

RICE UNIVERSITY

**Experimental and Analytical Evaluation of
Multi-User Beamforming in Wireless LANs**

by

Ehsan Aryafar

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Doctor of Philosophy

APPROVED, THESIS COMMITTEE:

Edward W. Knightly, Chair
Professor of Electrical and Computer
Engineering

David B. Jonhson
Professor of Computer Science and
Electrical and Computer Engineering

Ashutosh Sabharwal
Associate Professor of Electrical and
Computer Engineering

Houston, Texas

June, 2011

ABSTRACT

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

by

Ehsan Aryafar

Adaptive beamforming is a powerful approach to receive or transmit signals of interest in a spatially selective way in the presence of interference and noise. Recently, there has been renewed interest in adaptive beamforming driven by applications in wireless communications, where multiple-input multiple-output (MIMO) techniques have emerged as one of the key technologies to accommodate the high number of users as well as the increasing demand for new high data rate services.

Beamforming techniques promise to increase the spectral efficiency of next generation wireless systems and are 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 thesis, I experimentally and analytically evaluate the performance of adaptive beamforming techniques on the downlink channel of a wireless LAN.

To this end, I present the design and implementation of the first multi-user beamforming system and experimental framework for wireless LANs. Next, I evaluate the

benefits of such system in two applications.

First, I investigate the potential of beamforming to increase the unicast throughput through spatial multiplexing. Using extensive measurements in an indoor environment, I evaluate the impact of user separation distance, user selection, and user population size on the multiplexing gains of multi-user beamforming. I also evaluate the impact of outdated channel information due to mobility and environmental variation on the multiplexing gains of multi-user beamforming. Further, I investigate the potential of beamforming to eliminate interference at unwanted locations and thus increase spatial reuse.

Second, I investigate the potential of adaptive beamforming for efficient wireless multicasting. I address the joint problem of adaptive beamformer design at the PHY layer and client scheduling at the MAC layer by proposing efficient algorithms that are amenable to practical implementation. Next, I present the implementation of the beamforming based multicast system on the WARP platform and compare its performance against that of omni-directional and switched beamforming based multicast. Finally, I evaluate the performance of multicast beamforming under client mobility and infrequent channel feedback, and propose solutions that increase its robustness to channel dynamics.

Acknowledgments

I would like to start by thanking my advisor and mentor, Prof. Edward W. Knightly for his thoughtful comments, support and guidance during my graduate studies at Rice University. Prof. knightly has been an example of excellent professor, wonderful advisor, and an efficient manager, and I am grateful for having the chance of being his student. He has shown great care and concern for his students, and has always helped them advance and become better not only in their research, but also in their personal and professional lives.

Also, I would like to thank Ashu Sabharwal and Dave Johnson for serving on my PhD, MS, and/or 599 committees and providing novel perspectives to my work. Additionally, I would like to express my appreciation to the department chair, Prof. Behnaam Aazhang, for leading this department during my graduate studies at Rice.

I am thankful for the friendships and collaborations within our research group including Theodoros, Omer, Jingpu, Joe, Joshua, Misko, Eugenio, Tasos, Ahmed, Bruno, Naren, Ryan, Cen, and Oscar. For my work on the Multi-User Beamforming I am thankful for the collaboration with Theodoros Salonidis and Narendra Anand, and for the work on multicast beamforming, I am thankful for the collaboration with Karthik Sunderasan, Mohammad (Amir) Khojastepour, and Sampath Rangarajan

from NEC Labs.

I am thankful for Prof. Aazhang and Prof. Sabharwal efforts in creating and maintaining the CMC lab where I have had the opportunity to learn and experiment, some of which is presented in this thesis. Further, I would like to thank the entire WARP team especially Patrick Murphy, Melissa Duarte, Chris Hunter, and Siddharth Gupta for all their guidance and support on the WARP platform.

My sincere thanks go to my parents, Afssaneh and Majid, for being helpful throughout all these years, and being supportive in all the aspects of my life. They dedicated their lives to the education and advancement of their children, and I am, and will forever be, grateful for that. I am also grateful for having great sisters, Kamelia and Parisa. I would not have envisioned achieving my academic goals without my family's support.

Finally, I would like to thank the friends that I have had the chance of interacting and collaborating with. I can not name all of them; however, I have to mention a few: Farbod, Marjan, Nahal, Ahmad, Kiarash, Mehrdad, Shirin, Mina, Alidad, Amir, Alireza, Borghan, Vincenzo, Mona, Yang, Paul,

Contents

Abstract	ii
Acknowledgments	iv
List of Illustrations	x
List of Tables	xii
1 Introduction	1
1.1 Summary of Thesis Contributions	4
1.2 Thesis Overview	6
2 Background	7
2.1 System Model	8
2.2 WARPLab Research Framework	9
3 Multi-User Beamforming	12
3.1 Introduction	13
3.2 Preliminaries	16
3.2.1 Single-User Scheme	16
3.2.2 Multi-User Beamforming	18

3.3	Experimental Setup	19
3.3.1	System Implementation	20
3.3.2	Multi-User Beamforming Implementation	20
3.3.3	Measurement Setup	23
3.4	Spatial Multiplexing Gains of ZFBF	27
3.4.1	Impact of Receiver Separation Distance	27
3.4.2	Impact of User Selection	30
3.4.3	Impact of User Population Size	35
3.5	Effects of Channel Variation	37
3.5.1	Impact of Environmental Variation	40
3.5.2	Impact of User Mobility	44
3.6	Impact of Beamforming on Spatial Reuse	46
3.6.1	Interference Reduction as a Function of Location	46
3.6.2	Multi-Point Interference Reduction	49
3.6.3	Impact of Multi-User Beamforming on Network Throughput	52
3.7	Incorporation of Overhead	55
3.8	Summary	58
4	Multicast Beamforming	60
4.1	Introduction	61
4.2	Background	65
4.3	Motivation	67

4.4	Design Challenges	70
4.4.1	Determination of Adaptive Beamformers	70
4.4.2	Scheduling	72
4.4.3	Channel Dynamics and Feedback Rate	73
4.5	Design of ADAM	75
4.5.1	Components of ADAM	75
4.5.2	Problem Formulation	77
4.5.3	User Partitioning	79
4.5.4	Multicast Beamformer Design	80
4.6	System Implementation	86
4.6.1	Hardware and Software	86
4.6.2	Multicasting Framework	88
4.6.3	Implementation	90
4.6.4	Performance Metrics	92
4.7	Gains of Adaptive Beamforming	93
4.7.1	Impact of discrete rates	94
4.7.2	Algorithm Evaluation	97
4.7.3	Adaptive vs. Switched beamforming	100
4.8	Impact of Channel Dynamics	102
4.8.1	Feedback Rate and Coherence Time	103
4.8.2	Reduced Feedback and Mobility	106

4.9 Summary	110
5 Related Work	111
5.1 Multi-User Beamforming	112
5.2 Multicast Beamforming	114
6 Conclusion	117
Bibliography	121

Illustrations

2.1	System model.	8
2.2	Transmitter platform.	11
3.1	Preamble structure.	22
3.2	Channel variation.	24
3.3	Experimental evaluation of spatial multiplexing as a function of receiver separation distance.	28
3.4	Capacity as a function of location.	29
3.5	Impact of concurrent user selection.	32
3.6	Impact of population size on aggregate capacity.	35
3.7	Impact of user population size on per-link SNR difference.	36
3.8	Channel emulator setup.	38
3.9	Operations inside a forward emulator module.	40
3.10	Impact of environmental variation.	42
3.11	Impact of mobility.	45
3.12	Experimental Scenario.	47

3.13	Interference reduction as a function of location.	48
3.14	Multi-point interference reduction	50
3.15	Maximum Capacity of two flows.	53
3.16	An example MAC protocol for MUBF.	55
3.17	MUBF Throughput comparison with 802.11.	57
4.1	Adaptive vs. switched beamforming	68
4.2	Impact of user size on adaptive gain.	72
4.3	Channel variations.	74
4.4	Switched beam patterns	89
4.5	WARP board SNR-rate relation.	92
4.6	Gains of ADAM.	95
4.7	Algorithm evaluation.	97
4.8	Evaluation of switched beamforming.	99
4.9	Impact of coherence time and feedback rate on ADAM.	104
4.10	SNR-Rate for $s = \frac{50}{8}$	108
4.11	Impact of training on throughput.	109

Tables

3.1	Multi-User Beamforming Physical layer parameters.	20
3.2	Channel model parameters.	41
3.3	Model parameters.	56
4.1	WarpLab Physical layer parameters.	87
4.2	Channel model parameters.	91

Chapter 1

Introduction

Multiple-Input Multiple-Output (MIMO) communication systems have received significant attention over the past several years and are already implemented by numerous companies on commercial products. For example, over the past years, ArrayComm [1], Quantenna [2], and Xirrus [3] have developed multi-antenna access points with as many as 48 antennas on a single access point. Such multi-antenna access points potentially allow higher throughput, increased diversity, and reduced interference as they communicate with multiple wireless users.

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 multiple clients each with one or more antenna. The downlink channel of such a system has received significant attention; MU-MIMO techniques are already being adopted by the next generation of wireless standards such as 802.11 ac [4], LTE [5], and WiMAX [6] and are planned to be included in future access points [2]. In this thesis, I focus on the downlink channel of a MU-MIMO system and investigate how MU-MIMO techniques can benefit the unicast and multicast applications in a wireless LAN.

Unicast transmit beamforming is a relatively young and dynamically developing research field. In classical beamforming, a single unicast vector of interest is matched to a beamforming vector and its goal is to ensure that the inner product of the beamforming weight vector and the unicast vector of interest is large, while the inner

product of the beamforming weight vector and all other vectors is small (to minimize interference). This paradigm applies to both transmit and receive beamforming corresponding to one user.

An alternative but relative case is that of multi-user beamforming, which arises in the downlink channel of a cellular or wireless LAN network. In this case, 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 sum streams in parallel over the base station's antenna array. The beamforming weight vector designed for a given user is such that it has a large inner product with the steering vector of its user, and small inner product with the steering vector of all the other users (such that inner-user interference is minimized or eliminated).

Now, what if the transmitter intends to transmit a common information to many users? The traditional way of doing this is blind, in the sense that little or no information is available regarding the spatial distribution of users that are listening to the transmitter. This is for example true in traditional radio and TV broadcasting where the signal is emitted isotropically. Another example is today's wireless LAN deployments in which all multicast packets are transmitted in an omni-directional way. In future wireless LAN deployments, some level of feedback is available at the transmitter from different clients in the form of channel state information.

This can be utilized to boost network coverage, quality of service, and spectral

efficiency as well as to minimize interference to other systems. This is the origin of a recent line of work on multicast beamforming. As a result, multicast beamforming is now part of the current universal mobile telecommunications system long-term evolution service (UMTS-LTE) draft for next generation cellular wireless services. Similar ideas are also making their way for future wireless LAN standards such as 802.11ac [4].

1.1 Summary of Thesis Contributions

Multi-User beamforming allows for a multi-antenna enabled access point to transmit different co-channel unicast transmissions, each meant to reach the receiver of one user. 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 thesis, I 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, I 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, my results reveal that two receivers achieve close to maximum performance with a minimum separation distance of a half of a wavelength. I 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. My 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.

Beamforming techniques can also help increase the throughput when an access point wishes to transmit a common information to a group of users (multicasting). In such a scenario, adaptive beamforming techniques can help significantly reduce the length of multicast transmission and hence increase multicasting throughput. However, all prior work has considered only the beamformer design problem without considering the medium access layer into account. Further, no prior work has been implemented to show the performance of such algorithms with real channel conditions.

Towards addressing these issues, I present the design and implementation of ADAM, the first adaptive beamforming based multicast system and experimental framework for wireless LANs. ADAM addresses the joint problem of adaptive beamformer design at the PHY layer and client scheduling at the MAC layer by proposing efficient greedy algorithms that are amenable to practical implementation.

ADAM is implemented on an FPGA platform and its performance is compared against that of omni-directional and switched beamforming based multicast. Using extensive measurements in an indoor environment, ADAM's performance is also evaluated under several practical considerations including discrete transmission rates,

infrequent channel information feedback, and client mobility.

My experimental results reveal that (i) switched multicast beamforming has fundamental limitations in indoor multi-path environments, whose deficiencies can be effectively overcome by ADAM to yield gains as high as 280%; (ii) the higher the dynamic range of the discrete transmission rates employed by the MAC hardware, the higher are the gains in ADAM's performance, yielding upto nine folds improvement over omni with the 802.11 rate table; and (iii) finally, ADAM's performance is susceptible to channel variations resulting from user mobility and infrequent channel information feedback. However, I show that training ADAM's SNR-rate relation to incorporate channel feedback rate and coherence time increases its robustness to channel dynamics.

1.2 Thesis Overview

The thesis proceeds as follows. In Chapter 2, I present the system model and the WARPLab research framework. In Chapter 3, I present the detailed evaluation of multi-user beamforming in controlled indoor environments as well as with emulated and repeatable channel conditions. In Chapter 4, I propose beamforming and user scheduling algorithms for multicasting along with extensive indoor measurements. In Chapter 5, related work is discussed on both topics. Finally, in Chapter 6, I conclude by discussing the implications and future directions that result from the thesis.

Chapter 2

Background

In this chapter, I present the system model and describe the hardware platform that is used in my experiments.

2.1 System Model

I 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. 2.1.

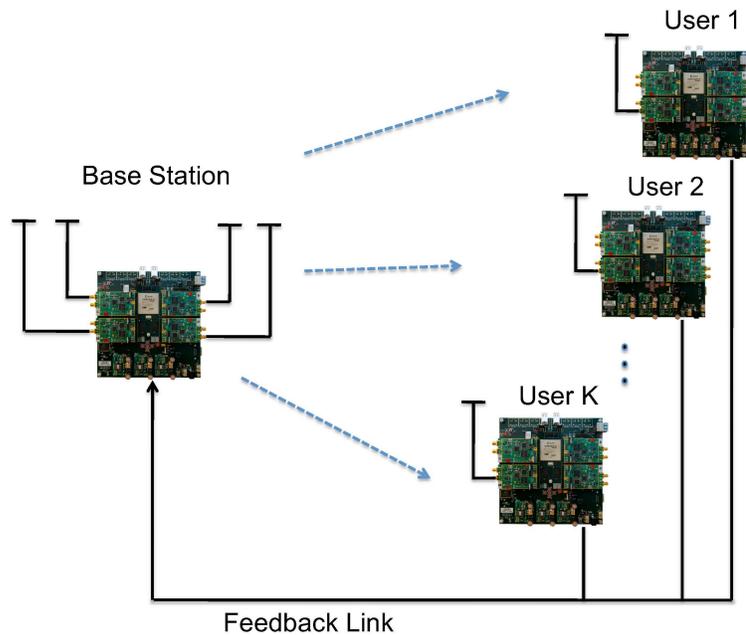


Figure 2.1 : System model.

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

$$y_k = \mathbf{h}_k \mathbf{x} + z_k, \quad k = 1, \dots, K \quad (2.1)$$

where \mathbf{x} is the transmitted symbol from the base station antennas, $\mathbf{h}_k = [h_{1k}, h_{2k}, \dots, h_{Nk}]^T$ 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., $\mathbf{x}^* \mathbf{x} \leq P$, *. The total transmit power does not depend on the number of transmit antennas and remains the same for all schemes studied in this thesis.

2.2 WARPLab Research Framework

I performed experiments using WARPLab [7], 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 on MATLAB workspace to facilitate interaction with the WARP nodes. In this framework up to 16 nodes and a host PC running MATLAB are connected to an Ethernet switch. The host PC constructs the baseband waveforms (samples) in MATLAB and stores them in the buffers of the transmitting WARP boards through the Ethernet links.

* x^* is the conjugate transpose of the transmitted symbol x .

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. 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.

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.

I used four radio boards at the base station transmitter to build a multi-antenna

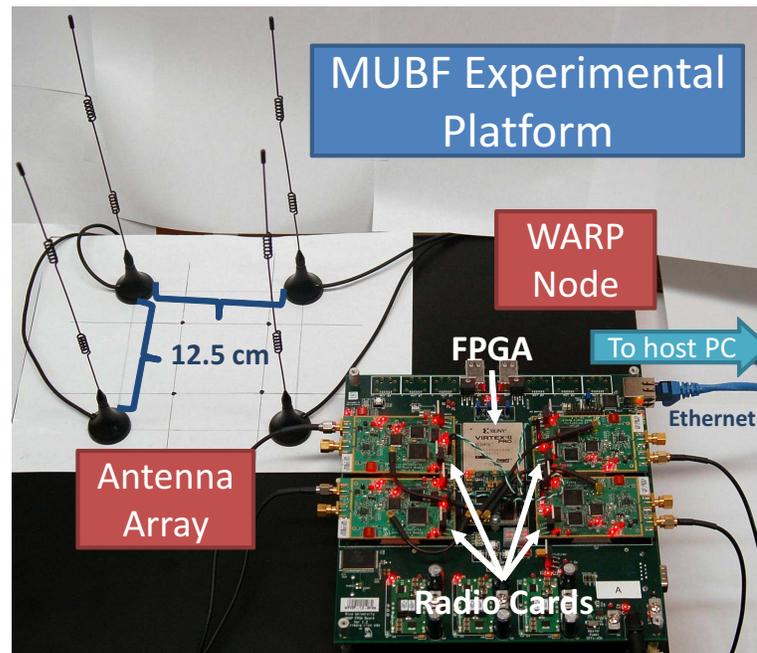


Figure 2.2 : Transmitter platform.

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. 2.2 depicts the antenna array at the transmitter connected to a WARP board. Each receiver only uses one radio board.

Chapter 3

Multi-User Beamforming

3.1 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 [8] 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 [5] and WiMAX [6].

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 [9], and several theoretical multi-user schemes have been proposed [10, 11, 12] for such systems. The optimal solution involves a theoretical pre-interference cancellation technique known as Dirty Paper Coding

(DPC) [13, 11]; however, DPC is difficult to implement due to its high computational complexity.

Multi-user beamforming (MUBF) [12] 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 chapter is to evaluate the performance of such downlink schemes in real-world deployments. To accomplish this, I design and implement a custom, FPGA-based, hardware framework that enables the evaluation of MUBF algorithms under real channel conditions. Specifically, I investigate a MUBF algorithm known as Zero Forcing Beamforming (ZFBF) [12]. I measure the performance of ZFBF as a function of receiver separation distance, concurrent user selection, and user population size. I also perform channel emulator experiments with controlled and repeatable channels to address the impact of outdated channel information due to mobility and environmental variation. I further investigate the potential of ZFBF to reduce interference at unwanted locations and increase spatial reuse. In all of my experiments I also perform TDMA-based single-user beamforming (SUBF) as the baseline.

My measurement study has the following main contributions: First, I design and implement a custom framework that allows for evaluation of different MUBF algorithms. To the best of my 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, I 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, I 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, I show that this allows for the simultaneous transmission of different data streams to users that are down to a half of a wavelength from one another.

Third, with controlled experiments performed with a channel emulator, I investigate the impact of outdated channel information due to environmental variation and user mobility on the performance of ZFBF. I 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, I find that a channel update rate of 10 ms is required for a mobile receiver with an average pedestrian speed of 3 mph.

Fourth, I investigate the potential benefits of ZFBF in reducing interference and thus increasing spatial reuse. My experimental results reveal that a user can obtain an interference-free channel by sending its channel information to a ZFBF-enabled transmitter. I 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 chapter is organized as follows: Section 3.2 provides a background of MUBF. Section 3.3 describes the design and implementation of the schemes studied in this paper. Section 3.4 describes the multiplexing gains of ZFBF. Section 3.5 investigates the impact of outdated channel information. Section 3.6 investigates the potential of ZFBF to increase spatial reuse. Section 3.7 investigates the impact of overhead on the aggregate throughput. Finally, I conclude this chapter in Section 3.8.

3.2 Preliminaries

In this section, I present background on the techniques I implemented using the WARPLab research framework.

3.2.1 Single-User Scheme

In a Single-User scheme, the base station transmits to only one user at a time in a TDMA fashion. I 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:

$$C_k = \log_2(1 + SNR_k) \quad (3.1)$$

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, I consider a round robin 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 $\sum_{k=1}^{k=K} \frac{C_k}{K}$.

3.2.2 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. (2.1) the resulting received signal vector for user k is:

$$y_k = (\sqrt{P_k} \mathbf{h}_k \mathbf{w}_k) s_k + \sum_{j \neq k} \sqrt{P_j} \mathbf{h}_j \mathbf{w}_j s_j + z_k \quad (3.2)$$

In Eq. (3.2), the first term represents the desired signal, the second term represents the multi-user interference and the third term is the noise. The receiver detects the transmitted symbol s_k by simply treating the interference terms as an additive Gaussian noise.

Finding the optimal \mathbf{w}_k s and P_k s that maximize the aggregate capacity is a difficult, non-convex optimization problem [10].

In this thesis I implement a simpler strategy known as Zero-Forcing Beamforming (ZFBF) [12]. 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. (3.2) 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 $M \subset \{1, \dots, K\}$, $|M| \leq N$ be the subset of users that the base station intends to serve concurrently, and $\mathbf{H}(M)$ and $\mathbf{W}(M)$ be the corresponding submatrices of

$\mathbf{H} = [\mathbf{h}_1^T, \mathbf{h}_2^T \dots \mathbf{h}_K^T]^T$ and $\mathbf{W} = [\mathbf{w}_1 \dots \mathbf{w}_K]$, respectively. In [14], Wiesel et al. show that the optimal choice of \mathbf{W}_M that gives zero-interference is the pseudo-inverse of $\mathbf{H}(\mathbf{M})$.

$$\mathbf{W}(\mathbf{M}) = \mathbf{H}(\mathbf{M})^\dagger = \mathbf{H}(\mathbf{M})^*(\mathbf{H}(\mathbf{M})\mathbf{H}(\mathbf{M})^*)^{-1} \quad (3.3)$$

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 thesis, I investigate two power allocation approaches with ZFBF. First, I consider the maximum throughput approach (ZFBF-MT), where the power allocation problem becomes:

$$\begin{aligned} \max_{\mathbf{p} \geq 0} \quad & \sum_k \log\left(1 + \frac{P_k}{\sigma^2}\right) \\ \text{s.t.} \quad & \sum_k P_k [(\mathbf{H}\mathbf{H}^*)^{-1}]_{k,k} \leq P \end{aligned} \quad (3.4)$$

This problem can be easily solved by using the well-known water filling solution [12]. Second, I consider a scheme that I call ZFBF-EP where the base station transmitter allocates equal power to its users. I use ZFBF-EP for a fair comparison with the round robin-based, Single-User schemes.

3.3 Experimental Setup

In this section, I describe the design and implementation of the multi-user beamforming testbed along with the conditions under which the measurements in this study

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

Table 3.1 : Multi-User Beamforming Physical layer parameters.

were performed.

3.3.1 System Implementation

3.3.2 Multi-User Beamforming Implementation

My implementation is based on the WARPLab research framework as described in the background of Chapter 2. Table 3.1 specifies the physical layer parameters used for the experiments in this chapter. 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, the 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, I will detail the three main components of the aforementioned system: Channel Training, Channel Estimation, and Beam Weight Calculation.

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. 3.1. 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° above or below room temperature. Such drifts could cause significant distortion between received and transmitted signal phase to the point where 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.

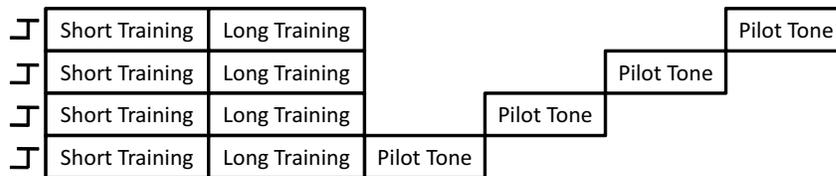


Figure 3.1 : Preamble structure.

The only difference between the antennas' preambles is the structure as is apparent in Fig. 3.1. 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. (3.3) 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.

3.3.3 Measurement Setup

In this section, I describe the conditions under which OTA transmissions were performed. First, I show that the feedback delay of the 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, I describe the metrics used to evaluate the performance of different schemes.

Channel Coherence Time

The total feedback delay in my 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 my evaluation testbed should be within the channel coherence time.

To measure channel coherence time, I studied the channel variation behavior of several randomly selected links for node deployments considered in this evaluation. For each of these links, I 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.

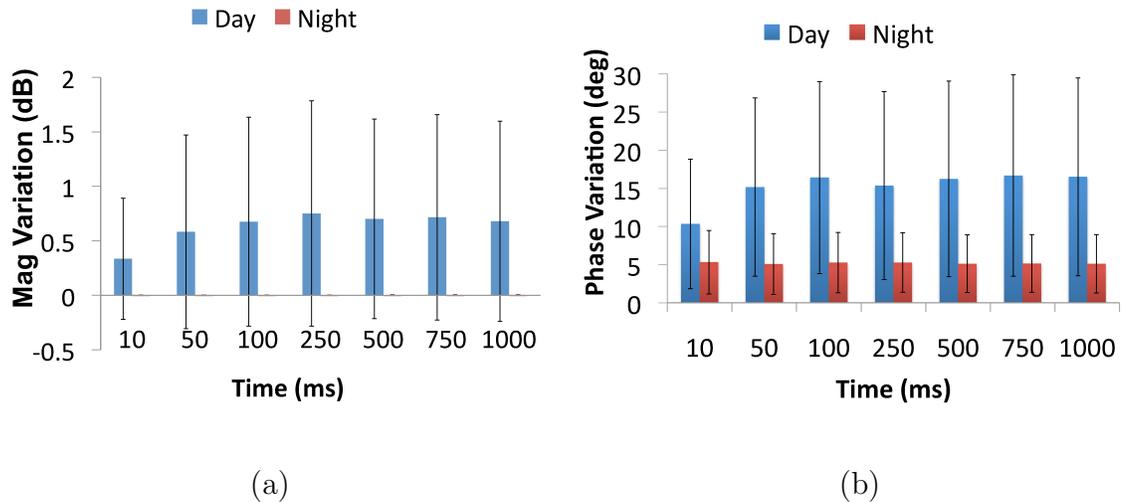


Figure 3.2 : Channel variation.

The experiments were conducted in an interference-free channel * and under two environmental conditions: late at night when no movement was happening in the

*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.

environment, and during office hours on an average day with normal human traffic around the nodes under study. I next calculate the channel's magnitude and phase variation from the measured data sets as a function of the time interval.

Fig. 3.2 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, I observed similar nighttime performance but varying daytime performance.

For the link studied in Fig. 3.2, 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. (3.2) 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. 3.2, 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. I perform all of the OTA experiments

in such conditions.

Performance Metrics

I use the *received signal strength* (RSS in dBm) value reported by the radio boards for performance comparison of different schemes studied in this paper. I 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. (3.2). Thus, in order to correctly measure the signal strength, this interference should be subtracted from the received signal in addition to noise power. I use the *signal to interference plus noise ratio* (SINR) as the metric for MUBF schemes.

For a given user k , I take the following approach to measure SINR. First, I perform multi-user beamforming and measure the RSS value. Next, I redo the multi-user beamforming measurement but this time I set the power allocated to user k to zero without changing the power allocated to the rest of the users. According to Eq. (3.2), 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, I obtain the SINR at k . In all OTA experiments, I take 10 SINR measurements and report the average and standard deviation for each data point. For the channel emulator experiments, I take 1000 SINR measurements for each data point.

In addition to SNR and SINR measurements, I also use the corresponding Shannon capacity in Eq. (3.1) 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.

3.4 Spatial Multiplexing Gains of ZFBF

In this section, I experimentally characterize the spatial multiplexing gains of ZFBF in indoor wireless networks. I 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, I study the impact of user selection based on link quality difference on ZFBF. Finally, I investigate the behavior of ZFBF as the number of concurrently served users increases.

3.4.1 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 [15], 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.

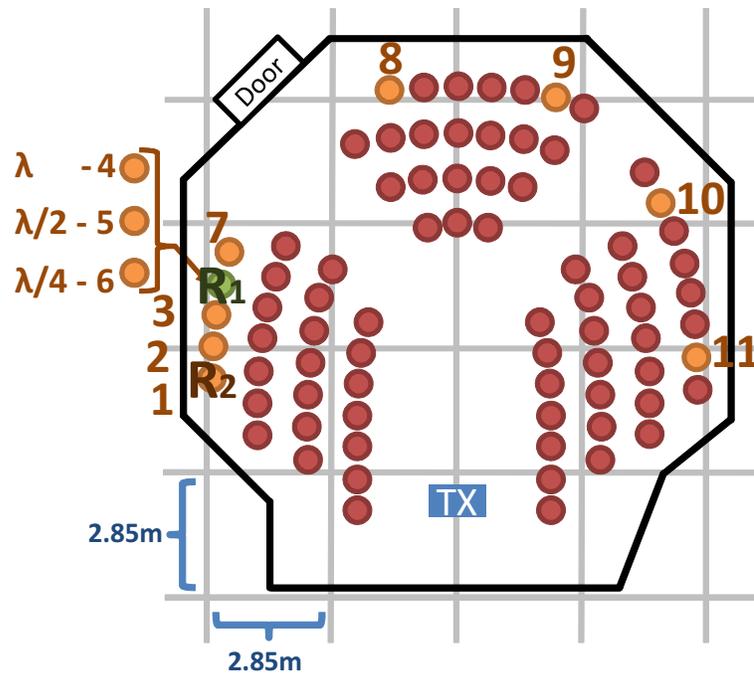


Figure 3.3 : Experimental evaluation of spatial multiplexing as a function of receiver separation distance.

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, I designed an experiment shown in Fig. 3.3 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. 3.3, I 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.

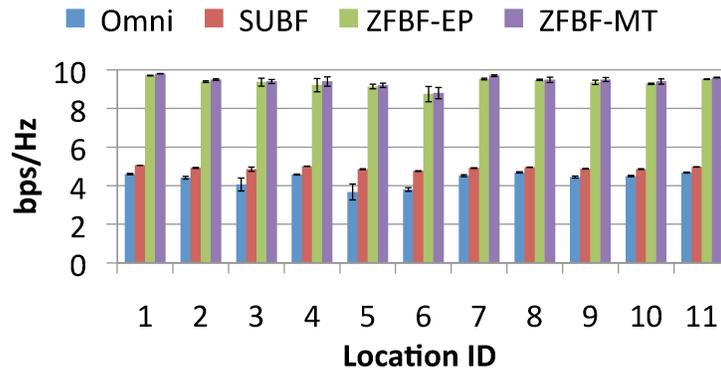


Figure 3.4 : Capacity as a function of location.

Fig. 3.4, 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. 3.4 reveals that the performance of the ZFBF scheme does not depend on the separation distance between the two receivers. This 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 only a small decrease in aggregate capacity. I 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 half 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 half 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.

3.4.2 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. (3.3), a particular receiver's SINR could vary depending on its partnered receivers. In this section, I investigate the performance of a single link's behavior as

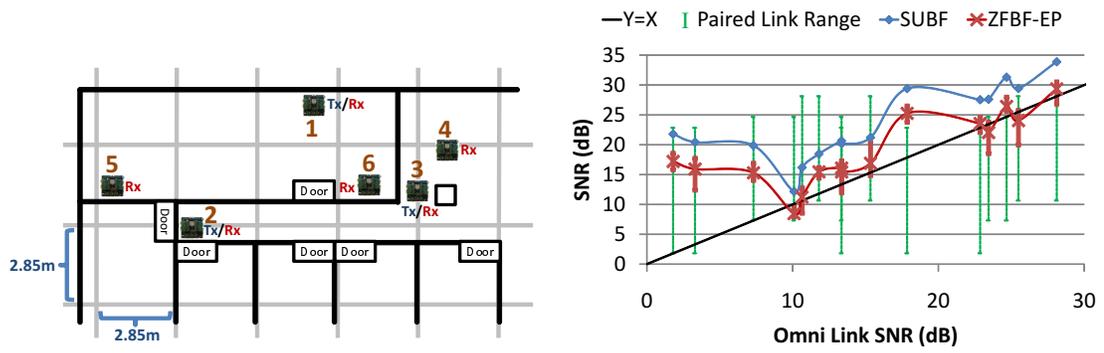
it is scheduled with different users with heterogeneous link qualities.

Scenario. Fig. 3.5(a), depicts the experimental setup in which I 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. I 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, I measure the SNR at each receiver from Omni and SUBF transmissions, and the SINR at each receiver from a jointly beamformed transmission. I repeat this experiment for all possible transmitter-receiver pairs.

Fig. 3.5(b) shows the SNR variation of each link in Fig. 3.5(a), when the link is scheduled with any other link in the network simultaneously. The x-axis of Fig. 3.5(b) represents a given link's measured Omni SNR. The y-axis shows the SNR value of the same link for the indicated schemes.

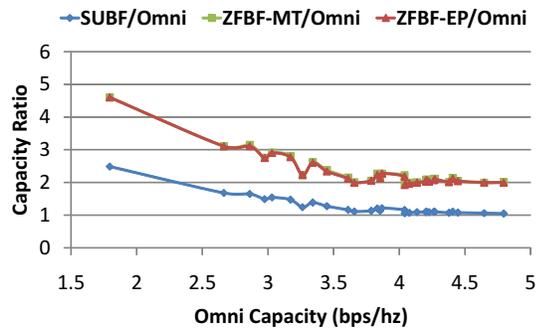
For a selected link l , there are four remaining links that can be scheduled simultaneously with l when using the ZFBF-EP scheme. Thus, for the ZFBF-EP results, I plotted the average SNR of l , when combined with each of the four other links. The thicker red bars indicate full range of l 's SNR when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

Fig. 3.5(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 trans-



(a) Map of the office environment.

(b) Per link SNR variation



(c) Aggregate capacity ratio

Figure 3.5 : Impact of concurrent user selection.

missions, 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. 3.5(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. 3.5(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. 3.5(b) also reveals information about concurrent user selection. The 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.

I now investigate the impact of link quality on the aggregate performance of ZFBF. Fig. 3.5(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. 3.5(a). I 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. 3.5(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 3.5(c) when Omni capacity is above 4 bps/Hz.

Fig. 3.5(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, I measured the SNR difference between the two schemes for all two-link sub-topologies of Fig. 3.5(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 the achieved capacity. This behavior is observed in the aggregate capacity results of Fig. 3.5(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.

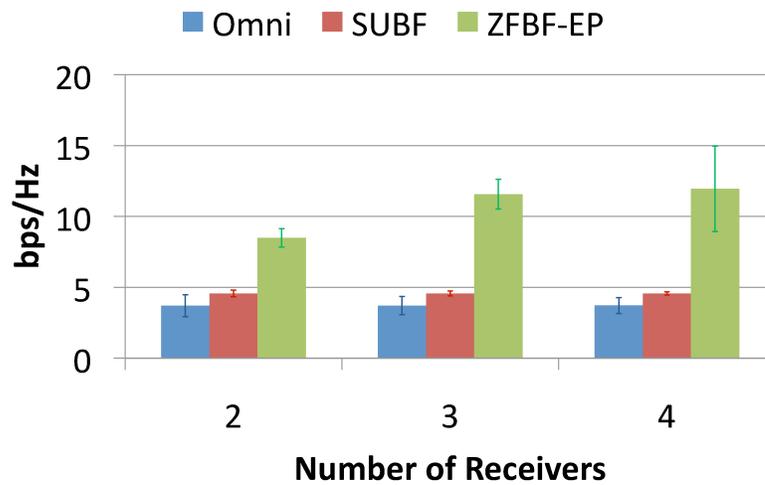


Figure 3.6 : Impact of population size on aggregate capacity.

3.4.3 Impact of User Population Size

I now investigate the performance of ZFBF as the number of served users approaches the number of transmitter antennas. I use the same node deployment setup of

Fig. 3.5(a) and perform the same set of experiments as in the previous section. However, instead of serving two users, I 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, I compute each sub-topology's aggregate capacity. Next, I group the sub-topologies based on receiver population size and calculate the average capacity for each group in Fig. 3.6. In addition, I find the average per-link SNR difference between ZFBF and Omni for each user population size as shown in Fig. 3.7.

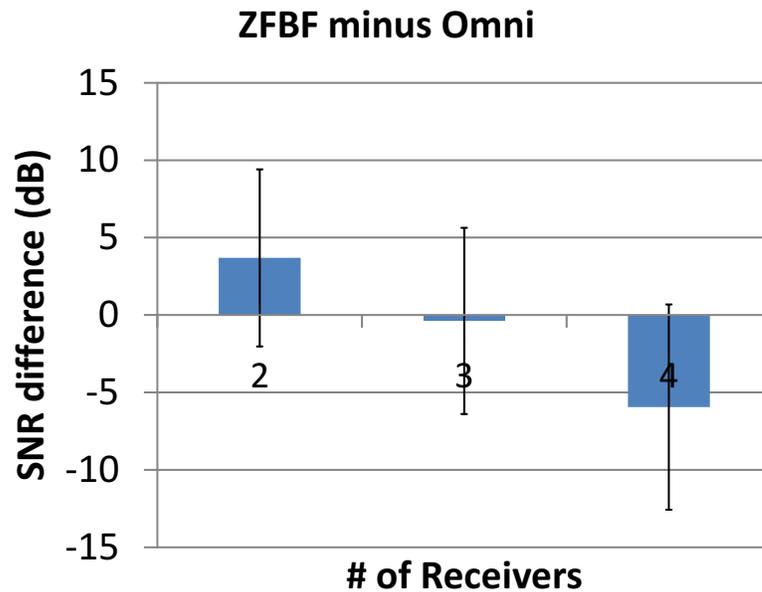


Figure 3.7 : Impact of user population size on per-link SNR difference.

Fig. 3.6 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. 3.7 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.

3.5 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. 3.1, 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, I explore the effects of channel variation on ZFBF performance.

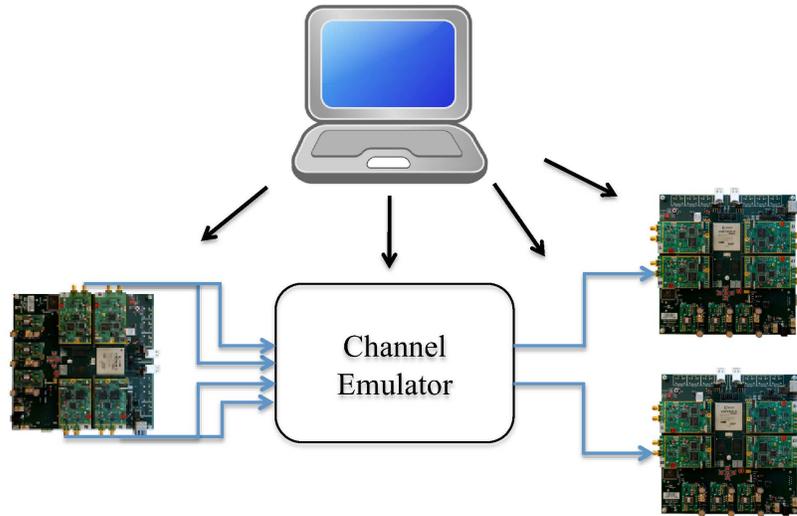


Figure 3.8 : Channel emulator setup.

Scenario. In order to have consistent and precise control over the channel and its variability, I use a channel emulator. Fig. 3.8 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 [16].

The ACE 400WB is a fully bidirectional and reciprocal 4x4 MIMO channel emulator. Internally, there are two emulator modules. In the forward module, the four input ports are connected to the four output ports. In the reverse module, the four output ports are connected to the four input ports. The inputs accept signals from a

transmitter and route them to the emulator outputs through sixteen possible paths. Each of these paths is referred to as a MIMO path.

The operations inside a single emulator module are depicted in Fig. 3.9. Each MIMO path is implemented as a tapped delay line filter. The filter coefficients can be constant (static), or time-varying, driven by the output of fading generators. The fading generators are random processes designed to emulate a particular Doppler spectrum. The fading generators may also be correlated to produce spatially-correlated fading.

The characteristics of each path can be modeled according to the movement velocity. If a user is mobile, all of the paths that are associated to the mobile client will have fading properties. On the other hand, the channel emulator allows for static assignment of some paths and fading assignment to some other paths. This can be potentially used to model variations in the environment while the clients is static.

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 3.2. The channel model is adapted from 802.11n task group (TGn) models used to evaluate the performance of MIMO in indoor environments [8]. This channel model is composed of nine Non-Line-of-Sight (NLOS) Rayleigh fading 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.

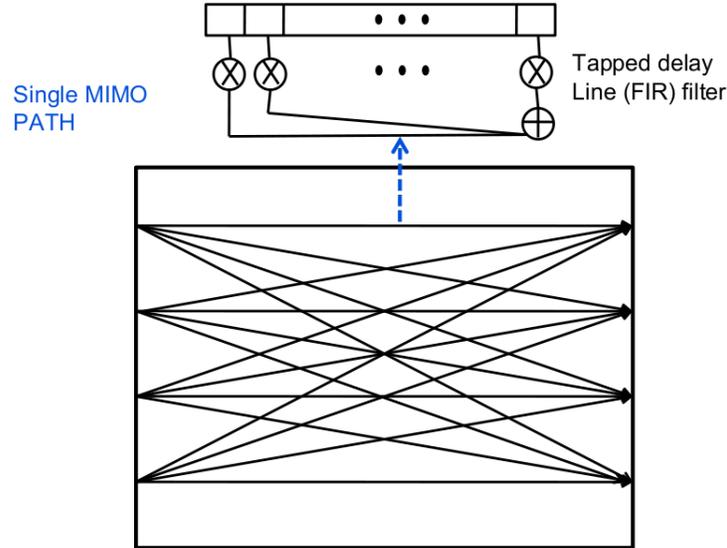


Figure 3.9 : Operations inside a forward emulator module.

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

3.5.1 Impact of Environmental Variation

In this section, I 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 [8]. I performed two sets of experiments using Doppler fading rates of

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 3.2 : Channel model parameters.

1.157 and 2.778 Hz to emulate typical (T) and rapidly (R) varying environments respectively. For each of these experiments, I varied the time interval between the channel estimate measurement and actual data transmission.

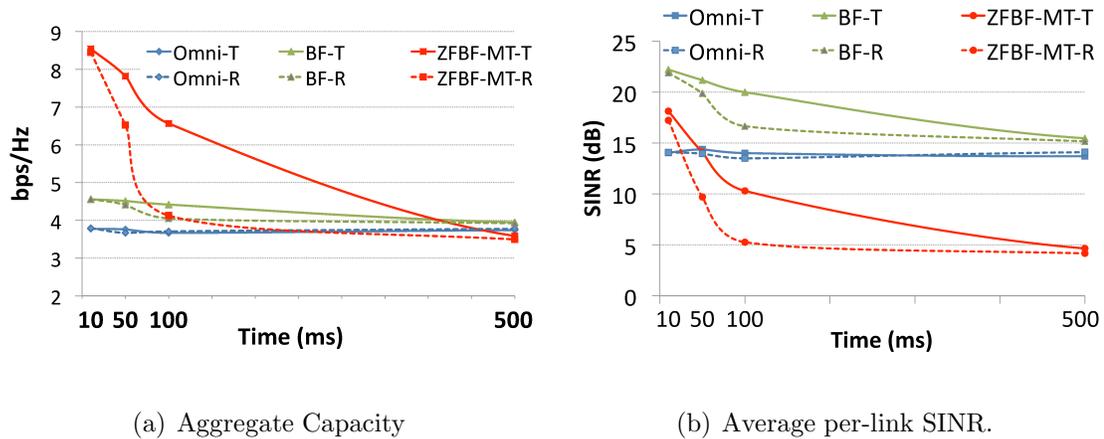


Figure 3.10 : Impact of environmental variation.

Fig. 3.10(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.

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. 3.10(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. 3.10(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[†]. In Fig. 3.10(b), I measured the average per-link SINR value for all of these schemes. Fig. 3.10(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

[†]From Section 3.2, recall that instead of multiplexing users, Single-User schemes link Omni and SUBF schedule users sequentially according to a TDMA schedule

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.

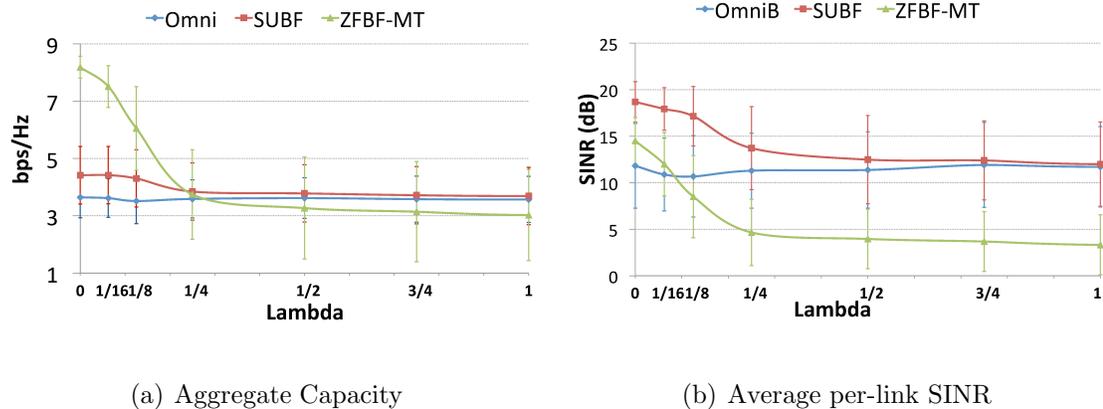
3.5.2 Impact of User Mobility

I 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.

I perform controlled experiments to quantify the drop in throughput as a function of user mobility. I use the same experiment setup as shown in Fig. 3.8; however, I 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. 3.11(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. 3.11(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



(a) Aggregate Capacity

(b) Average per-link SINR

Figure 3.11 : Impact of mobility.

would further decrease the performance of SUBF and ZFBF. However, Fig. 3.11(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. I 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 I 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 approx-

imately 10 ms.

3.6 Impact of Beamforming on Spatial Reuse

I now investigate the increase in spatial reuse opportunities offered by MUBF. In Section 3.6.1, I consider a single sender/receiver pair and a third node, W , at which I attempt to minimize the interference caused by the initial pair's transmission. I quantify the reduction in interference as a function of W 's location. Next, in Section 3.6.2, I investigate the ability for a sender to reduce its transmission footprint by minimizing interference at multiple unintended receivers simultaneously. Finally, in Section 3.6.3, I 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.

3.6.1 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 . I investigate

the ability of ZFBF to reduce interference as a function of W 's location.

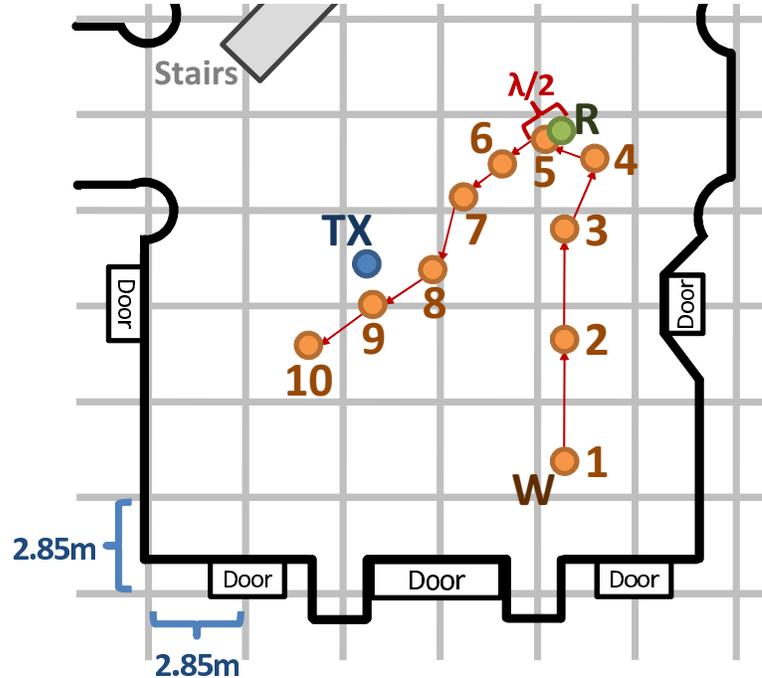


Figure 3.12 : Experimental Scenario.

Scenario. the experimental scenario is depicted in Fig. 3.12. The transmitter, TX , sends data to its receiver, R , such that the resulting interference at W is minimized. I investigate three different movement patterns of W . First, I 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, I place W and R adjacent to one another and move W along the line connecting the three nodes (location IDs 5 to 7). Finally, I 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, I take the following measurements: First, I perform an Omni transmission from TX to R and record the

received signal strength at W . Next, I perform joint beamforming with the objective of zero interference at W and measure the resulting signal strength at W .

Fig. 3.13 shows the resulting interference at W for each of the location IDs. In Omni mode, I 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. 3.13 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 .

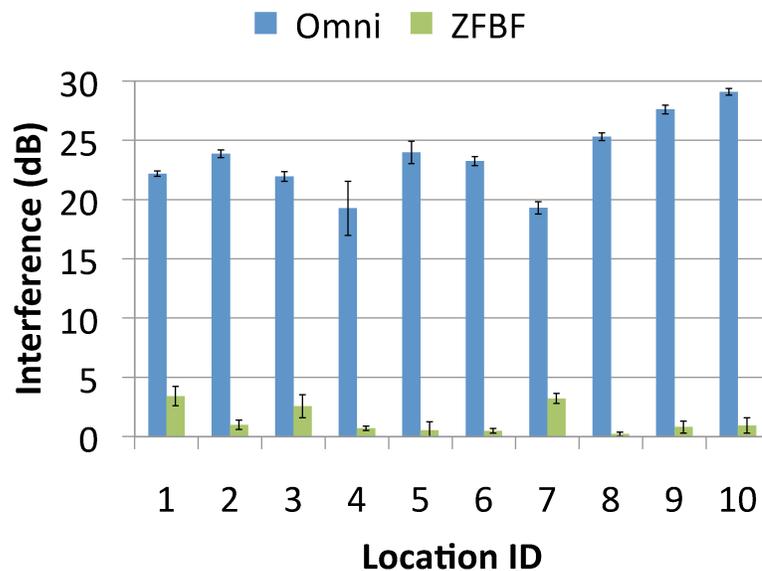


Figure 3.13 : Interference reduction as a function of location.

Finding: A user can obtain an interference-free channel by sharing its channel

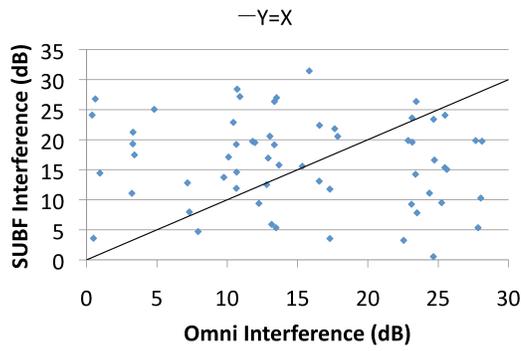
information to a ZFBF-enabled transmitter. The interference-free channel is obtained irrespective of the distance between the user and either the transmitter or the receiver.

3.6.2 Multi-Point Interference Reduction

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

I consider the node location setup described in Fig. 3.5(a). Nodes 1, 2, and 3 each have four antennas and thus can be used as four-antenna transmitters or single-antenna receivers. I select one of these nodes as the transmitter and one of the remaining nodes as the intended receiver. Then, I consider all possible combinations of 1, 2, or 3 nodes among the remaining nodes as locations at which I plan to minimize interference. I repeat this experiment for all possible transmitter-receiver pairs leading to 210 different sub-topologies. I perform Omni, SUBF, and ZFBF transmissions, and measure the resulting signal strength at the intended receiver as well as unintended receivers.

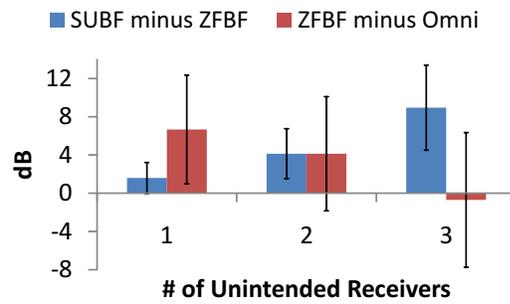
Fig. 3.14 shows the interference footprint for the three schemes. I first investigate the performance of Omni and SUBF. Fig. 3.14(a) shows the scatter plot of interference at unintended receivers with Omni and SUBF schemes. Each point in this graph corresponds to a sender-unintended receiver pair. From this plot, similar performance is observed between the Omni and SUBF schemes. For half of these locations, the



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

(a) Omni vs. SUBF interference

(b) Table 3: ZFBF interference (dB)



(c) SINR difference at the receiver

Figure 3.14 : Multi-point interference reduction

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 I present the measured mean and standard deviation of interference caused at the unintended receivers. Similar to the results of Fig. 3.13, I observe that the resulting interference is close to the noise floor power. However, unlike Fig. 3.13, these results are obtained as the transmitter used up all of its DoF. Thus, I 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. I investigate this behavior in Fig. 4.4(c). Here, I compare the SINR of ZFBF to Omni and SUBF schemes at the intended receiver as I 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 the measurements, I 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, I 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.

3.6.3 Impact of Multi-User Beamforming on Network Throughput

I now investigate the potential of ZFBF to increase network capacity by minimizing interference between concurrent links. I create 36 different sub-topologies consisting

of two sender-receiver pairs for the node setup shown in Fig. 3.5(a). For each of these sub-topologies, I 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.

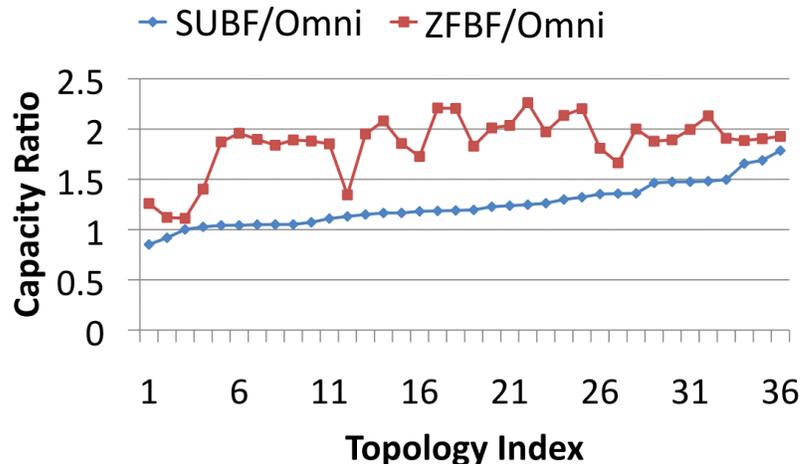


Figure 3.15 : Maximum Capacity of two flows.

Fig. 3.15, shows the relative capacity improvement of SUBF and ZFBF compared to Omni. I sort the sub-topologies based on increasing SUBF capacity ratio. For the first and last three sub-topologies, ZFBF performs close to SUBF. Careful investigation of these sub-topologies revealed that for the first three topology indices, 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 outper-

forms 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.

3.7 Incorporation of Overhead

The throughput and discussions so far, did not include the impact of overhead due to channel estimation. In this section, I address the impact of overhead on the overall system performance. To this end, I consider an example MAC protocol that has been proposed for the medium access protocol of the 802.11ac [4] standard.

The protocol components are depicted in Fig. 3.16. The medium access is composed of three main parts. In the first part, a training sequence is transmitted by the transmitter for the purpose of channel estimation by the clients. This training sequence also includes the addresses of the clients which the access points intends to transmit unicast packets.

After the transmission of the training sequence, clients that are included in the training sequentially send back their channel estimates. Finally the access point sends parallel unicast packets. Note that the transmission of unicast packets can happen for the length of channel coherence time before a new set of channel estimates are obtained by the access point.

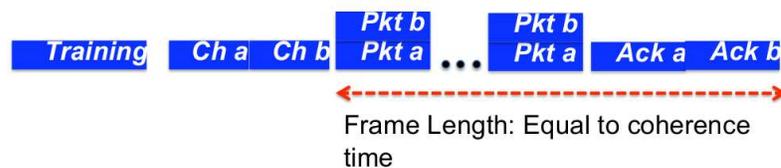


Figure 3.16 : An example MAC protocol for MUBF.

In the above protocol, the overhead is made of the training sequence and the

Parameter	Value
Spatial multiplexing gain	2
Packet size	1000 Byte
PLCP header length	192 bits
Data transmission rate	54 Mbps
Header transmission rate	1 Mbps
Base rate	2 Mbps
Training size	20 Byte
Channel estimate size	20 Byte

Table 3.3 : Model parameters.

channel estimates that are sent back by the clients. I next proceed to provide a worst case analysis on the impact of overhead on the benefits of multi-user beamforming. I assume that data transmission rate is equal to 54 Mbps so that the amount of overhead is large compared to the packet transmission time. Table 3.3 provides the parameters that are used for comparison of sequential transmission without overhead, to multi-user beamforming with training overhead.

Fig. 3.17 plots the throughput ratio between multi-user beamforming with over

head to sequential transmission without overhead as a function of number of back to back transmissions in a frame. Based on this graph, even when channel estimation is done on a per-packet basis, multi-user beamforming still provides gains compared to sequential transmission. The gains of beamforming increase as the number of back-to-back transmissions in a single frame increases. According to Fig. 3.17, with 12 transmissions the amount of overhead is less than 5% of the total transmission time and the gains of beamforming remain close to two.

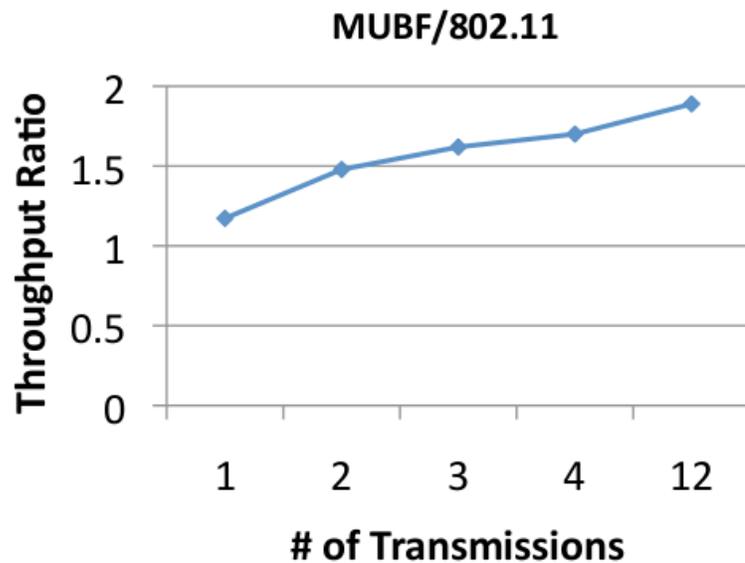


Figure 3.17 : MUBF Throughput comparison with 802.11.

As a result, when channel estimation is done for a set of back-to-back transmissions as opposed to on a packet by packet basis, the amount of overhead would be negligible. Note further that the above analysis is considering only the worst case scenario in

which only two packets are being multiplexed and the data transmission rate is high. When the two multiplexed packets support high data rates, multi-user beamforming only provides spatial multiplexing gains. On the other hand, if the omni transmission rate is low, beamforming has the potential to significantly increase the transmission rate as observed in Section 3.4.

3.8 Summary

In this chapter, I designed and implemented a custom MUBF platform that allows for the experimental evaluation of different beamforming strategies. Using this platform, I experimentally evaluated the multiplexing gains of ZFBF as a function of receiver separation distance, user selection, and user population size. I experimentally showed that a four-antenna, ZFBF-enabled transmitter is able to simultaneously transmit to two users that are within a half of a wavelength of one another. I 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. I also evaluated the impact of user mobility and environmental variation on the performance of ZFBF. I 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, I

investigated the potential of ZFBF to reduce interference at unwanted locations and increase spatial reuse. The 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.

Chapter 4

Multicast Beamforming

4.1 Introduction

The proliferation of mobile computing devices as well the rapid growth in applications and services involving group communication (network management and software updates, electronic class/conference rooms, MobiTV, etc.) has made wireless multicasting an important component in next generation wireless standards such as 802.11ac, LTE, and WiMAX. This coupled with the increasing demand from users to be able to stream media content on their wireless devices ubiquitously has prompted the need for designing efficient solutions for wireless multicasting.

While the inherent broadcast nature of the wireless medium allows for a single multicast transmission to cover a group of users simultaneously, its performance is determined by the client with the weakest channel (SNR). On a parallel front, beamforming antennas have recently gained a lot of attention in indoor wireless networks [17, 18, 19]. These are multiple-element arrays that are able to focus their signal energy in specific directions and hence form a natural solution to improve the channel to the weakest client and hence the multicast system performance. Beamforming could be either adaptive where the beam patterns are computed on the fly based on channel feedback from clients, or switched, where precomputed beams that cover the azimuth of 360° are used. Recent works [20, 21, 22] have advocated the use of switched beamforming to improve multicasting. However, the beamforming gain (from restricted signal footprint) comes at the cost of reduced broadcast advantage, thereby requiring multiple beamformed transmissions to cover all the clients unlike an omni-directional

transmission. Addressing this tradeoff in turn requires the use of composite beams that are generated by combining individual beams so as to effectively balance between beamforming gain and coverage [20].

In this work, I experimentally show that switched beamforming is ineffective for multicasting in indoor multipath environments. The reasons are two fold: (i) using a pre-determined set of beam patterns limits performance when simultaneously catering to a multitude of clients; (ii) since the resulting SNR on a composite beam is not available a priori, it is modeled based on the measured SNR from its constituent beams; however, such modeling is highly inaccurate in multipath environments, resulting in inefficient performance when a composite beam is actually applied. To address these deficiencies, I advocate the use of adaptive beamforming for multicasting in indoor wireless networks.

While adaptive beamforming incurs more overhead in the form of channel estimates from clients compared to switched beamforming, I show that its significant potential to improve indoor multicast system performance outweighs the overhead.

Translating the potential of adaptive beamforming into practically realizable benefits for multicasting is a highly challenging task. Specifically, (i) given the channel information of clients, determining an optimal adaptive beamformer for multicasting is a hard quadratically constrained quadratic optimization problem; (ii) it may be detrimental to beamform to all clients simultaneously and hence it becomes important to identify *when* and *how* a set of clients must be partitioned into separate groups

(scheduling) and beamforming executed within each of these groups sequentially; and (iii) in practical scenarios the rate of channel feedback available from a client may not be sufficient compared to the coherence time of its channel either due to limited feedback (for reducing overhead) or small coherence times (due to client mobility). In such cases, the adaptive beamformer must incorporate robust mechanisms to compensate for the lack of timely channel feedback not only to retain its benefits, but also to avoid degrading to worse than omni.

Towards addressing these challenges, I present ADAM- the first adaptive beamforming based system for improving multicast performance in indoor wireless networks. The two main contributions of ADAM are,

- incorporation of efficient yet practically amenable solutions to address each of the aforementioned challenges.
- implementation and comprehensive evaluation of the practical benefits of adaptive beamforming for multicasting, its superiority over switched beamforming and its robustness to channel dynamics.

Briefly, ADAM decouples the joint client scheduling and beamformer design problem into two individual sub-problems in a manner that allows their solutions to reinforce each other. For a given number of groups, it first partitions the clients into groups based on the “closeness” of their channels. This allows ADAM to later determine an efficient adaptive beamformer for the clients within the same group. It then employs a greedy, one-shot algorithm to provide a near-optimal multicast beamformer within

each group. This is in contrast to other sub-optimal beamformers in related work that are iterative in nature and exhibit slow convergence [23]. Finally, ADAM combines the solutions from these two components to arrive at the appropriate number of groups, along with the client membership and beamformer within each group.

ADAM is implemented on the WARP platform and its performance is extensively evaluated in indoor environments. The experimental results reveal that (i) while switched beamforming has fundamental limitations for multicasting in indoor multipath environments, ADAM is able to address these deficiencies to yield gains as high as three folds; (ii) ADAM's gains are more with a higher dynamic range of the (discrete) transmission rates employed by the MAC, yielding gains as high as nine folds over omni with the 802.11 rate table; and (iii) the rate tables employed for beamformed transmissions are strongly dependent on both the coherence time (t_c) of the channel as well as the channel feedback time scale (t_f) and more specifically on the s -ratio, where $s = \frac{t_f}{t_c}$. Hence, ADAM categorizes the clients based on their s parameter and employs client-specific rate tables in determining the beamformed transmission rate, thereby increasing its robustness to both client mobility and limited channel feedback.

The rest of this chapter is organized as follows: Section 4.2 provides a background on beamforming. Sections 4.3 and 4.4 describe the motivation and challenges of adaptive beamforming for multicasting. Section 4.5 describes the components of ADAM. Section 4.6 describes its implementation followed by detailed evaluation in

Sections 4.7 and 4.8. Finally, I summarize this chapter in Section 4.9.

4.2 Background

Beamforming: Beamforming antennas consist of an array of omni-directional elements, with sophisticated signal processing capabilities. The signals that are transmitted to each of these antenna elements can be weighted in both amplitude and phase to produce a desired beam pattern that increases the SNR at the receiver. These weights applied at the Tx antenna array can be written as $\mathbf{w} = [w_0 w_1 \dots w_{K-1}]$. Depending on the level of sophistication in adapting these weights, there are two main types of beamforming namely, switched and adaptive.

In switched beamforming, a set of pre-determined beam patterns covering the entire azimuth of 360° is made available. Each of these beam patterns has a main lobe of maximum gain and some side lobes representing leakage of energy. The larger the number of elements, the thinner the main lobe that can be formed and higher is the array gain. However, for a K element antenna, the thinnest beam that can be formed is $\frac{360^\circ}{K}$, thereby requiring K beams to cover 360° .

As switched beamforming is normally implemented as an open-loop procedure without fine-grained channel feedback from the Rx, a Tx will tend to choose a pattern that provides the strongest signal strength to the client. Such a beam may not coincide with the physical direction of the Rx depending on the multipath scattering in the environment.

In adaptive beamforming, channel estimation and feedback from the Rx is used to adapt the beam pattern in the signal domain at the Tx. The resulting beam pattern may not have the single main lobe structure (pointing in the direction of the Rx) of a switched beam, but is optimized to reinforce the multipath components of the signals arriving at the Rx from the different Tx antenna elements. Its versatility in indoor multipath environments comes at the cost of channel feedback overhead from the clients.

Multicast and Beamforming: Given that multicast performance of a group depends on the client with the weakest channel in the group, beamforming provides a natural solution to improve the SNR of the weakest client and hence the multicast group as a whole. However, as previous works [20, 21] have pointed out, the beamforming gain comes at the cost of spatially restricted transmissions, which in turn limits its broadcast advantage that is required to cater to multiple clients simultaneously.

The solution to address this tradeoff in the case of switched beamforming is to form a composite beam pattern from individual beam patterns that can be used to cover multiple clients simultaneously [20]. However, since the energy is conserved, the net power is distributed among the constituent beams and hence the resulting beamformed SNR at the clients due to the composite beam is reduced. Hence, it becomes important to intelligently choose composite beam patterns that tradeoff user coverage and beamforming gain.

Unlike switched beamforming, in the case of adaptive beamforming, the channel to each of the clients is estimated and fed back to the AP. With the complete channel information, the AP can directly solve (albeit challenging) to obtain a joint beamforming weight vector that maximizes the minimum SNR to all the clients in the group.

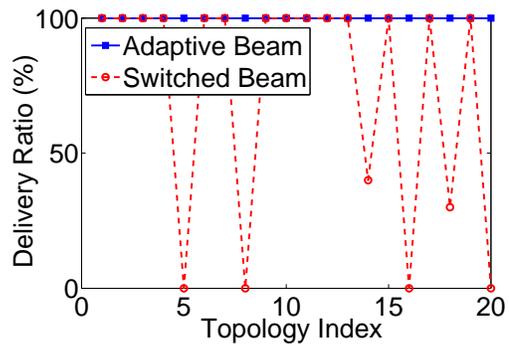
4.3 Motivation

Current beamforming solutions for improving the multicast performance [20, 21, 22] advocate the use of switched beamforming. Hence, in order to motivate the need for adaptive beamforming, we address the following two questions.

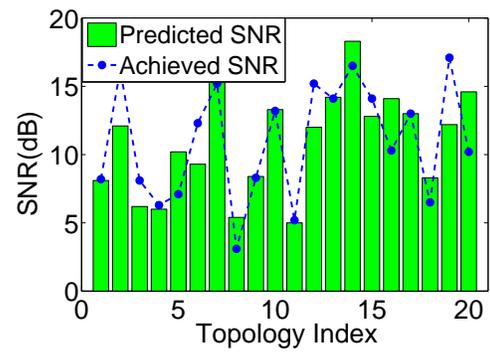
(i) Is switched beamforming a practical solution for improving multicast performance?

Given that the existing switched beamforming solutions are mostly theoretical solutions without a practical implementation, it remains to be seen if switched beamforming can indeed deliver the promised multicast gains in practice. Further, this will also unveil the validity of the modeling assumptions behind the switched beamforming solutions.

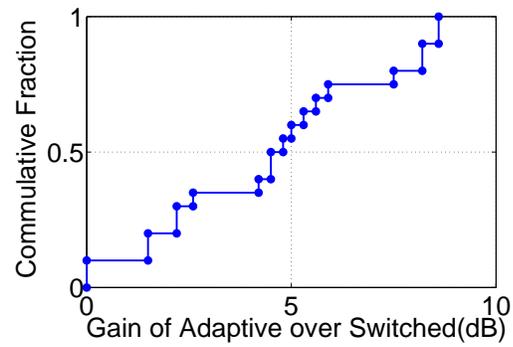
We conduct an experiment in the topology shown in Fig. 4.6(a), by considering three clients in a multicast group. A circular array of four antennas with four pre-determined beams is used for switched beamforming. Based on the beam with the best SNR reported by each client, the AP determines a composite beam pattern to cater to all the three clients simultaneously [20].



(a) PDR



(b) Modeling Accuracy in Switched Beam



(c) SNR Gain

Figure 4.1 : Adaptive vs. switched beamforming

However, the SNR at the clients for composite beams cannot be known a priori. Hence, the inherent modeling assumption made is that when a composite beam is formed from k individual beam patterns, the resulting SNR at the clients are reduced by $10 \log_{10}(k)$ dB (compared to individual beam SNR) due to the equal distribution of power across constituent beams. Thus, the AP selects a transmission rate according to the predicted resulting SNR of the weakest client.

By varying our clients, we generate multiple topologies and obtain the optimal switched beamforming solution, apply it and measure the resulting packet delivery ratio (PDR). The PDR should be close to 100% if the modeling assumption is accurate. However, the results in Fig. 4.1(a) are quite the contrary, where the PDR could be significantly less, thereby indicating that the switched beamforming solution applied is not an efficient one. In verifying the reason behind the poor performance, we plot the predicted multicast group SNR of the composite beams against the actual measured values in Fig. 4.1(b). It is clear that the modeling assumption behind the predicted SNR, which may hold in line-of-sight environments, does not hold good in many of our indoor topologies, where it either under-estimates or over-estimates the actual SNR. This in turn can be attributed to the multipath nature of the indoor environment, which makes it hard to predict the effect of composite beams needed for multicasting.

(ii) Given that switched beamforming cannot address multicasting efficiently in in-

door wireless environments, the next question to understand is whether *adaptive beamforming (designed to handle multipath) can address the deficiencies of switched beamforming for multicasting?* Towards this end, we estimate the channel to all the three clients and compute an adaptive beamformer that maximizes the minimum SNR for the multicast group (details deferred to Section 4.5). The resulting PDR for each of the topologies is compared against switched beamforming in Fig. 4.1(a). It can be clearly seen that adaptive beamforming is capable of delivering the predicted performance in contrast to switched beamforming. Further, the CDF of the SNR gain (in dB) of adaptive over switched beamforming over all the topologies, depicted in Fig. 4.1(c), clearly indicates the large potential of adaptive beamforming for improving multicast performance.

4.4 Design Challenges

In this section, I describe the challenges in realizing a practical adaptive beamforming system.

4.4.1 Determination of Adaptive Beamformers

Determining an adaptive beamformer that caters to all users in the multicast group is a challenge in itself. To see this, consider the objective of maximizing the minimum rate of the users in the multicast group under constant transmit power constraint. The problem can alternatively be posed as minimizing the transmit power while satisfying

a minimum rate requirement for all users [24].

The rate of the k^{th} user can be written as

$$R_k = \log_2(1 + \mathbf{h}_k \mathbf{w} \mathbf{w}^* \mathbf{h}_k^*) \quad (4.1)$$

The multicast beamforming problem is then

$$\begin{aligned} \max_{\mathbf{w}} \quad & \min_k \{ \log_2(1 + \mathbf{h}_k \mathbf{w} \mathbf{w}^* \mathbf{h}_k^*) \} \\ \text{s.t.} \quad & \mathbf{w}^* \mathbf{w} \leq P \end{aligned}$$

where P is the total power constraint. Without loss of generality I assume $\|s\|^2 = 1$. Also, we note that the power constraint $\mathbf{w}^* \mathbf{w} \leq P$ has to be satisfied with equality, since otherwise \mathbf{w}_{opt} (the value of w in the optimal solution) can be scaled up such that the power constraint is satisfied with equality and this increases the objective function and is therefore a contradiction.

Here, optimizing the rate is equivalent to optimizing the minimum SNR of the multicast group. Hence, the problem can be alternatively presented as the maximization of the minimum received SNR of all users, i.e.

$$\begin{aligned} \mathcal{P}_1 : \quad & \max_{\mathbf{w}} \quad \min_k \{ \mathbf{w}^* \mathbf{h}_k^* \mathbf{h}_k \mathbf{w} \} \\ \text{s.t.} \quad & \mathbf{w}^* \mathbf{w} \leq P \end{aligned}$$

The problem formulation in \mathcal{P}_1 , is a quadratically constrained quadratic optimization program (QCQP), which is a non-convex problem and its discrete version

is NP-hard as well. This makes it challenging to design an efficient algorithms to compute an adaptive multicast beamformer.

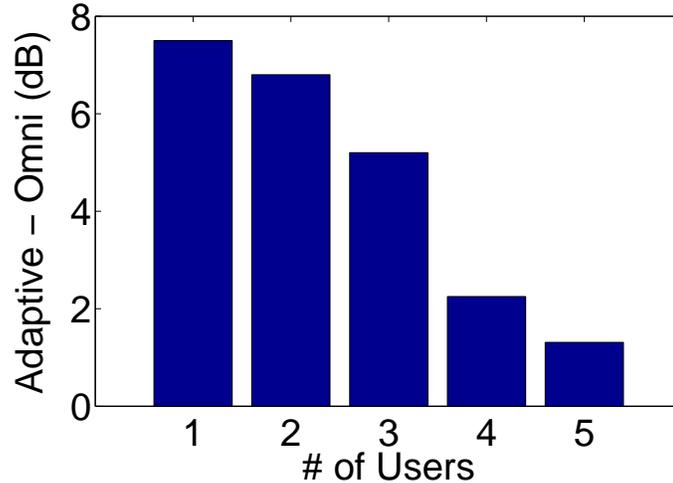


Figure 4.2 : Impact of user size on adaptive gain.

4.4.2 Scheduling

While the above challenge pertains to finding an adaptive beamformer for a group of users, the next aspect to understand is whether all users should be jointly beamformed to. I perform an experiment, where I increase the number of users in the multicast group from one to five in the topology of Fig. 4.6(a). The adaptive beamformer is determined for each group and the gain of the resulting minimum SNR of the beamformed transmission over omni-directional transmission is plotted in Fig. 4.2. It can be seen that as the size of the group increases, the adaptive beamforming benefits tend to decrease with its performance tending to that of an omni transmission. This

is because as the size of the group increases, the randomness of the channel vectors of different users makes the beamformed vector tend to that of an omni-directional transmission so as to cater to all the users. This in turn advocates the partitioning of users in a large multicast group into sub-groups of smaller size and enabling beamforming to improve transmissions in each of the sub-groups. The need for such partitioning (scheduling) is exacerbated in the presence of discrete rate tables. For example, consider two users that each achieves a 5 dB SNR when jointly beamformed to. With 802.11 rate table of Fig. 4.6(d), the transmission rate would be 1Mbps. Now, if sequential serving of the users increases each user's SNR by 3 dB, the resulting data rate of each client would be 9 Mbps. Thus, if the transmission time of transmitting L bytes with joint serving is $\frac{L}{1}$, the required time with sequential serving would be $\frac{L}{9} + \frac{L}{9} = \frac{L}{4.5}$, which is a gain of 450%.

Introducing scheduling complicates the beamforming problem further. Note that when users are partitioned into sub-groups, there is a (time) multiplexing loss with different sub-groups receiving transmissions sequentially. Hence, there is a tradeoff between operating on low rates (low min SNR) by beamforming to all the users in one shot or operate on higher rates in each sub-group but incur the multiplexing loss.

4.4.3 Channel Dynamics and Feedback Rate

The above two challenges are with respect to determination of a solution under the assumption of instantaneous channel information from clients. However, in any prac-

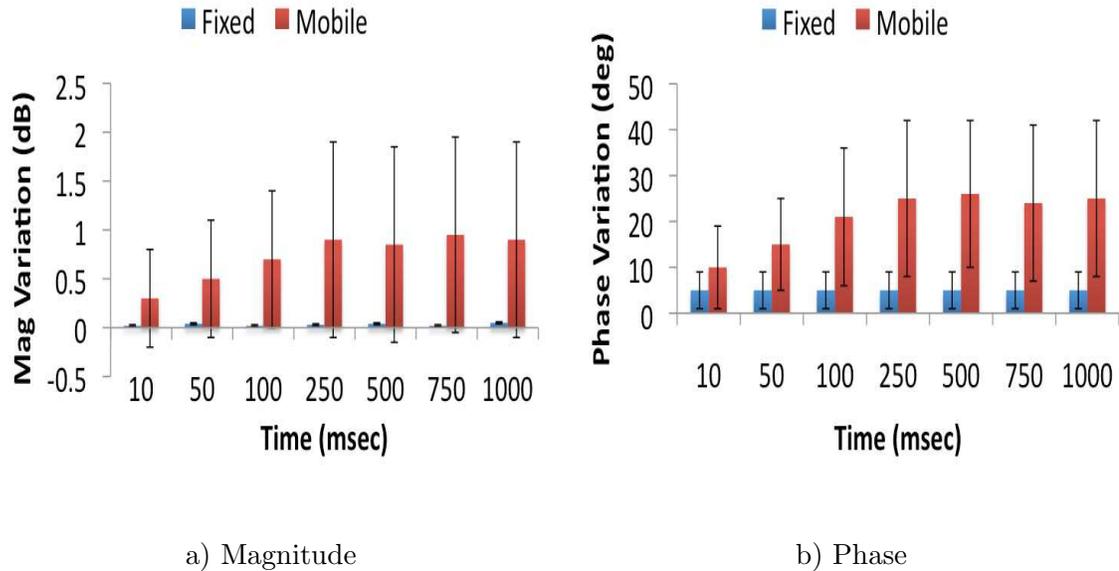


Figure 4.3 : Channel variations.

tical system, channel state feedback constitutes overhead and may not be available for every single packet. The mobility of clients further reduces the coherence time of the channel, thereby requiring increased feedback frequencies, the absence of which could render the feedback both outdated and inaccurate.

I conduct an experiment, where the AP transmits 100 pkts/sec to a static client at night. The client estimates the channel from the decoded preambles. The variation in the channel magnitude and phase for the measured samples in an interval is plotted as a function of the interval size in Fig. 4.3. The experiment is then repeated for a mobile client and the corresponding results are also indicated. It can be seen that the channel dynamics are almost negligible for a static client, indicating a large coherence time for the channel as well as its ability to withstand reduced feedback frequencies.

However, with a mobile client, the situation is quite the contrary, where the mean channel magnitude and phase variations are around 1 dB and 20-30° respectively. Note that the corresponding large standard deviation especially in the channel phase (critical for adaptive beamforming) indicates the small coherence time of the channel, thereby requiring high feedback frequencies on the order of few milliseconds.

Hence, it becomes important to understand the sensitivity of the adaptive beamforming solution for multicast to such channel dynamics as well as feedback frequency, and hence incorporate robustness into its design.

4.5 Design of ADAM

In this section, I describe the design of ADAM, the adaptive beamforming based multicast system that addresses the identified challenges. I first propose a joint user scheduling and beamformer design problem with the objective of minimizing the time that it takes to disseminate data to the multicast clients. Next, I propose efficient algorithms that are implemented in ADAM and are suitable for a practical system design. I address the impact of channel dynamics and ADAM's solutions to increase robustness in Section 4.8.

4.5.1 Components of ADAM

Once the AP receives data to be disseminated for a multicast session, ADAM operates as follows:

- **Step 1:** AP sequentially transmits training symbols on each of its antennas.
- **Step 2:** Each client measures the channel amplitude and phase for each of the transmitting antennas.
- **Step 3:** Clients sequentially feedback channel information to the AP.
- **Step 4:** AP runs its algorithms which partition the clients to different groups and find the beamformer for each group.
- **Step 5:** AP selects the appropriate rate for each group based on a rate table, and transmits multicast data.

The main algorithmic component of ADAM is to design efficient user partitioning and multicast beamformer for Step 4. To evaluate this, I use the notion of schedule length (delay) required for multicast data transfer to the entire group as the metric of optimization. I assume a packet delivery ratio (PDR) requirement of 100% for all of the clients. If some of the clients can tolerate a lower PDR, it can be incorporated in the solution. Furthermore, it is possible for an AP to send multiple multicast packets in each schedule in order to reduce overhead. The periodicity of channel estimation procedure can be determined based on its incurring overhead, the required PDR for each client, and the dynamics of channel due to user mobility or variations in the environment.

4.5.2 Problem Formulation

Assume K users in the system, and a multicast data size of L bytes. The objective is to partition the users into J groups and transmit L bytes sequentially on each group, such that the total schedule length to deliver L bytes to all users is minimized. The problem can be formally stated as:

$$\begin{aligned} \mathcal{P}_2 : \quad & \min \sum_{j=1}^J \frac{L}{R(\text{SNR}_j)} \\ & s.t. \quad \mathbf{w}_j^* \mathbf{w}_j \leq P \\ & \text{SNR}_j = \min_{k \in S_j} (\mathbf{h}_k \mathbf{w}_j \mathbf{w}_j^* \mathbf{h}_k^*) \end{aligned}$$

where j is the number of partitions, S_j is the set of user indices and \mathbf{w}_j is the beamforming vector for each partition j . The rate function $R(\text{SNR})$ maps SNR into the appropriate rate and it may be a continuous (e.g., log based capacity) or a discrete function. In practical systems there are only a finite set of modulation-coding schemes, which result in discrete rate functions. Hence, the emphasis of the work is on discrete rates.

As described in Section 4.4.1, finding a multicast beamformer is NP-Hard even for a single multicast group. The above problem formulation is further complicated as the optimal grouping depends on the rate of each group, which itself is dependent on the beamformer vector for that group and has a discrete nature for practical purposes.

To address these issues, I adopt a decomposition approach which divides the problem into two sub-problems in a manner that allows the two sub-problems to re-inforce each other. For a given number of groups, I first partition the users into groups based on the “closeness” of their channels. This allows us later to determine an efficient adaptive beamformer for the clients within the same group. I then employ a greedy, one-shot algorithm to provide a near-optimal multicast beamformer within each group.

By combining the above two sub-problems, I have developed two algorithms to solve the joint partitioning and beamformer (JPB) design problem of \mathcal{P}_2 . The algorithms are as follows:

JPB-A (All): This algorithm considers up to K number of partitions. Given the number of partitions (groups) j , it determines the client membership to the groups as well as the beamformer for each group, and calculates the resulting schedule length. Finally, JPB-A selects the number of partition j^* along with the corresponding beamformers and client membership that yield the minimum schedule length among all.

JPB-S (Successive): This algorithm increases the number of partitions one by one only if additional partitioning of the clients decreases the schedule length.

The above two algorithms need to address two sub-problems: given a number of partitions, *how* to assign the clients to the given number of partitions; next, *design* an appropriate beamformer for the clients within each group. These two components are discussed next.

4.5.3 User Partitioning

In order to optimize the overall performance, the users that are grouped together would be selected such that a beamformer that is appropriate for one is also desirable for the rest of the users in the group. This can significantly increase the minimum SNR of the group and the resulting transmission rate. I use the notion of chordal distance [25] between two vectors as the metric for closeness of user channels. Given two users with channels \mathbf{h}_i and \mathbf{h}_j , the chordal distance between the channels is defined as:

$$d_c(\mathbf{h}_i, \mathbf{h}_j) = \sqrt{1 - \frac{|\mathbf{h}_i \mathbf{h}_j^*|^2}{|\mathbf{h}_i|^2 |\mathbf{h}_j|^2}} \quad (4.2)$$

The multicast beamformer can be efficiently designed for a group of channels with low chordal distance between each other. This is because of two reasons. First, a beamformer \mathbf{w} that has a low chordal distance from one channel in such a group, would have a low chordal distance from any other channel in the group due to the following property of chordal distance

$$|d_c(\mathbf{h}_i, \mathbf{w}) - d_c(\mathbf{h}_j, \mathbf{w})| \leq d_c(\mathbf{h}_i, \mathbf{h}_j) \quad (4.3)$$

Second, based on Eq. 4.2, minimizing $d_c(\mathbf{h}_j, \mathbf{w})$ is equivalent to maximizing $\tilde{\mathbf{w}} \tilde{\mathbf{h}}_j \tilde{\mathbf{h}}_j^* \tilde{\mathbf{w}}^*$ (SNR) where $\tilde{\mathbf{w}}$, and $\tilde{\mathbf{h}}_j$ are the normalized beamforming vector and normalized channel vector.

Hence, when I later design a beamformer for clients that are grouped together based on their chordal distance, the beamformer would efficiently increase the SNR across all the clients.

Algorithm 1 summarizes the procedure for grouping of users into a given number of partitions. The algorithm is mainly composed of two steps:

Step 1: (Line 11) Partitioning: during this users are assigned to partitions which have the least chordal distance from the centroid or mean of the partition.

Step 2: (Line 13) Finding the centroid: the new mean of each partition is calculated.

Algorithm 1 takes the number of iterations as an input and converges to a partitioning in a small number of iterations.

4.5.4 Multicast Beamformer Design

The remaining component in algorithms JPB-A and JPB-S is that for a given set of users that are grouped together, how to design a beamformer that maximizes the minimum SNR of the users (problem \mathcal{P}_1). The solution to the optimization problem in \mathcal{P}_1 is equivalent (up to a scaling constant) to the solution to the following problem

$$\begin{aligned} \mathcal{P}_3 : \quad & \min_{\mathbf{w}} \quad \mathbf{w}^* \mathbf{w} \\ & s.t. \quad \min_k \mathbf{w}^* \mathbf{h}_k^* \mathbf{h}_k \mathbf{w} \geq \alpha, \quad \forall k \in [1, K] \end{aligned}$$

Algorithm 1 Multicast user partitioning GM-UP.

1: **Input:**

2: Channel vectors \mathbf{h}_k , $1 \leq k \leq K$

3: Number of partitions J and number of iterations Q

4: **Output:**

5: A partitioning of K clients into J sets (S_1, \dots, S_J)

6: Normalize the channel vectors $\mathbf{h}_k = \frac{\mathbf{h}_k}{\|\mathbf{h}_k\|}$, $1 \leq k \leq K$

7: Randomly assign clients to partitions s.t. $|S_i^{(0)}| \neq 0$

8: Let $\mathbf{M}_i^{(0)} = \frac{1}{|S_i^{(0)}|} \sum_{k \in S_i^{(0)}} \mathbf{h}_k \mathbf{h}_k^*$

9: Find partition centroid: $m_i^{(0)} =$ largest eigenvector $\mathbf{M}_i^{(0)}$

10: **for** $t = 1$ to Q **do**

11: $\forall j = 1, \dots, J$: Let $S_j^t = \{k : d_c(\mathbf{h}_k, m_j^{(t-1)}) \leq d_c(\mathbf{h}_k, m_i^{(t-1)}), \forall k = 1, \dots, K, \forall i = 1, \dots, J, j \neq i\}$

12: Let $\mathbf{M}_i^{(t)} = \frac{1}{|S_i^{(t)}|} \sum_{k \in S_i^{(t)}} \mathbf{h}_k \mathbf{h}_k^*$

13: Find partition centroid: $m_i^{(t)} =$ largest eigenvector $\mathbf{M}_i^{(t)}$

14: **end for**

15: $S_i = S_i^Q \quad \forall i \in \{1, \dots, J\}$

This is because the optimal solution to \mathcal{P}_1 will be given by the product of α and a scaling constant.

The Lagrangian and the KKT (Karush-Kuhn-Tucker) conditions for the optimality of \mathcal{P}_3 can be written as:

$$L(\mathbf{w}, \underline{\lambda}) = \mathbf{w}^* \mathbf{w} + \sum_{k=1}^K \lambda_k (\alpha - \mathbf{w}^* \mathbf{h}_k^* \mathbf{h}_k \mathbf{w})$$

$$\nabla_{\mathbf{w}} L(\mathbf{w}, \underline{\lambda}) = 2\mathbf{w} - 2 \sum_{k=1}^K \lambda_k \mathbf{h}_k^* \mathbf{h}_k \mathbf{w} = 0 \quad (4.4)$$

$$\lambda_k (\alpha - \mathbf{w}^* \mathbf{h}_k^* \mathbf{h}_k \mathbf{w}) = 0 \quad (4.5)$$

where $\underline{\lambda} = [\lambda_1, \dots, \lambda_k, \dots, \lambda_K]$ and $\lambda_k \geq 0, \forall k \in [1, K]$.

Based on the optimality conditions in 4.4 and 4.5, I make the following two observations, which serve as the basis for the beamformer design algorithm.

Observation 1: *The multicast beamformer \mathbf{w} is a linear combination of \mathbf{h}_k^* s.*

This can be inferred from 4.4. The reason is that 4.4 can be written as

$$\mathbf{w} = \sum_{k=1}^K \lambda_k \mathbf{h}_k^* \mathbf{h}_k \mathbf{w} = \sum_{k=1}^K \lambda_k \mathbf{h}_k^* a_k = \sum_{k=1}^K \beta_k \mathbf{h}_k^* \quad (4.6)$$

where $a_k = \mathbf{h}_k \mathbf{w}$ and $\beta_k = \lambda_k a_k$ are scalar values.

Observation 2: *Given a permutation of the users, the optimal solution can be represented as a function of the orthogonalized channels of each user with respect to the channels of users preceding it in the permutation.*

This can be inferred from (4.5). Motivated by the above observation, I provide a greedy approach to solve problem \mathcal{P}_3 . Suppose that M out of K values of λ_k are non-zero and the rest are zero. Assume an ordering $\pi(k)$ of users where $\lambda_{\pi(k)} \neq 0$ for $1 \leq k \leq M$. For a given permutation π and for all k from 1 to K , let $\mathbf{h}_{\pi,\pi(k)}^*$ be the vector obtained by successively orthogonalizing $\mathbf{h}_{\pi(k)}^*$ to all prior $\mathbf{h}_{\pi(i)}^*$ for $i = 1, \dots, k-1$. I can rewrite (4.6) as

$$\mathbf{w} = \sum_{k=1}^K \beta_{\pi,k} \mathbf{h}_{\pi,\pi(k)}^* \quad (4.7)$$

Note that by using the KKT condition and the assumption that $\lambda_{\pi(k)} \neq 0$ for $1 \leq k \leq M$, the constraint $\alpha - \mathbf{w}^* \mathbf{h}_{\pi(k)}^* \mathbf{h}_{\pi(k)} \mathbf{w} \leq 0$ has to be satisfied with equality for indices $\pi(k), 1 \leq k \leq M$. By using (4.7) and orthogonal construction of $\mathbf{h}_{\pi,k}^*$ I have the following for $j = 1, \dots, M$

$$\begin{aligned} \alpha &= \left(\sum_{k=1}^K \beta_{\pi,k}^* \mathbf{h}_{\pi,\pi(k)} \right) \mathbf{h}_{\pi(j)}^* \mathbf{h}_{\pi(j)} \left(\sum_{l=1}^K h_{\pi,\pi(l)}^* \beta_{\pi,l} \right) \\ &= \left(\sum_{k=1}^j \beta_{\pi,k}^* \mathbf{h}_{\pi,\pi(k)} \right) \mathbf{h}_{\pi(j)}^* \mathbf{h}_{\pi(j)} \left(\sum_{l=1}^j h_{\pi,\pi(l)}^* \beta_{\pi,l} \right) \end{aligned} \quad (4.8)$$

The expression (4.8) has the following interpretation, which can be used to build a greedy solution. When $j = 1$, I have

$$\alpha = \beta_{\pi,1}^* \mathbf{h}_{\pi,\pi(1)} \mathbf{h}_{\pi(1)}^* \mathbf{h}_{\pi(1)} \mathbf{h}_{\pi,\pi(1)}^* \beta_{\pi,1} \quad (4.9)$$

In this case $\beta_{\pi,1}$ can be found easily to satisfy the condition. Next for $j = 2$, I have

$$\begin{aligned} \alpha &= \beta_{\pi,1}^* \mathbf{h}_{\pi,\pi(1)} \mathbf{h}_{\pi(2)}^* \mathbf{h}_{\pi(2)} \mathbf{h}_{\pi,\pi(1)}^* \beta_{\pi,1} \\ &+ \beta_{\pi,2}^* \mathbf{h}_{\pi,\pi(2)} \mathbf{h}_{\pi(2)}^* \mathbf{h}_{\pi(2)} \mathbf{h}_{\pi,\pi(2)}^* \beta_{\pi,2} \\ &+ 2\Re\{\beta_{\pi,1}^* \mathbf{h}_{\pi,\pi(1)} \mathbf{h}_{\pi(2)}^* \mathbf{h}_{\pi(2)} \mathbf{h}_{\pi,\pi(2)}^* \beta_{\pi,2}\} \end{aligned} \quad (4.10)$$

Now, $\beta_{\pi,2}$ can be found to satisfy this condition given that $\beta_{\pi,1}$ from the previous step is used. Note that the successive orthogonality of $\mathbf{h}_{\pi,\pi(k)}$ with respect to k ensures that the conditions that are met before still remain intact as I find the values for the next $\beta_{\pi,k}$. However, at each step the value for $\beta_{\pi,k}$ that satisfies condition (4.8) might not be unique and hence it should be chosen so as to minimize the norm of the multicasting beamformer \mathbf{w} at the final step.

Thus, the key steps of the greedy algorithm are as follows.

Step 1: For a given permutation of users, orthogonalize the user channels with respect to the channels of users preceding it in the permutation (steps 6-8).

Step 2: With the help of the orthogonalized channels determined, each weight $\beta_{\pi,k}$ is obtained successively as a function of the orthogonalized channels of users $[1, k]$ such that they minimize the norm of \mathbf{w} (steps 9-20).

Step 3: Steps 1 and 2 are repeated for every permutation π to obtain the corresponding beamforming vector \mathbf{w}_π . The final beamforming vector is obtained as the one that has the minimum norm over all permutations (step 21) and the corresponding permutation $\hat{\pi}$ is obtained similarly as the one for which the beamforming vector

Algorithm 2 Greedy algorithm for multicast beamformer design GM-BF.

- 1: **Input:**
 - 2: Channel vectors \mathbf{h}_k , $1 \leq k \leq K$; SNR threshold $\alpha > 0$; Set of user permutations Π
 - 3: **Output:**
 - 4: A permutation $\hat{\pi}$ of K users; The beamforming vector $\mathbf{w} = \sum_{k=1}^K \beta_{\pi,k} \mathbf{h}_k^*$
 - 5: **for all** $\pi \in \Pi$ **do**
 - 6: **for** $k = 1$ to K **do**
 - 7: $\mathbf{h}_{\pi,\pi(k)} \leftarrow \mathbf{h}_{\pi(k)} - \sum_{l=1}^{k-1} \frac{\mathbf{h}_{\pi(k)} \mathbf{h}_{\pi,\pi(l)}^*}{\|\mathbf{h}_{\pi(k)}\| \|\mathbf{h}_{\pi,\pi(l)}\|} \mathbf{h}_{\pi,\pi(l)}$
 - 8: **end for**
 - 9: $\beta_{\pi,1} \leftarrow \sqrt{\alpha} |\mathbf{h}_{\pi,\pi(1)} \mathbf{h}_{\pi(1)}^*|^{-1}$
 - 10: **for** $j = 2$ to K **do**
 - 11: $A \leftarrow |\mathbf{h}_{\pi,\pi(j)} \mathbf{h}_{\pi(j)}^*|^2$; $B \leftarrow (\sum_{k=1}^{j-1} \beta_{\pi,k} \mathbf{h}_{\pi,\pi(k)}) \mathbf{h}_{\pi(j)}^* \mathbf{h}_{\pi(j)} \mathbf{h}_{\pi,\pi(j)}^*$
 - 12: $C \leftarrow |\mathbf{h}_{\pi(j)} \sum_{k=1}^{j-1} \mathbf{h}_{\pi,\pi(k)}^* \beta_{\pi,k}|$
 - 13: **if** $C \geq \alpha$ **then**
 - 14: $\beta_{\pi,j} = 0$
 - 15: **else**
 - 16: $\beta_{\pi,j} = \frac{1}{A} \left(|B| + \sqrt{|B|^2 + A(\alpha - C)} \right) e^{-j\varphi(B)}$, where $\varphi(B)$ is the phase of B .
 - 17: **end if**
 - 18: **end for**
 - 19: $\mathbf{w}_\pi = \sum_{k=1}^K \beta_{\pi,k} \mathbf{h}_k^*$
 - 20: **end for**
 - 21: $\mathbf{w} = \arg \min_{\mathbf{w}_\pi} \|\mathbf{w}_\pi\|$
 - 22: $\hat{\pi} = \arg \min_{\pi} \|\mathbf{w}_\pi\|$
-

has the minimum norm (step 22).

As we can see, the greedy approach tries to minimize the norm at each step although the optimal solution might not necessarily abide by this structure. For any given permutation function $\pi(k)$, if I perform this greedy algorithm for all values of k for $k = 1, \dots, K$, we can always find a vector \mathbf{w} such that all the constraints $\alpha - \mathbf{w}^* \mathbf{h}_{\pi(k)}^* \mathbf{h}_{\pi(k)} \mathbf{w} \leq 0$ are at least satisfied with equality.

The key advantage of the proposed algorithm is that there is no need for an iterative approach as in prior works [23]; such iterative approaches require fine adjustments to the solution parameters to obtain fast convergence and avoid divergence and are not amenable to practical implementations.

Note that in case that a large number of users are grouped together, considering all possible permutation of the users can become intractable. In this case, consider a small number of randomly selected permutations such that the overall algorithm is computationally tractable.

4.6 System Implementation

4.6.1 Hardware and Software

My implementation is based on the WARPLab framework as described in Chapter 2. In this framework, all WARP boards are connected to a host PC through an Ethernet switch. The host PC is responsible for baseband PHY signal processing, while WARP boards act as RF front-ends to send/receive packets over the air. Table 4.1 specifies

Frequency	2.4 GHz - Channel 14
Bandwidth	625 KHz
Payload size	100 bits
Modulation	BPSK, QPSK 16-QAM, 64 QAM
Convolutional Codes	Generator polynomials: $g_0(133)$ $g_1(171)$, Rate = 1/2, 2/3, 3/4
Base data rate	156 Kbps at 1/2 rate

Table 4.1 : WarpLab Physical layer parameters.

the PHY parameters used in the evaluation. The APs use four radio boards which are connected to 3 dBi antennas, and are mounted on a circular array structure with a half-wavelength ($\frac{\lambda}{2}$) distance between adjacent antennas (6.25 cm at 2.4 GHz).

Note that the implementation uses a channel bandwidth of 652 KHz. This channel bandwidth is smaller than the 20 MHz channel bandwidth used in 802.11 a/b/g. I emphasize 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 with OFDM modulation based standards (e.g., 802.11 a/g) where the channel is divided into many subcarriers, per subcarrier (or a group of subcarrier) channel information provides accurate channel information.

4.6.2 Multicasting Framework

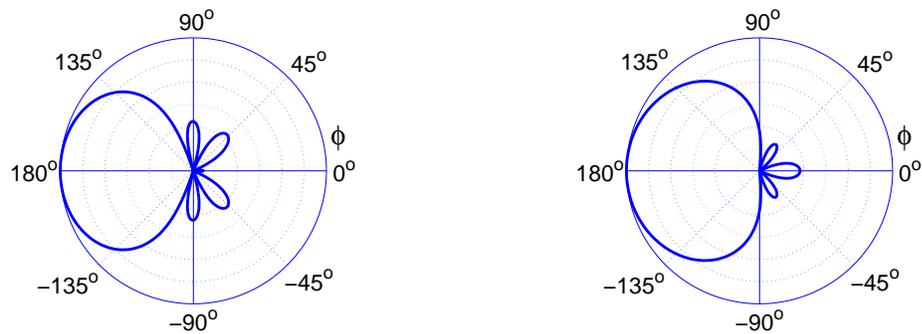
I implemented three multicast mechanisms on the testbed.

Omni. This mechanism obtains periodic SNR feedbacks from all of the clients in the multicast group. Next, it transmits multicast packets with the rate that is supported by the weakest client. This mechanism always uses the first (fixed) antenna for transmission.

Multicasting with Switched Beam Antennas. I have considered Linear and Circular arrays for switched beamforming with 3 and 4 orthogonal beams respectively. Fig. 4.4 depicts the beam patterns that are created for these antenna arrays. In a linear array with antenna separation distance of $\frac{15*\lambda}{100}$, 3 orthogonal beams can be created [26]. Fig. 4.4(a), and 4.4(b) depict two of these beam patterns. With appropriate shifting of the phase across the antennas, a third beam can be generated that is similar to Fig. 4.4(a), which however will point towards the 0° direction. In circular arrays, antenna elements are placed in a circle with equal distance between each two neighbor antennas. Fig. 4.4(c) depicts the resulting beam pattern for separation distance of $\frac{\lambda}{2}$ [26]. With appropriate shifting of the phase across the antennas, the beam pattern of Fig. 4.4(c) can be rotated to point towards the -90° , 0° , and 90° directions, thus, providing 4 orthogonal beams.

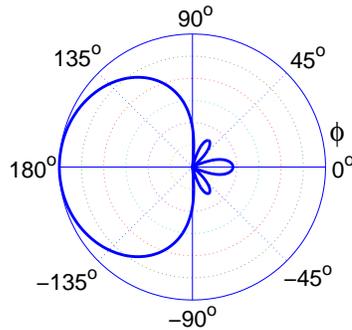
I have implemented switched multicast beamforming according to [20], whose solutions search over beam patterns that are a superset of those considered in [21], and shows considerable gains compared to [21]. In this approach, the base station

transmits training symbols for each of its beams sequentially. Next, the clients feed-back the beam index on which the strongest signal was received, together with the corresponding beam index. The base station then constructs a set of optimal beams to cover all of the clients. However, when a composite beam is used, the total power is equally distributed among its constituent beams. In such cases, the algorithm predicts the resulting SNR of the clients that are associated to a composite beam and selects a rate that is supported by the client with the lowest SNR.



(a) Linear beam pattern 1

(b) Linear beam pattern 2



(c) Circular beam pattern 1

Figure 4.4 : Switched beam patterns

ADAM. I have implemented the components of ADAM based on the discussion in the previous section.

4.6.3 Implementation

I now describe the components of the implementation.

Channel Training During the channel training, the transmitter sends a known preamble. The preamble is composed of a training sequence and a pilot tone. The training sequence is used to achieve frequency and phase synchronization between the transmitter and receiver. The pilot is used for actual channel estimation. In omni, the preamble is sent over the fixed antenna. For each of the beam patterns in switched beamforming, the preamble is multiplied by the corresponding beam weight. The weighted preambles are next transmitted sequentially. In adaptive beamforming, the base station transmits the preamble sequentially on each of its antennas. Thus, clients can correctly measure the channel for each transmitting antenna.

Channel Estimation. During the channel estimation, each client measures the \mathbf{h} or SNR for each of the preambles and sends it to the host PC. In omni, each of the clients measure the preamble's SNR and feeds back its value. In switched beamforming, each beam pattern's SNR is measured and the value of the highest SNR together with its beam index is fed back. In adaptive beamforming, \mathbf{h} is measured and fed back by each of the clients. The feedback delay of the implementation is approximately 60 *ms*.

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 4.2 : Channel model parameters.

Modulation and Coding Scheme (MCS) Selection. All of the studied protocols in this paper, select a MCS according to the resulting SNR. Thus, I need to quantify the SNR-rate relation for the WARP boards. I have used the Azimuth ACE 400WB channel emulator [16] to find the WARP board's rate table. I connect one single antenna transmitter and one single-antenna receiver to the emulator and vary the SNR accross the full range of allowable received power for the WARP radio board. The channel profile parameters used by the channel emulator are adapted from the 802.11n task group (TGn) models for a small office environment. The channel emulator parameters are summarized in Table 4.2.

Fig. 4.5 shows the packet delivery ratio (PDR) as a function of received power for various MCSs. I select the rate of an SNR value, as the highest MCS such that the given SNR achieves 100% PDR.

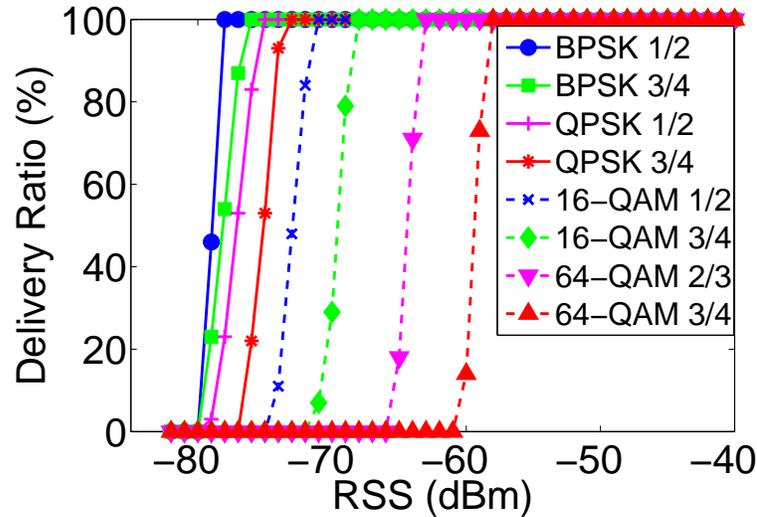


Figure 4.5 : WARP board SNR-rate relation.

Multicast Packet Transmission. In this step, the AP obtains the appropriate channel information (SNR or \mathbf{h}) by all of the clients. It then sends the multicast packet with the parameters according to the corresponding protocol.

4.6.4 Performance Metrics

All of the indoor experiments are conducted during night in an interference free environment and with static nodes. Experiments were conducted on the 802.11 2.4 GHz channel 14, which consumer devices are not allowed to use in the USA. As observed in Fig. 4.3, the variations in channel amplitude and phase in such conditions

are such that the channel remains coherent during the experiments. This allows for valid comparison among multiple multicasting schemes that are studied in this paper. Each data point in the indoor over-the-air experiments is an average of fifty samples. In the channel emulator based experiments, I take 1000 SNR measurements for each data point. I consider the received signal strength (dBm), schedule length (delay), packet delivery ratio (PDR), and throughput as the metrics for comparison of different schemes studied in this paper. I define PDR and throughput for a client, based on the number of packets that are received correctly by that client over all the transmitted packets. Next, I define the multicast PDR and the multicast throughput as the average of PDRs and throughputs over all of the clients.

4.7 Gains of Adaptive Beamforming

In this section, I compare the performance of ADAM to omni and switched beamform multicasting. I also evaluate the algorithmic components of ADAM.

Scenario. Fig. 4.6(a) depicts the experimental setup in which I deployed 6 nodes in an office environment. Nodes 1 and 2 each have four antennas and thus, can be used as transmitters or single-antenna receivers. I first consider node one as our transmitter, and amongst the remaining five nodes, consider all subsets of two, three, four, and five nodes as the different client sets for generating different topologies. I repeat the experiment with node 2 as the transmitter, leading to a total of 52 topologies. For each of these topologies, I measure the schedule length for the multicasting systems

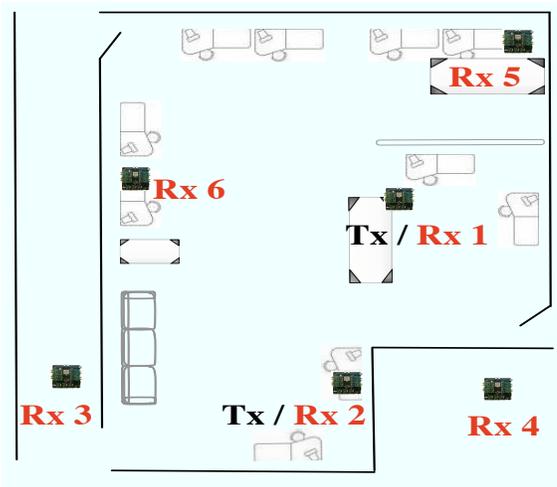
considered in this paper.

4.7.1 Impact of discrete rates

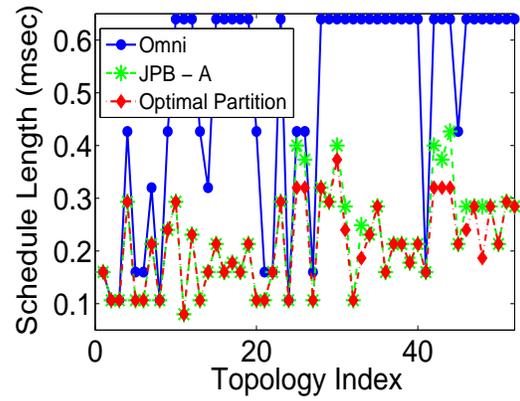
Performance Gains: Fig. 4.6(b) shows the schedule length of ADAM when the rate is selected according to the WARP SNR-rate relation of Fig. 4.5. Topology indices 1-10, 21-30, 41-45, and 51 are respectively 2,3,4, and 5 client topologies with node 1 as the transmitter. Topology indices 11-20, 31-40, 46-50, and 52 correspond to node 2 as the transmitter.

Fig. 4.6(b) shows that for some of the topologies with node 1 as the transmitter, ADAM provides negligible gains compared to omni. For these topologies, the minimum rate that is supported by omni is high. Thus, the increase in SNR due to adaptive beamforming does not provide high throughput gains. However, in topologies where at least one client has a weak channel, the gains of adaptive beamforming are much higher. In such topologies, omni would choose the lowest rate such that all clients can successfully receive the packet. A similar increase in the SNR would then result in high gains due to the nonlinear mapping of SNR-rate of WARP boards. On average, in this experiment ADAM reduces the schedule length by a factor of 2.8 compared to omni.

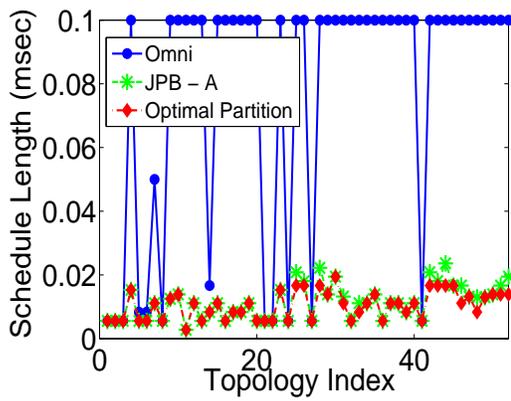
Sub-optimality of Partitioning: Fig. 4.6(b) also compares the performance of ADAM's user partitioning (JPB-A) to the optimal partition. I find the optimal partition of a given topology, by considering all possible partitions of its correspond-



a) Map of the environment.



b) Schedule length with WARP



c) Schedule length with 802.11 rates

SNR (dB)	Rate (Mbps)	SNR (dB)	Rate (Mbps)
802.11a			
≥ 24.56	54	[10.79, 17.04]	18
[24.05, 24.56)	48	[9.03, 10.79)	12
[18.8, 24.05)	36	[7.78, 9.03)	9
[17.04, 18.8)	24	[6.02, 7.78)	6
802.11 Base Rates			
[5.03, 6.02)	2	[2.01, 5.03)	1

d) 802.11 rate table

Figure 4.6 : Gains of ADAM.

ing client set and selecting the one with the minimum schedule length. According to Fig. 4.6(b), JPB-A has a performance that is very close to that of the optimal partition. On average, JPB-A increases the schedule length only by 7% compared to that of the optimal partition.

Dynamic Range of Rate Tables: ADAM's user partitioning and its overall schedule length is dependent on the SNR-rate mapping of its hardware. I now explore ADAM's performance when I select the rates according to 802.11's rate table. The SNR-rate mapping of 802.11a is shown in Fig. 4.6(d). Fig. 4.6(c) depicts the schedule length of ADAM as well as omni. In order to measure the schedule length, I measure the beamformed multicast packet's SNR at the corresponding clients. Next, I map the measured SNR to 802.11 rate table of Fig. 4.6(d) and calculate the resulting schedule length for each of the schemes.

Fig. 4.6(c) shows that ADAM has significantly reduced the schedule length with an average reduction factor of 9. 802.11a uses OFDM modulation with rates of 6 to 54 Mbps. It also supports basic rates of 1 and 2 Mbps with DSSS modulation. Thus ADAM has the potential to provide gains as high as 54. This in turn results in additional decrease in schedule length as compared to WARP board's SNR-rate table.

Finding: ADAM with four antennas can reduce the schedule length by about 2.8 times compared to omni. As the SNR of the weakest client increases, ADAM's gain decreases. ADAM's gains are also highly dependent on the SNR-rate table used by

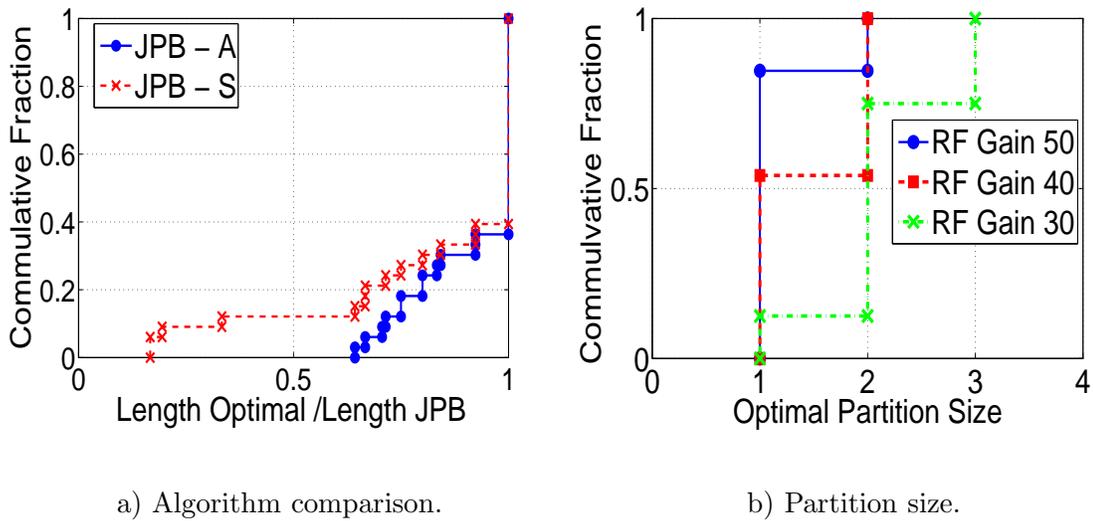


Figure 4.7 : Algorithm evaluation.

the specific hardware and can significantly increase when the dynamic range of a rate table is high.

4.7.2 Algorithm Evaluation

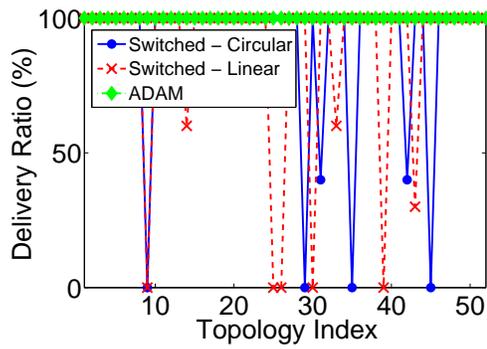
I now evaluate the algorithmic components of ADAM. I start by comparing the performance of JPB-A and JPB-S. JPB-A considers all possible number of partitions ([1 to K]) for K clients, whereas JPB-S successively increases the number of partitions (details in Section 4.5.2).

Performance vs. Complexity: Fig. 4.7(a) depicts the cumulative fraction (CDF) of the ratio between schedule length of the optimal user partitioning to that of the proposed partitioning algorithms. Observe that JPB-A achieves a schedule length that is close to that of optimal user partitioning. However, the performance

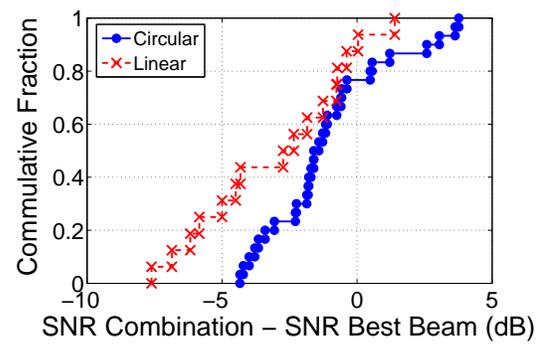
of JPB-S could be significantly lower than JPB-A. Our results show that JPB-S can converge to a local minimum, while JPB-A considers a higher number of partitions and thus can achieve a better performance.

Optimal Partition Size: Fig. 4.7(b) shows the CDF of the optimal partition sizes for three different transmission powers. For high transmission powers (RF transmit gain 50), up to 85% of topologies do not require partitioning. As I reduce the transmission power, the need for partitioning increases. Fig. 4.7(b) shows that with 10 dB reduction in transmission power (20 step reduction in RF gain), only 10% of the topologies would not require partitioning, while 70% would require at least two partitions. The need for partitioning with low power is due to two reasons: First, with a low transmission power, it may not be feasible to serve all of the clients in the same group. Second, with low transmission power, a higher number of clients would have low quality links. Due to the discrete nature of SNR-rate mapping and the fact that SNR increase in lower rates results in higher throughput gains, beamforming to a smaller group size provides a higher gain compared to serving all users together.

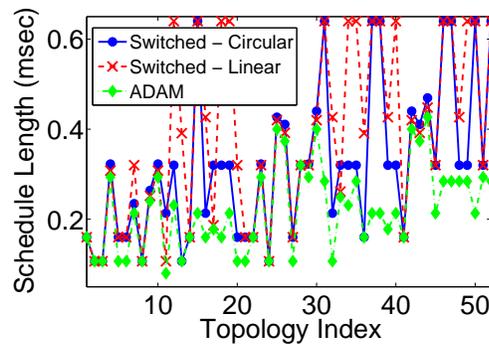
Finding: In general, the optimal partition size of K clients should be exhaustively found by considering up to K partitions. However, the experimental results show that the typical number of optimal partitions is low. Thus, in order to reduce the computational complexity one can limit the number of partitions to a small constant, independent of K .



(a) PDR



(b) Impact of beam combining



(c) Predicted schedule length of switched-beam

Figure 4.8 : Evaluation of switched beamforming.

4.7.3 Adaptive vs. Switched beamforming

In this section, I compare the performance of ADAM to that of switched beamforming. I have used the same experimental setup of Fig. 4.6(a). For each topology, I first perform adaptive beamforming. Next, without changing the antenna array, I perform switched multicast beamforming by using the pre-determined beams for the circular array. Finally, I change the antenna array to a linear array and perform switched multicast beamforming with its corresponding beam weights. While changing the antenna array, I keep the first antenna at its former location. Since the performance of omni is only dependent on the first antenna, its schedule length remains similar to that of Fig. 4.6(b)).

Drawback of Switched Beamforming: Fig. 4.8(a) shows the main drawback of switched beamforming for multicasting. The resulting PDR of switched beamforming could be a lot lower than the predicted 100%, and could be equal to zero for many topologies. This is due to the composite beam construction of switched beamforming. In particular, switched beam algorithms predict the resulting SNR of a composite beam according to the SNR of the constituent beams. According to [20], when two beams are combined together, the power allocated to each beam is distributed in half (with resulting SNR on each beam being 3 dB less) so that the total power remains unchanged, and selects a MCS accordingly.

I have performed an experiment to show the inaccuracy of such a modeling assumption. For each of the clients in the topology of Fig. 4.6(a), I find the beam that

achieves the highest SNR for both linear and circular array structures. Next, for each client I construct a two-lobe composite beam by combining its best beam, with every other beam of that particular antenna array. Finally, I measure the resulting SNR of the constructed composite beam, and subtract it from the SNR obtained by using the best beam alone. Fig. 4.8(b) shows the cumulative results (over beam combinations) for both linear and circular arrays. While [20] assumes that combining two beams would result in a 3 dB SNR drop, the experimental results reveal quite the contrary, indicating that the resulting SNR could be significantly higher or lower than the predicted SNR. This is because, even when the constituent beams are orthogonal, when a composite beam is used in an indoor multipath environment, the resulting energy at each client not only depends on its chosen constituent beam but also on other beams due to reflections and multipath scattering. Depending on whether the resulting effect is constructive or destructive, the resulting SNR could be higher or lower and hence highly inaccurate compared to that predicted from its chosen constituent beam.

Relative Gains: I now compare the gains of ADAM and switched beamforming with respect to the omni baseline of Fig. 4.6(b). ADAM provides an average gain of 2.8 over omni. Switched beamforming provides an average gain of 1.6 (with circular array), and 1.3 (with linear array) over omni. It is clear that even without considering PDR, switched beamforming provides limited gains compared to ADAM, which consistently outperforms switched beamforming in every topology. This can

be attributed to the fact that switched beam uses only a finite set of pre-determined beams which might even have a lower gain compared to an omni transmission in the presence of multipath. Indeed, by comparing Fig. 4.6(b) and Fig. 4.8(c) we observe that in many scenarios switched beamforming would not be used and instead the switched beam algorithm would end up using 802.11's omni transmission.

Finding: Switched beamforming for multicasting has fundamental limitations in indoor multipath environments. First, ADAM benefits from indoor multipath by choosing appropriate weights that reinforce the multipath components at the receiver, whereas switched beamforming has limited gains due to pre-determined beam patterns. Second, the resulting SNR of a composite beam could be significantly lower than the predicted SNR, thus lowering the PDR and increasing the actual schedule length.

4.8 Impact of Channel Dynamics

The experiments so far were conducted with perfect channel information at the transmitter. However, in any practical system the rate of channel feedback that is available from a client may not be sufficient compared to the coherence time of its channel. The channel feedback time scale could be inherently limited in the system for overhead reduction, and/or the channel coherence time could be small due to high variations in the environment or client mobility. This would cause inaccurate channel information at the transmitter which can significantly reduce the gains of ADAM and may even degrade its performance to worse than omni. In this section, I first explore the rela-

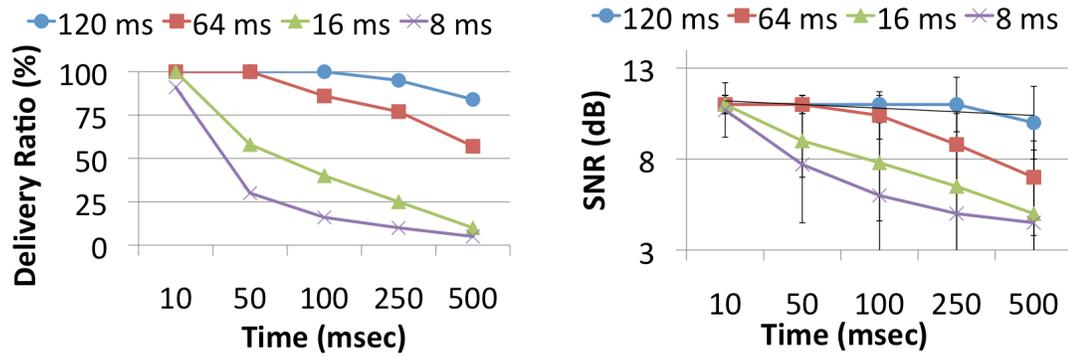
tion between channel feedback rate and channel coherence time on the performance of ADAM. Next, I propose solutions to compensate for the lack of timely channel feedback, such that the benefits of ADAM are retained.

Scenario. In order to have precise and repeatable channel conditions, I use a channel emulator for the experiments within this section. I use the same channel emulator configuration setup of Section 4.6. However, the topology is composed of a four-antenna transmitter, and three single-antenna receivers. The three receivers constitute a single multicast group to whom the the transmitter jointly beamforms.

4.8.1 Feedback Rate and Coherence Time

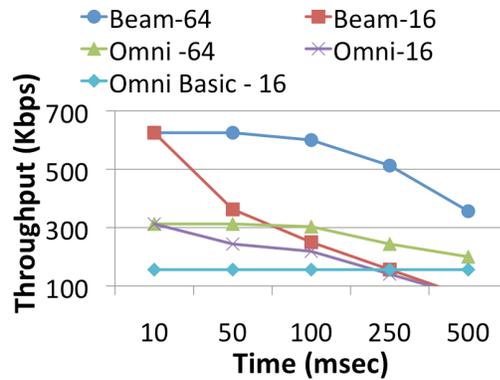
I now evaluate the gains of beamforming in changing channel conditions as a function of feedback rate. Specifically, I vary the time scale of channel information feedback (t_f) that is available at the transmitter. Once the transmitter obtains the channel information, it jointly beamforms towards the clients and transmits back-to-back multicast packets until the next channel information feedback is available. I repeat this experiment for four coherence time (t_c) values of 120, 64, 16, and 8 *ms*. The 120 and 64 *ms* t_c values are associated with a fixed wireless endpoint in slowly and highly varying environments, respectively. The 16 and 8 *ms* t_c values are associated with a typical pedestrian client in slowly and highly varying environments.

Coupling between t_f and t_c : Fig. 4.9(a) shows the average PDR as a function of channel feedback time scale for different coherence times. We observe that the



(a) PDR

(b) SNR



(c) Throughput

Figure 4.9 : Impact of coherence time and feedback rate on ADAM.

PDR of multicast beamforming drops as the time scale of channel feedback increases for a given coherence time, or as the coherence time decreases for a fixed feedback time scale. This drop in PDR is significant for smaller coherence times (16 and 8 *ms*) associated with user mobility. We also observe that for 8 *ms* coherence time, the time scale of 10 *ms* for channel feedback results in approximately 8% drop in PDR, whereas 100% PDR is achieved for all of the other t_c .

To understand the reason for the drop in PDR, I evaluate the variation in the received average SNR of clients in the multicast group in Fig. 4.9(b) as a function of channel feedback time scale. In these experiments, I measure the SNR value for every packet over all of the clients and plot the average SNR and its standard deviation. We observe that the average SNR drops as the time scale of channel feedback (t_c) increases for a given coherence time (t_f), or the coherence time decreases for a fixed feedback rate, thereby corroborating the corresponding trend observed in PDR. This also indicates the strong coupling between t_f and t_c (specifically the ratio of $s = \frac{t_f}{t_c}$) that keeps track of channel dynamics and hence impacts the multicast performance of a group.

Finding: Channel variations reduce the effective SNR of a multicast group, which in turn depends on both t_f and t_c , and more specifically on $s = \frac{t_f}{t_c}$

Impact on Performance: I next compare the performance of ADAM to omni. In omni, the transmitter selects a rate that is supported by the weakest client. This rate is used for all of the multicast packets until the next SNR feedback is available.

Omni with base rate uses the lowest MCS without any feedback requirement from the clients. This approach is currently implemented in 802.11 for multicasting.

Fig. 4.9(c) depicts the throughput results for 16 and 64 *ms* coherence times. While both ADAM and omni (denoted as omni FB) are highly sensitive to accurate channel information, the sensitivity is higher in ADAM as expected due to its stronger dependence on channel information. What is interesting is that even in the presence of increased channel dynamics, ADAM continues to provide gains over 802.11 with feedback. However, at extremely reduced feedback rate ($t_f = 500$ ms) and small coherence time ($t_c = 16$ ms), i.e. large s values, both the schemes degrade to perform even worse than omni with base rate.

Finding: In order to realize the benefits of ADAM, channel information must be obtained in relation to the clients' coherence times. Inaccurate channel information, characterized by large s values, can significantly reduce the multicast throughput to even lower than omni with base rate.

4.8.2 Reduced Feedback and Mobility

In any multicast system, the required PDR is dependent on the application. As seen in Fig. 4.9(a), for a given PDR requirement, clients with smaller coherence times require more frequent feedback. This could result in significant training and feedback overhead especially with a high number of clients and/or transmit antennas. Also, when clients in a multicast system have different coherence times, a single client with a small coherence time is sufficient to significantly increase the training overhead.

This is because the frequency at which the AP should transmit training symbols on each of its antennas depends on the client with the smallest coherence time. Thus, for any practical system it is desirable to reduce the feedback rate and hence the overhead.

Since there is no control over t_c of clients and would like to keep t_f fixed to a desired value to minimize the overhead, the resulting infrequent feedback (for clients with small t_c) reduces the effective SNR of the multicast system as seen in Fig. 4.9(b). Hence, to account for the reduced effective SNRs, I propose to train ADAM's operational SNRs based on both t_f and t_c . Since the inaccuracy in channel information is directly related to $s = \frac{t_f}{t_c}$, training here refers to obtaining the SNR-rate profiles that are specific to different s values. ADAM then categorizes clients based on their s value and applies the appropriate s -rate table for each client in determining the effective multicast rate. Thus, accounting for t_f and t_c of each client helps build robustness into ADAM's operation against infrequent feedback and client mobility.

s -valued Rate Tables: To train a rate table corresponding to a given $s = \frac{t_f}{t_c}$, I perform an experiment with channel emulator with one sender and one receiver. For each SNR value, the transmitter sends back-to-back packets to the receiver for a duration of t_f , measures the PDR and repeats this experiment for a thousand trials. The emulator uses the same configuration parameters of Section 4.6. However, instead of using a static channel ($t_c = \infty$), its t_c value is based on the s parameter.

Fig. 4.10 shows the achieved PDR as a function of the SNR (dBm) for each of the

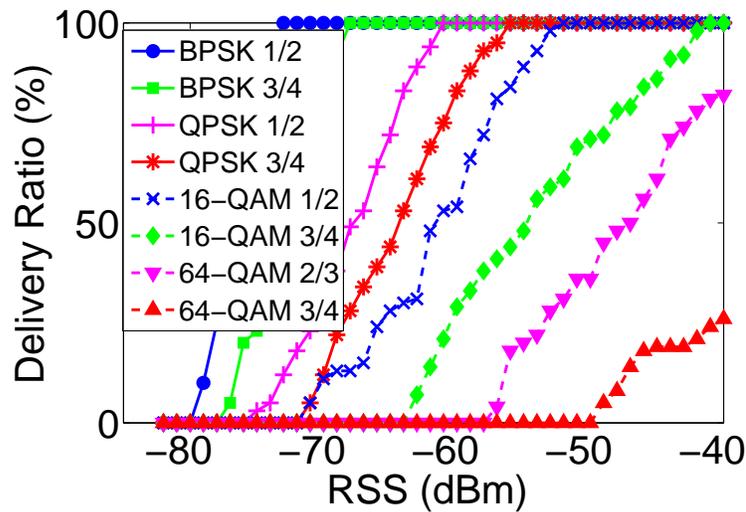


Figure 4.10 : SNR-Rate for $s = \frac{50}{8}$.

WARP MCSs for an $s = \frac{50}{8}$ ($t_f = 50, t_c = 8$ ms). Comparing Fig. 4.10 with Fig. 4.5, we observe that the required SNR for 100% PDR is now increased. In other words, a higher average SNR is required to sustain a given MCS so as to compensate for the infrequent feedback available to track the channel dynamics.

Impact on Robustness: I now quantify the gains of training ADAM based on s -rate tables. To achieve this, I use the same experimental setup of Fig. 4.9. However, I obtain our rate table according to Fig. 4.10 for $s = \frac{50}{8}$. Fig. 4.11 shows the performance of ADAM both with and without training for coherence times of 8 and 16 ms .

It can be seen that the gains of training are dependent on the time scale of channel update. With a 10 ms update rate, the untrained system is capable of tracking channel dynamics to yield high throughput. However, training becomes critical to

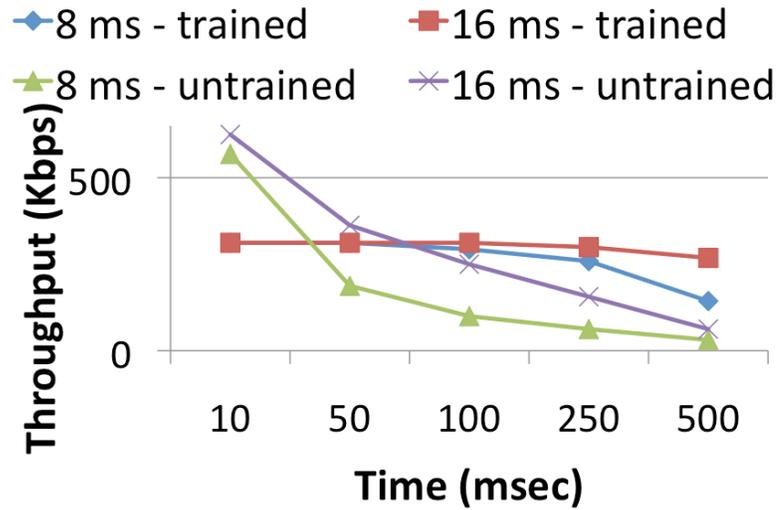


Figure 4.11 : Impact of training on throughput.

sustain high throughput when channel update rates are equal or higher than t_f for the corresponding s . Since a trained multicast system selects a lower MCS to account for channel variations, its resulting throughput compared to an untrained system would be lower for feedback time scales smaller than t_f , and higher for the time scales larger than t_f . Note that apart from throughput, PDR is another metric that should be considered in selecting between a trained vs. untrained rate table. In the above experiment, 100% PDR is achieved by the trained system for two data points, whose (t_c, t_f) is (8,50) ms and (16,100) ms respectively. However, their s value is the same ($s = \frac{50}{8}$), thereby indicating the performance dependence on the s value as opposed to the individual t_f and t_c values.

Finding: Training a rate table based on coherence time and feedback rate allows ADAM to effectively accommodate clients with varied (t_c) values. The client specific

SNR-rate mapping can be incorporated in the user scheduling optimization problem to further reduce the overall schedule length, which is an interesting avenue for future research.

4.9 Summary

In this chapter, I presented the design and implementation of ADAM, an adaptive beamforming system for multicasting in indoor wireless environments. I proposed efficient algorithms to solve the joint scheduling and beamformer design problem. I also implemented ADAM on the WARP platform, and through extensive indoor measurements showed significant gains compared to switched beamforming and omni. I also investigated the performance of ADAM as a function of channel feedback rate and user mobility, and proposed solutions to increase its robustness to channel dynamics.

Chapter 5

Related Work

5.1 Multi-User Beamforming

Single-User MIMO. Single-User MIMO systems, for example 802.11n [8] and BLAST [27], 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 [28], authors have implemented an SUBF platform to evaluate its performance in indoor wireless networks. As a baseline for comparison, I have also implemented the SUBF scheme. In [29], authors have performed EVM-based analysis for accurate SNR measurement of SUBF. Such approach can also be done for MUBF to provide accurate system performance analysis.

In [30], authors have developed a channel estimation scheme based on the received power for each of the access point's transmitting antennas. This approach could be used by legacy devices for channel estimation in multi-user beamforming.

Multiple Antennas on Mobile Devices. Multiple antenna devices such as wireless handsets are currently being produced. These devices use multiple (often two) omni-directional antennas for antenna diversity [31, 32, 33, 34]. However, client devices are typically constrained by cost and size and thus accommodate a lower number of antennas as compared to access points. The focus of my thesis is in such scenarios which are typical in current wireless LAN deployments.

Spatial Reuse. Prior research has used directional [35, 36, 37] or sectorized [38] antennas to increase SNR at intended receivers while increasing spatial reuse. [37] performs experimental characterization of multi-antenna arrays in outdoor environments, while [35, 36, 38] investigate the spatial reuse capabilities of the aforementioned antenna technology in indoor wireless networks. Similarly, I have investigated the benefits of ZFBF at reducing interference. However, in contrast to all prior work, I 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 [9, 39, 40, 41, 11]. Information theory results [9] have shown that DPC [13, 11] is the optimal strategy in MIMO downlink channels. However, DPC is difficult to implement due to extensive computational complexity. ZFBF [12] 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 [14]). The performance of these schemes are usually calculated under simulated channel conditions with uncorrelated channel gains. In [15], Kaltenberger et al. use measured channel gains to evaluate the aggregate performance of ZFBF in outdoor environments. In contrast to all, I have designed a platform to experimentally evaluate the performance of ZFBF in indoor wireless networks.

Practical MU-MIMO Protocols. Arraycomm [1] has built outdoor, cellu-

lar base stations with twelve antennas that can create up to four spatial channels. However, I have designed an open experimental framework for the prototyping and implementation of various MUBF algorithms. In addition, I have measured the performance of ZFBF in indoor wireless networks and have explored various factors that affect its performance.

Recent work [42, 43, 44] has proposed MU-MIMO protocols to increase network capacity. IAC [42] 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 my work, the system bottleneck is the number of antennas at the receiver. Thus, I consider the issue of using an AP's antennas to serve multiple users simultaneously. SAM [43] addresses the problem of serving multiple users with a single AP; however, this work considers the uplink channel problem. In contrast, I consider the downlink channel problem. In [44], Zhang et al. propose algorithms to solve the scheduling problem for a ZFBF-enabled transmitter. In my work, I identified the factors that affect the performance of ZFBF and evaluated their impact through indoor experiments.

5.2 Multicast Beamforming

Omni Antennas and Multicast: The growing demand for mobile applications involving group communication has resulted in increased focus on designing efficient link layer multicasting solutions. Link layer multicast solutions with omni-directional an-

tennas have been proposed in [45, 46, 47, 48]. While these solutions are restricted to theory, recently [49] proposed a practical multicast system for WiFi to alleviate its known problems of low data rate and high loss. However, by virtue of being designed for omni-directional antennas, these solutions cannot be directly applied for use with beamforming antennas.

Beamforming and Multicast: Beamforming has received a lot of attention recently in unicast [50, 51, 52, 53, 54] and multicast [55, 23, 56, 21, 20] applications. For unicast applications, these include both theoretical [50, 52] and practical [51, 53, 54] systems that leverage switched beam antennas. The joint problem of multicasting and (adaptive) beamforming has received significant attention in the physical layer community [55, 23, 56, 57, 58, 59] from a theoretical perspective. While these works target the continuous (power, rate) version of the problem without addressing the scheduling aspect, I consider both in this work, which makes the problem different. More importantly, I also build a practical system that realizes the benefits of adaptive beamforming for multicast. On the other hand, the joint problem of scheduling and beamforming has been considered in theory with respect to switched beamforming antennas [21, 20]. In addition to these solutions being less effective in practical indoor multipath environments (shown experimentally later), the problem formulation and hence solutions are significantly different when it comes to adaptive beamforming.

MU-MIMO Protocols: Multi-user MIMO has been recently explored in [19, 42] for unicast. In unicast, the different streams cause mutual interference to one another.

On the contrary, in multicasting a common stream needs to be optimized for all of the clients. Thus, MU-MIMO techniques for unicast do not apply to the multicast problem, necessitating complete redesign of the beamforming algorithms along with scheduling for multicast.

Chapter 6

Conclusion

In conclusion, I evaluated the benefits of multiple-antenna enabled access points in wireless LANs for unicast and multicast applications. I designed a custom, cross layer framework to evaluate the performance of multiple and previously unimplemented beamforming techniques.

By performing experiments in both in-lab controlled environments and repeatable emulated channels, I showed that a ZFBBF-enabled access point is able to simultaneously transmit to two users that are within a half a wavelength of one another. I also showed that the required channel information update rate is dependent on the channel coherence time as well as a per-link SNR requirement. Assuming that a link can tolerate an SNR decrease of 3 dB compared to Omni, I showed that the required channel update rate is equal to 100 and 10 *ms* for typical non-mobile receivers and mobile pedestrian speed of 3 mph respectively. Further, I demonstrated the potential of ZFBBF to reduce interference at unwanted locations and increase spatial reuse.

I also developed beamforming and scheduling algorithms when an access point wishes to transmit a common information to a group of users (multicasting). I designed and implemented the first adaptive-beamforming based multicast system and extensively evaluated its performance in indoor environments. I also evaluated the impact of channel variations and proposed solutions to increase the system tolerance to channel dynamics.

There are, however, remained challenges prior to widespread deployment of beamforming techniques in multiple-antenna enabled wireless LANs. First, as demon-

strated in the thesis, implementation of adaptive beamforming techniques requires fine-grained channel estimates in the form of channel amplitude and phase. There are a couple of solutions to this problem. With a more transparent interaction between physical layer and medium access layer, a client device can obtain channel information through the physical layer. However, even without a physical layer cooperation, an application layer solution could exist in the form of channel estimation based on the received power for each of the transmitting antennas. Second, as shown in the thesis, there is a need for knowledge of the coherence time of the channel for each client not only to optimally schedule users, but also to have an appropriate feedback rate. One solution to this problem could be to have an intelligent physical layer design in which a client obtains historic coherence time information through the physical layer. Without physical layer cooperation, another solution could exist in the form of using location information combined with motion changes to infer the coherence time.

Finally, there are several directions for future work. First, the scope of this thesis was limited to indoor wireless LANs. However, the work in this thesis can also be extended to outdoor mesh and sensor networks. Second, in this thesis I focused only on the downlink channel. Receive beamforming techniques coupled with appropriate MAC protocols could also be used in the uplink channel. Finally, the issue of robustness is an important issue of future research, due to delayed/erroneous/quantized channel feedback of practical systems, which is especially important for network beam-

forming.

Bibliography

- [1] “ArrayComm,” Available at: <http://www.arraycomm.com>.
- [2] “Quantenna communications,” Available at: <http://www.quantenna.com>.
- [3] “Xirrus Networks,” Available at: <http://www.xirrus.com>.
- [4] “Ieee 802.11ac work group,” Available at: <http://www.ieee802.org/11>.
- [5] D. Gesbert, F. Tosato, C. Rensburg, and F. Kaltenberger, *LTE, The UMTS Long Term Evolution: From Theory to Practice*, ch. Multiple Antenna techniques. Wiley and Sons, 2009.
- [6] J. Andrews, A. Ghosh, and R. Muhamed, *Fundamentals of WiMAX: Understanding Broadband Wireless Networking*, ch. Multiple-Antenna Techniques. Prentice Hall, 2007.
- [7] “Rice University WARP project,” Available at: <http://warp.rice.edu>.
- [8] “Antenna selection and RF processing for MIMO systems,” *IEEE 802.11-04/0713r0*, 2004.
- [9] G. Caire and S. Shamai, “On the achievable throughput of a multiantenna gaussian broadcast channel,” *IEEE Transactions on Information Theory*, vol. 49,

- pp. 1691–1706, July 2003.
- [10] S. Venkatesan and H. Huang, “System capacity evaluation of multiple antenna systems using beamforming and dirty paper coding,” *Bell Labs*, 2003.
- [11] H. Weingarten, Y. Steinberg, and S. Shamai, “The capacity region of the Gaussian multiple-input multiple-output broadcast channel,” *IEEE Transactions on Information Theory*, vol. 52, pp. 3936–3964, September 2006.
- [12] T. Yoo and A. Goldsmith, “On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming,” *IEEE Journal on Selected Areas in Communications*, vol. 24, pp. 528–541, March 2006.
- [13] M. Costa, “Writing on dirty paper (corresp.),” *IEEE Transactions on Information Theory*, vol. 29, pp. 439–441, May 1983.
- [14] A. Wiesel, Y. Eldar, and S. Shamai, “Zero-forcing precoding and generalized inverses,” *IEEE Transactions on Signal Processing*, vol. 56, pp. 4409–4418, September 2008.
- [15] 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*, vol. 8, pp. 4866–4875, September 2009.
- [16] “Azimuth Systems,” Available at: <http://www.azimuthsystems.com/>.

- [17] S. Lakshmanan, K. Sundaresan, R. Kokku, A. Khojastepour, and S. Rangarajan, "Towards adaptive beamforming in indoor wireless networks: an experimental approach," in *IEEE Infocom MiniConference*, 2009.
- [18] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannak, S. Seshan, and P. Steenkiste, "Dirc: Increasing indoor wireless capacity using directional antennas," in *SIGCOMM*, 2009.
- [19] E. Aryafar, N. Anand, T. Salonidis, and E. Knightly, "Design and Experimental Evaluation of Multi-User Beamforming in Wireless LANs," in *Proceedings of ACM MOBICOM*, Sep 2010.
- [20] K. Sundaresan, K. Ramachandran, and S. Rangarajan, "Optimal beam scheduling for multicasting in wireless networks," in *ACM MOBICOM*, Sep 2009.
- [21] S. Sen, J. Xiong, R. Ghosh, and R. Choudhury, "Link layer multicasting with smart antennas: No client left behind," in *IEEE ICNP*, Nov 2008.
- [22] H. Zhang, Y. Jiang, S. Rangarajan, and B. Zhao, "Wireless data multicasting with switched beamforming antennas," in *Proceedings of IEEE INFOCOM*, 2011.
- [23] A. Lozano, "Long-term transmit beamforming for wireless multicasting," in *IEEE ICASSP*, Apr 2007.
- [24] N. Sidiropoulos and T. Davidson, "Broadcasting with channel state information,"

- [25] R. H. H. J. H. Conway and N. J. Sloane, "Packing lines, planes, etc.: Packings in grassmannian spaces," in *Experimental Math.* 5:2, 1996.
- [26] C. A. Balanis, "Antenna theory: Analysis and design,"
- [27] G. Golden, R. Valenzuela, P. Wolniansky, and G. Foschini, "V-BLAST: an architecture for realizing very high data rates over the rich-scattering wireless channel," in *Proceedings of URSI International Symposium on Signals, Systems, and Electronics*, (Pisa, Italy), September 1998.
- [28] 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.
- [29] M. Duarte, A. Sabharwal, C. Dick, and R. Rao, "Beamforming in miso systems: empirical results and evm-based analysis," *IEEE Transactions on Wireless Communications*, vol. 9, October 2010.
- [30] S. Lakshmanan, K. Sundaresan, S. Rangarajan, and R. Sivakumar, "Practical Beamforming based on RSSI Measurements using Off-the-shelf Wireless Clients," in *Proceedings of ACM IMC*, Nov 2009.
- [31] O. Norklit, P.D.Teal, and R. Vaughan, "Measurement and evaluation of multi-antenna handsets in indoor mobile communication," *IEEE Transactions on Antennas and Propagation*, 2001.

- [32] J. Colburn, Y. Rahmat-Samii, M. Jensen, and G. Pottie, "Evaluation of personal communications dual-antenna handset diversity performance," *IEEE Transactions on Vehicular Technology*, 1998.
- [33] A. A. Sani, H. Dumanli, L. Zhong, and A. Sabharwal, "Power efficient directional wireless communication on small form-factor mobile devices," August 2010.
- [34] A. A. Sani, L. Zhong, and A. Sabharwal, "Directional antenna diversity for mobile devices: characterizations and solutions," in *Proceedings of ACM MobiCom*, (Chicago, IL), September 2010.
- [35] 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.
- [36] 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.
- [37] V. Navda, A. Subramanian, K. Dhanasekeran, A. Timm-Giel, and S. Das, "MobiSteer: Using steerable beam directional antenna for vehicular network access," in *Proceedings of ACM MobiSys*, (San Juan, Puerto Rico), June 2007.
- [38] 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.
- [39] F. Rashid-Farrokhi, K. Liu, and L. Tassiulas, “Transmit beamforming and power control for cellular wireless systems,” vol. 16, October 1998.
- [40] M. Bengtsson and B. Ottersten, *Optimal and suboptimal transmit beamforming*, ch. 18. L. Godara, Ed. Boca Raton, FL, 2001.
- [41] C. Farsakh and J. Nosssek, “Spatial covariance-based downlink beamforming in an sdma mobile radio system,” *IEEE Transactions on Communications*, vol. 46, November 1998.
- [42] S. Gollakota, S. Perli, and D. Katabi, “Interference alignment and cancellation,” in *Proceedings of ACM SIGCOMM*, (Barcelona, Spain), August 2009.
- [43] 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.
- [44] 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.
- [45] Y. Park, Y. Seok, N. Choi, Y. Choi, and J. Bonnin, “Rate-adaptive multimedia multicasting over ieee 802.11 wlans,” in *IEEE CCNC*, Jan 2006.

- [46] A. Chen, G. Chandrasekaran, D. Lee, and P. Sinha, "Himac: High throughput mac layer multicasting in wireless networks," in *IEEE MASS*, Oct 2006.
- [47] S. Jain and S. R. Das, "Mac layer multicast in wireless multihop networks," in *IEEE COMSWARE*, Jan 2006.
- [48] P. Chaporkar, A. Bhat, and S. Sarkar, "An adaptive strategy for maximizing throughput in mac layer wireless multicast," in *ACM MOBIHOC*, May 2004.
- [49] R. Chandra, S. Karanth, and T. M. et. al., "Dircast: A practical and efficient wi-fi multicast system," in *IEEE ICNP*.
- [50] R. Ramanathan, "On the Performance of Ad Hoc Networks with Beamforming Antennas," in *Proc. of ACM MOBIHOC*, Oct. 2001.
- [51] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan, "On the effectiveness of switched beam antennas in indoor environments," in *Passive and Active Measurements*.
- [52] K. Sundaresan, W. Wang, and S. Eidenbenz, "Algorithmic aspects of communication in ad-hoc networks with smart antennas," in *ACM MOBIHOC*, May 2006.
- [53] V. Navda, A. P. Subramanian, K. Dhansekaran, A. Timm-Giel, and S. Das, "Mobisteer: Using steerable beam directional antenna for vehicular network access," in *ACM MOBISYS*, Jun 2007.

- [54] A. Subramanian, P. Deshpande, J. Gao, and S. Das, “Drive-by localization of roadside wifi networks,” in *IEEE INFOCOM*.
- [55] N. Sidiropoulos, T. Davidson, and Z. Luo, “Transmit Beamforming for Physical-Layer Multicasting,” *IEEE Transactions on Signal Processing*, vol. 54, pp. 2239–2251, June 2006.
- [56] E. Matakani, N. Sidiropoulos, and L. Tassiulas, “On multicast beamforming and admission control for umts-lte,” in *IEEE ICASSP*.
- [57] Z. Luo, W. Ma, M. So, Y. Ye, and S. Zhang, “Semidefinite relaxation of quadratic optimization problems,” *IEEE Transactions on Signal Processing*, 2010.
- [58] Z. Luo, N. Sidiropoulos, P. Tseng, and S. Zhang, “Approximation bounds for quadratic optimization with homogeneous quadratic constraints,” *SIAM Journal of Optimization*, vol. 18, pp. 1–28, February 2007.
- [59] T. Chang and Z. L. C. Chi, “Approximation bounds for semidefinite relaxation of max-min fair multicast transmit beamforming problem,” *IEEE Transactions on Signal Processing*, vol. 56, August 2008.