

ADAM: An Adaptive Beamforming System for Multicasting in Wireless LANs

Ehsan Aryafar¹, Mohammad (Amir) Khojastepour², Karthikeyan Sundaresan²
Sampath Rangarajan², and Edward W. Knightly¹

¹Rice University, Houston, TX ²NEC Laboratories America, Princeton, NJ

Abstract—We present the design and implementation of ADAM, the first adaptive beamforming based multicast system and experimental framework for indoor wireless environments. ADAM addresses 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. ADAM is implemented on an FPGA platform and its performance is compared against that of omni-directional and switched beamforming based multicast.

Our experimental results reveal that (i) switched multicast beamforming has limited gains in indoor multi-path environments, whose deficiencies can be effectively overcome by ADAM to yield an average gain of three-fold; (ii) the higher the dynamic range of the discrete transmission rates employed by the MAC hardware, the higher the gains in ADAM’s performance, yielding upto nine-fold improvement over omni with the 802.11 rate table; and (iii) finally, ADAM’s performance is susceptible to channel variations due to user mobility and infrequent channel information feedback. However, we show that training ADAM’s SNR-rate mapping to incorporate feedback rate and coherence time significantly increases its robustness to channel dynamics.

I. INTRODUCTION

Wireless multicasting is becoming increasingly important for applications such as video/audio streaming. 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 [1], [2]. 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 work advocated the use of switched beamforming to improve multicasting [3], [4], [5]. 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 [3]. Be it individual or composite beams, we experimentally show that the perfor-

mance of switched beamforming is limited for multicasting in indoor wireless networks. The reasons are two fold: (i) using a pre-determined set of beam patterns limits performance when simultaneously catering to a multitude of clients in multipath environments; (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 inaccurate in multipath environments, resulting in limited performance when a composite beam is applied. To address these deficiencies, we advocate the use of adaptive beamforming for multicasting in indoor wireless networks.

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, an optimal solution needs to identify *if* and *how* a set of clients must be partitioned into separate groups (scheduling) and *how* to design an adaptive beamformer that simultaneously caters to all clients within the same group; (ii) if such a solution can be realized and implemented in practice to overcome the deficiencies of switched beamforming and provide gains in indoor multipath environments, and *what* are the factors affecting its performance; and (iii) in practical scenarios the rate of channel feedback 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, we present ADAM—the first adaptive beamforming based system for multicasting in indoor wireless networks. 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, while also making them amenable to practical implementation. It first partitions the set of clients into groups based on the “closeness” of their channels. This allows it to later determine an efficient adaptive beamformer for clients within the same group, wherein a greedy, one-shot algorithm yielding near-optimal performance is employed.

ADAM is implemented on the WARP platform and its performance is extensively evaluated in indoor environments. Our experimental results reveal that (i) while switched beamforming provides limited gains for multicasting in indoor multipath

environments, ADAM is able to address these deficiencies to yield a three-fold average gain; (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-fold over omni with the 802.11 rate table.

Finally, with controlled experiments performed with a channel emulator, we show that the performance of ADAM is 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 paper is organized as follows: Section II provides a background on beamforming along with related work. Sections III describe the challenges in realizing a practical adaptive beamforming system for multicast. Section IV describes the components of ADAM. Section V describes its implementation followed by detailed evaluation in Sections VI and VII. Finally, we conclude the paper in Section VIII.

II. BACKGROUND AND RELATED WORK

A. Preliminaries

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

In switched beamforming, a set of pre-determined beam patterns is available. Each of these beams has a main lobe of maximum gain and some side lobes representing leakage of energy. In switched beamforming, a Tx normally chooses a beam that provides the strongest signal strength at the client, without requiring fine-grained channel information. Such a beam may not coincide with the physical direction of the Rx depending on the multipath scattering in the environment [6].

In adaptive beamforming, channel estimation 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.

Multicast and Beamforming: The solution to address the beamforming-coverage tradeoff with switched beamforming is to employ individual beams sequentially or form a composite beam by combining multiple individual beam patterns so as to cover multiple clients simultaneously [3]. However, since the energy is conserved, the net power in a composite beam is distributed among its constituent beams, and hence the

resulting beamformed SNR at the clients is reduced. Hence, it becomes important to intelligently choose composite beam patterns that tradeoff user coverage and beamforming gain [3].

In adaptive beamforming, the channel to each of the clients is estimated and fed back to the AP. With the complete channel information, the AP determines and applies a beamformer that maximizes the minimum SNR among all the clients.

B. Related Work

Beamforming and Multicast: Beamforming has received a lot of attention recently in unicast [7], [8] and multicast [9], [10], [11], [4], [3] applications. The problem of multicasting with adaptive beamforming has received significant attention in the physical layer community [9], [10], [11] from a theoretical perspective. While these works target the continuous (power, rate) version of the problem without addressing the scheduling aspect, we consider both. Further, our focus is on building a practical system that realizes the benefits of adaptive beamforming for multicast. The joint problem of scheduling and beamforming has been considered in theory with switched beamforming antennas [4], [3]. In addition to switched beamforming 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: MU-MIMO implementations have been explored in [2], [12] for unicast. In unicast, different streams cause mutual interference to one another. However, in multicasting a common stream needs to be optimized for all the clients. Thus, MU-MIMO techniques for unicast do not apply to the multicast problem, necessitating redesign of the beamforming algorithms along with scheduling for multicast.

III. DESIGN CHALLENGES

In this section, we describe the system model and challenges in realizing a practical adaptive-beam multicast system.

A. System Model

We consider a single-cell environment, where a smart antenna AP is equipped with N antennas and transmits to K clients each equipped with a single antenna. Once a multicast session has been selected, our goal is to determine: (i) how to group (schedule) the clients that belong to a multicast session, into one or multiple transmissions, (ii) how to calculate the adaptive beamformer for each of the transmissions, and (iii) the transmission rate for each of the groups.

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

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

here \mathbf{x} is the transmitted symbol from the base station antennas, $\mathbf{h}_k = [h_{1k} \ h_{2k} \dots \ h_{Nk}]$ is the channel gain for the k^{th} user, and z_k represents the circularly symmetric additive white Gaussian noise at the receiver. In this model, the base station transmitter is subject to a total power constraint P , i.e., $\mathbf{x}^* \mathbf{x} \leq P$ (\mathbf{x}^* is the conjugate transpose of \mathbf{x}). The total transmit power does not depend on the number of transmit

antennas and remains the same for all the schemes studied in this paper. With beamforming, the transmitted signal \mathbf{x} is given by $\mathbf{x} = \mathbf{w}s$, where \mathbf{w} is the beamformer vector and s is the intended symbol. When beamforming is applied, the resulting SNR at a client k is equal to $\mathbf{h}_k \mathbf{w} \mathbf{w}^* \mathbf{h}_k^*$.

B. Addressing the Beamforming-Multicasting Problem

Solving the joint beamforming-multicasting problem is challenging at two levels: (i) determining an adaptive beamformer catering to a set of users simultaneously; (ii) determining *if* and *how* a set of users must be partitioned into sub-groups, where beamforming is executed separately in each group.

Adaptive Beamformer Design: Consider the objective of maximizing the minimum rate of the users in a single multicast group under constant transmit power constraint. The rate of the k^{th} user can be written as

$$R_k = \log_2(1 + \mathbf{h}_k \mathbf{w} \mathbf{w}^* \mathbf{h}_k^*) \quad (2)$$

The multicast beamforming problem is then

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

Without loss of generality we assume $\|s\|^2 = 1$. 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 [9]. Designing an efficient algorithm to address this problem, while being amenable to implementation is all the more challenging.

User Partitioning: We perform an experiment in which we increase the number of users in the multicast group from one to five in the topology of Fig. 3(a). The adaptive beamformer is determined for each group and the gain of the resulting minimum SNR of the beamformed transmission over omni is plotted in Fig. 1(a). 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 in turn advocates the need to restrict beamforming to small multicast groups and hence to partition users in a large multicast group into smaller groups, where beamforming is executed separately within each group. The need for such partitioning 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. 2(b), the transmission rate would be 1Mbps. Now, if sequential beamforming of the users increases each user's SNR by 3 dB, the resulting data rate of each client would be 9 Mbps. Since multiple transmissions are required, the performance metric shifts to latency (schedule length) -

total time required to transmit L bytes of multicast data to all users. Thus, while joint beamforming incurs a transmission time of $\frac{L}{1}$, sequential beamforming would incur $\frac{L}{9} + \frac{L}{9} = \frac{L}{4.5}$, which is latency reduction of over four-fold.

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.

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 using adaptive beamforming, such that the total schedule length to deliver L bytes to all users is minimized. The problem can now be formally stated as:

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

where J is the number of partitions, S_j is the set of user indices and \mathbf{w}_j is the beamformer for partition j . The rate function $R(\text{SNR})$ maps SNR into the appropriate rate based on a coding-modulation scheme and is discrete in practical systems. J , S_j , and \mathbf{w}_j are the outputs of the problem.

C. Robustness to 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 practical 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.

We conduct an experiment in which the AP transmits 100 pkts/sec to a static client at night. The client estimates the channel from the preambles. The variation in the channel magnitude and phase is plotted as a function of the interval size in Fig. 1(b),(c). The experiment is then repeated for a mobile client and the 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 30° , indicating the small coherence time of the channel and thereby the need for high feedback frequency.

Hence, it is important to understand the sensitivity of the adaptive beamforming multicast solution to such channel dynamics, and hence incorporate robustness into its design.

IV. DESIGN OF ADAM

We now present our adaptive-beam multicast system (ADAM) that addresses the identified challenges. We present our solution to the beamforming-multicasting problem that is

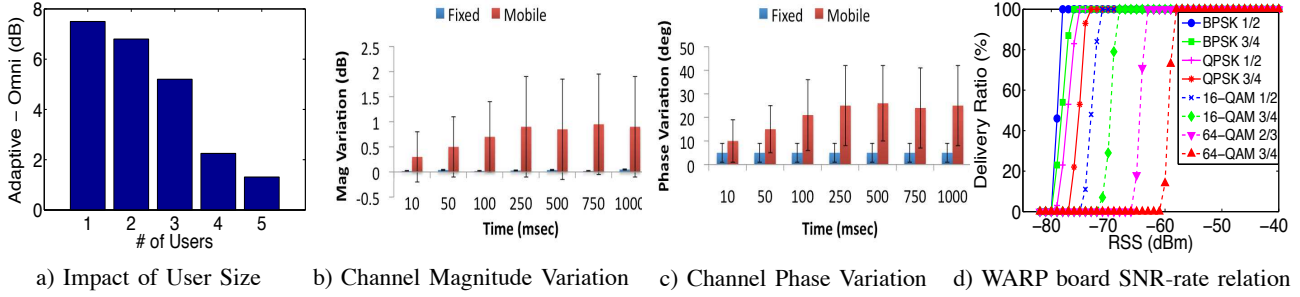


Fig. 1. Adaptive beamformer challenges ((a), (b), (c)), and WARP board SNR-rate relation (d).

amenable to a practical implementation, while its extension to handle channel dynamics is deferred to Section VII.

A. Operations in ADAM

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

(i) AP sequentially transmits training symbols on each of its antennas; (ii) Each client measures the channel amplitude and phase for each of the transmitting antennas; (iii) Clients sequentially feedback channel information to the AP; (iv) AP runs algorithms which partition the clients to different groups and find the beamformer for each group; (v) AP selects the appropriate rate for each group based on a rate table, and transmits multicast data. The algorithmic component of ADAM (step iv) is responsible for designing an efficient user partitioning and multicast beamformer algorithm.

B. Algorithm Overview

The optimal user partitioning in problem \mathcal{P}_2 depends on the rate achieved by each group (partition), which depends on the beamformer used for that group (NP-hard), which in turn depends on the set of users grouped together. To address this cyclic dependency, we decompose the problem into two sub-problems in a manner that they re-inforce each other and propose the following algorithm (**JPB-A**). For a given number of partitions, we first partition the users based on the “closeness” of their channels. This allows us to employ a greedy, one-shot algorithm to provide a near-optimal multicast beamformer within each partition. Based on the client membership and beamformer employed, the resulting rate in each partition is determined to obtain the net schedule length. The procedure is repeated for all number of partitions (up to K) and the one (j^*) yielding the minimum schedule length is chosen along with its corresponding beamformers and client memberships.

The above algorithm needs to address two sub-problems: (i) given a number of partitions, *how* to assign the clients to the given number of partitions; and (ii) *design* an appropriate beamformer for the clients within each group. These two components are discussed next.

C. User Partitioning

We use the notion of chordal distance [13] between two vectors as our metric for user partitioning. Given two users

with channels \mathbf{h}_i and \mathbf{h}_j , the chordal distance between them is defined as:

$$d_c(\mathbf{h}_i, \mathbf{h}_j) = \sqrt{1 - \frac{|\mathbf{h}_i \mathbf{h}_j^*|^2}{|\mathbf{h}_i|^2 |\mathbf{h}_j|^2}} \quad (3)$$

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

$$|d_c(\mathbf{h}_i, \mathbf{w}) - d_c(\mathbf{h}_j, \mathbf{w})| \leq d_c(\mathbf{h}_i, \mathbf{h}_j) \quad (4)$$

Second, based on Eq. 3, minimizing $d_c(\mathbf{h}_j, \mathbf{w})$ is equivalent to maximizing $\tilde{\mathbf{w}} \mathbf{h}_j \tilde{\mathbf{h}}_j^* \tilde{\mathbf{w}}^*$ (SNR) where $\tilde{\mathbf{w}}$, and $\tilde{\mathbf{h}}_j$ are the normalized beamforming and channel vectors. Hence, when we subsequently 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 Multicast user partitioning GM-UP.

- 1: **Input:**
 - 2: Channel vectors $\mathbf{h}_k, 1 \leq k \leq K$
 - 3: Number of partitions n and number of iterations Q
 - 4: **Output:**
 - 5: A partitioning of K clients into n sets (S_1, \dots, S_n)
 - 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: $\forall j = 1, \dots, n$: Let $S_j^t = \{k : d_c(\mathbf{h}_k, m_j^{(t-1)}) \leq d_c(\mathbf{h}_k, m_i^{(t-1)}), \forall k = 1, \dots, K, \forall i = 1, \dots, n, j \neq i\}$
 - 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: $S_i = S_i^Q \quad \forall i \in \{1, \dots, n\}$
-

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

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

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

D. Multicast Beamformer Design

The remaining component in algorithm JPB-A 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

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

This is because the optimal solution to \mathcal{P}_1 will be given by the product of α and a scaling constant. Based on the KKT optimality conditions of \mathcal{P}_3 , we make the following two observations which serve as the basis for our beamformer design algorithm.

Observation 1: The multicast beamformer \mathbf{w} is a linear combination of \mathbf{h}_k^* , $k \in [1, K]$.

Observation 2: Given a permutation of the users (π), 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 ($\mathbf{h}_{\pi, \pi(k)}^*$), i.e. $\mathbf{w} = \sum_{k=1}^K \beta_{\pi, k} \mathbf{h}_{\pi, \pi(k)}^*$.

Leveraging these observations, the key steps of our proposed greedy algorithm can be summarized 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.

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} .

Step 3: Steps 1 and 2 are repeated for every permutation π (or a randomly selected number of permutations to reduce complexity) 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.

Note that since the main focus of this paper is on the implementation and experimental evaluation of adaptive beamforming for multicasting, we have omitted details of the beamformer algorithm (derivation, pseudo-code, etc.) and presented them in [14].

V. EXPERIMENTAL SETUP

A. Hardware and Software

Our implementation is based on the WARPLab framework [15]. In this framework, all WARP boards are connected to a host PC through an Ethernet switch. The host PC is responsible

Freq	2.4 GHz – Ch 14	SNR (dB)	Rate (Mbps)	SNR (dB)	Rate (Mbps)
Load	100 bits	802.11a			
Modulation	BPSK, QPSK, 16-32 QAM,	≥ 24.56	54	[10.79, 17.04]	18
Coding	$=\frac{1}{2}, \frac{2}{3}, \frac{3}{4}$	[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
Base Rate	156 Kbps	802.11 Base Rates			
		[5.03, 6.02]	2	[2.01, 5.03]	1

a) WARPLab PHY Parameters.

b) 802.11 SNR-Rate table.

Fig. 2. WARPLab (a) and 802.11 SNR-Rate Table.

for baseband PHY signal processing, while WARP boards act as RF front-ends to send/receive packets over the air. Fig. 2(a) specifies the PHY parameters used in our evaluation. Our APs use four radios connected to 3 dBi antennas. The antennas 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).

B. Multicasting Framework

We implemented three multicast mechanisms on our testbed.

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

Multicasting with Switched Beam Antennas. We have considered Linear and Circular arrays for switched-beam multicasting with 3 and 4 orthogonal beams respectively. Our implementation is according to [3], whose solutions considers both individual and composite beams, and shows considerable reduction in schedule length compared to schemes using individual beams alone. In this approach, the base station transmits training symbols for each of its beams sequentially. Next, the clients feedback 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 composite 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 associated to a composite beam and selects a rate supported by the client with the lowest SNR.

ADAM. Our proposed adaptive beam multicast system.

C. System Implementation

We now describe the components of our 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 ADAM, 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 and sends it to the host PC. In omni, each of the clients measures the preamble's SNR and feeds back its value. In switched beamforming, each beam's SNR is measured and the value of the highest SNR together with its beam index is fed back. In ADAM, \mathbf{h} is measured and fed back by each of the clients.

Modulation and Coding Scheme (MCS) Selection. All of the studied protocols in this paper select a MCS according to the resulting SNR. Thus, we need to quantify the SNR-rate relation for the WARP boards. We have used the Azimuth ACE 400WB channel emulator [16] to find the WARP board's rate table. We connect one single antenna transmitter and one single-antenna receiver to the emulator and vary the SNR across 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 profile is composed of 14 Rayleigh fading channels with multipath RMS delay spread of 30 ns, and maximum delay of 200 ns. Fig. 1(d) shows the packet delivery ratio (PDR) as a function of received power for various MCSs. We select the rate of an SNR value, as the highest MCS such that the given SNR achieves 100% PDR.

Multicast Packet Transmission. In this step, the AP sends the multicast packet with the appropriate parameters.

D. Performance Metrics

All of our indoor experiments are conducted during night 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. 1(b),(c), the variations in channel in such conditions are such that the channel remains coherent during the experiments. This allows for valid comparison among multiple multicasting schemes. Each data point in our indoor over-the-air experiments is an average of fifty samples. In our channel emulator based experiments, we take 1000 SNR measurements for each data point. We consider the received signal strength (dBm), schedule length (delay), packet delivery ratio (PDR), and throughput as our metrics of comparison. We 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, we define the multicast PDR and throughput as the average of PDRs and throughputs over all of the clients.

VI. GAINS OF ADAPTIVE BEAMFORMING

In this section, we compare the performance of ADAM to omni and switched beamform multicasting.

Scenario. Fig. 3(a) depicts our experimental setup in which we 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. We first consider node 1 as our transmitter, and amongst the remaining five nodes, consider all subsets of two, three, four, and five nodes as our different client sets for generating different topologies. We repeat the experiment with node 2 as our transmitter, leading to a total

of 52 topologies. For each of these topologies, we measure the schedule length for all of the multicasting systems.

A. ADAM's performance characterization

Performance Gains: Fig. 3(b) shows the schedule length of ADAM when the rate is selected according to the WARP SNR-rate relation of Fig. 1(d). 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. 3(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. 3(b) also compares the performance of ADAM's user partitioning (JPB-A) to the optimal partition. We find the optimal partition of a given topology, by considering all possible partitions of its corresponding client set and selecting the one with the minimum schedule length. According to Fig. 3(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. We now explore ADAM's performance when we select the rates according to 802.11's rate table. The SNR-rate mapping of 802.11a is shown in Fig. 2(b). Fig. 3(c) depicts the schedule length of ADAM as well as omni. In order to measure the schedule length, we measure the beamformed multicast packet's SNR at the corresponding clients. Next, we map the measured SNR to 802.11 rate table of Fig. 2(b) and calculate the resulting schedule length for each of the schemes.

Fig. 3(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 the specific hardware and can significantly increase when the dynamic range of a rate table is high.



a) Map of the environment.



b) Schedule length with WARP



c) Schedule length with 802.11 rates

Fig. 3. Gains of ADAM.

B. Adaptive vs. Switched beamforming

In this section, we compare the performance of ADAM to that of switched beamforming. We have used the same experimental setup of Fig. 3(a). For each topology, we first perform adaptive beamforming. Next, without changing the antenna array, we perform switched multicast beamforming by using the pre-determined beams for the circular array. Finally, we change the antenna array to a linear array and perform switched multicast beamforming with its corresponding beam weights. While changing the antenna array, we 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. 3(b).

Relative Gains: We now compare the schedule length of switched beamforming to that of adaptive beamforming. Fig. 4(a) shows that ADAM provides an average gain of 1.8 and 2.1 over switched beamforming with circular and linear arrays respectively. Further, ADAM 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. 3(b) and Fig. 4(a) we observe that in many scenarios switched beamforming would not be used and instead the switched beam algorithm would end up using omni transmission.

Drawback of Switched Beamforming: Fig. 4(b) shows the drawback of switched beamforming when employing composite beams. 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. For example when two beams are combined and the power allocated to each beam is divided in half (so that total power is conserved), the inherent assumption is that the resulting SNR in each beam reduces by 3 dB and a MCS is selected accordingly.

We have performed an experiment to show the inaccuracy of such a modeling assumption in indoor multipath environments. For each of the clients in the topology of Fig. 3(a), we find the beam that achieves the highest SNR for both linear and circular array structures. Next, for each client we construct a two-lobe composite beam by combining its best beam, with every other beam of that particular antenna array. Finally, we measure the resulting SNR of the constructed composite beam, and subtract

it from the SNR obtained by using the best beam alone. Fig. 4(c) shows that when combining two beams, 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, making it hard to leverage composite beams in indoor multipath environments.

Finding: Switched beamforming has limited performance for multicasting in indoor multipath environments, while ADAM benefits from indoor multipath by choosing appropriate weights that reinforce the multipath components at the receiver.

VII. 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, we first explore the relation between channel feedback rate and coherence time on the performance of ADAM. Next, we 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, we use a channel emulator for the experiments within this section. We use the same channel emulator configuration setup of Section V. However, our topology is composed of a four-antenna transmitter, and three single-antenna receivers. The three receivers constitute a single multicast group to whom the transmitter jointly beamforms.

A. Feedback Rate and Coherence Time

We now evaluate the gains of beamforming in changing channel conditions as a function of feedback rate. Specifically, we vary the time scale of channel information feedback (t_f)

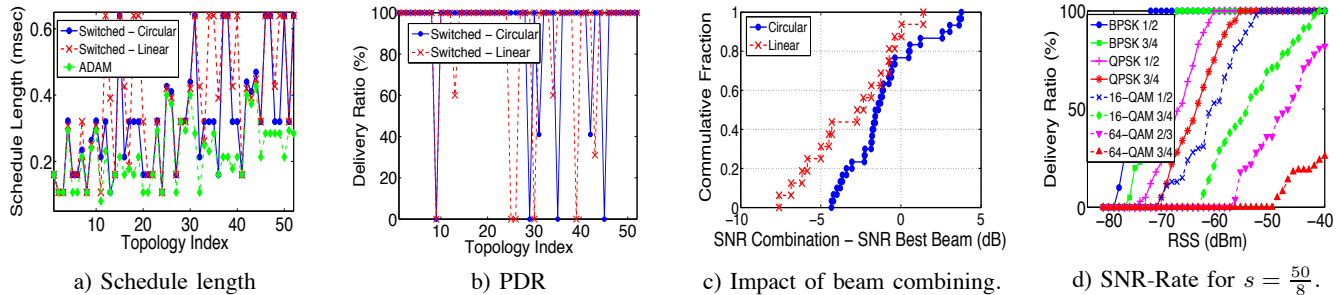


Fig. 4. Evaluation of switched beamforming ((a), (b), (c)), and WARP SNR-rate for $s = \frac{50}{8}$.

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. We 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. 5(a) shows the average PDR as a function of channel feedback time scale for different coherence times. We observe that the 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, we evaluate the variation in the received average SNR of clients in the multicast group in Fig. 5(b) as a function of channel feedback time scale. In these experiments, we 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.

Impact on Performance: We 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. 5(c) depicts the throughput results for 16 and 64 ms coherence times. While both ADAM and omni are highly sensitive to accurate channel information, the sensitivity is higher in ADAM as expected due to its stronger dependence on

channel information. Further, 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: 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 $s = \frac{t_f}{t_c}$. Inaccurate channel information, characterized by large s values, can significantly reduce the multicast throughput to even lower than omni with base rate.

B. Reduced Feedback and Mobility

In any multicast system, the required PDR is dependent on the application. As seen in Fig. 5(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 we have 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. 5(b). Hence, to account for the reduced effective SNRs, we 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 $s = \frac{t_f}{t_c}$, we 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 V. However, instead of

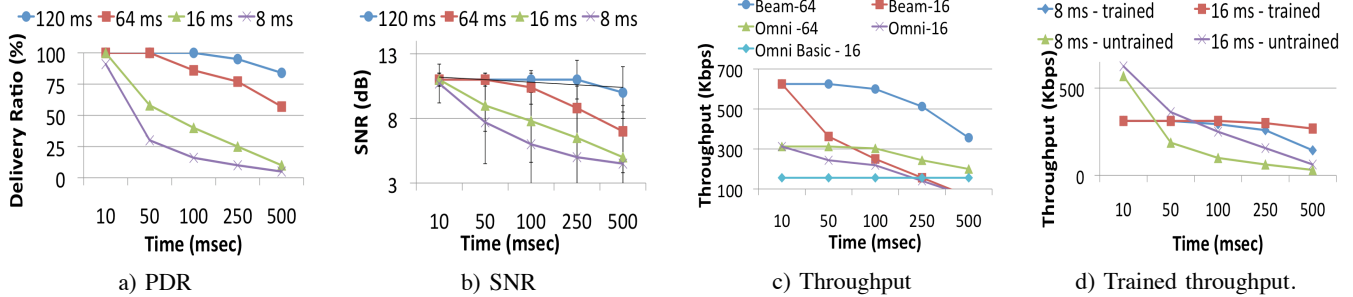


Fig. 5. Impact of coherence time and feedback rate on ADAM ((a), (b), (c)), and training impact on throughput (d).

using a static channel ($t_c = \infty$), its t_c value is based on the s parameter. Fig. 4(d) shows the achieved PDR as a function of the SNR (dBm) for each of the WARP MCSs for an $s = \frac{50}{8}$ ($t_f = 50, t_c = 8$ ms). Comparing Fig. 4(d) with Fig. 1(d), 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: We now quantify the gains of training ADAM based on s -rate tables. To achieve this, we use the same experimental setup of Fig. 5. However, we obtain our rate table according to Fig. 4(d) for $s = \frac{50}{8}$. Fig. 5(d) 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 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.

VIII. CONCLUSIONS

In this paper, we presented the design and implementation of ADAM, an adaptive beamforming system for multicasting in wireless LANs. We proposed efficient algorithms to solve the joint scheduling and beamforming problem. We also implemented ADAM on the WARP platform, and through extensive indoor measurements showed significant gains compared to

switched-beam and omni. We also evaluated the performance of ADAM with respect to feedback rate and user mobility, and proposed solutions to increase its robustness to channel dynamics.

IX. ACKNOWLEDGEMENTS

This research was supported in part by NSF grants CNS-0923479, CNS-1012921, and CNS-1012831.

REFERENCES

- [1] 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.
- [2] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly, "Design and Experimental Evaluation of Multi-User Beamforming in Wireless LANs," in *Proceedings of ACM MOBICOM*, Chicago, IL, September 2010.
- [3] K. Sundaresan, K. Ramachandran, and S. Rangarajan, "Optimal beam scheduling for multicasting in wireless networks," in *Proceedings of ACM MOBICOM*, Beijing, China, September 2009.
- [4] S. Sen, J. Xiong, R. Ghosh, and R. Choudhury, "Link layer multicasting with smart antennas: No client left behind," in *Proceedings of IEEE ICNP*, Orlando, FL, October 2008.
- [5] H. Zhang, Y. Jiang, S. Rangarajan, and B. Zhao, "Wireless data multicasting with switched beamforming antennas," in *Proceedings of IEEE INFOCOM*, Shanghai, China, April 2011.
- [6] Anand Prabhu, Henrik Lundgren, and Theodoros Salonidis, "Experimental characterization of sectorized antennas in dense 802.11 wireless mesh networks," in *Proceedings of ACM MobiHoc*, New Orleans, LA, May 2009.
- [7] V. Navda, A. P. Subramanian, K. Dhasekaran, 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.
- [8] A.P. Subramanian, P. Deshpande, J. Gao, and S. Das, "Drive-by localization of roadside wifi networks," in *Proceedings of IEEE INFOCOM*, Phoenix, AZ, April 2008.
- [9] N.D. Sidiropoulos, T.N. Davidson, and Z.Q. Luo, "Transmit Beamforming for Physical-Layer Multicasting," *IEEE Transactions on Signal Processing*, vol. 54, no. 6, pp. 2239–2251, June 2006.
- [10] A. Lozano, "Long-term transmit beamforming for wireless multicasting," in *Proceedings of IEEE ICASSP*, Honolulu, Hawaii, April 2007.
- [11] E. Matakani, N.D. Sidiropoulos, and L. Tassiulas, "On multicast beamforming and admission control for UMTS-LTE," in *Proceedings of IEEE ICASSP*, Las Vegas, NV, April 2008.
- [12] S. Gollakota, S.D. Perli, and D. Katabi, "Interference alignment and cancellation," in *Proceedings of ACM SIGCOMM*, Barcelona, Spain, August 2009.
- [13] R. H. Hardin J. H. Conway and N. J.A. Sloane, "Packing lines, planes, etc.: Packings in grassmannian spaces," in *Experimental Mathematics*, 1996, vol. 5, pp. 139–159.
- [14] E. Aryafar, M.A. Khojastepour, K. Sundaresan, S. Rangarajan, and E.W. Knightly, "ADAM: An adaptive beamforming system for multicasting in wireless LANs," <http://www.nec-labs.com/~karthiks/adam.pdf>.
- [15] "Rice University WARP project," Available at: <http://warp.rice.edu/>.
- [16] "Azimuth Systems," Available at: <http://www.azimuthsystems.com/>.