significant overhead to not only consider a large number of possible combinations of a clients group, but it is also a complex mathematical operation to check the orthogonality among the clients. The required processing time of an exhaustive search for a four antenna AP and 20 clients exceeds 100 ms with an Intel Xeon 16 cores machine [38]. To reduce such computational workload, sub-optimal user grouping algorithms have been adopted a greedy approach [39,40] or graph theory [41] to select user group candidates in low complexity. Recent proposals [4,38] pre-divide the user groups based on the capabilities and functionalities of client devices. They also utilize additional high computational resource. Nonetheless, simply the decompression of V-matrix for 10 clients took 23.4 ms (10 $\mu$s × 10 users × 234 sub-carriers) [4], which exceeds the channel coherence time of a mobile (1.5 m/s) user [38].

Our experiments showed that the AP not being able to accurately estimate channel capacity is the main culprit of poor MU-MIMO performance. To achieve high throughput, the APs must find the optimal settings for the clients group, the transmission mode, and the number of streams; all these selections require an accurate estimation of the achievable data rate based on channel state. Existing SNR-based approaches should not be applied in MU-MIMO due to inaccuracy. Thus, a new channel metric is needed.

Another issue that must be solved is the difficulty in obtaining the state of the entire wireless channel and the correlation between the clients. Note that Orthogonal Frequency-Division Multiple Access (OFDMA) that allocates specific sub-carrier sets to each client, is a new feature in the 802.11ax amendment, and could be an alternative to increasing user throughput by improving channel efficiency.

## 6. Conclusion

We conducted an empirical study on the latest WLAN specification and evaluated the user throughput achieved with commercial MU-MIMO devices in various network environments. We showed that current implementation of MU-MIMO in commercial devices does not increase the user throughput compared with SU-MIMO mainly due to (i) the inability of measuring channel state information within a channel coherence time and (ii) the difficulty of accurately estimating the channel capacity from individual clients' channel feedback. We also found that channel sounding overhead can be increased up to 50% of the total channel occupancy time, and that MU-MIMO performance is greatly influenced by environmental factors. We believe this study helps researchers and practitioners to understand the main issues of current performance of MU-MIMO and together solve research challenges that could be applied to practical deployments.

### References

[1] IEEE Standards Association IEEE 802.11–2016, Part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications, IEEE Standards Association.

[2] W. Sun, O. Lee, Y. Shin, et al., Future of Wi-Fi: Wi-Fi could be much more, IEEE Communications Magazine (2014).

[3] MIERCOM report, IEEE 802.11ac wave 2 access points comparative performance, https://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise-networks/802-11ac-solution/miercom-report-wave-2-aps-cisco-aruba-ruckus.pdf, 2016

[4] S. Sur, I. Pefkianakis, X. Zhang, K.-H. Kim, Practical MU-MIMO user selection on 802.11ac commodity networks, IEEE Communications Magazine (2014).

[5] Y. Zeng, P.H. Pathak, P. Mohapatra, Throughput, Energy Efficiency and Interference Characterisation of 802.11ac, Wiley Online Libraty, 2015.

[6] L. Kriara, E.C. Molero, T.R. Gross, Evaluating 802.11ac features in indoor WLAN: an empirical study of performance and fairness, Proceeding of the ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH), 2016.

[7] N. Anand, J. Lee, S.J. Lee, E.W. Knightly, Mode and user selection for multi-user MIMO WLANs without CSI, in: Proceedings of the IEEE International Conference on Computer Communications (INFOCOM), 2015.

[8] X. Xie, X. Zhang, K. Sundaresan, Adaptive feedback compression for MIMO networks, in: Proceedings of the Annual ACM International Conference on Mobile Computing and Networking (MobiCom), 2013.

[9] X. Xie, X. Zhang, Scalable user selection for MU-MIMO networks, in: Proceedings of the IEEE International Conference on Computer Communications (INFOCOM), 2014.

[10] E. Aryafar, N. Anand, T. Salonidis, E.W. Knightly, Design and experimental evaluation of multi-user beamforming in wireless LANs, in: Proceedings of the Annual ACM International Conference on Mobile Computing and Networking (MobiCom), 2010.

[11] W.L. Shen, Y.C. Tung, K.C. Lee, K.C. Lin, S. Gollakota, D. Katabi, M.S. Chen, Rate adaptation for 802.11 multiuser MIMO networks, in: Proceedings of the Annual ACM International Conference on Mobile Computing and Networking (MobiCom), 2012.

[12] Q. Yang, X. Li, H. Yao, J. Fang, K. Tan, W. Hu, J. Zhang, Y. Zhang, Bigstation: enabling scalable real-time signal processing in large MU-MIMO systems, In: Proceedings of the ACM Conference on Data Communication (SIGCOMM), 2013.

[13] Rice University WARP project, Wireless Open-Access Research Platform, https://warpproject.org/trac

[14] Ettus Research USRP, Universal software radio peripheral, https://www.ettus.com/.

[15] K. Tan, J. Zhang, J. Fang, H. Liu, Y. Ye, S. Wang, Y. Zhang, H. Wu, W. Wang, G.M. Voelker, Sora: high performance software radio using general purpose multi-core processors, Proceedings of the Sixth USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2009.

[16] MU-MIMO and the User Experience: Quantifying the benefits of more advanced 802.11ac features, Signals Research Group, https://www.qualcomm.com/documents/mu-mimo-and-user-experience-srg-whitepaper, 2016.

[17] Smallnetbuilder Smallnetbuilder: Network Equipment Review Community, https://www.smallnetbuilder.com/

[18] Ixia VeriWave equipment, https://www.ixiacom.com/products/ixveriwave

[19] Aruba Networks White Paper: 802.11ac In-Depth, https://www.arubanetworks.com/assets/wp/WP_80211acInDepth.pdf, Aruba Networks, 2014.

[20] J.S. Park, A. Nandan, M. Gerla, H. Lee, SPACE-MAC: Enabling spatial reuse using MIMO channel-aware MAC, Proceedings of the IEEE International Conference On Communications (ICC), 2005.

[21] E. Perahia, R. Stacey, Next generation wireless LANs: 802.11n and 802.11ac, Cambridge university press, 2013.

[22] M. Costa, Writing on dirty paper (corresp.), IEEE Transactions on Information Theory (1983).

[23] V. Veen, D. Barry, K.M. Buckley, Beamforming: a versatile approach to spatial filtering, IEEE ASSP magazine (1988).

[24] T. Yoo, A. Goldsmith, On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming, IEEE Journal on Selected Areas in Communications (2006).

[25] A. Zanella, M. Chiani, M.Z. Win, MMSE Reception and successive interference cancellation for MIMO systems with high spectral efficiency, IEEE Transactions on Wireless Communications (2005).

[26] H. Yu, T. Kim, Beamforming transmission in IEEE 802.11 ac under time-varying channels, The scientific world Journal (2014).

[27] P. Zhao, B. Daneshrad, M. Gerla, G. Pei, J.H. Kim, Concurrent link capacity of MIMO ad-hoc networks with qos constraint, Proceedings of the IEEE Military Communications Conference (MILCOM), 2008.

[28] J. Zhang, R.W. Heath Jr., M. Kountouris, J.G. Andrews, Mode switching for MIMO broadcast channel based on delay and channel quantization, Proceedings of the European Association for Signal Processing (EURASIP), 2018.

[29] Q. Wang, L.J. Cimini, L.J. Greenstein, D.S. Chan, Single-user and multi-user mode selection for MIMO broadcast channels with imperfect channel state information, Proceedings of the IEEE Conference on Information Sciences and Systems (CISS), 2012.

[30] TGac Channel Model Addendum Supporting Material, https://mentor.ieee.org/802.11/dcn/09/11-09-0569-00-00ac-tgac-channel-model-addendum-supporting-material.doc, 2009

[31] A. Bhartia, B. Chen, F. Wang, et al., Measurement-based, practical techniques to improve 802.11ac performance, ACM Internet Measurement Conference (IMC), 2017.

[32] Tamosoft Wi-Fi Packet Sniffer Tool, Commview for Wifi, https://www.tamos.com/products/commwifi/

[33] J. Chan, C. Zheng, X. Zhou, 3D printing your wireless coverage, Proceedings of the ACM Workshop on Hot Topics in Wireless (HotWireless), 2015.

[34] R. Porat, E. Ojard, N. Jindal, et al., Improved MU-MIMO performance for future 802.11 systems using differential feedback, Proceedings of the Information Theory and Applications Workshop (ITA), 2013.

[35] R. Crepaldi, J. Lee, R. Etkin, et al., CSI-SF: Estimating wireless channel state using CSI sampling & fusion, Proceedings of the IEEE International Conference on Computer Communications (INFOCOM), 2012.

[36] F. Boccardi, et al., Five disruptive technology directions for 5g, IEEE Communications Magazine 52.2 (2014) 74–80.

[37] B. Bellalta, IEEE 802.11 Ax: high-efficiency WLANs, IEEE Wireless Communications Magazine 23.1 (2016) 38–46.

[38] W.L. Shen, K.C.J. Lin, M.S. Chen, T. Kun, SIEVE: Scalable user grouping for large MU-MIMO systems, In: Proceedings of the IEEE International Conference on Computer Communications (INFOCOM), 2015.

[39] J. Wang, D. Love, M.D. Zoltowski, User selection with zero-forcing beamforming achieves the asymptotically optimal sum rate, IEEE Transactions on Signal Processing (2008).

[40] S. Huang, H. Yin, J. Wu, V.C.M. Leung, User selection for multiuser MIMO downlink with zero-forcing beamforming, IEEE Transactions on Vehicular Technology (2013).

[41] X. Ma, Q. Gao, V. Marojevic, J.H. Reed, Hypergraph matching for MU-MIMO user grouping in wireless LANs, Ad Hoc Networks (2016).
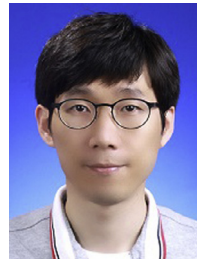
**Hyunwoo Choi** received his B.S. degree in Mechatronics Engineering from the Chungnam National University in 2007. He has been working on development for mobile devices in the Mobile Communication division at Samsung Electronics since 2007. He is currently pursuing his M.S. degree in Computer Science in Korea Advanced Institute of Science and Technology (KAIST). His research interests are in wireless communication, human-computer interaction, and mobile computing.

**Taesik Gong** received his Bachelors of Science in Computer Science from Yonsei University in 2016 (Summa Cum Laude). He received his Master of Science in Computer Science from Korea Advanced Institute of Science and Technology (KAIST) in 2017. He is currently pursuing his Doctor of Philosophy (Ph.D.) in Computer Science in Korea Advanced Institute of Science and Technology (KAIST). His research interests are in mobile computing, ubiquitous computing, and machine learning on mobile devices.

**Jaehun Kim** is currently pursuing his B.S. degree in Computer Science in Korea Advanced Institute of Science and Technology (KAIST). His research interests are in the areas of wireless networks, and moblie computing.

**Jaemin Shin** received his B.S. degree in Computer Science from Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea, in 2018. He is currently pursuing his M.S. degree in Computer Science in Korea Advanced Institute of Science and Technology (KAIST). He received the Best Paper Award from the Cloud-Assisted Networking (CAN) workshop at ACM CoNEXT 2017. His research interests are in the areas of wireless networks, and mobile computing.

**Sung-Ju Lee** is a Professor and KAIST Endowed Chair Professor at KAIST. He received his Ph.D. in computer science from the University of California, Los Angeles in 2000, and spent 15 years in the industry in Silicon Valley before joining KAIST. His research interests include computer networks, mobile computing, network security, and HCI. He is the winner of the HP CEO Innovation Award, the Best Paper Award at IEEE ICDCS 2016, and the Test-of-Time Paper Award at ACM WINTECH 2016. He is an IEEE Fellow and ACM Distinguished Scientist.