Consistently High MIMO Rates via Switched-beam Antennas

Yonglong Zhang and Konstantinos Psounis, *Fellow, IEEE*

Abstract—The demand for wireless bandwidth is rising to unprecedented levels. The industry has responded with the inclusion of advanced PHY techniques, most notably multi-user (MU) MIMO, in the most recent WiFi and LTE standards. However, despite the theoretical promise for large multiplexing gains, in practice the rate gains are modest due to a combination of large overhead to collect channel state information and not-so-well-conditioned channel matrices.

In this paper we propose to replace omni-directional antennas with inexpensive switched-beam antennas to produce well-conditioned channel matrices for MU-MIMO purposes with very low overhead. Remarkably, experimental results with both software defined radios and commercial WiFi chipsets show that, when appropriate antenna modes are used, this leads to a 3.5x-5x average throughput improvement in indoor environments. What is more, our backward compatible protocol extension coupled with an efficient algorithm to select appropriate antenna modes, achieve the aforementioned gains with almost zero overhead.

I. INTRODUCTION

The rapid increase in the quantity and capability of consumer mobile wireless devices has accelerated the growth in demand for wireless bandwidth tremendously. The US government has responded to this wireless bandwidth crunch with an effort to release new spectrum for wireless communications and promote spectrum sharing technologies. The industry has responded with including in the latest standards novel, highly promising PHY techniques. Most notably, MU-MIMO was included in the latest WiFi and LTE standards, 802.11ac and LTE-Advanced, and implemented in the 2nd wave of 802.11ac chipsets.

In theory, MU-MIMO offers significant spatial multiplexing gains: a transmitter with n antennas may concurrently transmit to n users, yielding an n times performance gain. In practice however, this requires the collection of instantaneous channel state information (CSI) from the n users and a corresponding channel matrix which is well-conditioned [1], [2]. Since it is rarely the case that a randomly selected set of n users will yield a well-conditioned channel matrix, more than n users need to be probed [3], [4]. But, unfortunately, the overhead to collect instantaneous CSI is so large that the industry has settled with probing n users only, and, in the likely event that

the resulting channel matrix is not well-conditioned, go for a smaller than n multiplexing gain.

In this paper we propose to use inexpensive switched-beam antennas in place of typical omni-directional antennas at transmitters to pre-condition the channel and ensure that the channel matrix of *any* random selection of users is well conditioned. Pre-conditioning the channel with analog front ends is not a new idea, see, for example, [5], [6]. However, we do so using inexpensive antennas that can find their way in commercial WiFi access points (APs), and, more importantly, not only conduct extensive experiments using software defined radios (SDRs) and commercial chipsets showcasing under real world conditions that we can achieve surprisingly large gains thanks to this pre-conditioning (3.5x-5x on average in indoors environments depending on the number of antennas) but also present (i) protocol extensions which are compatible with the 802.11ac standard to collect long-term channel statistics with zero overhead, and (ii) efficient algorithms which use those long-term channel statistics to appropriately configure the switched-beam antennas such that gains are materialized in practice.

The rest of the paper is organized as follows. Section II discusses prior work, Section III motivates this work by showing how large the gap can be when using omnidirectional versus directional antennas for indoors MU-MIMO, Section IV discusses how to efficiently select the antenna configurations to achieve the 3.5x-5x gains, Section V presents extensive experiments with SDRs and commercial devices, and Section VII discusses how to collect long-term channel statistics, used by the antenna configuration algorithms, in an 802.11ac backward compatible manner and with zero overhead, and what is the effect of mobility on performance.

II. PRIOR WORK

There is a large body of work on switched-beam antennas, ranging from traditional Butler matrix antennas [7], [8] to more compact ones [9], [10] with varying directionality properties and price.

Researchers first focused on the use of switched-beam antennas in cellular networks [11], [12] to improve the signal to noise ratio (SNR) of SISO transmissions. Recently, researchers have used the directionality of switched-beam antennas to allow concurrent nearby transmissions, e.g. [13] uses switched-beam antennas to minimize inter-cell interference in a multi-cell setup via coordination. Also, switched-beam antennas have found their way to commercial WiFi products [14]. In all these prior instances, switched-beam antennas are used to

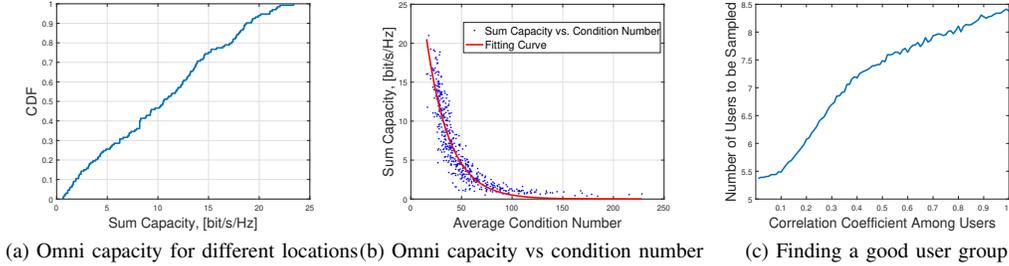(a) Omni capacity for different locations (b) Omni capacity vs condition number (c) Finding a good user group

Fig. 1: MIMO capacity heavily depends on channel matrix $\mathbf{H}$.

improve the SNR thanks to their directionality. In contrast, in our work we use switched-beam antennas to improve the channel matrix in the context of MU-MIMO transmissions, and, as a matter of fact, the line-of-sight directions are not the ones that yield the best improvement.

MU-MIMO can increase wireless capacity sizably thanks to its multiplexing gain [15] and a large body of experimental work has showcased MIMO benefits in various real world setups [16], [17], [18], [19], [20], [21], [22]. To achieve the large multiplexing gain in practice, the MU-MIMO channel matrix needs to have a small condition number [1], [2], [23], [24]. Motivated by this, in recent years researchers have proposed various approaches to achieve a channel matrix with a small condition number. The most classic approach is to choose the best subset from a large group of users in each round of MU-MIMO transmission such that the resulting channel matrix is well conditioned and the sum data rate of users is maximized, see, for example, [3], [25], [4]. However, not only is it NP-hard to select the best user group, but also the overhead from collecting instantaneous CSI from all these users is prohibitively large. Thus, while greedy approaches have been proposed to select a good enough user group [25], [26], the fact of the matter is commercial chipsets collect CSI from the minimum possible number of users and it is very unlikely that a large number of users will ever be sampled for CSI purposes in practice, due to the significant overhead.

Another approach to get a good channel matrix may be to pre-condition the channel using directional antennas. For example, the authors in [17] have observed that directional antennas may yield higher MIMO performance not thanks to higher signal strength, but rather thanks to a better ensuing channel. However, they didn't proceed to design a scheme that would collect CSI information and then select the best directions using this information. Other recent work on directional antennas has considered the use of reconfigurable phased-array antennas in conjunction with MIMO to either reduce the number of RF chains [5] or to reduce inter-cell interference in the context of a multi-cell environment [27], [6] with [6] also making the observation that directionality helps to decorrelate users within a cell and thus achieve a better channel matrix. However, reconfigurable phased-array antennas are too expensive and large, especially for commercial WiFi systems which is the focus of our work, and, we consider instead inexpensive switched-beam antennas which leads to a fundamentally different antenna configuration

problem which does not require any signal processing due to the small number of predetermined modes that switched-beam antennas offer. Finally, in [28] the researchers work with a theoretical model for a "switched-beam based" antenna array which assumes very sharp beams and no inter-beam interference. In practice this requires an antenna array with a very large number of antenna elements and it is completely different from real world switched-beam antennas of the type that we consider. Other major differences from this work is that it focuses on millimeter wave massive MIMO systems (whereas our algorithm is designed for 802.11 devices which have less tightly focused beam pattern and much less number of Tx antennas), and it does not consider the condition number as a factor.

Last, even though no prior work has attempted to build a 802.11-compatible system to configure and use switched-beam antennas for MIMO channel pre-conditioning, there are works that have discussed metrics to select "good" directions of directional antennas or "good" user groups. We do consider such metrics, e.g. the signal-to-leakage ratio and the user orthogonality, and establish that they lead to suboptimal selection of directions in our setup, further motivating our work.

III. MOTIVATION

In this section, we present experimental results where a single AP transmits to a number of users using MU-MIMO over OFDM. We highlight scenarios where omni-directional antennas yield bad performance whereas switched-beam antennas can significantly increase the throughput thanks to channel pre-conditioning.

Recall that in theory and under appropriate conditions an AP with n antennas can transmit to up to n users' antennas concurrently on the same frequency band by precoding the transmitting signal based on CSI. Each user antenna receives its own signal while interference (from signals for other user antennas) is cancelled, yielding a spatial diversity multiplexing gain of n.

A. Omni-mode can be bad

Today's MU-MIMO-enabled APs are mostly equipped with omni-directional antennas. Due to size limitations of commercial APs, those omni-directional antennas are tightly placed in a small area and end up with high spatial correlation [29]. As a result, the correlation of the eventual channel matrix $\mathbf{H}$ will

be significantly increased (see, for example, the widely used Kronecker model), especially when users are located close to each other, severely affecting the capacity of the MIMO channel.

To see this experimentally, we conduct 4x4 MU-MIMO communication between one AP with 4 omni-directional antennas and 4 users (we use software radio platform WARP v3 to act as both AP and user) in a typical office environment, each equipped with one omni-directional antenna. (For more details on experiment topology and settings see Section V.) We measure the channel capacity based on the receivers' effective SNRs [1] and constantly change their locations while maintaining the same distance towards the AP in such a way that all users maintain LOS and constant received signal strength (RSS) from the AP. As shown in Figure 1a, the system capacity has a huge variation when users are on different locations even though they are under the same RSS level. For example, when users are located closely in the room, we measured a sum capacity of only 1.7 bits/s/Hz, which is nearly one fifth of the average capacity over all tested topologies.

B. Condition number matters

The condition number of an n-by-n channel matrix $\mathbf{H}$ is defined as the ratio of the largest to the smallest singular value [15]. For a channel matrix $\mathbf{H}$, its condition number is directly determined by the channel correlation coefficient [30], and it is a good indication of the multipath richness of the channel [1].

A channel matrix that has a low condition number (often referred as "well-conditioned") benefits MIMO transmissions in two ways: (i) it implies a higher effective channel gain for each user [15], [3], and (ii) it reduces the effect of the error caused by using a noisy estimate of the channel to perform precoding, as illustrated in the next paragraph. (Note that this is directly related to the notion of condition number in the context of numerical analysis, which is used to measure how sensitive a system is to the imprecision in the input and how much imprecision in the output results from it [31]). As a result of these two reasons, the lower the condition number, the higher the capacity that the channel can support [2], [23]. [2]

To better illustrate the effect of the condition number on error propagation, we consider two matrices:

$$\mathbf{A} = \begin{bmatrix} 1 & 4 \\ 2.1 & 2 \end{bmatrix} \qquad \mathbf{B} = \begin{bmatrix} 1 & 2 \\ 2.1 & 4 \end{bmatrix},$$

which stand for two ground-truth 2-by-2 MIMO channel matrices (for simplicity we use real rather than imaginary numbers). Note that both matrices have the same power gain,

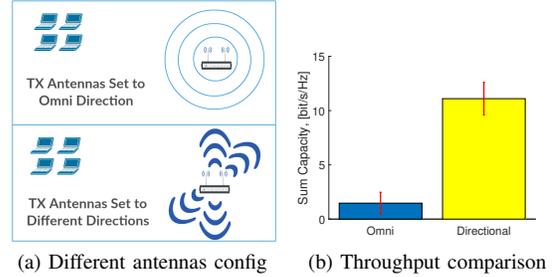

(a) Different antennas config (b) Throughput comparison

Fig. 2: Changing antenna direction yields higher capacity.

while matrix $\mathbf{A}$ has a lower condition number than $\mathbf{B}$. Now, we add the same amount of noise to both matrices and obtain:

$$\mathbf{A_{noise}} = \begin{bmatrix} 1.1 & 4.1 \\ 2.2 & 2.1 \end{bmatrix} \qquad \mathbf{B_{noise}} = \begin{bmatrix} 1.1 & 2.1 \\ 2.2 & 4.1 \end{bmatrix},$$

which represent the measured, noisy channel matrices. Recall that in Zero-forcing Beamforming (ZFBF, the most popular MU-MIMO precoding scheme), we compute the pseudo-inverse of the measured channel matrix and use it to precode the transmission signal, hoping to obtain an identity matrix after the signal propagates through the actual channel. We compute the result of this operation for both matrices and get:

$$\mathbf{A}\mathbf{A_{noise}^{-1}} = \begin{bmatrix} 0.999 & -0.044 \\ -0.001 & 0.955 \end{bmatrix} \mathbf{B}\mathbf{B_{noise}^{-1}} = \begin{bmatrix} 2.727 & -0.909 \\ 1.727 & 0.091 \end{bmatrix}.$$

Observe that while $\mathbf{A}\mathbf{A_{noise}^{-1}}$ is very close to the identity matrix and both users may achieve a good effective SNR. (From the definition of the effective SNR in the first footnote[1] and assuming no background noise, we get $10 \log_{10} \frac{0.999^2}{(-0.044)^2} = 27dB$ and $10 \log_{10} \frac{0.955^2}{(-0.001)^2} = 56dB$ respectively.) $\mathbf{B}\mathbf{B_{noise}^{-1}}$ is very far from the identity matrix and the MIMO transmission would fail due to too much cross interference (e.g. the achieved SNRs of the two users are only 9dB and -26dB respectively). It is evident that the effect of the same level of noise in the case of the matrix with larger condition number is disproportionally larger.

We next perform a total of 500 MU-MIMO communication measurements to show the relationship between the condition number of $\mathbf{H}$ and the channel capacity. We take the average of all the 64 OFDM sub-carriers' condition numbers [3] and plot them in Figure 1b where we use an exponential function to fit the individual measurements. It is evident that the condition number has a drastic effect on the channel capacity.

C. Finding a good user group is expensive

To get a good condition number an AP needs to collect CSI from a number of users and select a good user group. Consider an AP with n omni-directional antennas. The AP may randomly pick n users and examine their CSI to form the n-by-n channel matrix $\mathbf{H}$. If $\mathbf{H}$ is ill-conditioned, the AP may pick another user, examine its CSI and try all the $\binom{n+1}{n}$ combinations to see if it is possible to get a well-conditioned

---

[1] In the context of a MIMO transmission, the *effective SNR* of a specific user is defined as the user's "useful" signal strength over noise which consists of 1) the signal for other users that has not been entirely canceled though MIMO due to real world imperfections and 2) the usual background noise. The former would be much larger than the latter and will greatly affect the transmission when the channel is ill-conditioned as will be illustrated later in Sec. III-B.

[2] To accommodate non-square matrices, the condition number of $\mathbf{W} = \mathbf{H}\mathbf{H}^\dagger$ is sometimes used where $\dagger$ denotes the conjugate-transpose operation. When $\mathbf{H}$ is a square matrix the condition number of $\mathbf{W}$ equals the square of that of $\mathbf{H}$, and, minimizing the non-negative condition number of $\mathbf{W}$ is equivalent to minimizing that of $\mathbf{H}$.

[3] Unless otherwise specified, in this paper we assume MU-MIMO transmissions on a typical 20MHz channel under 802.11ac standard. Like in 802.11a/g/n, each 20MHz channel in 802.11ac contains 64 OFDM subcarriers.

channel matrix with this new user and $n-1$ of the old users. If not, the AP may pick yet another user and so on and so forth.

Finding the best user group among many users is NP-hard and a number of greedy approaches have been suggested. But the main issue is not the computational cost, it is the overhead of collecting the CSI for all those users. As a concrete example, today's WiFi chipsets implement the so-called "explicit feedback" mechanism which involves the AP transmitting channel sounding symbols to the users, followed by each user sending one by one its channel estimation back to the AP under the lowest data rate. This is so expensive that 4x4 MU-MIMO WiFi chipsets randomly select 4 users and settle with a user group of cardinality 3 or less if the resulting channel matrix is not full rank. This means the AP would rather not take advantage of the maximum multiplexing gain than engage in collecting additional CSI from more users. [4]

We use a modified simulator based on [33], where we employ the Kronecker model to generate channel matrices using the formula $\mathbf{H} = \mathbf{R}_R^{\frac{1}{2}} \widetilde{\mathbf{H}} \mathbf{R}_T^{\frac{1}{2}}$ ($\mathbf{R}_R$ and $\mathbf{R}_T$ are correlation matrices determined by the spatial correlations among transmitter antennas and receiver antennas, respectively, while $\widetilde{\mathbf{H}}$ has $\mathcal{CN}(0,1)$ entries, see [33] for more details), to compute the expected number of users to be examined before a well-conditioned subset can be found in the context of 4x4 MU-MIMO. We use measured data from [29] to set the correlation coefficient of $\mathbf{R}_T$ to a fixed value and vary the correlation coefficient of $\mathbf{R}_R$, effectively simulating a wide range of real world conditions. Figure 1c plots the number of users to be sampled to achieve a condition number of 25 or less (which, based on Figure 1b yields good rates), as a function of the correlation coefficient among users. It is evident that even with low user correlation the AP has to examine a handful of additional users before it can find a good subset, which increases the overhead sizably.

Last, note that there are cases where improving the condition number by sampling more users is not even an option as there might not be any more users to transmit to. With the trend to design wireless networks with smaller and smaller cells, this case will become even more likely. Concluding, attempting to correct a channel matrix by sampling additional users is very expensive, and sometimes it may not even be an option.

D. Directionality to the rescue

Consider the scenario illustrated in Figure 2. We replace the 4 antennas with compact switched-beam antennas with 9 modes (8 directional modes, and an omin-directional one, which, as a matter of fact, is the mode we used for the omni-directional experiments above) in the same experimental setting as the one used in Section III-A and select the directional modes illustrated in the lower Figure 2a. Figure 2b shows that we get 6x the performance of the omni mode by using a carefully selected antenna direction configuration while maintaining the same level of RSS for all users.

---

[4]Instead of explicit feedback, in theory one may use the so called implicit feedback mechanism which has less overhead. But, because in this case the system needs to be calibrated [32], the WiFi industry has rejected its use for practical reasons.

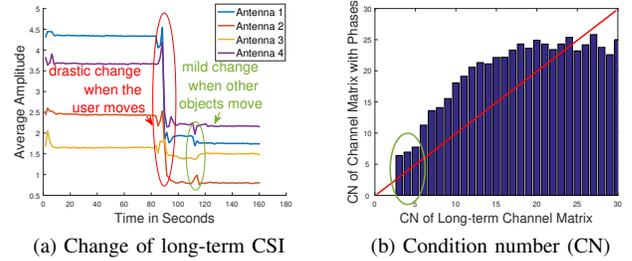(a) Change of long-term CSI   (b) Condition number (CN)

Fig. 3: Properties of long-term CSI.

While the rest of the paper discusses important issues such as how and at what cost one may select the right directional modes, it is interesting to note that the best modes are not necessarily the ones corresponding to a line-of-sight. Also, the channel pre-conditioning achieved by the antennas is apparently enough to yield the maximum multiplexing gain even in the most correlated scenarios. Further, as we will establish later, any random set of 4 users with appropriate directional modes can achieve near-optimal rates, the results hold even in the case of 8x8 MIMO channels (which are of interest given that the maximum number of antennas in the 802.11ac standard is 8), and, nomadic user mobility is not a problem either as the timescale of indoors mobility is much slower than the timescale of updating the long term CSI information of a user that moves to a new location. Last, note that due to the nature of MU-MIMO, the beam patterns do not need to be super directional or "sharp", as the purpose of changing antenna direction is to get a better channel when the channel is ill-conditioned, thus not only we have some room when selecting the right modes (a sizable number of sets of modes would do the trick) but also there is no performance benefit from using more directional and thus much more expensive antennas, such as phased-array antennas.

IV. SYSTEM DESIGN

Consider an MU-MIMO enabled AP with n switched-beam antennas, each having d modes (directions). Without loss of generality, we assume that there are a total of m users in the network ($m \geq n$), and each user is equipped with one omni-directional antenna. At each transmission the AP may serve n users concurrently.

With n antennas each having d modes, the AP has a total of d^n possible antenna configurations to choose from. This gives us d^n possible channel realizations for every n users being served. To improve channel capacity, our goal is to find the antenna configuration that results in the lowest channel condition number for a given group of n users with as little overhead as possible.

A. The long-term CSI matrix

We use long term channel statistics, that is, MIMO channel statistics that are relatively steady over a long period of time, to determine the best antenna configuration. (See later in this

section for the precise long-term statistics that we use.) This way we keep the overhead low while ensuring a high mode selection accuracy. Denote by $\mathbf{G}$ the m by $(d \times n)$ matrix where the ith row contains the long-term statistics of user i towards every transmit antenna for every direction:

$$\mathbf{G} = \overbrace{\begin{bmatrix} g^1_{1,1} & g^2_{1,1} & \cdots & g^d_{1,1} & g^1_{1,2} & \cdots & g^d_{1,2} & \cdots & g^1_{1,n} & \cdots & g^d_{1,n} \\ g^1_{2,1} & g^2_{2,1} & \cdots & g^d_{2,1} & g^1_{2,2} & \cdots & g^d_{2,2} & \cdots & g^1_{2,n} & \cdots & g^d_{2,n} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\ g^1_{m,1} & g^2_{m,1} & \cdots & g^d_{m,1} & g^1_{m,2} & \cdots & g^d_{m,2} & \cdots & g^1_{m,n} & \cdots & g^d_{m,n} \end{bmatrix}}^{\text{Antenna 1} \quad \text{Antenna 2} \quad \text{Antenna } n} \quad (1)$$

where $g^k_{i,j}$ represents the long-term channel gain between user i and antenna j, given that antenna is set to direction k.

We define and compute the long-term channel gain as follows: let $h^{k,c}_{i,j} = |a^{k,c}_{i,j}|e^{j\omega}$ represent the instantaneous channel gain between user i and antenna j on subcarrier c, given that the transmit antenna is set to direction k. This contains both amplitude and phase information of the channel. We ignore the phase ω since it changes very fast and focus on the amplitude $|a^{k,c}_{i,j}|$ which remains relatively steady as long as neither of the antennas change their physical location. Thus, we take the average amplitude of all OFDM subcarriers (as defined 802.11ac standard) between user i and antenna j on direction k, and save it as $g^k_{i,j}$. Without loss of generality consider 20MHz channels where 52 of the total 64 are data subcarriers and we have $g^k_{i,j} = \frac{\sum_{c=1}^{52} |h^{k,c}_{i,j}|}{52}$.

To verify the time-invariant nature of the long-term gain, we perform an experiment where we constantly measure the elements of $\mathbf{G}$ for 3 minutes in a typical office room. In the first minute and a half the environment is relatively steady: about 6 people work on their desks. In the middle of the experiment we change the location of the user by about 1 meter. In the second half of the experiment, the 6 people begin to walk around in the office, opening and closing doors. As shown in Figure 3a, the long-term channel gains are quite steady during the first half of the experiment. In the second half, the long-term gains vary a little bit due to moving objects, but within a small magnitude (less than 10%). Last, as expected, there is a sizable change when the user/antenna moves. Concluding, we confirm that in a typical office environment the entries in $\mathbf{G}$ stay valid for a relatively long time since the location of laptops, tablets, etc. changes in the order of minutes [34], [35], [36].

B. Formulating the problem of finding the best antenna configuration

We address how to populate and keep on updating the $\mathbf{G}$ matrix in Section VII. Here we assume that the AP has access to the $\mathbf{G}$ matrix and solve the following problem. Let $\mathbf{G}_{\boldsymbol{u},\boldsymbol{d}}$ denote an n by n submatrix of $\mathbf{G}$ formed by the rows whose indices are dictated by vector $\boldsymbol{u}$ and the columns whose indices are dictated by vector $\boldsymbol{d}$. For a given subset of users with user indices $\boldsymbol{u} = (u_1, u_2, \ldots, u_n)$, find a vector of column indices $\boldsymbol{d} = (d_1, d_2, \ldots, d_n)$ where $d_i \in [(i-1)n+1, in]$ for all $i \in [1, n]$, such that the condition number of $\mathbf{G}_{\boldsymbol{u},\boldsymbol{d}}$ is minimized. That is,

$$\boldsymbol{d}^* = \underset{(d_1,\ldots,d_n):d_i \in [(i-1)n+1,in],\forall i}{\arg\min} \kappa(\mathbf{G}_{(u_1,u_2,\ldots,u_n),(d_1,d_2,\ldots,d_n)}) \quad (2)$$

where $\kappa(\cdot)$ denotes the 2-norm condition number of a matrix. Note that the "feasibility" constraint $d_i \in [(i-1)n+1, in]$ ensures that we assign exactly one direction to each antenna and we will refer to vectors $\boldsymbol{d}$ that satisfy this as "feasible".

Such column selection problems are NP-hard even without the feasibility constraint [37]. In our setting, an AP with 4 antennas and 9 modes per antenna yields a modest 6561 possible antenna configurations, but, raising the antennas to 8 (max number of antennas on 802.11ac) raises the number of possible configurations to 43 million. Thus, using brute-force search algorithm to find $\boldsymbol{d}^*$ is not appealing.

C. Greedy algorithms for the best antenna configuration problem

As already discussed, there is a large body of work on greedy approaches to improve channel matrices. For example, multiple researchers including [3] proposed different greedy user selection algorithms to maximize users' orthogonality, while some other authors proposed to maximize the system's signal-to-leakage ratio (SLR). We apply these approaches to our problem and show (in Section V) that they fail to select very good modes. Motivated by this, we propose a novel algorithm based on singular value and orthogonal-triangular decomposition. We will refer to this algorithm as the Condition Number SVD (CN-SVD) algorithm.

Before we describe the CN-SVD algorithm, we briefly describe how to apply to our problem the two prior-work approaches mentioned above. First, the orthogonality of two vectors $\boldsymbol{h}$ and $\boldsymbol{g}$ is determined by the ratio $\frac{|\boldsymbol{h}\boldsymbol{g}^*|}{\|\boldsymbol{h}\|\|\boldsymbol{g}\|}$. (The smaller the ratio, the more orthogonal the two vectors are.) Using this approach, the goal is to find the feasible direction vector $\boldsymbol{d}$ such that for the given set of users $\boldsymbol{u}$, the columns of indices in $\boldsymbol{d}$ are as orthogonal to each other as possible. This approach, applied to the user grouping problem under a very large number of choices has been shown to work well, but when applied to our directional antenna mode selection problem which has much more limited number of feasible choices it does not perform well, see Section V. Second, the algorithm based on SLR maps each transmit antenna i to a unique user j. As a result, for user j, only the signal sent from its designated antenna i is considered useful while signals sent from other antennas are considered leakage, and the SLR of user j defined as $SLR_{i,j} = \frac{P_{i,j}}{\sum_{k \neq i} P_{k,j}}$, where $P_{i,j}$ is the received power from antenna i to user j. Then, the direction vector $\boldsymbol{d}$ that yields the maximum system-wide SLR is selected. Section V shows that this approach does not perform well in our setting either. We conjecture this is because in indoor environments MU-MIMO benefits more from a well-conditioned channel than from stronger signal strength, as we will show later in Section V.

The CN-SVD algorithm: The CN-SVD algorithm uses a subroutine, which is based on singular value and orthogonal-triangular decomposition, for selecting the "best" k columns from a large matrix in polynomial time. The subroutine starts by performing a singular value decomposition of a scaled (see below for more details) version of $\mathbf{G} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^*$ and lets $\mathbf{W}$ be the submatrix of $\mathbf{V}$'s first k columns. Then, an orthogonal-triangular (QR) decomposition of $\mathbf{W}^\intercal$ is performed such that

$\mathbf{W}^\mathsf{T}\mathbf{E} = \mathbf{QR}$, where $\mathbf{Q}$ is unitary, $\mathbf{R}$ is upper triangular and $\mathbf{E}$ is a permutation matrix designed to make the absolute value of diagonal elements of $\mathbf{R}$ decreasing. Finally, it outputs a vector consisting of the column indices of non-zero elements of $\mathbf{E}$'s first n rows. We denote the above subroutine by $\mathrm{ALG}(\mathbf{A}, k)$ where $\mathbf{A}$ is the input matrix and k is the number of columns we want to select from $\mathbf{A}$.

Algorithm 1 below presents the subroutine using pseudo-code, where $\mathrm{SCALE}(\mathbf{A})$ divides each entry of a matrix $\mathbf{A}$ by its column's root-sum-of-squares (i.e. by $\sqrt{\sum_i |a_{ij}|^2}$ for each column j), $\mathrm{SVD}(\mathbf{A})$ stands for the singular-value decomposition of $\mathbf{A}$, and $\mathrm{QR}(\mathbf{A}, k)$ represents the QR decomposition described above on the first k rows of $\mathbf{A}$. At last, the algorithm shall return a vector of the selected column indices $(c_1, c_2, \ldots, c_k)$. More details of how and why this subroutine works can be found in [38], [39].

Algorithm 1 The *ALG* Algorithm

---

procedure $\mathrm{ALG}(\mathbf{G}, k)$
    $\mathbf{A} \leftarrow \mathrm{SCALE}(\mathbf{G})$
    $(\mathbf{U}, \mathbf{S}, \mathbf{V}) \leftarrow \mathrm{SVD}(\mathbf{A})$
    $(\mathbf{Q}, \mathbf{R}, \mathbf{E}) \leftarrow \mathrm{QR}(\mathbf{V}^\mathsf{T}, k)$
    for $i = 1 \ldots k$ **do**
        $c_i \leftarrow \arg\max_j \mathbf{E}_{i,j}$
    end for
    return $(c_1, c_2, \ldots, c_k)$
end procedure

The output vector from the subroutine above may contain multiple columns (modes) corresponding to a single antenna and no columns corresponding to some of the antennas. For this reason, the CN-SVD algorithm uses a combination of a recursive application of the ALG subroutine to select at least one mode for each antenna, and a greedy step to select one of the multiple modes of those antennas which have multiple modes selected.

Look at the pseudo-code of Algorithm 2 below. Let $\boldsymbol{a}$ denote the set of all Tx antennas and $\boldsymbol{u}$ denote the set of users to which the AP is transmitting. Denote by $\boldsymbol{d}_f$ the set of Tx antennas' finalized directions/modes, which the algorithm populates as it settles on which mode to use for each antenna in $\boldsymbol{a}$. $\mathrm{GETCOLUMNS}(\boldsymbol{a})$ extracts the set of indices of columns (of matrix $\mathbf{G}$) belonging to the antennas in set $\boldsymbol{a}$. $\mathrm{CNSVD}(\boldsymbol{a}, \boldsymbol{u}, \mathbf{G}, \boldsymbol{d}_f)$ starts with calling subroutine ALG on the submatrix of $\mathbf{G}$ corresponding to the users $\boldsymbol{u}$ and the directions/columns of the antennas in $\boldsymbol{a}$. The subroutine selects columns and yields three sets of antennas, $\boldsymbol{a}_{us}$, $\boldsymbol{a}_{ms}$, and $\boldsymbol{a}_{ns}$, corresponding to the set of antennas with a unique selected mode, with multiple selected modes, and with no selected mode. (We use $\mathrm{GETUNIQUESELECTANT}(\boldsymbol{d})$ and $\mathrm{GETMULTISELECTANT}(\boldsymbol{d})$ to extract the sets $\boldsymbol{a}_{us}$, $\boldsymbol{a}_{ms}$ from $\boldsymbol{d}$.) For antennas in $\boldsymbol{a}_{us}$, their modes can thus be fixed and stored in $\boldsymbol{d}_f$, and if $\boldsymbol{a}_{ns}$ is not empty CNSVD is recursively called till every Tx antenna has been assigned to a mode. Last, getting back from the recursions, for each antenna in $\boldsymbol{a}_{ms}$ that have multiple selected modes, we greedily select the one which minimizes the condition number of the submatrix which consists of the already finalized modes in $\boldsymbol{d}_f$ and this mode.

Algorithm 2 The *CN-SVD* Algorithm

---

procedure $\mathrm{CNSVD}(\boldsymbol{a}, \boldsymbol{u}, \mathbf{G}, \boldsymbol{d}_f)$
    $\boldsymbol{c} \leftarrow \mathrm{GETCOLUMNS}(\boldsymbol{a})$
    $\boldsymbol{d} \leftarrow \mathrm{ALG}(\mathbf{G}_{\boldsymbol{u},\boldsymbol{c}}, |\boldsymbol{a}|)$
    $\boldsymbol{a}_{us} \leftarrow \mathrm{GETUNIQUESELECTANT}(\boldsymbol{d})$
    $\boldsymbol{a}_{ms} \leftarrow \mathrm{GETMULTISELECTANT}(\boldsymbol{d})$
    $\boldsymbol{a}_{ns} \leftarrow \boldsymbol{a} \setminus (\boldsymbol{a}_{us} \cup \boldsymbol{a}_{ms})$
    $\boldsymbol{d}_f \leftarrow \boldsymbol{d}_f \cup (\boldsymbol{d} \cap \mathrm{GETCOLUMNS}(\boldsymbol{a}_{us}))$
    if $\boldsymbol{a}_{ns} \neq \emptyset$ **then**
        $\boldsymbol{d}_f \leftarrow \mathrm{CNSVD}(\boldsymbol{a}_{ns}, \boldsymbol{u}, \mathbf{G}, \boldsymbol{d}_f)$
    else
        return $\boldsymbol{d}_f$
    end if
    for all $a \in \boldsymbol{a}_{ms}$ **do**
        $d^* \leftarrow \displaystyle\arg\min_{d \in \boldsymbol{d} \cap \mathrm{GETCOLUMNS}(\{a\})} \kappa(\mathbf{G}_{\boldsymbol{u}, \{d\} \cup \boldsymbol{d}_f})$
        $\boldsymbol{d}_f \leftarrow \boldsymbol{d}_f \cup \{d^*\}$
    end for
    return $\boldsymbol{d}_f$
end procedure

D. Long-term gains versus instantaneous CSI

We use long-term channel gains to select antenna modes which yield a low condition number, whereas it is the instantaneous CSI that determines the actual condition number. Thus, it is interesting to investigate the difference among the condition number of the actual channel matrix $\mathbf{H}$ (with phase information) versus that of the $\mathbf{G}_{\boldsymbol{u},\boldsymbol{d}}$ matrix. To do so we start with a measured long-term gain matrix $\mathbf{G}_{\boldsymbol{u},\boldsymbol{d}}$ and add a random phase [40] and some amplitude turbulence to create a corresponding channel matrix $\mathbf{H}$. We do this for all the OFDM sub-carriers and compute the resulting average condition number. Figure 3b plots the condition number of $\mathbf{G}_{\boldsymbol{u},\boldsymbol{d}}$ and $\mathbf{H}$ for 1000 different long-term gain matrices. As conjectured, it is evident from the plot that well-conditioned $\mathbf{G}_{\boldsymbol{u},\boldsymbol{d}}$ matrices are strongly correlated with well-conditioned $\mathbf{H}$ matrices, see the green circle on the plot.

V. EXPERIMENTAL RESULTS

A. Experiments with SDRs

We first conduct experiments with 2 WARPv3 boards, each with 4 RF ports. We use one WARP as a transmitter with 4 switched-beam antennas (Adant Star 160 [41]), while the other WARP acts as 4 independent users by using WARP's ability to separately process each RF chain and by positioning the antennas corresponding to each user at different locations using long cables. Each user is equipped with an omni-directional antenna. We conduct MU-MIMO downlink communication between the AP and the 4 users using ZFBF with explicit feedback (as in 802.11ac).

The experiments are done in a typical office room as shown in Figure 4a. We place users in different locations and measure their SNR and the resulting sum capacity under: (i)the omni-directional mode for all AP antennas (Omni), (ii) directional modes selected by minimizing the condition number using brute-force search (CN-BF), (iii) directions selected by the CN-SVD algorithm (CN-SVD), (iv) directions selected by maximizing the SLR using brute-force search (SLR) and (v) directions selected by maximizing the orthogonality using brute-force search (Orthogonal).
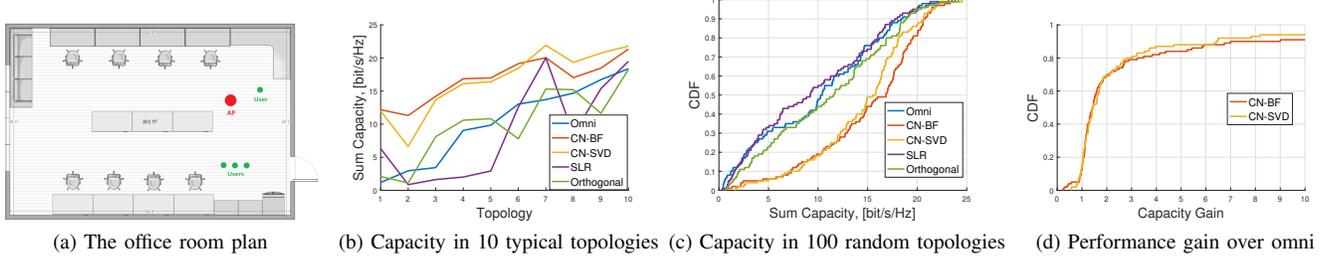
(a) The office room plan (b) Capacity in 10 typical topologies (c) Capacity in 100 random topologies (d) Performance gain over omni

Fig. 4: Results of 10 typical topologies and 100 random topologies.

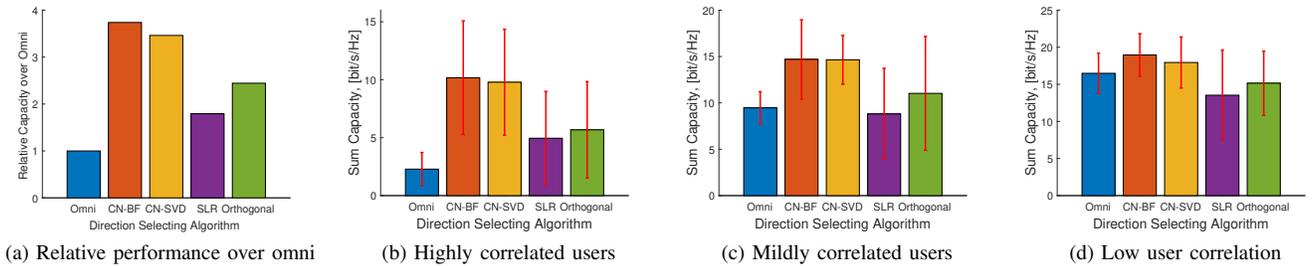(a) Relative performance over omni (b) Highly correlated users (c) Mildly correlated users (d) Low user correlation

Fig. 5: Averaged results of random topologies.

1) Typical topologies: We choose 10 typical topologies of varying spatial user correlation. These topologies range from cases where users are well-separated to cases where all four users are close to each other. Figure 4a shows one typical topology where three users are close to each other. For each topology, we get 10 measurements and report the average. Figure 4b plots the results, where topologies are numbered in decreasing level of spatial user correlation. It is evident that the algorithms which minimize the condition number (CN-BF and CN-SVD) outperform the others, especially when users are highly correlated. Note that the SLR and Orthogonal approaches do not have a steady performance. We discuss the reasons for this in the next subsection.

2) Randomly generated topologies: We compare the performance of the algorithms under 100 random topologies. For each topology we decide the location of each user by uniformly generating two values and using them as coordinates. Like before, for each topology we measure the channel capacity 10 times and compute the average. Figure 4c plots the empirical CDF of those averages, where, like before, the algorithms based on the condition number outperform the others, and, Figure 4d plots the CDF of the relative gain of CN-BF and CN-SVD over Omni. Note that we zoom in the interesting part and don't show the CDF when the gain is more than 10x, since in such scenarios Omni has such a poorly-conditioned matrix that in practice it makes sense to find a well-conditioned submatrix than insisting on transmitting to 4 users concurrently. Last, Figure 5a plots the relative performance of all 4 schemes over Omni, averaged over all 100 topologies. CN-BF and CN-SVD outperform Omni by about 3.5x, while SLR and Orthogonal outperform it by 1.7x and 2.4x respectively.

To better understand the results, we order the topologies in

increasing Omni performance and group them in three cases: (i) bottom third w.r.t Omni performance due to relatively high correlation of users, (ii) middle third due to mild correlation and (iii) top third due to almost no user correlation. Figure 5b plots the sum capacity in the first case. Omni yields a bad condition number but the CN-BF and CN-SVD algorithms search for antenna directions which minimize the channel condition number and achieve 5x the Omni's performance. Figure 5c shows the results when users have mild correlation. In this case Omni performs much better than before, CN-BF and CN-SVD achieve about 1.5x the performance of Omni, and SLR and Orthogonal perform similar to Omni. Figure 5d shows the results when users have almost no correlation. All algorithms perform well since the channel is well-conditioned. [5]

To explain the large fluctuations on the performance of SLR and Orthogonal as well as their lower-than-Omni performance when users are not spatially correlated, we compare the condition number of the $\mathbf{G}_{u,d}$ matrix resulting from selecting directions using the 4 selection algorithms. Figure 6a plots the results for 40 topologies under the 4x4 MU-MIMO setup and reports the average condition number, which is referred to as mean in the legend, for each algorithm. The spikes on the condition number for SLR and Orthogonal cause the performance fluctuations noticed in Figure 4b and the lower-than-Omni performance observed in Figure 5d. Figure 6b plots the results for the same 40 topologies under a 2x2 MU-MIMO setup. In this smaller scale problem, SLR and Orthogonal

---

[5]Note that the reported gains in Fig. 5a represent the average over all topologies of the throughput ratio of a scheme over Omni, as we aim to give equal weight to all topologies. That said, since Omni does particularly bad in some topologies, if one were to first compute the average throughput over all topologies for each scheme and then compute the ratio, the gains are still significant but lower.
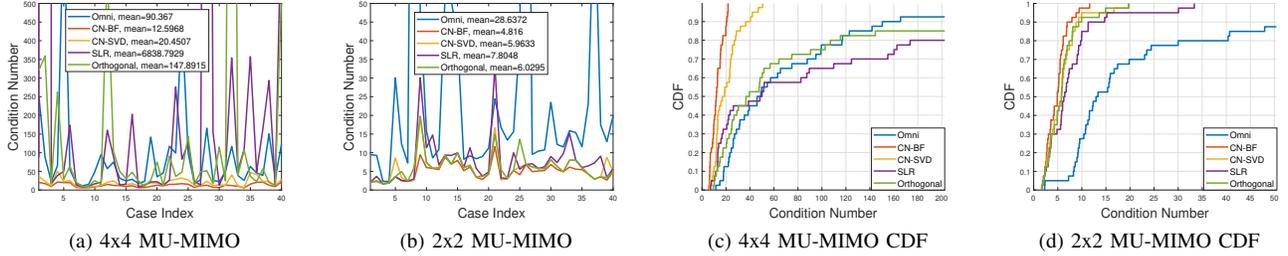
Fig. 6: Condition number of varying number of antennas.

make better directional mode selections and are expected to perform similar to the CN-BF and CN-SVD schemes, which is consistent with the 2x2 MU-MIMO results reported in [6] for SLR.

3) Multiplexing downgrading: An AP may choose to serve less than the maximum number users when the channel matrix is ill-conditioned, such that the extra transmit antennas can be used to increase the *diversity gain* and thus obtain a better-conditioned channel matrix (of lower dimension). For example, downgrading from a multiplexing gain of 4 to 3 causes a 25% multiplexing loss but if the 4x3 sub-matrix is well-conditioned while the original 4x4 is not we may end up with a better throughput overall. To investigate this idea we proceed as follows: if the condition number of the original 4x4 channel matrix is less than a threshold (equal to 30) then we transmit to the 4 users like before. If not, we select the best 3 users to transmit to, by computing the condition number of the ensuing four 4x3 sub-matrices and selecting the one with the smallest condition number, unless no 4x3 sub-matrix has a condition number smaller than the threshold, in which case we select the best 2 users to transmit to by computing the condition number of the ensuing six 4x2 sub-matrices. (We never had to downgrade to a single user only.)

We apply the above procedure to Omni, CN-BF, CN-SVD, and Orthogonal. Also, motivated by real-world chipsets which, in the interest of simplicity, may randomly choose 3 (or 2) out of the 4 users to serve when the original 4x4 channel matrix is ill conditioned, we also consider Omni with such random user selection for downgrading purposes and denote this as Omni-rand. Figure 7a shows the average capacity for all the five schemes over both typical and random topologies as before. We can see that both CN-BF and CN-SVD achieve a sizable improvement over Omni, thought the gains are smaller than when we enforce to serve 4 users utilizing the maximum multiplexing gain.

To better understand the tradeoff between the multiplexing gain and the ensuing condition number, we measure and sort the capacity of 4x4, 4x3 and 4x2 transmissions over all $9^4 = 6561$ antenna configurations for one specific topology in Figure 7b. As expected, 4x2 and 4x3 transmissions are more robust than 4x4 due to the diversity gain from the extra transmit antennas. However a 4x4 transmission may still achieve higher performance when the correct configuration is used thanks to the combination of a low condition number with the highest multiplexing gain. The plot also indicates the

performance of Omni for each of the three multiplexing gain cases for comparison purposes.

4) 8x8 MIMO experiments: We conduct the same set of experiments under an 8x8 MIMO setup, since 802.11ac allows up to 8 antennas at the transmit side. Note that CN-BF can no longer be computed easily so we do not report results for it. Figure 7c reports the throughput averaged over the same typical and random topologies used before. CN-SVD achieves a nearly 5x gain over Omni, higher than the 3.5x average gain in the 4x4 case. Note however that the absolute throughput of the CN-SVD algorithm in the 8x8 case does not increase significantly comparing to the 4x4 case. This is because the larger the multiplexing gain the harder it is to get a channel matrix with a low condition number [42]. The significant spikes in Figure 8a, which shows the condition numbers achieved by different direction-selecting algorithms under the 8x8 setup, is further evidence of the difficulty to get the maximum multiplexing gain in an 8x8 channel. Motivated by this we conjecture that if 8x8 chipsets become available in the future, some of the additional antennas will mostly be used for diversity than for multiplexing gains.

Like in the 4x4 case, we also consider multiplexing downgrading, and, in particular, 8x8, 8x7 and 8x6 MIMO transmissions according to a condition number threshold. As shown in Figure 7d, we observe that the CN-SVD algorithm has a sizable capacity gain over Omni, which, nevertheless, is smaller than when we enforce to serve the maximum possible number of users (8 in this case).

5) Channel pre-conditioning versus power gain: Last, to illustrate the fact that the capacity gains of CV-SVD are mainly caused by the pre-conditioning of the channel matrix and not by a higher Received Signal Strength (RSS), we conduct an 8x8 experiment with two different antenna configurations, one which achieves a low condition number and one which achieves a higher condition number. We perform 5 transmissions with both configurations and record (i) their total capacity and, (ii) their total power gain of the channel (computed by $\mathrm{Tr}[\mathbf{HH}^*] = \sum_{i,j} |h_{ij}|^2$). As shown in Figure 8c, although the second configuration has a slightly higher total power gain it achieves a much lower total MIMO throughput because of the higher condition number.

B. Experiments with commercial devices

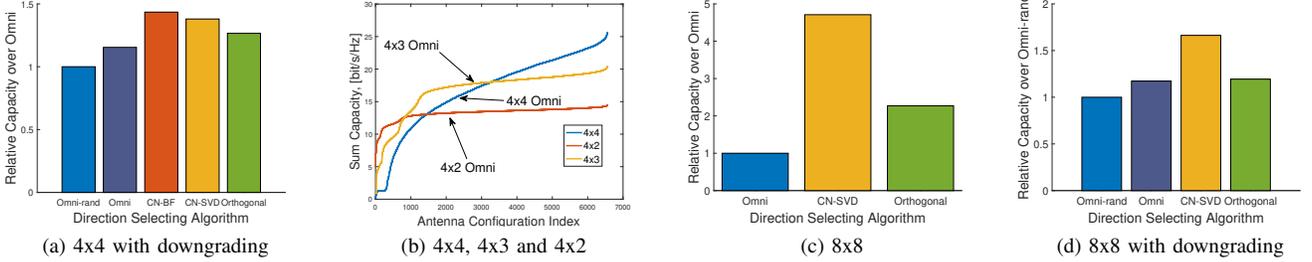To further validate our proposed approach, we conduct experiments using commercial 802.11ac wave 2 devices: a

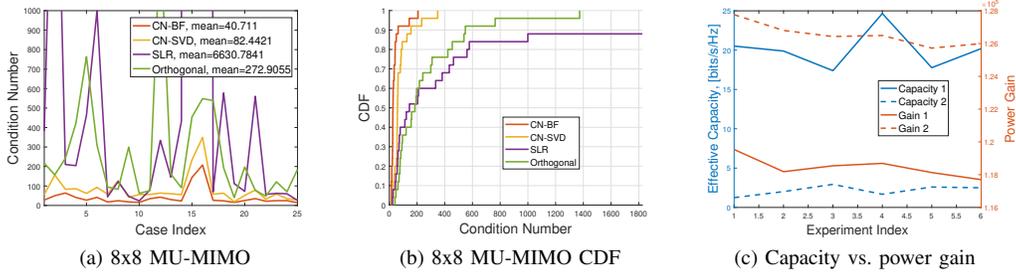Fig. 7: Results of experiments with and without downgrading.

(a) 4x4 with downgrading (b) 4x4, 4x3 and 4x2 (c) 8x8 (d) 8x8 with downgrading

Fig. 8: Condition number determines the outcome of 8x8 transmission.

(a) 8x8 MU-MIMO (b) 8x8 MU-MIMO CDF (c) Capacity vs. power gain

Netgear Nighthawk AP with a Quacomm chipset equipped with 4 Adant Star 160 switched-beam antennas and 3 Xiaomi Mi4i smartphones acting as users. (This chipset is configured to transmit to up to 3 users.) The experiments are performed in a typical office floor shown in Figure 9a. We present results from four typical topologies where users are sometimes located in different and sometimes at the same room. A number of typical antenna configurations are chosen including configurations where antennas point to the same direction, and configurations where they point to different directions. We let the AP perform backlogged MU-MIMO transmissions towards the users for 12 hours in each antenna configuration, and record the achieved throughput of each user. All transmissions are conducted under the generic 802.11ac wave2 standard that is built-in in the chipsets of the commercial devices that we are using.

In Figure 9b we compare the sum rate of the 3 users averaged over time using the omni mode against the highest throughput achieved by one of the preset directional configurations. Due to commercial hardware limitations, we only consider 36 out of the $8^4 = 4096$ total antenna configurations. Thus, it is unlikely that one of them is the one that maximizes the throughput. What is more, transmitting to 3 rather than 4 users protects the omni mode from very poorly-conditioned matrices since it operates on a 4x3 submatrix of the channel matrix (but, of course, also reduces the multiplexing gain from 4 to 3). Still, the system's throughput increases on average by 50%. Last, Figure 9c plots the performance of all 36 antenna configurations including the omni-mode and the configuration which uses the "line-of-sight" directions (main lobes pointing to users), for the topology in Figure 9a. Interestingly the line-of-sight configuration that has the highest antenna gain is not

the one that maximizes the throughput. This observation is consistent with the one we made earlier in Figure 8c and is something that occurs frequently in indoor environments.

VI. EXTENSIONS

A. Joint selection of antenna modes and users

Although we have shown that with the CN-SVD algorithm we can achieve a low condition number with almost any given group of users, it is worth investigating how to select antenna modes and users, which we refer as "cross selection" because essentially we are selecting a subset of columns and a subset of rows from the $\mathbf{G}$ matrix at the same time, such that the resulting submatrix has a low condition number (see Figure 12).

As shown below, our proposed "cross selection" algorithm works as follows: given a set of m users and their long-term channel gain matrix $\mathbf{G}$, we first use the ALG algorithm on $\mathbf{G}^\mathsf{T}$ to find the most independent k out of m users (assuming the AP transmits to k users concurrently). Then, we run the CN-SVD algorithm on those k users, and thus obtain the best antenna modes to minimize the condition number of those k users. For fairness purposes already selected users won't be selected again until all the m users are served once (else we would be always picking the best set of users).

Algorithm 3 The *Cross Selection* Algorithm

---

procedure CROSS($\boldsymbol{a}, \mathbf{G}$)
    $\boldsymbol{u} \leftarrow \text{ALG}(\mathbf{G}^\mathsf{T}, |\boldsymbol{a}|)$
    $\boldsymbol{d}_f \leftarrow \text{CNSVD}(\boldsymbol{a}, \boldsymbol{u}, \mathbf{G}_\mathbf{u}, \emptyset)$
    return $\boldsymbol{u}, \boldsymbol{d}_f$
end procedure

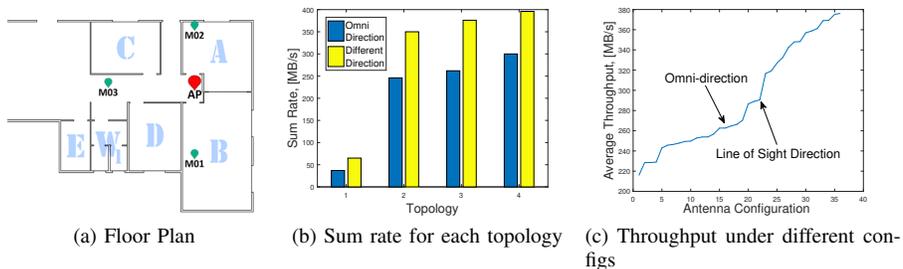We investigate the performance of the CROSS algorithm with simulations using actual channel data collected in the

(a) Floor Plan   (b) Sum rate for each topology   (c) Throughput under different configs

Fig. 9: Experiment performed with commercial hardware.

	Mean CN		Median CN	
Scheme	40	100	40	100
1	17.77	10.08	14.22	8.82
2	29.52	18.10	28.09	18.84
3	257.2	43.91	186.6	47.16
4	108.7	22.38	61.14	16.42
5	23.44	19.99	29.81	22.75
6	696.4	183.4	376.0	142.1

Fig. 10: Statistics of CROSS simulation

(a) 40 users   (b) 100 users   (c) 40 users CDF   (d) 100 users CDF
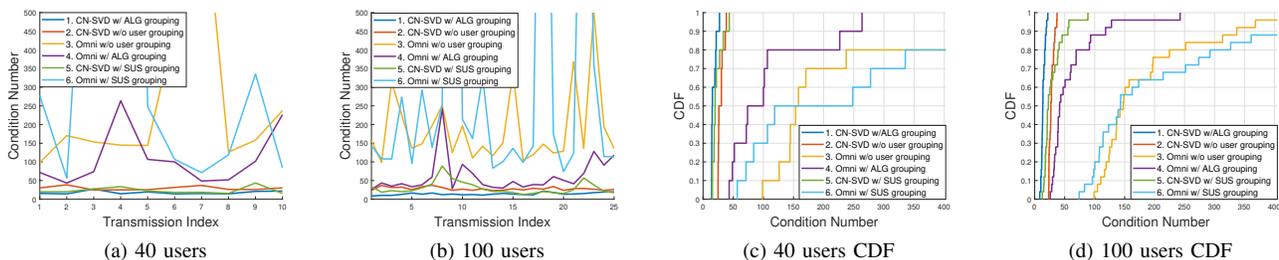
Fig. 11: Condition number achieved by different approaches.

Fig. 12: Example of cross selection: combining user grouping and antenna mode selection.

previous experiments assuming a total of 40 to 100 users and a 4x4 MU-MIMO setup. We compare the condition numbers achieved by 6 different approaches: 1) CN-SVD with ALG user-selection, where we perform both user grouping and antenna selection as proposed in Algorithm 3, 2) CN-SVD without user-selection, which means we randomly select a group of 4 users at each round, and transmit to them under the antenna modes found by the CN-SVD algorithm, 3) Omni without user-selection, where the AP randomly selects 4 users each time and then transmits to them in omni mode, 4) Omni with ALG user selection, in which case we first select the best 4 users from the not-yet-served users as described in the first step of the cross selection algorithm, and then transmit to them using omni mode, 5) CN-SVD with SUS user-selection, where we first do user grouping using the orthogonality method

proposed in [3] and then select antenna modes with the CN-SVD algorithm, and 6) Omni with SUS user grouping, where we first select users based on their orthogonality and then transmit to them using omni antenna mode. For each simulation run, we generate a $\mathbf{G}$ matrix containing 40 or 100 users's data randomly from more than 300 real channel measurements. We then simulate the 6 approaches mentioned above and record the condition numbers achieved by each one, at each transmission (for example, for a total of 40 users there would be 10 transmissions in each experiment to serve all of them 4 at a time, similarly with 100 users there would be 25 transmissions).

We repeat the experiment 1000 times for each case, and report the mean and median condition number over all transmissions in Figure 10. We also plot the average condition number in each transmission and their corresponding CDFs in Fig. 11. As we can see, the cross selection algorithm can further reduce the mean condition number by almost half on top of the plain antenna-mode selection method (CN-SVD without user grouping), while keeping the condition numbers low even for the last transmissions. In contrast, if we only select the users but not the antenna modes (as shown by Omni with ALG grouping and Omni with SUS grouping), the condition numbers are highly unstable, especially when there are fewer users left to select. Last, it is worth noting that the ALG user grouping method yields a lower condition number than SUS irrespectively of whether we perform antenna mode selection afterwards or not. This is expected since the ALG algorithm is specifically designed to reduce the condition number of the channel matrix.

B. Support for OFDMA

The latest 802.11 standard, 802.11ax, allows to use OFDMA, that is, the spectrum can be split in time-frequency resource units (RUs) and the AP can assign subsets of subcarriers to different users, achieving more flexible multiple access. This new feature complicates the antenna mode selection problem for the following reason: suppose the OFDM subcarriers are divided into two RU groups, each group having 4 users, and the AP transmits to all these 8 users concurrently using 802.11ax MU-MIMO. Now, it is likely that each group of users has a distinct optimal antenna configuration. However, the AP can only transmit under one specific antenna configuration, which means we need to find good antenna modes that work well for both RU groups.

To do so we start with a simple approach based on the CN-SVD algorithm: given multiple RUs $u_1, u_2, ..., u_n$ and their corresponding long-term channel gains $\mathbf{G}$, we first find the best antenna configuration d for each RU separately using the CN-SVD algorithm, then from those configurations we select the one that works best for all the RUs (for instance, the one that minimizes the maximum condition number of all RUs), as shown in Algorithm 4.

To investigate the performance of this approach we perform the following simulation: given 2, 4 or 8 RUs each containing 4 users, we compare the average condition number over 100 experiments (referred to as mean in plot legends) achieved by 1) our algorithm, 2) omni-directional antennas and 3) the optimal solution (found by brute-force search). Again we use actual channel measurements collected in the previous experiments to generate the channel matrices of the different RUs. As illustrated in Figure 13, our greedy algorithm performs significantly better than the Omni, and also closely to the optimal solution in all the cases of different number of RUs. Also, it is interesting to note that as the number of RU grows, the achieved average condition number (both the optimal one and the one of our algorithm) increases slightly. We conjecture this is because the more the groups the harder to "satisfy" all of them at the same time with a single antenna configuration.

Algorithm 4 The *CN-SVD for 802.11ax* Algorithm

---

procedure CN-SVDAX($a, u_1, u_2, \ldots, u_n, \mathbf{G}$)
  for $i = 1...n$ **do**
    $d_i \leftarrow \text{CNSVD}(a, u_i, \mathbf{G}, \emptyset)$
  end for
  $\kappa_{minmax} \leftarrow \infty$
  for $i = 1 \ldots n$ **do**
    $\kappa_i = 0$
    for $j = 1 \ldots n$ **do**
      $\kappa_i \leftarrow \max\left(\kappa_i, \kappa\left(\mathbf{G}_{u_j, d_i}\right)\right)$
    end for
    if $\kappa_i < \kappa_{minmax}$ **then**
      $\kappa_{minmax} \leftarrow \kappa_i$
      $d_f \leftarrow d_i$
    end if
  end for
  return d_f
end procedure

It is worth noting that the above approach is obviously not optimal. In the context of 802.11ax, one may envision

to jointly optimize the antenna configuration problem with the problem of user selection, or even with the problem of jointly dividing the channel into multiple RUs and assigning users to them so that the total throughput can be maximized, see [43] for a treatment of RU scheduling. However, the complexity of this problem increases sizably rendering any solution impractical in the context of real world APs.

VII. PROTOCOL SUPPORT FOR 802.11

We propose two protocol extensions by which a commercial 802.11 system can benefit from our direction selecting algorithm. The main purpose of these protocol extensions is to update the $\mathbf{G}$ matrix at minimum overhead. Unless otherwise stated, we consider a 4x4 MU-MIMO system, where the AP is equipped with switched-beam antennas with 8 directional modes. In this case, each user has a total of $8 \times 4 = 32$ long-term gain entries in $\mathbf{G}$.

A. Active feedback protocol extension

The first protocol extension, which we refer to as Active Feedback, allows the system to actively update the $\mathbf{G}$ matrix with almost zero overhead but requires minor support from end clients. The key idea is to inject training symbols into beacon frames. Beacon frames are used to announce the presence of a wireless network, and are typically broadcasted to the users every 100ms. We insert 4 training symbols into the beacon frame, one for each antenna. Then, beacons are sent while setting the antennas in different directions using a round robin fashion to cover all directions as fast as possible as illustrated in Table I. Upon reception of the training symbols, users compute their instantaneous CSI per subcarrier, and compute the average channel gain values as discussed in Section IV-A. Thus, each beacon frame informs 4 values and after 8 beacons (about 800ms) all 32 entries of the user are known. Once the user has obtained all its channel gain information, it injects the information into an ACK frame and sends it to the AP. Upon reception of the channel gain entries, the AP updates the row of $\mathbf{G}$ corresponding to this user.

TABLE I: Antenna directions at each beacon/transmission

Direction	beacon/Tx 1	beacon/Tx 2	beacon/Tx 3	...
ANT 1	↑	↗	→	...
ANT 2	→	↘	↓	...
ANT 3	↓	↙	←	...
ANT 4	←	↖	↑	...

One may worry that beacons may not be overheard by users if they are sent with directional modes, and/or that new users or mobile users that just changed their location won't be able to successfully receive packets during the warm up period of updating their row in the $\mathbf{G}$ matrix. Intuitively though, the main lobe of the inexpensive switched-beam antenna modes we consider is often more than 90 degrees, the main to side/back lobe gain ratio is often less than 6dB [41], and these two together with indoors multipathing imply we should not worry: the wireless coverage would be as good as before. To remove any doubt that this is not a real concern, we conduct

(a) 2 RUs, 8 users total

(b) 4 RUs, 16 users total
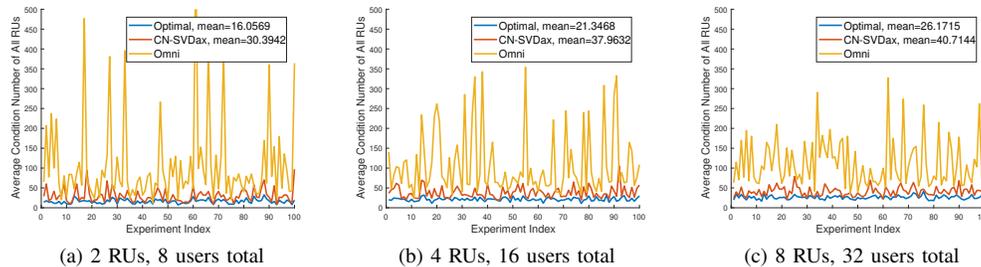
(c) 8 RUs, 32 users total

Fig. 13: Simulation of antenna mode selection for 802.11ax.

the following experiment with the WARP boards. We measure the SNR received by users when using the 8 training configurations of Table I and when using the omni mode under 20 randomly generated topologies. We compute the average sum capacity (over the different topologies) for both approaches, where the sum capacity of the training configurations are also averaged over the 8 different configurations. Figure 14a shows that, somewhat surprisingly, the 8 training configurations yield on average 1.2x higher rates than the omni mode, varying from 0.6x to 1.9x depending on the configuration and the topology.

Airtime overhead: We introduce overhead by increasing the size of the beacon frame and of some ACKs. We first compute the beacon overhead. A beacon frame is typically 50-100 bytes long and is modulated in BPSK. A training symbol is merely an OFDM symbol and thus equivalent to 8 bytes (for 20MHz channels). Therefore, we add $4 \times 8 = 32$ bytes to each beacon frame for all 4 antennas. The data rate for beacon frames is normally set to the lowest PHY rate which is 6.5 Mbps (for 20MHz channels), thus we introduce approximately 40μs of additional airtime per beacon. Since beacons are sent every 100ms, the overhead is only 40μs/100ms = 0.04%.

We now compute the overhead in the ACK frame. For each channel gain value, 2 bytes precision is more than enough. Therefore, to feedback all the long-term CSI of a given user we add $32 \times 2 = 64$ bytes into the ACK, which corresponds to about 80μs of additional airtime since the ACK is also modulated in the lowest PHY rate in order to be robust. Since channel gain updates occur every 800ms, and assuming say a network with 40 active users, on average every $\frac{800}{40} = 20$ms one user will need to update its long-term CSI to the AP, which results in a small overhead of 80μs/20ms = 0.4%. The total overhead is less than 0.5% of airtime, and even in an 8x8 MU-MIMO system it would be less than 1%.

B. Passive feedback protocol extension

The Passive Feedback protocol extension requires no support from end clients and relies on the channel sounding stage of a standard MU-MIMO transmission to update the **G** matrix. Similar to the active feedback extension, the AP sets its antennas to different directions in a round robin fashion so that all antenna modes can be eventually trained. The difference is that instead of changing directions when transmitting beacon frames, it changes directions during normal transmissions.

Clearly this protocol extension is fully compatible with the 802.11 standard and it introduces no airtime overhead

whatsoever. However, every time a new user joins the AP or an already associated user moves by a sizable amount, the AP has to use 8 data transmissions towards a user group that includes this user to update the row of the **G** matrix corresponding to this user. To do so, the 8 antenna configurations shown in Table I have to be used. As already discussed, the average performance of these 8 configurations is still 1.2x that of the omni mode, but clearly they perform far from the 3.5x gains achieved by the best configurations. The performance hit from this depends on the portion of time that transmissions take place using these 8 configurations versus using a configuration found by the CN-SVD algorithm. This, in turn, depends on user dynamics, e.g. how often new users join, and, more important, how mobile the users are, forcing the system to update elements of the **G** matrix. We investigate the effect of mobility and estimate the long-run performance gain in the presence of period **G** matrix updates below.

C. User mobility and long-run performance

When a user changes its location, the **G** matrix may need to be updated either with the active or the passive feedback protocol, and, there is a temporary performance hit. To see this, we perform MU-MIMO transmissions in a dynamic office environment and report the performance of Omni versus CN-SVD with passive feedback. Specifically, we conduct 20 transmissions with static users, change the location of some users at transmission 21, and conduct further transmissions after that. Figure 14b shows the higher throughput achieved by the CN-SVD algorithm during the first 20 transmissions, the throughput drop at transmission 21 at the levels of Omni, and the quick rebound after the update of the **G** matrix during the 8 transmissions following transmission 21. Note that at transmission 21 the system automatically detects the drastic change in some users' long-term CSI as shown in Figure 3a and triggers the use of the 8 antenna configurations shown in Table I to update **G** matrix.

How often does the system has to update the **G** matrix? Not too often, since wireless devices connected to a WiFi network change their location in minute time-scales [34], [35], [36], whereas the process of updating the row of a device/user in the **G** matrix takes much less. To make this precise and compute the long-run performance gain we resort to both back of the envelope calculations and to simulations.

In 802.11ac, the maximum transmission length is about 5ms when packet aggregation is used. A feedback report,

(a) Round-robin capacity VS Omni capacity

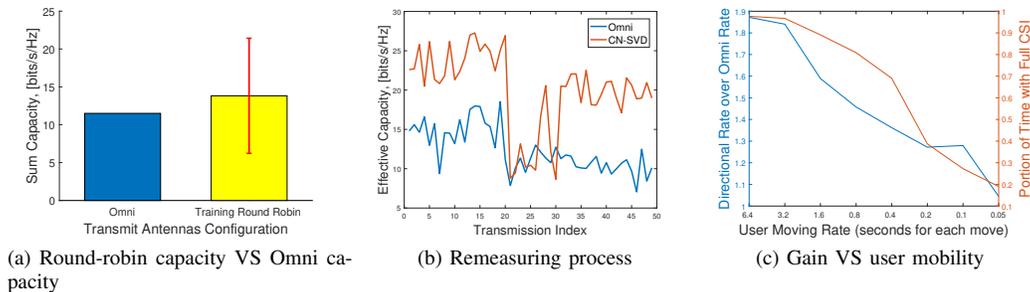(b) Remeasuring process

(c) Gain VS user mobility

Fig. 14: Experiments with mobile users.

which contains the instantaneous CSI of the users, is typically $16 \times 12 \times 52 = 9984$ bits long in a 4x4 scenario and takes up to 1.5ms to transmit [44]. Including the airtime for sounding and the ACK transmissions from the 4 users (each ACK frame takes 48μs to be transmitted [45]), it takes up to 7ms to complete an MU-MIMO transmission. An AP needs 8 transmissions to learn the long-term gain information of 4 users, or, equivalently, at most 56ms. (Clearly, it makes no sense to use packet aggregation when transmitting with the 8 training antenna configurations, but we choose to be conservative in our calculations.) Let's assume that a user changes location every say 10 seconds. Assuming there are 40 active users associated with the AP and the AP serves them equally, the user under consideration will be part of 10% of the user groups. Thus, during the 10 seconds that the user stays put, it is receiving data for 1 second (assuming saturation regime). Based on the previous analysis, for 56ms the user's gain is 1.2x and for the rest of time it is 3.5x, resulting in a $1.2 \times 56/1000 + 3.5 \times 944/1000 \approx 3.4$x average performance gain.

Last, we use simulations to determine the effect of varying mobility to erformance. An AP serves a total of 40 mobile users through 8x8 MU-MIMO transmissions. Users' inter-move times are exponentially distributed with an average of less than 10 seconds (see x-axis in Figure 14c for used values). Whenever a user moves, its long-term CSI changes and has to be remeasured. We run the simulation for 100 seconds and compare the average capacity obtained by Omni versus CN-SVD with passive feedback. We also record the proportion of time that the passive feedback mechanism spends on remeasuring the long-term CSI. As shown in Figure 14c, even if the users move as frequently as every 3 seconds we still achieve a significant gain over Omni.

VIII. IMPLEMENTATION OVERHEAD

In this section we are going to discuss the implementation overhead of our proposed algorithm in practice. Three types of overhead will be introduced to the WiFi system: 1) the airtime overhead for measuring the long-term channel gains, 2) the computation overhead for finding the best antenna configurations and 3) the additional power consumption from the use of switched-beam antennas. As we have already discussed the airtime overhead in Section VII, we will be focusing on the additional computation and power consumption here.

Assume the AP is equipped with n switched-beam antennas to serve n users concurrently, while each antenna has a total of d directional modes. In this case, we would have a $\mathbf{G}$ matrix of size n-by-nd. Recall that we use Algorithm 2 to recursively search for the best mode for each of the n antennas. The number of total recursions could vary from 1 to n: in the worst case, we will need to run the ALG algorithm (see Algorithm 1) n times and compute $\kappa(\mathbf{G}_{\boldsymbol{u}, \{d\} \cup \boldsymbol{d}_f})$ for $1 + 2 + \cdots + n = \mathcal{O}(n^2)$ times. Now, both ALG and κ operations are dominated by a singular-value decomposition (SVD), which costs $\mathcal{O}(\min(mn^2, m^2n))$ for an m-by-n matrix. Thus, the total worst-case time complexity of the CN-SVD algorithm would be $\mathcal{O}(n^4d^2 + n^5)$. Since the latest standard (802.11ax) only allows the AP to serve up to 8 users concurrently via MU-MIMO transmission, it is reasonable to assume that $d > n$ and thus the overall time complexity becomes $\mathcal{O}(n^4d^2)$, which is much smaller than the time complexity of brute-force search, $\mathcal{O}(n^3d^n)$.

Similarly, the worst-case time complexity of the extended algorithms can be found as follows: assuming a total of m users and other things remain unchanged, the *Cross Selection* algorithm (CROSS, see Algorithm 3) first selects n out of the m users and then performs CN-SVD on them. The first operation costs $\mathcal{O}(\min(m^2nd, n^2d^2m))$ and will only be executed once. Therefore, the total worst-case time is upper-bounded by either $\mathcal{O}(n^4d^2)$ or $\mathcal{O}(\min(m^2nd, n^2d^2m))$ depending on the value of m, as opposed to $\mathcal{O}(n^3m^nd^n)$ by brute-force search. On the other hand, the time complexity of *CN-SVD for 802.11ax* (CN-SVDax, see Algorithm 4) is upper-bounded by $\mathcal{O}(kn^4d^2)$ with k being the number of RU groups, where the brute-force search still has an exponential time complexity $\mathcal{O}(n^3d^n)$.

In practice, the CN-SVD algorithm takes a couple of millisecond to finish in a standard CPU in Matlab, and will take much shorter time if programmed in C. As a result, such time would be shorter than the total time of a typical MU-MIMO transmission, which allows the AP to find the best antenna configuration to be used in the next transmission in advance while performing the current transmission.

Finally, our proposed algorithm poses little additional power consumption to the system. Under the Active Feedback mechanism described in Section VII-A, the AP injects training symbols to its beacons, while the user sends back its long-term channel gain information by also injecting it into the

ACK frame. As a result, we have not introduced any new transmission into the system, but rather slightly extended some existing ones (beacon and ACK). Clearly, the extra power consumption is negligible in this case. Similarly, for the Passive Feedback mechanism (Section VII-B), all additional information is collected through ordinary data transmissions, thus no extra power consumption to the system is incurred.

IX. CONCLUSION/ACKNOWLEDGEMENTS

In this paper we use SDRs and commercial hardware to show that switched-beam antennas conjunction with MU-MIMO can achieve a 3.5x-5x performance gain over omni-mode MU-MIMO, with negligible overhead and while being fully compatible with the 802.11ac standard.

We would like to thank Adant Technologies, Inc. for their assistance in configuring the switched beam antennas used in this paper as well as with the experiments involving commercial devices.

REFERENCES

[1] V. Erceg, P. Soma, D. S. Baum, and A. J. Paulraj, "Capacity obtained from multiple-input multiple-output channel measurements in fixed wireless environments at 2.5 ghz," in *Proc. of ICC '02*.
[2] Agilent, "Mimo performance and condition number in lte test," 2009.
[3] T. Yoo and A. Goldsmith, "On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming," *IEEE J. Sel. Areas Commun.*, Mar 2006.
[4] T. Yoo, N. Jindal, and A. Goldsmith, "Multi-antenna downlink channels with limited feedback and user selection," *IEEE J. Sel. Areas Commun.*, Sept 2007.
[5] A. Adhikary, J. Nam, J. Y. Ahn, and G. Caire, "Joint spatial division and multiplexing; the large-scale array regime," *IEEE Trans. Inf. Theory*, Oct 2013.
[6] X. Xie, E. Chai, X. Zhang, K. Sundaresan, A. Khojastepour, and S. Rangarajan, "Hekaton: Efficient and practical large-scale mimo," in *Proc. of MobiCom '15*.
[7] D. Parker and D. C. Zimmermann, "Phased arrays-part ii: implementations, applications, and future trends," *IEEE Trans. Microw. Theory Techn.*, Mar 2002.
[8] W.-R. Li, C.-Y. Chu, K.-H. Lin, and S.-F. Chang, "Switched-beam antenna based on modified butler matrix with low sidelobe level," *Electronics Letters*, Mar 2004.
[9] G. Cerri, R. D. Leo, V. M. Primiani, C. Monteverde, and P. Russo, "Design and prototyping of a switching beam disc antenna for wideband communications," *IEEE Trans. Antennas Propag.*, Dec 2006.
[10] M. I. Lai, T. Y. Wu, J. C. Hsieh, C. H. Wang, and S. K. Jeng, "Compact switched-beam antenna employing a four-element slot antenna array for digital home applications," *IEEE Trans. Antennas Propag.*, Sept 2008.
[11] A. R. Lopez, "Performance predictions for cellular switched-beam intelligent antenna systems," *IEEE Commun. Mag.*, Oct 1996.
[12] M.-J. Ho, G. L. Stuber, and M. D. Austin, "Performance of switched-beam smart antennas for cellular radio systems," *IEEE Trans. Veh. Technol.*, Feb 1998.
[13] A. Michaloliakos, W. C. Ao, K. Psounis, and Y. Zhang, "Asynchronously coordinated multi-timescale beamforming architecture for multi-cell networks," *IEEE/ACM Trans. Netw.*, under submission.
[14] Ruckus Wireless, Inc., "White paper: Beamflex, 11ac wave 2 and mimo," 2015, https://www.ruckuswireless.com/technology/beamflex/.
[15] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
[16] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "Airsync: Enabling distributed multiuser mimo with full spatial multiplexing," *IEEE/ACM Trans. Netw.*, Dec 2013.
[17] T. H. Kim, T. Salonidis, and H. Lundgren, "Mimo wireless networks with directional antennas in indoor environments," in *Proc. of INFOCOM '12*.
[18] C. Kong, Z. Yuan, X. Han, F. Yang, X. Wang, T. Wang, and S. Lu, "Vsmc mimo: A spectral efficient scheme for cooperative relay in cognitive radio networks," in *Proc. of INFOCOM '15*.

[19] Y. Yan, P. Yang, X.-Y. Li, Y. Tao, L. Zhang, and L. You, "Zimo: Building cross-technology mimo to harmonize zigbee smog with wifi flash without intervention," in *Proc. of MobiCom '13*.
[20] S. Gollakota, S. D. Perli, and D. Katabi, "Interference alignment and cancellation," in *Proc. of SIGCOMM '09*.
[21] D. Bharadia and S. Katti, "Full duplex mimo radios," in *Proc. of USENIX NSDI '14*.
[22] H. Qiu, K. Psounis, G. Caire, K. C. M, and K. Wang, "High-rate wifi broadcasting in crowded scenarios via lightweight coordination of multiple access points," in *Proc. of MobiHoc '16*.
[23] M. Matthaiou, D. I. Laurenson, and C. X. Wang, "On analytical derivations of the condition number distributions of dual non-central wishart matrices," *IEEE Trans. Wireless Commun.*, March 2009.
[24] Y. Zhang and K. Psounis, "Efficient mu-mimo via switched-beam antennas," in *Proceedings of the 18th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. Mobihoc '17, 2017.
[25] G. Dimic and N. Sidiropoulos, "On downlink beamforming with greedy user selection: Performance analysis and a simple new algorithm," *IEEE Trans. Signal Process.*, Oct 2005.
[26] A. Michaloliakos, R. Rogalin, V. Balan, K. Psounis, and G. Caire, "Efficient mac for distributed multiuser mimo systems," in *Proc. of WONS '13*.
[27] A. Adhikary and G. Caire, "Jsdm and multi-cell networks: Handling inter-cell interference through long-term antenna statistics," in *Proc. of Asilomar Conf. '14*.
[28] J. Wang, H. Zhu, L. Dai, N. J. Gomes, and J. Wang, "Low-complexity beam allocation for switched-beam based multiuser massive mimo systems," *IEEE Transactions on Wireless Communications*, vol. 15, no. 12, pp. 8236–8248, Dec 2016.
[29] E. Bjornson, D. Hammarwall, and B. Ottersten, "Exploiting quantized channel norm feedback through conditional statistics in arbitrarily correlated mimo systems," *IEEE Trans. Signal Process.*, Oct 2009.
[30] M. Matthaiou, M. R. Mckay, P. J. Smith, and J. A. Nossek, "On the condition number distribution of complex wishart matrices," *IEEE Trans. Commun.*, June 2010.
[31] D. A. Belsley, E. Kuh, and R. E. Welsch, *Regression diagnostics: Identifying influential data and sources of collinearity*. John Wiley & Sons, 2005.
[32] R. Rogalin, O. Y. Bursalioglu, H. C. Papadopoulos, G. Caire, A. F. Molisch, A. Michaloliakos, H. V. Balan, and K. Psounis, "Scalable synchronization and reciprocity calibration for distributed multiuser MIMO," *IEEE Trans. Wireless Commun.*, April 2014.
[33] E. Bjrnson, M. Kountouris, M. Bengtsson, and B. Ottersten, "Receive combining vs. multi-stream multiplexing in downlink systems with multi-antenna users," *IEEE Trans. Signal Process.*, July 2013.
[34] A. Balachandran, M. Voelker, P. Bahl, and V. Rangan, "Characterizing user behavior and network performance in a public wireless lan," in *Proc. of SIGMETRICS '02*.
[35] D. Tang and M. Baker, "Analysis of a local-area wireless network," in *Proc. of MobiCom '00*.
[36] M. Balazinska and P. Castro, "Characterizing mobility and network usage in a corporate wireless local-area network," in *Proc. of MobiSys '03*.
[37] A. Çivril and M. Magdon-Ismail, "On selecting a maximum volume submatrix of a matrix and related problems," *Theor. Comput. Sci.*, 2009.
[38] G. Golub and C. V. Loan, *Matrix Computations (3rd Ed.)*. JHU Press, 1996.
[39] M. Kane and B. W. Lewis, "Svd subset selection," 2014, http://bwlewis.github.io/GLM/svdss.html.
[40] A. F. Molisch, *Wireless communications*. John Wiley & Sons, 2007.
[41] Adant Tech. Inc. (2016) Adant star160 (white paper). [Online]. Available: www.adant.com
[42] A. Pyzara, B. Bylina, and J. Bylina, "The influence of a matrix condition number on iterative methods' convergence," in *Proc. of FedCSIS '11*.
[43] K. Wang and K. Psounis, "Scheduling and resource allocation in 802.11ax," in *Proc. of INFOCOM '18*.
[44] Aruba Networks, Inc., "802.11ac in-depth (white paper)," 2014, http://www.arubanetworks.com/pdf/technology/whitepapers/WP_80211acInDepth.pdf.
[45] T. Vanhatupa, "Wi-fi capacity analysis for 802.11 ac and 802.11 n: Theory & practice," 2013.