

# Resilient Multi-User Beamforming WLANs: Mobility, Interference, and Imperfect CSI

Oscar Bejarano\*, Roger Pierre Fabris Hoefel<sup>†</sup>, and Edward W. Knightly<sup>‡</sup>

\*Cisco Systems, Inc.

<sup>†</sup>Federal University of Rio Grande do Sul

<sup>‡</sup>Rice University

**Abstract**—In this paper we present the design of **CHRoME**, a downlink multi-user beamforming (MUBF) protocol that addresses the inherent sensitivity of multi-stream systems to mobility, inter-stream interference, and imperfect channel state information. Our contributions are: (i) a technique for accurately selecting the downlink bit rate in the presence of inter-stream interference via a custom multi-user probe and feedback signal, immediately preceding data transmission, and (ii) a fast retransmission scheme that exploits liberated antenna resources to increase the expected per-user signal-to-interference-plus-noise ratio (SINR) and retransmit without having to re-sound the channel. We implement each mechanism and evaluate via a combination of indoor over-the-air experiments and trace-driven emulation. We demonstrate that **CHRoME** increases the resilience of MUBF systems to inter-stream interference and achieves multi-fold throughput gains compared to IEEE 802.11ac.

## I. INTRODUCTION

Downlink Multi-User Beamforming (MUBF) is a key technique to scale throughput in dense wireless local area networks (WLANs) as it enables an Access Point (AP) to simultaneously transmit multiple independent data streams to different users in the same frequency resource block.<sup>1</sup> Such multi-user transmission has been demonstrated in WLAN systems (e.g., [3,4]), massive MIMO systems (e.g., [23]), and is now standardized in IEEE 802.11ac [13] and commercialized.

To achieve concurrent transmission, the AP precodes the independent streams by multiplying them by a beam-steering weight matrix in a way that reduces or removes inter-user or inter-stream interference. Such precoding requires knowledge of the channels between the antenna array at the AP and each concurrently served user. In protocols such as IEEE 802.11ac, this Channel State Information (CSI) is obtained via a *sounding* process in which predefined pilots are transmitted by the AP so that channel state is estimated by the receiver and fed back to the transmitter.

Unfortunately, client mobility, environmental mobility, and any source of precoding error (e.g., due to CSI feedback compression/quantization) can vastly degrade performance. In particular, imperfect beam steering does not merely result in a poorer quality signal at the receiver due to energy being directed away from the receiver: in a multi-user system, imperfect beam steering also increases inter-stream and inter-user interference, i.e., influencing both the signal  $S$  and interference  $I$  components of the signal-to-interference-plus-noise ratio (SINR) (the argument can be made rigorous via capacity analysis [11]).

In this paper we present the design, implementation, and experimental evaluation of **CHRoME** and make the following contributions:

First, we propose **M<sup>3</sup>CS** (Multi-user Multi-stream MCS), a technique for “just-in-time” multi-user bit-rate selection. In contrast to single-stream systems, multi-stream modulation and coding scheme (MCS) selection introduces the challenge of selecting multiple and potentially unequal MCS instead of just a single one. Schemes where the AP selects the MCS based on collected CSI at the transmitter (CSIT) have been demonstrated to have strong performance in SU-MIMO systems [12]. However, we will show that in MUBF systems, if the AP selects the MCS for each stream based solely on collected CSIT [5,21], performance will rapidly degrade with increasing mobility and estimation error. In principle, as mobility and other uncontrollable factors degrade SINR, the AP can maintain successful frame reception at the users by sufficiently reducing the MCS.

Our key technique is to make the selection as late as possible (immediately prior to data transmission) and to use a beamformed probe so that clients can assess the actual SINR of the beamformed transmission vs. the predicted SINR due to measurements of the channel training sequences. In this way, the AP can re-tune its selections accordingly, just-in-time for the downlink data transmission. We will show that with mobile environments, mobile users, or imperfect CSIT, the additional overhead introduced by the MUBF probing and feedback is far outweighed by avoiding rate under-selection (unnecessarily low MCS that wastes airtime) or over-selection (excessively high MCS that yields frame loss).

Second, despite the aforementioned resilience mechanisms, frames will occasionally be non-decodable due to excessive co-stream interference or mobility. Unfortunately, current retransmission strategies, inherited from the original CSMA design, require re-contention after a doubled backoff window. Consequently, physical layer parameters such as beam-steering weights are likely stale by the time retransmission is feasible, and therefore the time and resource penalty of channel sounding must be incurred again. In contrast, we design a *soundless fast retransmission* strategy in which the AP triggers a one-time immediate retransmission using the same CSIT as in the original transmission. Yet, because the original transmission failed, it is clear that the retransmission strategy must be changed. Thus, because only a *subset* of the users’ transmissions will have failed, the retransmission will exploit the “liberated” degrees of freedom (DoF), e.g.,

<sup>1</sup>In this work we use MU-MIMO and MUBF interchangeably.

an 8 antenna AP transmitting to 8 users with 2 failed frames will lead to an additional 6 DoFs for the retransmission to 2 users. To avoid resounding the channel, we design a scheme in which the user's block acknowledgements (BA) that follow the failed transmission, piggyback a measurement of inter-stream interference obtained during the data transmission. With this hint, the AP can characterize the expected retransmission SINR, and reset beam-steering weights and bit rates such that they are sufficiently robust to enable reception *despite* the use of increasingly outdated CSIT.

Finally, we implement both components of CHRoME on the WARP platform [1], and perform an extensive set of over-the-air experiments combined with trace-driven emulation. Our evaluation reveals that the MCS selection mechanism in CHRoME can achieve between 7% and 280% throughput gains under mobility and dynamic channel scenarios, compared to CSIT-based MCS selection schemes. Likewise, under non-ideal quantization, CHRoME can reach between 9% and 600% throughput gains. Similarly, the fast retransmission scheme in CHRoME outperforms 802.11ac by at least 66% in terms of throughput.

## II. M<sup>3</sup>CS: BEAMFORMED PROBING FOR “JUST-IN-TIME” MCS SELECTION

Multi-user Multi-Stream MCS, M<sup>3</sup>CS, assesses the channel and inter-stream interference affecting each user, just prior to the data MUBF transmission, and adapts each stream's MCS accordingly. Modulation and coding scheme selection in MUBF is fundamentally different to the case of single-input single-output (SISO) systems. SISO transmitters typically rely on the SNR of previous packets as well as packet loss history to determine the best MCS to be used in the next transmission, i.e., SNR-based and packet loss-based algorithms [7]. Nevertheless, a MUBF AP cannot rely on individual SNR knowledge unless the channels to all users are completely orthogonal. Otherwise, any dependence among channel vectors to the multiple users would introduce an interference signal component. Similarly, using packet loss as a MUBF MCS indicator would require the set of concurrently served users to be the same for successive transmissions. Otherwise, the channel correlation between the different user groups would lead to different MCS requirements.

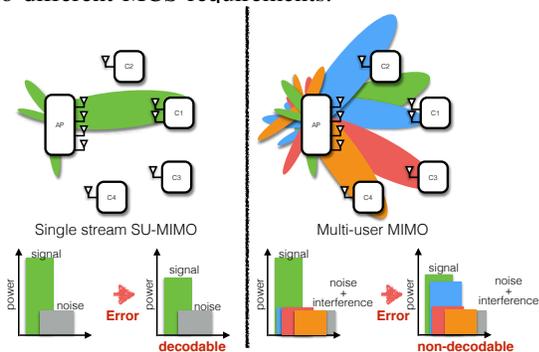


Fig. 1: Behavior of single-stream vs. multi-stream systems under beamforming errors. The figure considers only a single stream at a time in SU-MIMO.

Figure 1 illustrates the key difference between how single-stream and multi-stream systems are affected by errors in

channel estimation or by environmental/user mobility. In particular, the figure shows that beamforming errors in the single-stream case merely result in a decrease in signal strength whereas in the multi-stream case they can also lead to an increase in inter-stream interference.

### A. CSIT-Based MCS Selection

In order to accurately determine the most appropriate MCS for each individual user within a concurrent group, the AP needs to estimate the expected SINR with which the transmitted signals will arrive at each user during the data MUBF transmission. In the hypothetical case of the AP acquiring perfect CSIT (no quantization errors or channel variations), the SINR for every user can be directly calculated as follows. Consider a narrowband channel model and a network comprised of a single AP with  $M$  transmit antennas and  $K$  users. Let  $\Upsilon$  be the set of transmit antennas at the AP ( $|\Upsilon| = M$ ),<sup>2</sup> and  $S$  be the set of users selected by the AP to be served in the next MUBF transmission, i.e.,  $S \subseteq \{1, \dots, K\}$ ,  $|S| \leq M$ . Also, let  $\Gamma$  be the power matrix and each entry  $\gamma_{ji} \in \Gamma$  be the power from stream  $j$  measured at user  $i$ .  $\mathbf{W} = [\mathbf{w}_1 \dots \mathbf{w}_{|S|}]$  represents the precoding matrix applied by the AP to generate a single stream to each of the users in  $S$ . Thus, given the collected CSIT, the AP computes the SINR <sub>$i$</sub>  of the intended stream at user  $i$  (at pre-detection point) as follows:

$$\text{SINR}_i = \frac{\gamma_{ii}}{N_o + \sum_{j:j \neq i} \gamma_{ji}} \quad (1)$$

$$= \frac{|\sum_{m=1}^{|\Upsilon|} h_{im} w_{mi}|^2}{N_o + \sum_{j:j \neq i} |\sum_{m=1}^{|\Upsilon|} h_{im} w_{mj}|^2} \quad (2)$$

where  $h_{im}$  denotes the complex channel gain between AP antenna  $m$  and user  $i$ . Similarly,  $w_{mi}$  represents the complex weight applied to AP antenna  $m$  and  $N_o$  is the noise power at user  $i$ . More specifically, via the sounding process, the AP learns the complex channel vector representing the path between all transmitting antennas and the users, i.e., channel matrix  $\mathbf{H}$ , and uses this information to compute the corresponding precoding weight matrix  $\mathbf{W}$ . As shown in Equation (2), this information is sufficient to determine the MUBF SINR given the current CSIT and current channel conditions.

MCS selection based solely on CSIT has several drawbacks. First, Equation (2) assumes that channels remain static between sounding and MUBF data transmission, whereas channel variation will decrease the SINR. Because a decrease as low as 2 or 3 dB in SINR requires a reduction in MCS, throughput can be severely degraded with such an SINR decrease. Feedback compression or feedback reduction schemes in which sounding does not take place before every data transmission [6,28] are therefore particularly vulnerable since channel variation over multiple packet transmissions is more significant.

Second, inaccurate CSIT estimation due to quantization or inter-stream interference will have a similar effect. Quantization primarily affects *explicit* sounding in which users estimate the channel based on a training sequence transmitted

<sup>2</sup>Symbol  $|\cdot|$  indicates set cardinality.

by the AP, and then feed back a quantized version of these estimates utilizing a small number of bits to limit overhead. Likewise, *implicit* feedback, in which users transmit training sequences and the AP estimates the reciprocal channel based on this training, any other sources of interference (or noise) at the users will lead to inaccurate CSIT estimation because the channel measurements take place at the AP and not at the users. Consequently, this interference information is not considered in such estimates which in turn can lead to poor MCS selection.

### B. MCS Selection via MUBF Probing

The combination of quantization errors, channel dynamics, and inaccurate information with respect to the noise and interference observed at each user, affect the performance of MCS selection schemes in MUBF systems. Therefore, the amount of inter-stream interference affecting each user, as well as negative effects due to current channel conditions can only be known during the actual downlink MUBF transmission. We design a multi-user inter-stream interference probing mechanism that proactively evaluates the MCS selection resulting from predicting per-user SINR based on the acquired CSIT. In particular, CHRoME employs a multi-stage MCS selection scheme that probes the multi-user channels to evaluate the accuracy of the precoding scheme and adapts each stream's MCS, just-in-time for data transmission. While the first two stages are dedicated to acquiring CSIT and to probing the channel to adapt the MCS of all users, a third stage consists of reporting back this information to the AP. Figure 2 depicts the entire sounding, probing, and feedback process.

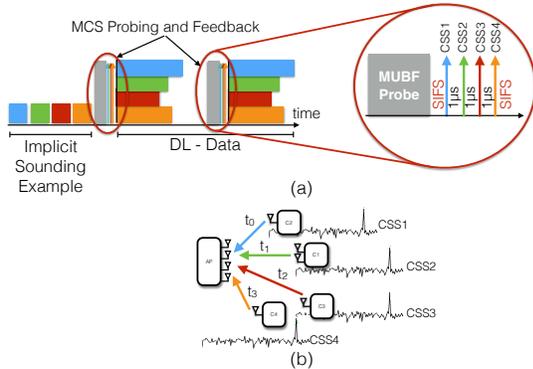


Fig. 2: MUBF probing and CSS feedback.

1) *Multi-User Inter-Stream Interference Probing*: Since CHRoME probes themselves are transmitted at a particular MCS, we use (necessarily sub-optimal) CSIT-based selection to set this initial MCS for each stream of the probe. Thus, using the most recent CSIT for each user  $s \in \mathcal{S}$ , the AP computes the beam-steering weight matrix and applies it to the independent data streams for the probe. The AP then triggers a multi-user probe by transmitting a minimum-length downlink multi-user frame at the rates determined using CSIT-based MCS selection. The multi-user probing frame enables each user to infer channel variations since sounding occurred, as well as the inter-stream interference affecting the transmission. Thus, upon reception of this probing frame, each user measures its effective SINR ( $\text{SINR}_{\text{eff}}$ ) and maps it to

the corresponding preferred MCS. Notice that in the ideal case that channels are completely static and CSIT is perfectly estimated, the measured MUBF SINR corresponds to the MCS previously estimated by the AP via the CSIT-based MCS selection scheme. In contrast, in non-ideal cases, the AP can now adapt to the true conditions. CHRoME is agnostic to the feedback mechanism implemented (i.e., it can operate with implicit or explicit systems) and does not make any assumption with respect to how frequently sounding occurs.

2) *Correlatable MCS Feedback*: In order for the AP to readjust the MCS according to the current channel conditions, each user needs to report back the computed MCS to the AP. Moreover, this feedback process needs to take place within the shortest time possible to minimize overhead. Since MCS is identified with an index (0 to 9 in 802.11ac), we can represent each MCS selection with only a few bits. CHRoME maps each MCS index to a predefined pseudo noise binary codeword (i.e., a correlatable symbol sequence or CSS [15]). The transmission length and processing required to identify these sequences is significantly lower than what is required for decoding a packet, thus making them ideal for this application. More specifically, upon MCS selection at the users, they reply with a corresponding CSS. Figure 2 depicts both (a) the timeline (not to scale) showing where the MCS probing and feedback take place within a given MUBF transmission, as well as (b) a simplified representation of CSS usage.

**Signaling MCS with a CSS.** Broadly, correlatable symbol sequences are BPSK sequences that are filtered, up-sampled, and transmitted via wideband techniques. While CSS preserve the statistical properties of sampled white noise, cross-correlation of any CSS with a matching copy will produce a spike indicating a positive match. The advantages of CSS over decodable packets include higher detection reliability, higher robustness to radio parameter imperfections, and substantial transmission time reduction. More specifically, as demonstrated in [15], 127-symbol Gold sequences can be reliably detected at low SINR (-6 dB) with 5.7% false negatives and no false positives. Consequently, these can be detected at 10 dB lower compared to 6 Mbps OFDM frames. Moreover, CSS do not require a preamble or data processing thus reducing the time needed for their transmission to only  $6.35 \mu\text{s}$  [15].

**CHRoME's dictionary.** 802.11ac features 10 different modulation and coding schemes indexed from 0 to 9. Given that 127-symbol Gold sequences allow 127 different sequences while retaining a low theoretical cross-correlation among them, these can easily support a mapping to 10 different MCS indexes. To support all 10 codewords the AP's hardware could either use 10 simultaneous correlators (higher design complexity and cost), or buffer the received sequences and evaluate them sequentially one at a time (longer processing).

**Feedback Processing.** CSS do not require data decoding. Out of the  $16 \mu\text{s}$  that 802.11 allocates for SIFS,  $14 \mu\text{s}$  are used for such processing and the rest for switching the radio modality (between TX and RX). Given that during CSS reception the AP does not need to switch from TX chain to RX chain, and vice versa, the transmission of consecutive CSS requires only up to  $1 \mu\text{s}$  in between to account for signal propagation delay. Notice that the order in which users reply follows the order established in the sounding frames (e.g., Null

Data Packet Announcement *NDPA*, in 802.11ac).

### C. Incurred Overhead

The additional time required to trigger the probing and feedback mechanism in CHRoME can be broken down into the following components.

$$T_{\text{overhead}} = \text{SIFS} + \text{Probe} + \text{SIFS} + |S| \cdot \text{CSS} + \dots \\ (|S| - 1) \cdot (1\mu\text{s})$$

where SIFS and RIFS take  $16 \mu\text{s}$  and  $2 \mu\text{s}$ , respectively. As previously mentioned, each CSS requires  $6.35 \mu\text{s}$ . Finally, the length of the probing frame depends on the minimum MCS used to transmit one A-MPDU to each user. That is, if four users are probed and three of them are served at MCS-3 (16-QAM,  $\frac{1}{2}$ ) but the remaining one is probed at MCS-0 (BPSK,  $\frac{1}{2}$ ), then the maximum time length is computed based on the MCS-0 transmission. We implement A-MPDUs as short as RTS frames. Therefore, the probing frame would take approximately  $52 \mu\text{s}$  including preambles and assuming 6 Mbps transmission rate. For a 4-user system employing the lowest transmission rate (worst case), the total overhead  $T_{\text{overhead}}$  reaches a maximum of  $112.4 \mu\text{s}$ . Notice however that increasing the transmission rate of the probing frame would *significantly decrease* the total overhead.

## III. MULTI-USER INTERFERENCE-AWARE FAST RECOVERY

In this section we describe the design and implementation factors for our multi-user fast recovery scheme.

### A. Overview

Compared to single-user systems, failed multi-user transmissions incur a lengthier recovery time thus reducing the system's efficiency. In particular, multi-user retransmissions typically involve not only a contention phase as in the case of 802.11 legacy retransmissions but also a re-sounding phase. Namely, in 802.11-based MUBF systems, upon a failed transmission, the AP triggers a binary exponential backoff process and begins contending for the medium. Once the AP gains access to the medium, it re-sounds the channel to generate the beam-steering weights needed for MUBF.

In contrast, we propose a multi-user retransmissions scheme that precludes the need to re-sound the channel by triggering a one-time immediate retransmission. That is, our scheme targets to realize a throughput gain by reducing the overhead incurred from repeated channel estimation. Nonetheless, by doing this, the AP faces the challenge of precisely determining the MCS which yields a successful retransmission even when the CSIT it possess for each user is increasingly outdated and inaccurate. Merely selecting MCS based on previously collected CSIT would likely lead to a failed retransmission, especially given the fact that the original transmission using this information has already failed. Similarly, arbitrarily decreasing the MCS to account for uncertainty in the current channel conditions might become overly conservative. Our joint retransmission and MCS selection scheme considers two key concepts, *decreased receiver-dimensionality* due to liberated antenna resources after successfully serving at least one user in the original user set, and *per-user inter-stream interference awareness* at the AP obtained via feedback during the acknowledgment process.

### B. Retransmission Overhead

The retransmission process in multi-user systems is less efficient than that in single-stream systems due to the need to re-sound the channel. Consider the 4x4 example in Figure 3; the top figure illustrates the MUBF retransmission process in 802.11-based networks. First, the AP beamforms four different streams to four different users but only two of them are completely decoded. Consequently, the AP initiates a contention phase after DIFS time and then triggers a sounding phase to acquire the channel estimates of the two remaining users as well as two other users. In contrast, CHRoME eliminates the need to re-sound the two remaining users by immediately attempting a retransmission (Figure 3 - bottom). The potential overhead reduction is due to both eliminating re-sounding as well as to avoiding any increase in the contention window (CW) that is readjusted (incremented) after a frame loss. While the legacy system can maximize the number of streams served in a particular MUBF transmission, we demonstrate that the availability of additional degrees of freedom provides the opportunity to use the same CSIT collected in the previous sounding phase and still attain gains compared to legacy retransmission.

In CHRoME, before the retransmission is triggered, the AP evaluates whether to serve all users in a multi-user MIMO fashion or via a TDMA MISO (Time Division Multiple Access Multiple-Input Single-Output) transmission. In particular, the AP assesses the time required to complete the transmissions in the two different modes and selects the configuration that minimizes the retransmission time (considering the respective MCS to each user). The AP evaluates the time it takes to serve all users one at a time (sequentially) vs. serving them concurrently. Notice that in the TDMA MISO case the MCS for each user is expected to be higher due to a higher expected SINR enabled by a power and diversity gain. Consequently, the increase in MCS could lead to a faster overall transmission.

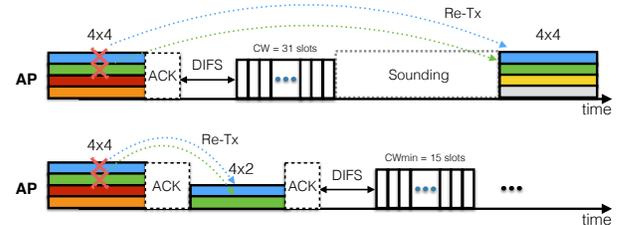


Fig. 3: Illustration of the retransmission strategy in legacy 802.11 (top) and in CHRoME (bottom)

### C. Receiver-Dimensionality Reduction

Since CHRoME avoids re-sounding the channel, we allow the AP to reuse the CSIT employed in the original transmission in order to generate the beam-steering weights needed in the beamformed retransmission.

**SINR Enhancement Due to Liberated Antennas.** Our rationale for allowing the reuse of possibly outdated channel information is based on the counteracting effect provided by the sudden availability of additional (liberated) antennas at the transmitter. That is, if any users were successfully served in the original transmission, every additional degree of freedom that becomes available at the transmitter (with respect to

the number of concurrent users) yields an increase in per-user SINR due to both antenna diversity gain and per-stream transmit power increase. CHRoME exploits these gains in order to counteract the SINR reduction that is due to the use of inaccurate or outdated channel estimates to generate the beam-steering weights. Without re-sounding the channel, the AP reshapes the channel matrix to account only for the remaining users and computes the beamforming weights for those users. Reducing the number of users to be served relative to the number of transmit antennas simplifies the construction of non-interfering streams, thus leading to a lower SINR penalty due to imperfect beamforming weights.

To show the potential SINR gains that can be attained by MUBF systems when the number of transmit antennas increases relative to the number of simultaneous users we simulate a scenario with one multi-antenna AP and 100 single-antenna users in an i.i.d Rayleigh MIMO channel. The AP employs a zero forcing precoding strategy and we assume perfect CSIT estimation (no quantization). Moreover, the channel input is subject to an average power constraint  $\mathbb{E}[\|\mathbf{x}\|_2^2] \leq \text{SNR}$ , where we let  $\text{SNR} = 10$  dB. The user group selected at each transmission is based on the individual channel norm for each user, i.e.,  $\|\mathbf{h}_s\|$ . Therefore, at every transmission the AP serves the  $M$  users with highest channel norm. Each data point consists of an average obtained over 10000 channel realizations. Figure 4 (left) shows the post-processing per-user SINR as a function of the number of transmit and receive antennas.

While the increase in SINR due to an additional transmit antenna varies depending on the overall configuration, the minimum increase we observed is roughly 2 dB (8x6 to 8x5 configuration). Moreover, these results demonstrate that the steepest increases occur at both extremes, that is, when the system approaches the maximum diversity gain, i.e.,  $M \times 1$ , as well as in the case where there is only one single additional antenna. More importantly, notice that the SINR increase observed in our simulations, closely match the expected value that is roughly approximated by Equation (3) [2] and plotted in Figure 4 (right). Therefore, the SINR of a signal transmitted with power  $P$  scales proportionally to  $\frac{M-|S|+1}{|S|}$ .

$$\mathcal{E}\{\text{SINR}_{BF}\} = 10 \cdot \log_{10} \left( \frac{M - |S| + 1}{|S|} \frac{P}{N_0} \right) \quad (3)$$

Although there is a clear scaling difference between both plots in Figure 4, the difference in SINR from increasing or decreasing the number of users relative to the number of transmit antennas remains the exact same. Based on these results, we can observe that in the case of the scenario presented in Figure 3, the per-user SINR in the retransmission can increase by close to 7 dB. Considering the required 802.11 receiver sensitivity this could mean an increase of more than two MCS indexes under the assumption of static channel conditions. Therefore, we expect that the decrease in SINR due to outdated CSIT can be significantly mitigated in CHRoME via receiver-dimensionality reduction.

**Inter-Stream Interference-Aware Retransmission.** While a failed transmission indicates that the channel cannot support the current MCS given the current transmission resources and

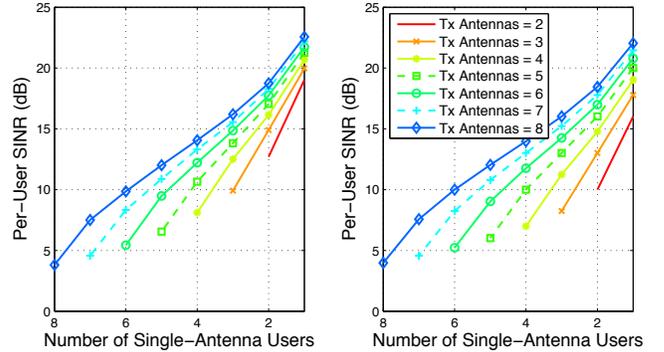


Fig. 4: Per-user SINR gain due to reduction in receiver-dimensionality (Left: Simulation; Right: Analytical).

conditions (i.e., transmit antennas and concurrent users), the AP has no other information to update its MCS selection according to current channel conditions. The default approach would be to let the AP select the MCS based on the CSIT for each user in the retransmission set. Nonetheless, as previously discussed, this would not consider the effects of inter-stream interference on each individual user. Similarly, the AP could conservatively select an MCS by simply decreasing the CSIT-based selection by one or two, e.g., MCS-4 to MCS-3 or to MCS-2. Notice however, that in this approach the AP merely relies on speculation that would possibly lead to an inaccurate selection (either by under- or over-selecting).

In CHRoME we enable users to *piggyback* information with respect to the SINR measured at each of these users, in the block acknowledgements (BA). More specifically, upon a failed transmission, users append the individual per-stream SINR components to their block acknowledgement. Each individual SINR corresponds to the individual components induced by each independent data stream. For instance, if the original transmission to user 2 failed, this user reports three individual SINR values based on the measured  $I_{1 \rightarrow 2}$ ,  $I_{3 \rightarrow 2}$ , and  $I_{4 \rightarrow 2}$  components.<sup>3</sup> Assuming that the retransmission user set contains both users 2 and 4 (reduced from four to two users), the AP considers only the SINR induced by user 4 onto user 2 in order to select the highest possible MCS according to the 802.11 receiver sensitivity specifications [13,17]. Notice that this requires extending the BA frame by  $|S| - 1$  fields, each consisting of only one octet. Recall  $S$  denotes the set of users selected by the AP to be served in the previous MUBF transmission. If the original transmission to a particular user was successful, a regular 802.11 BA is used.

Notice that recalculating the beam-steering matrix with fewer users would yield a higher per-user SINR compared to the SINR measured during the beamformed transmission. However, relying on such SINR increases robustness to outdated CSIT.

#### Discussion on TXOP and Channel Release Mechanism.

In the context of 802.11, a modification to the retransmission strategy would need to consider its effect on the transmit opportunity (TXOP) mechanisms. While we have not explicitly addressed this issue in our implementation, the TXOP can be

<sup>3</sup>We designed a Multi-Stream Interference Training structure (MSIT preamble) that allows each user to measure all the individual interference components.

adjusted so as to allow one fast retransmission at a minimum, for delay sensitive traffic such as voice and video. Similarly, in the case that no retransmission is required and the AP has no more data to transmit at a particular TXOP, channel release mechanisms such as [15] can be easily implemented.

#### IV. SYSTEM IMPLEMENTATION, MEASUREMENTS, AND EVALUATION

We validate CHRoME via an implementation and an extensive set of testbed and system emulation experiments. First, we describe our implementation and experimental methodology. Then, we investigate the performance of each individual technique in CHRoME using a combination of over-the-air transmissions as well as trace-driven emulation to accurately model 802.11 timings while transmitting over collected channel traces.

##### A. Implementation and Experimental Methodology

**Implementation and testbed.** We implemented CHRoME in WARP and WARPLab [1]. The WARPLab environment allows us to perform all the signal processing including encoding in a PC, and then transmit these signals over the air for decoding on the receiver side. Nonetheless, in WARPLab, reading (writing) from (to) the board’s buffers do not allow us to evaluate our protocol in real-time; therefore, to accurately represent the time-scales at which 802.11 operates, we also rely on trace-driven emulation where we first collect continuous channel samples and then use these to evaluate our scheme. More importantly, by doing this we ensure we can replay the same channels for the schemes to be compared, therefore achieving repeatability and a rigorous evaluation.

**Trace-driven emulation.** To accurately model the 802.11 time-scales we implement an 802.11ac-based MUBF trace-driven emulator featuring an entire OFDM transmit and receive RF chain. We collect a comprehensive set of channel traces with our testbed platform and use those as input to our emulator. Precoding consists of a zero-forcing scheme with equal power allocation. Transmit side error vector magnitude (EVM) is determined according to the highest MCS within a user group. We implement least squares channel estimation based on our MSIT preamble and the resulting per subcarrier SINR (post processing SINR at the MIMO detector output) is used to compute the effective SINR ( $\text{SINR}_{eff}$ ) which we then map to an MCS. Unless otherwise stated, we consider a 4-antenna AP serving up to 4 single antenna users at a time (single stream per user). Similarly, we consider 20 MHz transmissions over the 5 GHz band.

**From single-carrier to OFDM: effective SINR to MCS mapping.** In contrast to single carrier systems where SINR can be directly mapped to an MCS, in multi-carrier systems an intermediate step is necessary to map the per-subcarrier SINR to an effective SINR scalar metric, and in turn to a given MCS. The SINR in MIMO OFDM systems operating over frequency selective fading channels presents a highly dynamic range among the subcarriers. The performance of OFDM coded systems over these multi-carrier channels depends on the joint statistics of the SINR considering all data subcarriers, therefore, average SINR is not a useful metric to accurately estimate the system’s performance. The 802.11ax task group

(TGax) is considering the use of a mutual information based MCS mapping method to achieve this PHY abstraction [25]. This method uses the per-subcarrier SINR to compute a received bit information rate (RBIR) metric which is then mapped to an  $\text{SINR}_{eff}$ . Further, this  $\text{SINR}_{eff}$  is used to compute the packet error rate (PER) for different MCS.

In both plots in Figure 5 we present the curves we generated in order to map the per-subcarrier SINR to an MCS. That is, the figure on the left shows the relation between the RBIR metric and  $\text{SINR}_{eff}$  (where the scalar  $\text{SINR}_{eff}$  is obtained over the MIMO fading channel). This  $\text{SINR}_{eff}$  allows us to map to a PER obtained over an additive white Gaussian noise channel, as shown in the figure on the right.

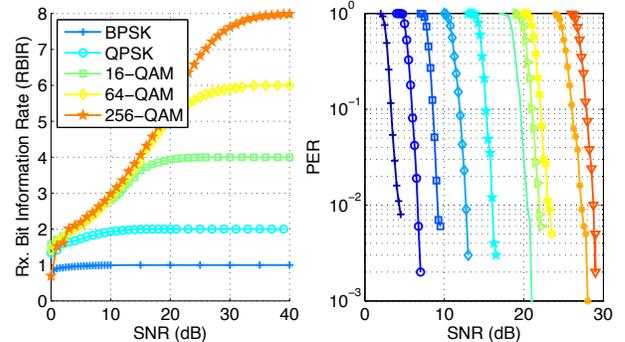


Fig. 5: (Left) RBIR as a function of SINR for least square channel estimation. (Right) PER vs. SINR with an MMSE receiver (binary convolutional code with soft-detection Viterbi decoder). From left to right: MCS0 to MCS9.

##### B. M<sup>3</sup>CS: Probing-Based MCS Selection

**MCS selection accuracy in explicit and implicit MUBF sounding systems.** While MCS over-estimation can lead to significant frame losses, under-estimation leads to an *opportunity loss* in which the current channel conditions could have supported higher rates thereby leading to a throughput increase. We investigate the MCS selection accuracy of our probing scheme compared to *baseline MCS selection* as well as a more conservative approach in which we decrease the baseline MCS by one. To this end, we evaluate the extent to which these schemes under- or over- select the MCS. Notice that the baseline MCS selection represents a scheme where each stream’s MCS is chosen according to the collected CSIT (i.e., a purely CSIT-based technique). We collect channel traces for over 28 different user locations and run 15,000 frame transmissions. For each transmission and MCS selection scheme we measure the number of frames in which the MCS was over-, under-, and accurately selected. Moreover, we consider two different feedback algorithms: *Explicit*: The AP sounds the channel before every packet transmission and follows the feedback process mandated in 802.11ac. *Implicit*: All users transmit a training pilot sequentially to allow the AP to estimate the channel. We modify our emulator to achieve perfect channel reciprocity (including transmit and receive RF chains) to eliminate calibration effects from our study.

Figure 6 depicts the MCS selection accuracy of each scheme. The top plots correspond to experiments where we generated *out-of-cell* interference from neighboring APs and

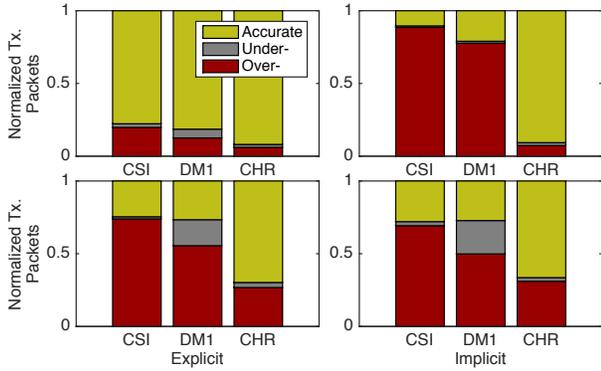


Fig. 6: Selection accuracy for all three MCS selection schemes. Top (Bottom) with (without) out-of-cell interference at users; CSI=Baseline CSIT-Based/DM1=Decrease-MCS-by-1 (Conservative)/CHR=CHRoME.

their users, i.e., interference that is *not* inter-stream interference. This out-of-cell yielded additional interference to the in-cell multi-user clients ranging from -70 to -90 dBm. Likewise, the left plots correspond to systems that obtain CSIT with explicit feedback measured by the clients and the right plots obtain CSIT with implicit measurements at the AP.

The results indicate that CHRoME is highly resilient to out-of-cell interference in both explicit and implicit systems. This is because CHRoME re-adjusts MCS according to the interference learned and observed by each user during the probe. In contrast, the baseline schemes perform poorly in *implicit* systems (right plots) because sounding does not take out-of-cell or inter-stream interference into consideration, therefore leading to substantial over-selection. On the other hand, the relatively fair performance of the two baseline schemes for explicit feedback systems with interference is due to the fact that this interference forces a dramatic drop in MCS to the lowest indexes thus avoiding significant over-selection. Consequently, CHRoME yields greater gains when the channel supports a wide range of MCS.

For the same experiments described above, we investigate the aggregate throughput of the multiple schemes (see Figure 7). That is, we consider the MAC/PHY overhead involved in the transmission process, including the additional overhead incurred by CHRoME. Notice that CHRoME outperforms the baseline schemes in all instances, achieving gains ranging from 16% to 280% in the case of implicit systems and between 16% and 42% in the case of explicit systems. Therefore, substantial gains outweigh the limited overhead necessary to enable MUBF probing and feedback in CHRoME.

**Adaptation response time.** To illustrate how well MCS selection in CHRoME follows the *best possible* MCS (i.e., the highest MCS that can be supported during the data MUBF transmission) compared to baseline schemes, we plot a timeline showing 50 samples in time and the MCS selected at each instance. Figure 8 (left) shows the MCS index as a function of time for the cases where there is interference at the users, and no interference, top and bottom respectively. Observe that the green (CHRoME) curve closely matches the best MCS (ground truth - measured at user during data transmission) whereas the basic baseline scheme frequently

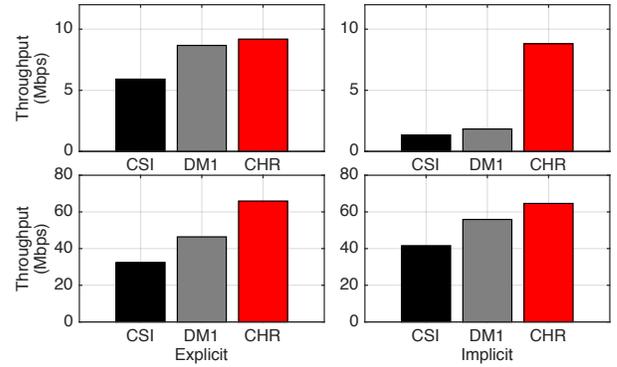


Fig. 7: Throughput of each feedback system and MCS scheme. One-to-one correspondence with plots in Figure 6.

over-selects. This demonstrates the capability of CHRoME to rapidly track the ideal selection even with drastic changes in channel conditions where the desired MCS jumps as high as eight MCS indexes.

**Robustness to suppression of channel sounding.** As shown in prior work [6,28], the overhead incurred by 802.11ac explicit feedback is can be a significant fraction of air-time. These same works have proposed suppression of channel sounding in order to reduce overhead by avoiding sounding before every packet transmission. Nonetheless, such schemes are susceptible to transient channel variations and stale CSIT. Therefore, we explore the ability of CHRoME to protect the system against these changes.

In Figure 8 (right) we present the throughput performance of CHRoME compared to the other MCS selection schemes as the time gap between sounding and the beamformed transmission is increased. Notice that this evaluation considers the overhead required to trigger CHRoME. While the slope for both baseline schemes is steep especially before reaching 50 ms, the slope of CHRoME's curve decreases at a much slower pace. Consequently, our probing scheme counteracts the degradation due to outdated beam-steering weights solely by allowing the AP to re-adjust MCS according to current channel conditions.

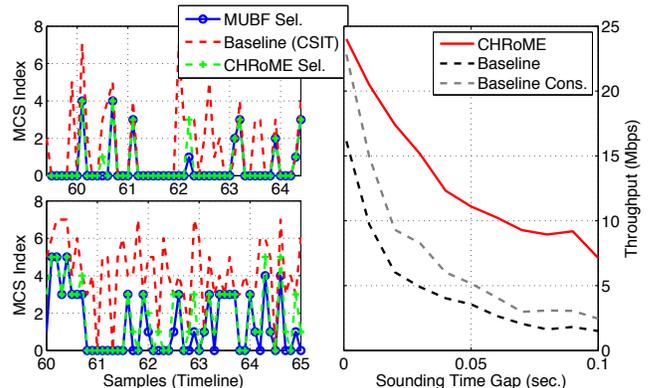


Fig. 8: (Left) MCS selection timeline; (Right) sum throughput for different time gaps between sounding and data MUBF transmission

**Resilience to feedback quantization error.** Similarly to changes in the environment and user mobility, errors due to

poor CSIT quantization will also hinder performance by inducing errors on the beam-steering weight calculations thereby increasing the inter-stream interference. We evaluate performance as a function of the number of bits used to quantize the channel estimates such that each user quantizes feedback using  $B$  bits. We perform scalar quantization [16,19] where the total number of bits are evenly allocated to magnitude and phase components. Following the scheme in [19], the elements of the channel vector  $\mathbf{h}_k = [h_1, \dots, h_M]^T$  – where  $M$  is the number of transmit antennas (and in this case, the number of concurrently served users) – are divided by the element  $h_1$  to yield  $M-1$  complex elements. Then, the  $M-1$  phases (relative) are quantized individually using uniform quantization in  $[-\pi, \pi]$ . On the other hand, the inverse tangents  $\tan^{-1}\left(\frac{|h_m|}{|h_1|}\right)$  of the relative magnitudes for  $m = 2, \dots, M$ , are quantized uniformly on the interval  $[0, \frac{\pi}{2}]$ .

Figure 9 depicts the MCS selection accuracy of all three selection schemes (i.e., (a) to (c)) as well as their corresponding per-user throughput performance (i.e., (d)). To guarantee that our results only reflect the effect of quantization, we perform this experiment under controlled scenarios where the channel obtained during sounding is the exact same as the one during the MUBF transmission. Observe that the selection accuracy of CHRoME is much higher than both baseline schemes in all cases. This is reflected in the throughput gain attained by CHRoME in Figure 9 (d). Since channels remain the same, the gap between schemes narrows as the number of bits needed to represent the actual channel vector increases. As shown in the figure, gains of CHRoME compared to the best of the baseline schemes range from 600% ( $B = 20$ ) to 9% ( $B = 100$ ).

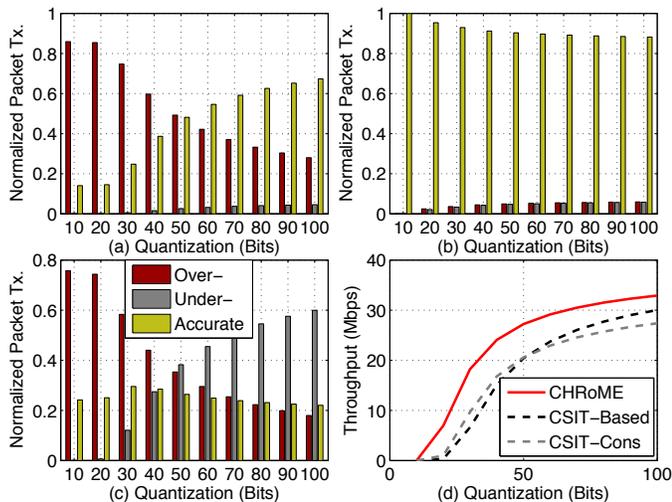


Fig. 9: Feedback quantization: throughput and MCS selection accuracy; (a) baseline, (b) CHRoME, and (c) conservative baseline.

### C. CHRoME's Fast Recovery

Throughput gain/loss of our retransmission system compared to an 802.11 MUBF approach depends on two main factors: incurred overhead and success rate of retransmission frames. Avoiding a re-sounding phase should decrease overhead yet could also decrease the likelihood of a successful

retransmission if the same channel information is used. In contrast, while 802.11 incurs higher overhead, it also uses more updated CSIT estimates to beamform. CHRoME attempts to increase the likelihood of a successful retransmissions while avoiding CSIT collection overhead.

We implement CHRoME's retransmission scheme and compare against 802.11 with re-sounding and against two other baselines. The first baseline consists of a MUBF retransmission scheme where  $R$  users are always served simultaneously, whereas the second one consists of a MISO TDMA scheme that serves all users sequentially (one at a time). We provide the 802.11 scheme an advantage by assuming that all the retransmissions use all available DoF to maximize the multiplexing gain. That is, in a system with 4 antennas at the AP, all retransmissions consider 4 concurrent users.

We transmit a total of 12,000 packets and plot the system's throughput in Figure 10. First, observe that the additional overhead incurred by the combination of doubling of the backoff window and channel re-sounding in the 802.11 scheme leads to a significant throughput penalty. Second, the difference in MCS due to higher number of concurrently served users causes a large throughput difference. Although the TDMA scheme serves each user at a higher rate compared to the MU-MIMO case, serving users sequentially leads to a similar drop in throughput. Therefore, with outdated CSIT and a small number of failed users, an MU-MIMO retransmission scheme has similar performance to a MISO TDMA scheme. Finally, observe that CHRoME performs at least as well as the best performing scheme (i.e., either MU-MIMO or MISO TDMA). That is, by selecting the configuration that minimizes the retransmission time we outperform the other strategies.

In addition, for each number of failed users, we compute the amount of times that a retransmission was 100% successful (i.e., no failed users during the retransmission). For the MU-MIMO scheme, there is a decrease in percentage as the number of failed users increases; nonetheless, this decrease only goes from 96.6% to 80.5%. Similarly we observe that for the TDMA scheme it ranges from 98% to 86% due to having a more aggressive MCS selection mechanism therefore incurring in over-selection. Furthermore, CHRoME achieved an overall overhead reduction of 64.6% compared to the 802.11 baseline. The significant reduction in overhead due to avoiding sounding, as well as the resilience provided by the liberated degrees of freedom which enable a high successful retransmission rate, shift the accuracy/overhead tradeoff in favor of CHRoME's fast retransmission scheme leading to high throughput gains.

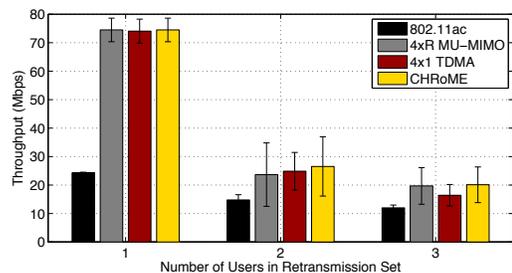


Fig. 10: System's throughput of CHRoME's retransmission scheme compared to 802.11.

## V. RELATED WORK

Prior work can be divided into *pre-* and *post-* transmission techniques, according to whether their protocol mechanisms are applied before data transmission (channel sounding, precoding, etc.) or after data transmission (retransmission).

**Pre-MUBF Transmission.** Both theoretical and practical work has focused on developing user, mode, and MCS selection strategies that attempt to maximize a rate or fairness metric [2,8,9,14,20,22]. More specifically, user and mode selection (or spatial scheduling) algorithms have been designed to group users based on their spatial correlation with the purpose of maximizing SINR at each user, consequently enabling the use of higher MCS. More importantly, these works rely on the CSIT measured during the sounding phase in order to allow the AP to determine the SINR of each user, and select the highest possible MCS based on the inferred SINR.

In contrast, our work focuses on enabling resilience regardless of the pre-chosen user set and their corresponding MCS. Therefore, CHRoME complements these protocols. In addition, other approaches aimed at eliminating inter-stream interference in MUBF focus on the implementation of accurate CSI estimation and quantization techniques [13], as well as a wide variety of precoding strategies [9,26,27]. CHRoME works in combination with any type of legacy CSI estimation and quantization, as well as precoding strategy thereby also complementing all of these downlink systems.

**Post-MUBF Transmission.** Upon a failed MUBF transmission, previous work on downlink MUBF WLANs either follows the retransmission mechanism proposed in 802.11 (via binary exponential backoff) or ignores the retransmission process [3,4,11,18,23,30]. On the other hand, cellular systems such as LTE Advanced employ mechanisms such as Hybrid Automatic Repeat-reQuest (HARQ) [10,24] to efficiently recover from losses. In contrast to MUBF WLAN systems, CHRoME avoids the costly overhead incurred by the combination of sounding and binary exponential backoff via a one-time immediate retransmission thus reusing the same CSIT but exploiting additional degrees of freedom. Moreover, CHRoME benefits from the added robustness that can be added via HARQ schemes such as incremental redundancy.

## VI. ACKNOWLEDGEMENTS

This research was supported by Cisco Systems, Intel, the Keck Foundation, and by NSF grants CNS-1514285, CNS-1444056, CNS-1126478 and CNS-1012831.

## VII. CONCLUSION

In this paper we present the design and implementation of CHRoME, a novel MUBF scheme that increases resilience against downlink inter-stream interference due to channel variations, user mobility, and poor feedback quantization. CHRoME features an inter-stream-interference-robust MCS selection technique and a fast retransmission scheme that obviates the need to re-sound the channel therefore minimizing overhead while guaranteeing robust retransmissions. We demonstrate that by obtaining and incorporating knowledge with respect to inter-stream interference into design decisions, our protocol can attain significant throughput gains compared to legacy systems.

## REFERENCES

- [1] Rice University WARP project. Available at: <http://warp.rice.edu>.
- [2] N. Anand, J.-K. Lee, S.-J. Lee, and E. Knightly. Mode and user selection for multi-user MIMO WLANs without CSI. In *Proc. of IEEE INFOCOM*, 2015.
- [3] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly. Design and experimental evaluation of multi-user beamforming in wireless LANs. In *Proc. of ACM MobiCom*, 2010.
- [4] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire. Achieving high data rates in a distributed MIMO system. In *Proc. of ACM MobiCom*, 2012.
- [5] H.-V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire. Airsync: Enabling distributed multiuser MIMO with full spatial multiplexing. *IEEE/ACM Trans. on Networking*, 21(6):1681–1695, 2013.
- [6] O. Bejarano, E. Magistretti, O. Gurewitz, and E. Knightly. MUTE: Sounding inhibition for MU-MIMO WLANs. In *Proc. of IEEE SECON*, 2014.
- [7] J. Camp and E. Knightly. Modulation rate adaptation in urban and vehicular environments: cross-layer implementation and experimental evaluation. *IEEE/ACM Trans. on Networking*, 18(6):1949–1962, 2010.
- [8] R. Chen, Z. Shen, J.-G. Andrews, and R.-W. Heath. Multimode transmission for multiuser MIMO systems with block diagonalization. *IEEE Transactions on Signal Processing*, 56(7):3294–3302, 2008.
- [9] Y. Cheng, S. Li, J. Zhang, F. Roemer, B. Song, M. Haardt, Y. Zhou, and M. Dong. An efficient transmission strategy for the multicarrier multiuser MIMO downlink. *IEEE Transactions on Vehicular Technology*, 63(2):628–642, 2014.
- [10] E. Dahlman, S. Parkvall, and J. Skold. *4G: LTE / LTE-advanced for mobile broadband*. Academic Press, 2013.
- [11] D. Gesbert, M. Kountouris, R. Heath, C.-B. Ch., and T. Salzer. Shifting the MIMO paradigm. *IEEE Signal Proc. Mag.*, 24(5):36–46, 2007.
- [12] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. Predictable 802.11 packet delivery from wireless channel measurements. *Proc. of ACM SIGCOMM*, 2011.
- [13] IEEE 802.11ac/D7.0, Enhancements for Very High Throughput for Operation in Bands Below 6 GHz, 2013.
- [14] V. Jungnickel, M. Schellmann, L. Thiele, T. Wirth, T. Haustein, O. Koch, W. Zirwas, and E. Schulz. Interference-aware scheduling in the multiuser MIMO-OFDM downlink. *IEEE Comm. Magazine*, 47(6):56–66, 2009.
- [15] E. Magistretti, O. Gurewitz, and E. Knightly. 802.11ec: collision avoidance without control messages. In *Proc. of ACM MobiCom*, 2012.
- [16] A. Narula, M.-J. Lopez, M.-D. Trott, and G.-W. Wornell. Efficient use of side information in multiple-antenna data transmission over fading channels. *IEEE JSAC*, 16(8):1423–1436, 1998.
- [17] E. Perahia and R. Stacey. *Next Generation Wireless LANs: 802.11 n and 802.11 ac*. Cambridge university press, 2013.
- [18] H. S. Rahul, S. Kumar, and D. Katabi. JMB: Scaling wireless capacity with user demands. In *Proc. of ACM SIGCOMM*, 2012.
- [19] N. Ravindran and N. Jindal. Multi-user diversity vs. accurate channel state information in MIMO downlink channels. *IEEE Transactions on Wireless Communications*, 11(9):3037–3046, 2012.
- [20] A. Rico-Alvarino and R.-W. Heath. Learning based link adaptation in multiuser MIMO-OFDM. In *Proc. of IEEE EUSIPCO*, 2013.
- [21] W. Shen, Y. Tung, K. Lee, K. Lin, S. Gollakota, D. Katabi, and M. Chen. Rate adaptation for 802.11 multiuser MIMO networks. In *Proc. of ACM MobiCom*, 2012.
- [22] Z. Shen, R. Chen, J. Andrews, R. Heath, and B. Evans. Low complexity user selection algorithms for multiuser MIMO systems with block diagonalization. *IEEE Trans. on Signal Proc.*, 54(9):3658–3663, 2006.
- [23] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong. Argos: practical many-antenna base stations. In *Proc. of ACM MobiCom*, 2012.
- [24] H. Shirani-Mehr, H. Papadopoulos, S.-A. Ramprasad, and G. Caire. Joint scheduling and ARQ for MU-MIMO downlink in the presence of inter-cell interference. *IEEE Transactions on Communications*, 59(2):578–589, 2011.
- [25] J. Sun and et.al. Overview on RBIR-based PHY abstraction. July 2014.
- [26] A. Tarighat, M. Sadek, and A.-H. Sayed. A multi user beamforming scheme for downlink MIMO channels based on maximizing signal-to-leakage ratios. In *Proc. of IEEE ICASSP*, 2005.
- [27] A. Wiesel, Y. Eldar, and S. Shamai. Zero-forcing precoding and generalized inverses. *IEEE Trans. Sig. Proc.*, 56(9):4409–4418, 2008.
- [28] X. Xie, X. Zhang, and K. Sundaresan. Adaptive feedback compression for MIMO networks. In *Proc. of ACM MobiCom*, 2013.
- [29] X. Zhang, K. Sundaresan, M. A. Khojastepour, S. Rangarajan, and K. G. Shin. NEMOx: Scalable network MIMO for wireless networks. In *Proc. of ACM MobiCom*, 2013.