Design and Experimental Evaluation of Multi-User Beamforming in Wireless LANs

Ehsan Aryafar[1], Narendra Anand[1], Theodoros Salonidis[2], and Edward W. Knightly[1]
[1]Rice University, Houston, TX, USA [2]Technicolor, Paris, France
[1]{ehsan, nanand, knightly}@rice.edu, [2]theodoros.salonidis@technicolor.com

ABSTRACT

Multi-User MIMO promises to increase the spectral efficiency of next generation wireless systems and is currently being incorporated in future industry standards. Although a significant amount of research has focused on theoretical capacity analysis, little is known about the performance of such systems in practice. In this paper, we present the design and implementation of the first multi-user beamforming system and experimental framework for wireless LANs. Using extensive measurements in an indoor environment, we evaluate the impact of receiver separation distance, outdated channel information due to mobility and environmental variation, and the potential for increasing spatial reuse. For the measured indoor environment, our results reveal that two receivers achieve close to maximum performance with a minimum separation distance of a quarter of a wavelength. We also show that the required channel information update rate is dependent on environmental variation and user mobility as well as a per-link SNR requirement. Assuming that a link can tolerate an SNR decrease of 3 dB, the required channel update rate is equal to 100 and 10 ms for non-mobile receivers and mobile receivers with a pedestrian speed of 3 mph respectively. Our results also show that spatial reuse can be increased by efficiently eliminating interference at any desired location; however, this may come at the expense of a significant drop in the quality of the served users.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless Communication*

General Terms

Measurement, Performance, Reliability, Experimentation, Design

Keywords

Multi-User MIMO, Beamforming, Wireless LAN, Channel Information, Mobility, Interference Suppression

MobiCom'10, September 20–24, 2010, Chicago, Illinois, USA.

1. INTRODUCTION

Multiple-Input Multiple-Output (MIMO) offers the potential to achieve high throughput in point-to-point wireless links. It is already included in several wireless standards such as IEEE 802.11n [4] and is implemented in commercially available devices.

Recently, there has been a growing interest in how to fully realize the benefits of MIMO in a multi-user scenario. In a Multi-User MIMO (MU-MIMO) system, the base station is equipped with several antennas and communicates simultaneously with several users each with one or more antennas. The downlink channel of such a system has received a great deal of attention; MU-MIMO techniques are already being adopted by the next generation of wireless standards such as LTE [9] and WiMAX [5].

In traditional single user systems, one user is served at a time with a mechanism such as time division multiple access (TDMA). However, the throughput of such a system would be limited by the minimum number of antennas at the base station and receiver. Typically, a base station could accommodate a large number of antennas, whereas a user device would have a small number of antennas. As a result, in such a system, the benefits of MIMO would be constrained by the number of user antennas.

Information theory results for downlink MIMO systems show that it is optimal to serve multiple users simultaneously [7], and several theoretical multi-user schemes have been proposed [18, 19, 21] for such systems. The optimal solution involves a theoretical pre-interference cancellation technique known as Dirty Paper Coding (DPC) [8, 19]; however, DPC is difficult to implement due to its high computational complexity.

Multi-user beamforming (MUBF) [21] is a sub-optimal yet simple method of serving multiple users. In MUBF, multiple users can be served simultaneously by multiplying each individual data stream by its appropriate beamforming weight vector, adding the resulting streams, and then transmitting the summed streams in parallel over the base station's antenna array. Careful selection of these beamforming weights can reduce or eliminate inter-user interference.

The performance of the aforementioned algorithms has been usually evaluated under the idealized case of uncorrelated, Gaussian channels. The primary goal of this paper is to evaluate the performance of such downlink schemes in real-world deployments. To accomplish this, we design and implement a custom, FPGA-based, hardware framework that enables the evaluation of MUBF algorithms under real channel conditions. Specifically, we investigate a MUBF algorithm known as Zero Forcing Beamforming (ZFBF) [21]. We measure the performance of ZFBF as a function of receiver separation distance, concurrent user selection, and user population size. We also perform channel emulator experiments with controlled and repeatable channels to address the impact of

outdated channel information due to mobility and environmental variation. We further investigate the potential of ZFBF to reduce interference at unwanted locations and increase spatial reuse. In all of our experiments we also perform TDMA-based single-user beamforming (SUBF) as the baseline.

Our measurement study has the following main contributions: First, we design and implement a custom framework that allows for evaluation of different MUBF algorithms. To the best of our knowledge, this is the first platform that allows for multi-antenna based simultaneous transmission of different data streams to different users while providing a framework for implementation of different MUBF strategies.

Second, we evaluate the multiplexing gain of ZFBF as a function of receiver separation distance, concurrent user selection, and user population size. Through extensive over-the-air (OTA) measurements, we find that when the number of selected users is smaller than the number of transmitting antennas, the multiplexing gains of ZFBF are not affected by the receiver separation distance. In fact, we show that this allows for the simultaneous transmission of different data streams to users that are down to a quarter of a wavelength from one another.

Third, with controlled experiments performed with a channel emulator, we investigate the impact of outdated channel information due to environmental variation and user mobility on the performance of ZFBF. We find that the necessary channel update rate is dependent on the environmental variation and user mobility as well as the link quality. Assuming that a link can tolerate SNR losses of up to 3 dB compared to an omni transmission, a maximum channel update rate of 100 ms is required to guarantee acceptable performance in a typical, indoor, non-mobile environment. However, we find that a channel update rate of 10 ms is required for a mobile receiver with an average pedestrian speed of 3 mph.

Fourth, we investigate the potential benefits of ZFBF in reducing interference and thus increasing spatial reuse. Our experimental results reveal that a user can obtain an interference-free channel by sending its channel information to a ZFBF-enabled transmitter. We show that the capability of ZFBF to eliminate interference is not affected by the location of an unintended receiver or the number of such unintended receivers; however, as the number of the unintended receivers increases, the link quality of the currently served receivers can drop significantly.

The rest of this paper is organized as follows: Section 2 provides a background of MUBF. Section 3 describes the design and implementation of the schemes studied in this paper. Section 4 describes the multiplexing gains of ZFBF. Section 5 investigates the impact of outdated channel information. Section 6 investigates the potential of ZFBF to increase spatial reuse. Finally, we discuss related work in Section 7 and conclude the paper in Section 8.

2. BACKGROUND

In this section, we describe the system model and present background on the techniques we implemented using our experimental platform.

2.1 System Model

We consider a multi-user, multi-antenna downlink channel in which a base station is equipped with N transmit antennas and transmits to K user terminals, each equipped with a single antenna. This scenario is typical in current WLAN systems and standards where base stations can afford to utilize sophisticated multi-antenna technologies while the clients, driven by cost and simplicity, use single-antenna technologies. An example of such a network is shown in Fig. 1.

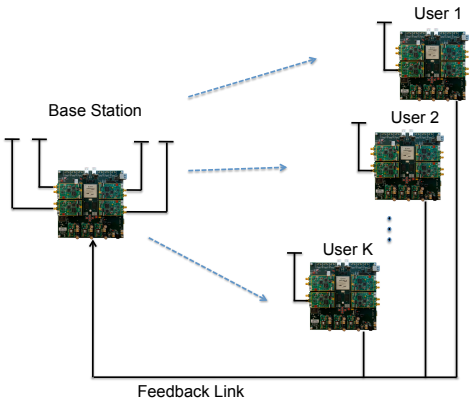

Figure 1: Multi-user beamforming system model.

We consider a narrowband system model, where the received baseband signal y_k of the k-th user is given by:

$$y_k = \mathbf{h_k}\mathbf{x} + z_k, \quad k = 1, ..., K \tag{1}$$

where $\mathbf{x}$ is the transmitted symbol from the base station antennas, $\mathbf{h_k} = [h_{1k}, h_{2k}, ..., h_{Nk}]^T$ is the channel gain matrix of the k^{th} user, and z_k represents the circularly symmetric additive white Gaussian noise at the receiver with zero mean and variance σ^2. In this model, the base station transmitter is subject to a total power constraint P, i.e., $\mathbf{x}^*\mathbf{x} \leq P$, [1]. The total transmit power does not depend on the number of transmit antennas and remains the same for all schemes that serve multiple users.

We proceed to describe the different schemes we implement and investigate in this paper.

2.2 Single-User Scheme

In a Single-User scheme, the base station transmits to only one user at a time in a TDMA fashion. We consider two such schemes: (i) In *Omni* transmission mode, no channel estimate feedback is available at the base station. Thus, the base station uses a fixed single antenna for all of its transmissions. (ii) In *Single-User Beamforming* (SUBF) mode, channel estimates are available at the base station through feedback. When the channel estimates are available at the base station, the signals fed to each of its antenna elements are weighted with suitable amplitude and phase components (beamforming weights) to increase SNR at the receivers.

In this scheme, the transmitted signal $\mathbf{x}$ is given by $\mathbf{x} = \mathbf{w}s$, where $\mathbf{w}$ is the beamforming vector and s is the intended symbol. The beamforming vector $\mathbf{w}$ is selected such that the transmit power of symbol s is not increased, i.e. $\|\mathbf{w}\|^2 = 1$. When serving only one user, the beamforming vector can be selected to maximize the SNR at the receiver. In this case, the SNR-maximizing weight vector equals $\frac{\mathbf{h}^*}{\|\mathbf{h}\|}$.

In both Omni and SUBF schemes, the aggregate throughput can be maximized by only serving the user with the largest single-user capacity, where the capacity of user k is given by:

$$C_k = log_2(1 + SNR_k) \tag{2}$$

Although aggregate throughput maximization is attractive, in practice, wireless providers must serve all their users. Thus, providing fairness among users is an important issue that can not be ignored by the service provider. Therefore, we consider a round robin

[1] x^* is the conjugate transpose of the transmitted symbol x.

scheduling scheme in which all of the users are provided with an equal amount of serving time. Thus, the sum rate of each Single-User scheme is equal to $\Sigma_{k=1}^{k=K} \frac{C_k}{K}$.

2.3 Multi-User Beamforming

An alternative approach to Single-User schemes is to serve multiple users simultaneously. Let s_k, $\mathbf{w_k} \in \mathbb{C}$, and $P_k \in \mathbb{R}$, be the data symbol, weight vector, and transmit power scaling factor for user k, respectively. In a Multi-User scheme with linear weights, the transmitted signal $\mathbf{x}$ equals $\sum_{k=1}^{K} \sqrt{P_k}\mathbf{h_k}\mathbf{w_k}$. Thus, from Eq. (1) the resulting received signal vector for user k is:

$$y_k = (\sqrt{P_k}\mathbf{h_k}\mathbf{w_k})s_k + \sum_{j \neq k} \sqrt{P_j}\mathbf{h_j}\mathbf{w_j}s_j + z_k \qquad (3)$$

In Eq. (3), the first term represents the desired signal, the second term represents the multi-user interference and the third term is the noise.

Finding the optimal $\mathbf{w_k}$s and P_ks that maximize the aggregate capacity is a difficult, non-convex optimization problem [18].

In this paper we implement a simpler strategy known as Zero-Forcing BeamForming (ZFBF) [21]. In ZFBF, weight vectors are selected with the goal of zero inter-user interference (i.e., $\mathbf{h_k}\mathbf{w_j} = 0$ for $j \neq k$), and thus the second term in Eq. (3) is equal to zero. With ZFBF, the maximum number of receivers that can be served simultaneously is equal to the number of transmitting antennas, N. Thus, the ZFBF scheme has N degrees of freedom (DoF).

Let $M \subset \{1, ..., K\}$, $|M| \leq N$ be the subset of users that the base station intends to serve concurrently, and $\mathbf{H}(\mathbf{M})$ and $\mathbf{W}(\mathbf{M})$ be the corresponding submatrices of $\mathbf{H} = [\mathbf{h_1^T}, \mathbf{h_2^T} ... \mathbf{h_K^T}]^T$ and $\mathbf{W} = [\mathbf{w_1} ... \mathbf{w_K}]$, respectively. In [20], Wiesel et al. show that the optimal choice of $\mathbf{W_M}$ that gives zero-interference is the pseudo-inverse of $\mathbf{H}(\mathbf{M})$.

$$\mathbf{W}(\mathbf{M}) = \mathbf{H}(\mathbf{M})^{\dagger} = \mathbf{H}(\mathbf{M})^{*}(\mathbf{H}(\mathbf{M})\mathbf{H}(\mathbf{M})^{*})^{-\mathbf{1}} \qquad (4)$$

Thus, the only remaining parameters that need to be specified are the power coefficients, P_k. These coefficients can be selected such that the aggregate throughput is maximized or different fairness objectives are achieved.

In this paper, we investigate two power allocation approaches with ZFBF. First, we consider the maximum throughput approach (ZFBF-MT), where the power allocation problem becomes:

$$\begin{aligned} &max_{\mathbf{p} \geq 0} \qquad \Sigma_k log(1 + \tfrac{P_k}{\sigma^2}) \\ &s.t. \qquad \Sigma_k P_k [(\mathbf{H}\mathbf{H}^{*})^{-\mathbf{1}}]_{k,k} \leq P \end{aligned} \qquad (5)$$

This problem can be easily solved by using the well-known water filling solution [21]. Second, we consider a scheme that we call ZFBF-EP where the base station transmitter allocates equal power to its users. We use ZFBF-EP for a fair comparison with our round robin-based, Single-User schemes.

3. EXPERIMENTAL SETUP

In this section, we describe the design and implementation of our multi-user beamforming testbed along with the conditions under which the measurements in this study were performed.

3.1 System Implementation

3.1.1 WARPLab Research Framework

We performed experiments using WARPLab [3], a framework that enables rapid implementation of physical layer algorithms in MATLAB and real-time, over-the-air (OTA) transmission of data using WARP boards. WARPLab provides a software interface between a host PC running MATLAB and up to 16 WARP boards through an Ethernet switch. The host PC is responsible for the construction and processing of baseband waveforms that are transmitted and received between the connected WARP nodes.

Parameter	Value
Carrier Frequency	2.4 GHz
Number of subcarriers	1
Bandwidth	625 KHz
ADC/DAC sampling frequency	40 MHz
Symbol time	3.2 μs
Modulation	16-QAM
Coding Rate	1 (No Correction Code)

Table 1: Warplab Physical layer parameters.

When used with WARPLab, WARP nodes are essentially large data buffers connected to wireless radio daughter cards that perform RF up/down conversion and amplification. The host PC creates baseband waveforms using a user-defined MATLAB script that implements a physical layer algorithm (i.e. MUBF). This baseband waveform is downloaded to the transmitting node's buffer via Ethernet and then sent OTA through the radio board. The receiving node streams this data into its own buffer after which the node uploads the received data back to the host PC for further baseband processing. To synchronize the transmission and reception of data, the host PC uses a trigger pulse sent to the connected WARP nodes.

Table 1 specifies the physical layer parameters used in WARPLab. The current reference design of the WARPLab framework supports a channel bandwidth of 625 KHz. This channel bandwidth is smaller than the channel used in standards such as 802.11a/b/g where a channel width of 20 MHz is used. However, we note that similar experimental results would be obtained with a higher channel width provided that either flat fading channel conditions exist or more accurate channel information is available. For example in an OFDM modulation system (e.g., 802.11a/g) in which the channel is divided into many subcarriers, per subcarrier channel information could be used to provide accurate channel information.

The main component of the WARP board is a Xilinx Virtex-II Pro FPGA. Each WARP node also has four daughter card slots which allow the FPGA to connect to up to four radio boards.

We used four radio boards at the base station transmitter to build a multi-antenna system. Four 3 dBi antennas are mounted in a circular array structure with a one-wavelength distance between adjacent antennas (12.5 cm at 2.4 GHz). Fig. 2 depicts the antenna array at the transmitter connected to a WARP board. Each receiver only uses one radio board.

3.1.2 Multi-User Beamforming Implementation

MUBF requires a feedback mechanism to allow the transmitter to obtain channel information in order to properly construct beam weights. In order to accomplish this goal, our system does the following: First, the transmitter sends a packet with a known training preamble. The clients receive this transmission and upload their received versions of the preamble to the host PC. Then, the host PC computes the $\mathbf{H}$ matrix from the received preamble and uses it to compute the beamforming weights. These weights are then downloaded to the transmitting node where they are used to beamform the second transmission. The receivers now measure the received

signal strength (RSS) value of this transmission and upload the data to the host PC for logging. In this section, we will detail the three main components of the aforementioned system: Channel Training, Channel Estimation, and Beam Weight Calculation.

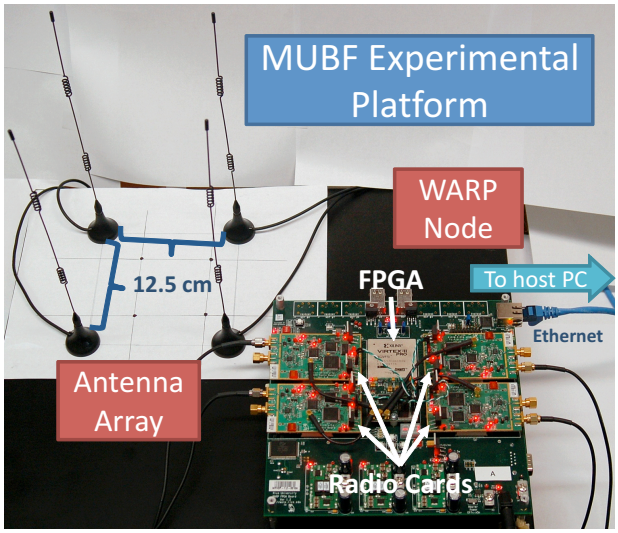

Figure 2: Transmitter platform.

Channel Training. During channel training, the base station simultaneously transmits a preamble sequence on all of its antennas. The structure of the preamble is shown in Fig. 3. Each preamble is composed of three main sections. The first is the Short Training sequence, which is a narrow-band tone used by the receiver's Automatic Gain Control (AGC) mechanism. The second is the Long Training sequence, a wide-band sequence from the 802.11a standard with strong autocorrelation properties that is used for timing synchronization at the receiver. This sequence is crucial to the system's performance because it helps eliminate the adverse effects of Carrier Frequency Offset (CFO) that are caused by oscillator drift between the transmitter and receiver. The CFO problem in a wireless system occurs due to differences between transmitter and receiver oscillators. The oscillator is responsible for generating the high frequency carrier signal. In today's hardware, oscillators drift on the order of parts per million (ppm) per C^o above or below room temperature. Such drifts could cause significant distortion between received and transmitted signal phase to the point were the correct signal can not be decoded. In a communication system, the preamble is used at the receiver to correct the CFO that exists between the transmitter and receiver. The third is the pilot tone, a narrow-band tone used for actual channel estimation. All three parts of the preamble have identical values for each antenna.

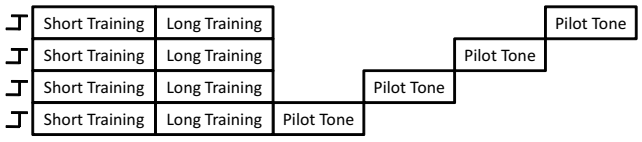

Figure 3: Preamble structure.

The only difference between the antennas' preambles is the structure as is apparent in Fig. 3. All four transmit antennas send the Long and Short training symbols in parallel because the receiver does not care which antenna the training symbols originated from. However, because channel estimates (and $\mathbf{H}$ matrices) need information for each antenna path, the transmitter sends them such that during the Pilot section of the preamble, only one antenna is transmitting a tone for channel estimation at a time.

Channel Estimation and Beam Weight Calculation. Channel estimation is accomplished by comparing the received Pilot Tones to the expected Pilot Tone. Once the $\mathbf{H}$ matrix is obtained, the beamforming weights can be found from the desired beam weight calculation algorithm (Eq. (4) in ZFBF). After this, the required power allocation scheme is applied to each of the selected beams. The resulting beam weights are then downloaded to the FPGA, which constructs the beamformed data and transmits it through the radio cards.

3.2 Measurement Setup

In this subsection, we describe the conditions under which OTA transmissions were performed. First, we show that the feedback delay of our system (i.e. the time interval between channel estimation at the receiver and beamformed data transmission at the transmitter) is within the channel coherence time. Then, we describe the metrics used to evaluate the performance of different schemes.

3.2.1 Channel Coherence Time

The total feedback delay in our implementation is equal to 60 ms due to the nature of the WARPLab framework. Because all baseband processing happens at the host PC in MATLAB, the system has the added latency of downloading and uploading data streams over Ethernet. If the channel varies significantly during this time interval, the initial channel estimate would become outdated. The resulting multi-user interference within the selected user group could be high enough to adversely affect system performance. Thus, for valid OTA transmissions the system feedback delay in our evaluation testbed should be within the channel coherence time.

To measure channel coherence time, we studied the channel variation behavior of several randomly selected links for node deployments considered in this paper. For each of these links, we studied the channel variation characteristics for a continuous duration of one hour by sending back-to-back preamble packets at a rate of 100 pkts/s (which is as fast as the testbed can transmit). As the receiving node receives the preamble packets, it uploads the received data to the host PC where each corresponding channel estimate is calculated and stored.

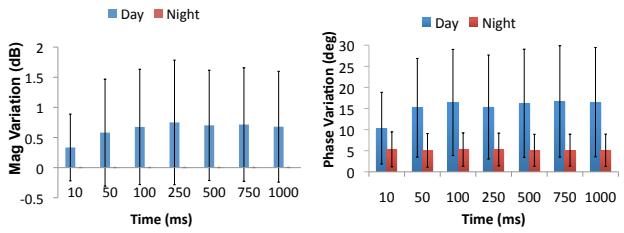

(a) (b)

Figure 4: Channel variation.

The experiments were conducted in an interference-free channel [2] and under two environmental conditions: late at night when

[2]The OTA experiments were conducted on the 802.11-2.4GHz channel 14, which consumer WiFi devices are not allowed to use in the USA.

no movement was happening in the environment, and during office hours on an average day with normal human traffic around the nodes under study. We next calculate the channel's magnitude and phase variation from our measured data sets as a function of the time interval. Fig. 4 depicts the mean and standard deviation of such changes in the two different environmental conditions for one of the links. For the rest of the links, we observed similar nighttime performance but varying daytime performance.

For the link studied in Fig. 4, during daytime experiments, a delay of 50 ms is enough to cause a mean channel magnitude variation of 0.7 dB and a phase variation of 15 degrees. Furthermore, the high standard deviation values demonstrate that there is a high unpredictability for both channel amplitude and phase estimation. Such channel variations would cause the interference term in Eq. (3) to be nonzero and would reduce the signal to interference plus noise (SINR) ratio.

On the other hand, the nighttime experiments show that the average channel magnitude change is almost zero, and the average phase change is close to 5 degrees. The standard deviations for both of these experiments are very low. As observed in Fig. 4, this behavior is independent of the time interval over which the channel estimates are calculated. Moreover, an average phase variation of 5 degrees is an inherent part of the system and exists among different packets due to the slight variations of the multiple hardware elements in the testbed. Thus, the above results guarantee that OTA measurements that are done in an interference-free channel and late at night are within the channel coherence time. We perform all of our OTA experiments in such conditions.

3.2.2 *Performance Metrics*

We use the *received signal strength* (RSS in dBm) value reported by the radio boards for performance comparison of different schemes studied in this paper. We observed that the reported RSS values among different cards can vary up to 1 dB for the same received power.

In all of our schemes, noise power is measured at the receiver prior to any packet reception. In Omni and SUBF schemes, this noise power is then subtracted from the RSS of the received packet and provides the *signal to noise ratio* (SNR) at the receiver.

In MUBF schemes, the recorded RSS value of each receiver contains the multi-user interference term in addition to the signal term as shown by Eq. (3). Thus, in order to correctly measure the signal strength, this interference should be subtracted from the received signal in addition to noise power. We use the *signal to interference plus noise ratio* (SINR) as the metric for MUBF schemes.

For a given user k, we take the following approach to measure SINR. First, we perform multi-user beamforming and measure the RSS value. Next, we redo the multi-user beamforming measurement but this time we set the power allocated to user k to zero without changing the power allocated to the rest of the users. According to Eq. (3), the measured RSS value at k is equal to the interference caused by other users plus noise power at k. By subtracting the two values, we obtain the SINR at k. In all OTA experiments, we take 10 SINR measurements and report the average and standard deviation for each data point. For the channel emulator experiments, we take 1000 SINR measurements for each data point.

In addition to SNR and SINR measurements, we also use the corresponding Shannon capacity in Eq. (2) for performance comparison. The overall end to end throughput of a system is dependent on the specific MAC protocol implementation and is an active research area. Shannon capacity is a measure of physical layer capacity and is also an upper bound on the throughput that would be achieved by any MAC protocol.

4. SPATIAL MULTIPLEXING GAINS OF ZFBF

In this section, we experimentally characterize the spatial multiplexing gains of ZFBF in indoor wireless networks. We first consider a two receiver scenario and investigate the capability of ZFBF to transmit independent data streams as a function of receiver separation distance. Next, we study the impact of user selection based on link quality difference on ZFBF. Finally, we investigate the behavior of ZFBF as the number of concurrently served users increases.

4.1 Impact of Receiver Separation Distance

The performance of ZFBF is highly dependent on the channel vectors from transmitter to receivers. When different users' channel vectors are uncorrelated with one another, we expect increased multiplexing performance. As users move closer to one another, the channel vectors could become increasingly correlated, which would cause a drop in received SINR for each receiver thus lowering multiplexing gains. In [12], the authors have shown that in outdoor environments, user separation distances of up to 70 m are required to achieve the full multiplexing gains of ZFBF with two receivers.

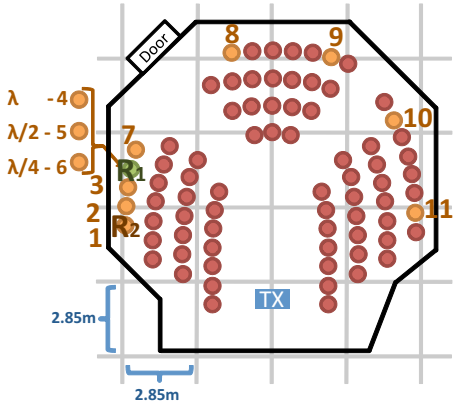

Figure 5: Experimental evaluation of spatial multiplexing as a function of receiver separation distance.

This conjecture raises the following important question: what receiver separation distance will result in a loss in multiplexing gain in indoor environments (measured in terms of aggregate capacity)?

Scenario. To answer this question, we designed an experiment shown in Fig. 5 consisting of a single transmitter and two receivers. The first receiver, R_1, is at a fixed location, while the second receiver, R_2, approaches R_1 and passes close by it before continuing around the room. For each of the location IDs in Fig. 5, we perform Omni, SUBF, ZFBF-EP, and ZFBF-MT transmissions toward the receivers. The experiment is conducted in a large classroom with many metallic chairs that cause reflections and multi-path scattering. The transmitted signal has a Line-of-Sight (LOS) component to both receivers.

Fig. 6, depicts the mean and standard deviation of the aggregate capacity as a function of R_2's location. For all locations, SUBF provides an average of 7 dB improvement over Omni. This results in a small capacity improvement for SUBF since both links have an average Omni SNR of 19 dB and thus an additional 7 dB does not increase capacity by much due to the logarithmic capacity function.

Fig. 6 reveals that the performance of the ZFBF scheme does not depend on the separation distance between the two receivers. This

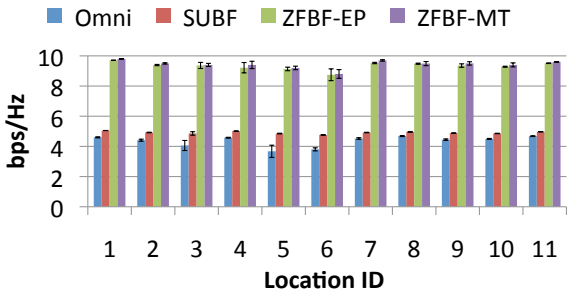

Figure 6: Capacity as a function of location.

is specifically observed at locations 4, 5, and 6, where the physical distances between the two receivers are equal to λ, $\frac{\lambda}{2}$, and $\frac{\lambda}{4}$ respectively. At the 6th location ($\frac{\lambda}{4}$), the bases of the two receivers' antennas are physically touching each other meaning that the nodes cannot be placed any closer. However, even with adjacent antennas, we still observe an unchanged aggregate capacity. We repeated this experiment in another indoor environment in which the transmitter lacks a LOS component to either receiver and measured the multiplexing gains of ZFBF as R_2 moves toward R_1 and passes close by it. For all of these experiments, we observed that the multiplexing gain does not change even when the receivers are placed at a quarter of wavelength from each other.

Finding: The spatial multiplexing gain of ZFBF with a four-antenna transmitter and two single-antenna receivers does not depend on the separation between the two receivers (down to a minimum of a quarter of a wavelength). The rich scattering characteristics of the indoor environment, the intrinsic randomness in each receiver's hardware implementation, and a higher number of antennas at the transmitter result in constant multiplexing gains irrespective of user separation distance.

4.2 Impact of User Selection

One of the key issues that is closely related to the performance of ZFBF is concurrent user selection. Because zero forcing beam weights are computed for a set of users as shown in Eq. (4), a particular receiver's SINR could vary depending on its partnered receivers. In this section, we investigate the performance of a single link's behavior as it is scheduled with different users with heterogeneous link qualities.

Scenario. Fig. 7(a), depicts our experimental setup in which we deployed six nodes in an office environment. Nodes 1, 2, and 3 are each equipped with four antennas and thus can be used as transmitters or single-antenna receivers. We select one of these three nodes as the transmitter and consider all possible two-receiver combinations from the remaining five nodes. For all of these sub-topologies, we measure the SNR at each receiver from Omni and SUBF transmissions, and the SINR at each receiver from a jointly beamformed transmission. We repeat this experiment for all possible transmitter-receiver pairs.

Fig. 7(b) shows the SNR variation of each link in Fig. 7(a), when the link is scheduled with any other link in the network simultaneously. The x-axis of Fig. 7(b) represents a given link's measured Omni SNR. The y-axis shows the SNR value of the same link for the indicated schemes.

For a selected link l, there are four remaining links that can be scheduled simultaneously with l when using the ZFBF-EP scheme. Thus, for the ZFBF-EP results, we plotted the average SNR of l, when combined with each of the four other links. The thicker red bars indicate full range of l's SNR when combined with different links. The dashed green bars show the full range of the other links' measured Omni SNRs.

Fig. 7(b) also indicates a single link's SNR value when SUBF is used. According to this graph, SUBF provides an average gain of 7.5 dB compared to Omni with minimum and maximum gains of 2 and 20 dB respectively. In all of the Omni transmissions, the transmitter always uses its first antenna for packet transmission. If the path from this antenna has a low gain compared to the other antennas, the Omni link SNR value will be low. On the other hand, SUBF uses all of the antennas at the transmitter and thus can leverage antennas with higher path gains while beamforming. This would significantly increase the SNR as is observed in the first data point of Fig. 7(b).

In the ZFBF-EP scheme, each link's SNR value is below that of SUBF and greater than or equal to that of Omni. In this scheme, power is allocated equally to each user resulting in each receiver being allocated half of the overall power at the transmitter. As a result, individual links served by ZFBF-EP will always have a lower received power than SUBF. However, the results of Fig. 7(b) demonstrate that the received power remains greater than or equal to that of Omni. This demonstrates how ZFBF's selected beam weights are able to compensate for the lower power allocation at the transmitter. Similar to SUBF, ZFBF-EP greatly enhances the per-link SNR value in the low SNR region revealing the potential of these schemes to enhance network connectivity.

Fig. 7(b) also reveals information about concurrent user selection. Our results show that each link's SNR remains the same irrespective of the user that it is paired with. This is demonstrated by the low SNR variation of a given link (thick red bars), even when it is combined with different links of highly variable quality as shown by the wide ranges of the green bars.

Finding: When the number of simultaneous users is fewer than the maximum DoF at the transmitter, different receiver pairing causes at most 3-4 dB difference on each link's SNR. For a system that can tolerate this loss, the performance would not be affected by different combinations of user scheduling.

We now investigate the impact of link quality on the aggregate performance of ZFBF. Fig. 7(c) plots the aggregate capacity ratio of SUBF to Omni and ZFBF (equal power and maximum throughput) to Omni for all two-receiver sub-topologies of Fig. 7(a). We consider equal time share for each receiver in the Omni and SUBF schemes. Low Omni capacity values correspond to low link qualities at the receivers. As observed in Fig. 7(b), both SUBF and ZFBF can significantly enhance SNR in this region and thus increase aggregate throughput. Furthermore, ZFBF serves two users simultaneously, thus benefiting from its ability to multiplex users.

When both links have high Omni SNR values, SNR gain over Omni due to SUBF and ZFBF would only slightly increase the capacity of each link due to the logarithmic capacity function. Thus, SUBF performs similarly to Omni, whereas ZFBF benefits from its ability to multiplex users. This behavior is observed in Fig 7(c) when Omni capacity is above 4 bps/Hz.

Fig. 7(c) also reveals that the capacity achieved by the two power allocation schemes is very close to one another. In order to quantify the difference between the equal power (EP) and maximum throughput (MT) schemes, we measured the SNR difference between the two schemes for all two-link sub-topoogies of Fig. 7(a). The average SNR difference and its standard deviation are equal to 1.53 and 0.42 dB respectively. With minimum ZFBF SNR values of 15 dB, such variations would cause a slight difference in

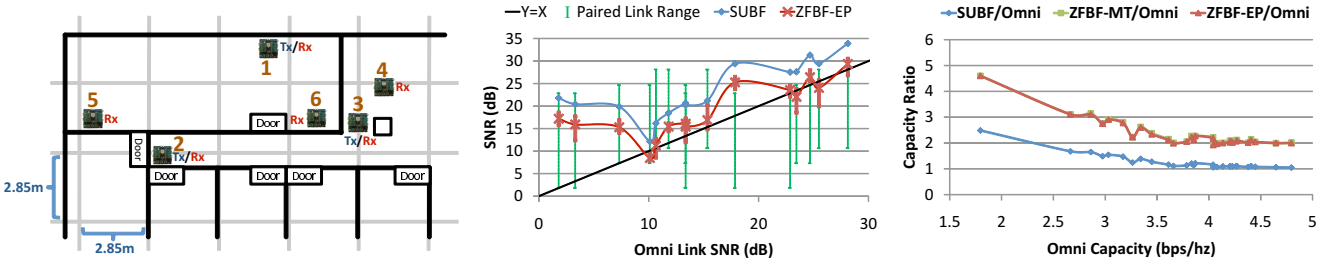

(a) Map of the office environment. (b) Per link SNR variation (c) Aggregate capacity ratio

Figure 7: Impact of concurrent user selection.

the achieved capacity. This behavior is observed in the aggregate capacity results of Fig. 7(c).

Finding: In a low SNR region, ZFBF and SUBF can significantly enhance the receiver's SNR resulting in large gains compared to Omni. With higher link qualities, SUBF only causes a small capacity improvement over Omni, whereas ZFBF benefits from user multiplexing and thus causes a 2x capacity improvement.

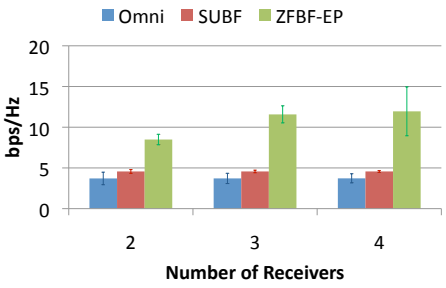

(a) Aggregate Capacity

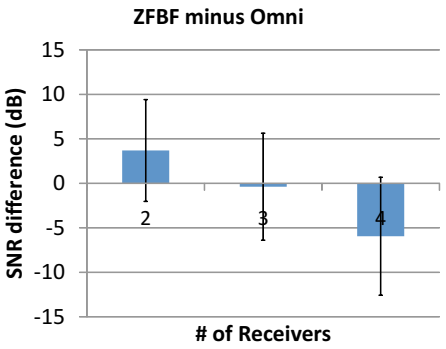

(b) Per link SNR difference

Figure 8: Impact of user population size.

4.3 Impact of User Population Size

We now investigate the performance of ZFBF as the number of served users approaches the number of transmitter antennas. We use the same node deployment setup of Fig. 7(a) and perform the same set of experiments as in the previous subsection. However, instead of serving two users, we evaluate the performance of Omni, SUBF, and ZFBF-EP as the transmitter serves two, three, or four users.

Using the measured SNRs of each link for the Omni, SUBF, and ZFBF schemes, we compute each sub-topology's aggregate capacity. Next, we group the sub-topologies based on receiver population size and calculate the average capacity for each group in Fig. 8(a). In addition, we find the average per-link SNR difference between ZFBF and Omni for each user population size as shown in Fig. 8(b).

Fig. 8(a) shows that Omni and SUBF capacities remain constant regardless of user population size because the net capacity is simply the average of each per-link SNR. Therefore even if user population size increases, the average of all possible topologies will remain the same. In ZFBF, we observe a considerable capacity improvement from 2 to 3 concurrent users, however only a marginal improvement from 3 to 4 users.

On the other hand, Fig. 8(b) reveals that as we increase the number of receivers, ZFBF's relative per-link SINR gain over Omni decreases. ZFBF's per-link SINR is several dB greater than Omni for the two-receiver case. However, for the three receiver case, the per-link SINR gain over Omni is essentially 0 while the SINR for the four receiver case is almost 6 dB below that of Omni.

Finding: The aggregate capacity of ZFBF saturates as the number of served users approaches the DoF at the expense of a significant drop in per-link SINR. Thus, the number of users ZFBF can serve depends on the link quality constraints of the individual user.

5. EFFECTS OF CHANNEL VARIATION

Thus far, the experiments were conducted with perfect channel information at the transmitter. However, in practice, channel information can become outdated for multiple reasons. For example, as observed in Fig. 3, even with fixed wireless endpoints, the mobility of objects or people in the environment can cause significant channel variation. Furthermore, a device's mobility can outdate a channel estimate by the time it is used to transmit beamformed data. Inaccurate channel information can destroy the zero-interference condition of the selected beams, potentially rendering the packets undecodable. Therefore, it is crucial to understand the effects of channel update rate and variation on overall performance. In this section, we explore the effects of channel variation on ZFBF performance.

Scenario. In order to have consistent and precise control over the channel and its variability, we use a channel emulator. Fig. 9 depicts the setup over which the experiments were conducted. The four-antenna transmitter and two single-antenna receivers are connected to the Azimuth ACE 400WB Channel Emulator [2]. The boards and channel emulator are connected to the host PC that manages the transmission of the boards and channel profile used by the channel emulator. The channel profile parameters used by the channel emulator are shown in Table 2. The channel model is adapted from 802.11n task group (TGn) models used to evaluate the performance of MIMO in indoor environments [4]. This channel model is composed of nine Non-Line-of-Sight (NLOS) Rayleigh fading

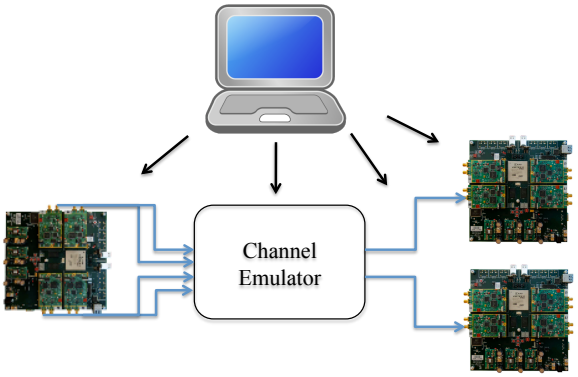

Figure 9: Channel emulator setup.

Parameter	Value
Number of multi-paths	9
Fading model per path	Rayleigh
Delay per path (ns)	0, 10, 20, 30
	40, 50, 60, 70, 80
Path loss per path (dB)	0, 5.428, 2.516, 5.890, 9.160
	12.510, 15.612, 18.714, 21.816

Table 2: Channel model parameters.

paths and is used to emulate a typical, residential environment. The channel emulator is configured to output an average SNR value for each receiver while varying the instantaneous SNR according to environmental variation or user mobility.

We investigate two issues with this setup. First, we consider static nodes to characterize the performance of ZFBF as a function of environmental variation. Next, we emulate mobile receivers in order to characterize the impact of user mobility on ZFBF's performance.

5.1 Impact of Environmental Variation

In this section, we quantify the performance of ZFBF as a function of environmental variation and channel estimation delay. The 802.11n taskgroup uses the Doppler fading rate interval of [0.028 2.778] Hz as the quantitative metric for environmental variation [4]. We performed two sets of experiments using Doppler fading rates of 1.157 and 2.778 Hz to emulate typical (T) and rapidly (R) varying environments respectively. For each of these experiments, we varied the time interval between the channel estimate measurement and actual data transmission.

Fig. 10(a) depicts the sum-rate performance of Omni, SUBF, and ZFBF for the two fading rates as a function of channel estimation delay. The solid lines in this figure correspond to a typically varying environment while the dashed lines correspond to a rapidly varying environment. We observe that Omni's capacity remains similar irrespective of environmental variation or channel estimation delay. Omni does not require channel information and thus its performance does not change with channel estimation delay. Furthermore, when run for a long time, the average output Omni SNR would remain the same regardless of environmental variation or user mobility.

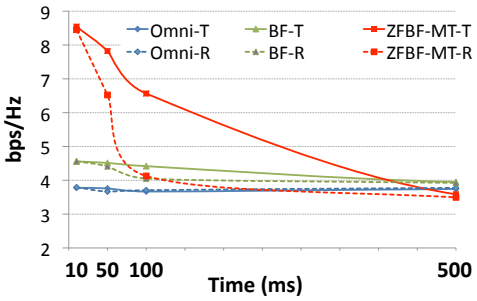

(a) Aggregate Capacity

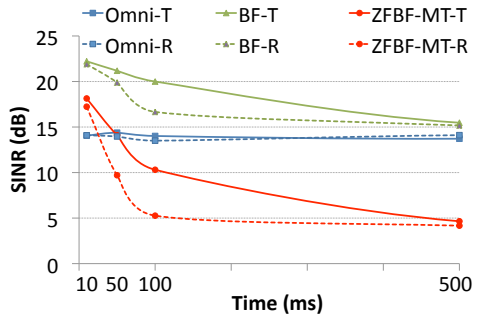

(b) Average per-link SINR.

Figure 10: Impact of environmental variation.

On the other hand, the SUBF scheme is vulnerable to inaccurate channel estimate information. SUBF requires accurate channel information at the transmitter to form a beam that maximizes SNR at its receiver. According to Fig. 10(a), the performance of SUBF becomes equivalent to Omni with a time interval of 500 ms. Additional increases in the time interval further decreases the performance of SUBF compared to Omni.

Fig. 10(a) indicates that the ZFBF scheme is highly dependent on accurate channel information. In the rapidly varying environment, the aggregate capacity decreases sharply, while both environments demonstrate an aggregate capacity equivalent to Omni at a 500 ms update rate.

Note that in the ZFBF scheme, both receivers are served at the same time. As a result the capacity of this scheme benefits from multiplexing the two users. Thus, while aggregate capacity of this scheme could be equal to or higher than Omni, per-link SINR values could be significantly lower [3]. In Fig. 10(b), we measured the average per-link SINR value for all of these schemes. Fig. 10(b) reveals that per-link SINR value is 10 dB less than Omni at a channel update rate of 500 ms. Thus, a link's SINR region must be considered to identify the necessary channel update rate. In a high SNR region, such power reduction due to environmental variation could be tolerated by the system, whereas with lower link qualities such variation would not.

Finding: The necessary channel update rate with static devices depends on environmental variation as well as link quality. Assuming links can tolerate an SNR decrease of up to 3 dB compared to Omni, a maximum channel update rate of 100 ms is required to guarantee acceptable performance in a typical indoor environment.

[3]From Section 2, recall that instead of multiplexing users, Single-User schemes link Omni and SUBF schedule users sequentially according to a TDMA schedule

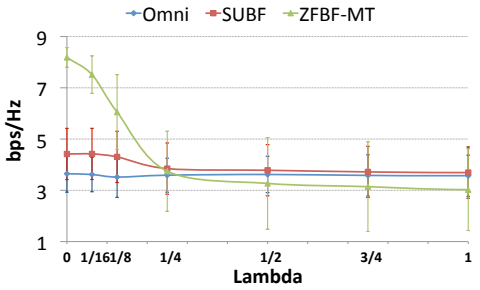

(a) Aggregate Capacity

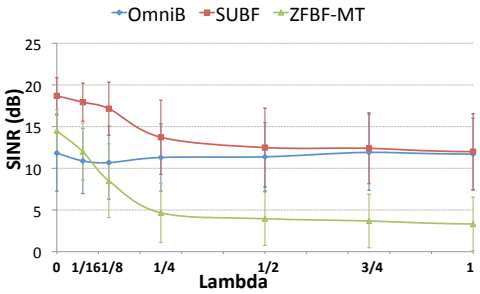

(b) Average per-link SINR

Figure 11: Impact of mobility.

5.2 Impact of User Mobility

We now investigate the effects of channel variation due to user mobility. Mobile users would travel some distance between the time a transmitter obtains a channel estimate and actually transmits beamformed data, thus causing channel variation. The channel variation due to user mobility can significantly increase the multi-user interference and reduce the effectiveness of spatial multiplexing.

We perform controlled experiments to quantify the drop in throughput as a function of user mobility. We use the same experiment setup as shown in Fig. 9; however, we instruct the channel emulator to change the channel for both receivers as a function of the distance that the users have moved. The channel emulator is configured such that users have equivalent speeds although their movement direction is random and independent from one other.

Fig. 11(a) plots the aggregate capacity of different schemes as a function of movement distance in number of wavelengths by the receivers. Omni remains robust irrespective of user mobility; however, SUBF and ZFBF are both highly dependent on receiver movement distance.

Fig. 11(a) shows that a user movement of $\frac{\lambda}{4}$ drops the aggregate capacity of SUBF and ZFBF to that of Omni. Additional increases in the movement distance would further decrease the performance of SUBF and ZFBF. However, Fig. 11(b) shows how the implications of this drop are different for the per-link SNR. For ZFBF, at $\frac{\lambda}{4}$, the average SNR of each link drops 6 dB below that of Omni; whereas, for SUBF, the average SNR of each link remains 3 dB above that of Omni. Thus, in a low SNR region, ZFBF's per-user capacity would be significantly lower than Omni and SUBF.

Finally, the channel model considered in these experiments has been restricted to NLOS environment. We have also investigated the impact of having a LOS component, where a user may be able to move a greater distance before a change in the channel occurs. In these experiments we observed the same behavior as NLOS experiments.

Finding: ZFBF is vulnerable to channel changes due to user mobility. Assuming links can tolerate SINR losses of up to 3 dB compared to Omni, user movement distance of up to $\frac{\lambda}{8}$ is acceptable. At 2.4 GHz, this is equivalent to 1.56 cm. With a typical pedestrian speed of 3 mph, this is equivalent to channel update rate of approximately 10 ms.

6. IMPACT OF BEAMFORMING ON SPATIAL REUSE

We now investigate the increase in spatial reuse opportunities offered by MUBF. In Section 6.1, we consider a single sender/receiver pair and a third node, W, at which we attempt to minimize the interference caused by the initial pair's transmission. We quantify the reduction in interference as a function of W's location. Next, in Section 6.2, we investigate the ability for a sender to reduce its transmission footprint by minimizing interference at multiple unintended receivers simultaneously. Finally, in Section 6.3, we consider a scenario with multiple sender/receiver pairs and investigate the impact of the senders' cooperation on reducing interference at each other's receivers compared to Omni-mode transmission.

6.1 Interference Reduction as a Function of Location

The multi-element antenna array at the transmitter can be used to increase SNR at the receiver(s), while suppressing interference at multiple other users (unintended receivers). In ZFBF, this is achieved by obtaining channel information from all receivers and calculating the appropriate beam weights; however, zero power is allocated to the unintended receivers' beams while the total power budget is given to the intended receiver(s). With one intended receiver R, and one unintended receiver W, the resulting beam would point toward R while causing no interference at W. We investigate the ability of ZFBF to reduce interference as a function of W's location.

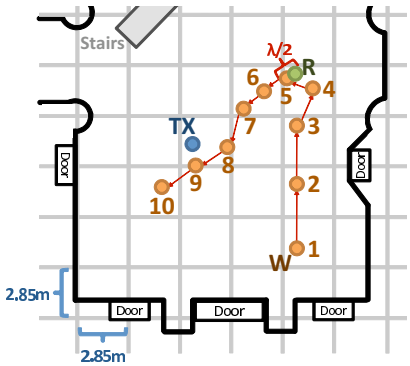

Figure 12: Experimental Scenario.

Scenario. Our experimental scenario is depicted in Fig. 12. The transmitter, TX, sends data to its receiver, R, such that the resulting interference at W is minimized. We investigate three different movement patterns of W. First, we start with a fixed distance between W and R, and move toward R along the line connecting the two points (location IDs 1 to 4). Second, we place W and R adjacent to one another and move W along the line connecting the three nodes (location IDs 5 to 7). Finally, we investigate the ability of ZFBF to cancel interference at W as it is moved closer to the transmitter (location IDs 8 to 10). For each of these locations,

we take the following measurements: First, we perform an Omni transmission from TX to R and record the received signal strength at W. Next, we perform joint beamforming with the objective of zero interference at W and measure the resulting signal strength at W.

Fig. 13 shows the resulting interference at W for each of the location IDs. In Omni mode, we observe high interference values at locations 1 to 7. As W moves closer to the transmitter, the amount of interference increases.

The ZFBF scheme causes far less interference than Omni. The resulting interference caused by ZFBF has an average of 1.1 dB above the noise floor for all of the location IDs. Fig. 13 also shows that even when TX, W, and R are on the same line, or as W approaches TX, the ZFBF scheme is still able to cancel interference at W.

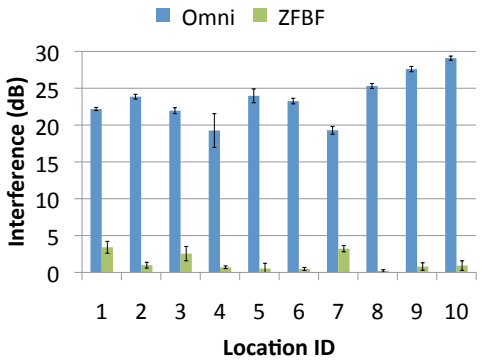

Figure 13: Interference reduction as a function of location.

Finding: A user can obtain an interference-free channel by sharing its channel information to a ZFBF-enabled transmitter. The interference-free channel is obtained irrespective of the distance between the user and either the transmitter or the receiver.

6.2 Multi-Point Interference Reduction

In this section, we evaluate MUBF's interference suppression performance when the transmitter communicates with an intended receiver while attempting to minimize interference at multiple unintended receivers.

We consider the node location setup described in Fig. 7(a). Nodes 1, 2, and 3 each have four antennas and thus can be used as four-antenna transmitters or single-antenna receivers. We select one of these nodes as the transmitter and one of the remaining nodes as the intended receiver. Then, we consider all possible combinations of 1, 2, or 3 nodes among the remaining nodes as locations at which we plan to minimize interference. We repeat this experiment for all possible transmitter-receiver pairs leading to 210 different sub-topologies. We perform Omni, SUBF, and ZFBF transmissions, and measure the resulting signal strength at the intended receiver as well as unintended receivers.

Fig. 14 shows the interference footprint for the three schemes. We first investigate the performance of Omni and SUBF. Fig. 14(a) shows the scatter plot of interference at unintended receivers with Omni and SUBF schemes. Each point in this graph corresponds to a sender-unintended receiver pair. From this plot, similar performance is observed between the Omni and SUBF schemes. For half of these locations, the resulting interference of SUBF is higher than that of Omni, whereas, for the other half, the Omni interference is higher.

Finding: SUBF obtains channel information from its intended receiver without regard to any other user. The corresponding beam pattern would cause a high SNR at the intended receiver, while the resulting interference would be location dependent. This interference could be significantly higher or lower than an Omni transmission and is dependent on the environment and location of the unintended receivers.

The interference reduction performance of ZFBF is shown in Table 3, where we present the measured mean and standard deviation of interference caused at the unintended receivers. Similar to the results of Fig. 13, we observe that the resulting interference is close to the noise floor power. However, unlike Fig. 13, these results are obtained as the transmitter used up all of its DoF. Thus, we conclude that the interference suppression capabilities of ZFBF are not constrained by the number of DoF used. The transmitter can efficiently construct beamforming weights that cause minimal interference at unintended receivers.

Although ZFBF's interference cancellation ability does not depend on the number of DoF used, there is a potential impact on the received signal strength at the intended receiver. We investigate this behavior in Fig. 14(c). Here, we compare the SINR of ZFBF to Omni and SUBF schemes at the intended receiver as we increase the number of unintended receivers. Note that in this case, the SINR of SUBF and Omni remains constant since the receiver's SINR does not depend on the number of unintended receivers, whereas ZFBF's SINR does.

From our measurements, we present the average and standard deviation of $SINR_{SUBF} - SINR_{ZFBF}$ along with $SINR_{ZFBF} - SINR_{Omni}$. With only one unintended receiver, the performance of ZFBF is close to that of SUBF and higher than that of Omni. As the number of unintended receivers increases, the SINR of ZFBF decreases at the intended receiver. When all DoF of the ZFBF scheme are used, we observe that ZFBF's SINR is on average 0.5 dB lower than Omni. The high standard deviations indicate that the SINR could decrease up to 8 dB below Omni as the ZFBF scheme uses all of its DoF. The resulting drop in capacity of the served links depends on the Omni SNR value. In a high SINR region, such a drop in signal strength would result in a small decrease in link capacity, whereas in a lower SNR region, the link capacity decrease would be more significant.

Finding: ZFBF's interference reduction capabilities do not depend on the location of the receivers nor the number of DoF used. However, the increase in the number of unintended receivers decreases the link quality of the intended user(s). When all DoF are used, the performance of a given user can significantly drop below that of an Omni transmission.

6.3 Impact of Multi-User Beamforming on Network Throughput

We now investigate the potential of ZFBF to increase network capacity by minimizing interference between concurrent links. We create 36 different sub-topologies consisting of two sender-receiver pairs for the node setup shown in Fig. 7(a). For each of these sub-topologies, we first calculate the overall maximum capacity of the SUBF and Omni schemes. This overall maximum capacity is the maximum of the single-link capacities and the sum capacity of the two links when the two links are active simultaneously. With ZFBF, both flows are active simultaneously and thus the transmitters jointly beamform such that the resulting interference at the other flow's receiver is minimized.

Fig. 15, shows the relative capacity improvement of SUBF and ZFBF compared to Omni. We sort the sub-topologies based on increasing SUBF capacity ratio. For the first and last three sub-

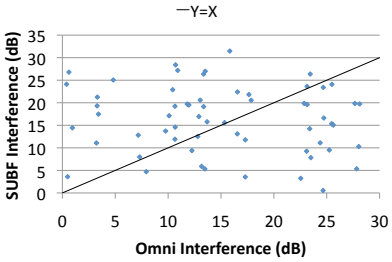

(a) Omni vs. SUBF interference

# of ZF Points	Interference (μ, σ)
1	(0.93, 0.85)
2	(0.73, 0.64)
3	(0.84, 0.92)

(b) Table 3: ZFBF interference (dB)

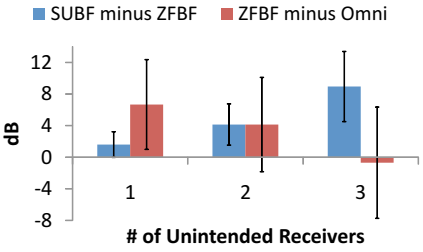

(c) SINR difference at the receiver

Figure 14: Multi-point interference reduction

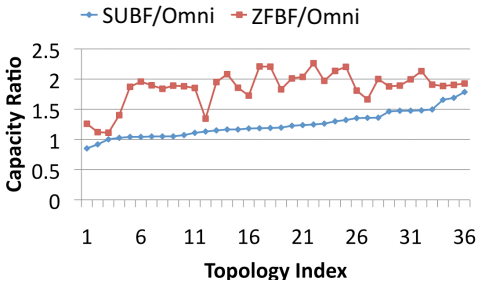

Figure 15: Maximum Capacity of two flows.

topologies, ZFBF performs close to SUBF. Careful investigation of these sub-topologies revealed that for the first three topology indicies, a high Omni capacity is achieved when both links are active simultaneously. However, SUBF causes significant interference at the other flow's receiver and thus achieves its maximum capacity when only the highest capacity link is active. Thus, Omni outperforms SUBF for these sub-topologies. On the other hand, for these sub-topologies, Omni causes less interference at the other flow's receiver and therefore ZFBF does not benefit from its interference reduction capabilities and achieves a performance close to Omni.

For the last three sub-topologies, Omni achieves its maximum capacity when only one link is active. In these topology indicies, SUBF causes less interference at the other flow's receiver and achieves its maximum throughput when both links are active at the same time. This results in a high capacity ratio of SUBF compared to Omni. In these sub-topologies, ZFBF reduces the remaining interference thus slightly increasing the capacity. For the rest of the sub-topologies, a high mutual interference exists among the flows for the Omni or SUBF schemes. As a result, ZFBF is able to benefit by reducing mutual interference resulting in a high performance gain.

Finding: In a network with multiple sender-receiver pairs, ZFBF can reduce mutual interference allowing for sender-receiver pairs to transmit simultaneously thus increasing the overall throughput. As the amount of mutual interference for the Omni or SUBF schemes decreases, the performance gain of ZFBF decreases compared to these other schemes. With SUBF, the overall network capacity could decrease compared to Omni due to increases in mutual interference.

7. RELATED WORK

Single-User MIMO. Single-User MIMO systems, for example 802.11n [4] and BLAST [10], enhance the capacity of a point-to-point communication link. When users have a smaller number of antennas than the base station, the system capacity is constrained by the receiver antennas. However, MU-MIMO schemes can benefit from the full number of antennas at the transmitter with a high number of users. In SUBF, multiple antenna elements are used at the transmitter to increase a single link's SNR values. In [13], Lakshmanan et al. have implemented an SUBF platform to evaluate its performance in indoor wireless networks. As a baseline for comparison, we have also implemented the SUBF scheme.

Spatial Reuse. Prior research has used directional [6, 14, 15] or sectorized [16] antennas to increase SNR at intended receivers while increasing spatial reuse. [15] performs experimental characterization of multi-antenna arrays in outdoor environments, while [6, 14, 16] investigate the spatial reuse capabilities of the aforementioned antenna technology in indoor wireless networks. Similarly, we have investigated the benefits of ZFBF at reducing interference. However, in contrast to all prior work, we experimentally show that ZFBF is able to increase SNR at intended receivers while eliminating interference at any undesired location.

Theoretical Work on Multi-User MIMO. Extensive theoretical research exists on the subject of MU-MIMO. Information theory results [7] have shown that DPC [8, 19] is the optimal strategy in MIMO downlink channels. However, DPC is difficult to implement due to extensive computational complexity. ZFBF [21] is a simple strategy to serve multiple users simultaneously and it achieves a large fraction of DPC capacity. There are several papers on ZFBF focusing on different design criteria (for a comprehensive survey refer to [20]). The performance of these schemes are usually calculated under simulated channel conditions with uncorrelated channel gains. In [12], Kaltenberger et al. use measured channel gains to evaluate the aggregate performance of ZFBF in outdoor environments. In contrast to all, we have designed a platform to experimentally evaluate the performance of ZFBF in indoor wireless networks.

Practical MU-MIMO Protocols. Arraycomm [1] has built outdoor, cellular base stations with twelve antennas that can create up to four spatial channels. However, we have designed an open experimental framework for the prototyping and implementation of various MUBF algorithms. In addition, we have measured the performance of ZFBF in indoor wireless networks and have explored various factors that affect its performance.

Recent work [11, 17, 22] has proposed MU-MIMO protocols to increase network capacity. IAC [11] improves the capacity of wireless LANs assuming that the access point's (AP) number of antennas is the bottleneck. IAC allows collaboration between APs such that multiple AP-client pairs can concurrently transmit. In our work, the system bottleneck is the number of antennas at the receiver. Thus, we consider the issue of using an AP's antennas to serve multiple users simultaneously. SAM [17] addresses the

problem of serving multiple users with a single AP; however, this work considers the uplink channel problem. In contrast, we consider the downlink channel problem. In [22], Zhang et al. propose algorithms to solve the scheduling problem for a ZFBF-enabled transmitter. In our work, we identified the factors that affect the performance of ZFBF and evaluated their impact through indoor experiments.

8. CONCLUSION

In this work, we designed and implemented a custom MUBF platform that allows for the experimental evaluation of different beamforming strategies. Using this platform, we experimentally evaluated the multiplexing gains of ZFBF as a function of receiver separation distance, user selection, and user population size. We experimentally showed that a four-antenna, ZFBF-enabled transmitter is able to simultaneously transmit to two users that are within a quarter of a wavelength of one another. We also showed that when the number of scheduled users is fewer than the maximum DoF at the transmitter, different receiver pairings cause at most 3-4 dB difference on each link's SNR. We also evaluated the impact of user mobility and environmental variation on the performance of ZFBF. We showed that the required channel information update rate is dependent on environmental variation and user mobility as well as a per-link SNR requirement. Assuming that a link can tolerate an SNR decrease of 3 dB compared to Omni, the required channel update rate is equal to 100 and 10 ms for typical non-mobile receivers and mobile pedestrian speeds of 3 mph respectively. Finally, we investigated the potential of ZFBF to reduce interference at unwanted locations and increase spatial reuse. Our results showed that a ZFBF-enabled transmitter is able to minimize interference at any undesired location(s); however, this may come at the expense of a significant drop in the quality of the served users.

9. ACKNOWLEDGEMENTS

We would like to thank Marios Kountouris and Melissa Duarte for their help and fruitful discussions on MIMO physical layer issues. We would also like to thank our shepherd, Kyle Jamieson, and the anonymous reviewers for their valuable feedback, which helped in improving the presentation of the paper. This work was supported by NSF grants CNS-0751173 and CNS-0832025, by Texas Instruments, and by the OPNEX project of the European Community Seventh Framework Programme (FP7-ICT-224218).

10. REFERENCES

[1] ArrayComm. Available at: `http://www.arraycomm.com`.

[2] Azimuth Systems. Available at: `http://www.azimuthsystems.com/`.

[3] Rice University WARP project. Available at: `http://warp.rice.edu`.

[4] Antenna selection and RF processing for MIMO systems. *IEEE 802.11-04/0713r0*, 2004.

[5] J.G. Andrews, A. Ghosh, and R. Muhamed. *Fundamentals of WiMAX: Understanding Broadband Wireless Networking*, chapter Multiple-Antenna Techniques. Prentice Hall, 2007.

[6] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sunderesan. On the effectiveness of switched beam antennas in indoor environments. In *Proceedings of the 9th international conference on Passive and active network measurement*, Cleveland, Ohio, April 2008.

[7] G. Caire and S. Shamai. On the achievable throughput of a multiantenna gaussian broadcast channel. *IEEE Transactions on Information Theory*, 49(7):1691–1706, July 2003.

[8] M. Costa. Writing on dirty paper (corresp.). *IEEE Transactions on Information Theory*, 29(3):439–441, May 1983.

[9] D. Gesbert, F. Tosato, C. Rensburg, and F. Kaltenberger. *LTE, The UMTS Long Term Evolution: From Theory to Practice*, chapter Multiple Antenna techniques. Wiley and Sons, 2009.

[10] G.D. Golden, R.A Valenzuela, P.W. Wolniansky, and G.J. Foschini. V-BLAST: an architecture for realizing very high data rates over the rich-scattering wireless channel. In *Proccedings of URSI International Symposium on Signals, Systems, and Electronics*, Pisa, Italy, September 1998.

[11] S. Gollakota, S. Perli, and D. Katabi. Interference alignment and cancellation. In *Proceedings of ACM SIGCOMM*, Barcelona, Spain, August 2009.

[12] F. Kaltenberger, M. Kountouris, D. Gesbert, and R. Knopp. On the trade-off between feedback and capacity in measured MU-MIMO channels. *IEEE Transactions on Wireless Communications*, 8(9):4866–4875, September 2009.

[13] S. Lakshmanan, K. Sundaresan, R. Kokku, A. Khojastepour, and S. Rangarajan. Towards adaptive beamforming in indoor wireless networks: an experimental approach. In *Proceedings of IEEE INFOCOM Mini-Conference*, Rio de Janeiro, Brazil, April 2009.

[14] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste. DIRC: increasing indoor wireless capacity using directional antennas. In *Proceedings of ACM SIGCOMM*, Barcelona, Spain, August 2009.

[15] V. Navda, A.P. Subramanian, K. Dhanasekeran, A. Timm-Giel, and S.R. Das. Mobisteer: Using steerable beam directional antenna for vehicular network access. In *Proceedings of ACM MobiSys*, San Juan, Puerto Rico, June 2007.

[16] A. Prabhu, H. Lundgren, and T. Salonidis. Experimental characterization of sectorized antennas in dense 802.11 wireless mesh networks. In *Proceedings of ACM MobiHoc*, New Orleans, Louisiana, May 2009.

[17] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. Voelker. SAM: enabling practical spatial multiple access in wireless LAN. In *Proceedings of ACM MobiCom*, Beijing, China, September 2009.

[18] S. Venkatesan and H. Huang. System capacity evaluation of multiple antenna systems using beamforming and dirty paper coding. *Bell Labs*, 2003.

[19] H. Weingarten, Y. Steinberg, and S. Shamai. The capacity region of the Gaussian multiple-input multiple-output broadcast channel. *IEEE Transactions on Information Theory*, 52(9):3936–3964, September 2006.

[20] A. Wiesel, Y.C. Eldar, and S. Shamai. Zero-forcing precoding and generalized inverses. *IEEE Transactions on Signal Processing*, 56(9):4409–4418, September 2008.

[21] T. Yoo and A. Goldsmith. On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming. *IEEE Journal on Selected Areas in Communications*, 24(3):528–541, March 2006.

[22] Z. Zhang, S. Bronson, J. Xie, and H. Wei. Employing the one-sender-multiple-receiver technique in wireless LANs. In *Proceedings of IEEE INFOCOM*, San Diego, California, March 2010.