ADAM: An Adaptive Beamforming System for Multicasting in Wireless LANs

Ehsan Aryafar, *Member, IEEE*, Mohammad Ali Khojastepour, *Member, IEEE*, Karthik Sundaresan, *Senior Member, IEEE*, Sampath Rangarajan, *Senior Member, IEEE*, and Edward Knightly, *Fellow, IEEE*

***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 a field programmable gate array (FPGA) platform, and its performance is compared against that of omnidirectional and switched beamforming based multicast. Our experimental results reveal that: 1) switched multicast beamforming has limited gains in indoor multipath environments, whose deficiencies can be effectively overcome by ADAM to yield an average gain of threefold; 2) the higher the dynamic range of the discrete transmission rates employed by the MAC hardware, the higher the gains in ADAM's performance, yielding up to ninefold improvement over omni with the 802.11 rate table; and 3) 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 signal-to-noise ratio (SNR)-rate mapping to incorporate feedback rate and coherence time significantly increases its robustness to channel dynamics.**

***Index Terms*—Adaptive beamforming, channel dynamics, channel feedback rate, scheduling, switched beamforming, wWireless multicast.**

I. 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 the next generation of wireless standards such as 802.11ac [1], LTE [2], and WiMAX [3].

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 [signal-to-noise ratio (SNR)]. On a parallel front, beamforming antennas have recently gained a lot of attention in indoor wireless networks [4]–[6]. These are multiple-element antenna 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 entire azimuth of 360° are used. Recent works [7]–[9] 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 omnidirectional 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 [7].

In this paper, we experimentally show that switched beamforming has limited gains for multicasting in indoor multipath environments. The reasons are twofold: 1) using a predetermined set of beam patterns limits performance when simultaneously catering to a multitude of clients; 2) 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, 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: 1) 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; 2) 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 3) 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 solution 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.

Manuscript received April 23, 2012; revised November 08, 2012; accepted November 08, 2012; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor A. Capone. Date of publication January 22, 2013; date of current version October 11, 2013.

E. Aryafar is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: earyafar@princeton.edu).

M. A. Khojastepour, K. Sundaresan, and S. Rangarajan are with NEC Labs America, Princeton, NJ 08540 USA (e-mail: amir@nec-labs.com; karthiks@nec-labs.com; sampath@nec-labs.com).

E. Knightly is with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005 USA (e-mail: knightly@rice.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2012.2228501

Toward 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 subproblems 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: 1) while switched beamforming provides limited gains for multicasting in indoor multipath environments, ADAM is able to address these deficiencies to yield a threefold average gain; 2) ADAM's gains are more with a higher dynamic range of the (discrete) transmission rates employed by the MAC, yielding gains as high as ninefold 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 timescale (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 and IV describe the motivation and challenges of adaptive beamforming for multicasting. Section V describes the components of ADAM. Section VI describes its implementation followed by detailed evaluation in Sections VII and VIII. Discussion and future work is presented in Section IX. Finally, we conclude the paper in Section X.

II. Background and Related Work

A. Preliminaries

Beamforming: Beamforming antennas consist of an array of omnidirectional elements, with sophisticated signal processing capabilities. 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 $\mathbf{w} = [w_0\, w_1 \ldots\, w_{N-1}]^T$. 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 predetermined beam patterns is available. A transmitter normally chooses a beam pattern 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 [10].

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 is such that it 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 a fine-grained channel estimation feedback.

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 [7], [8] 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 the beamforming-coverage tradeoff with switched beamforming is to form a composite beam by combining multiple individual beams so as to cover multiple clients simultaneously [7]. However, since the energy is conserved, the net power is distributed among the constituent beams of the composite beam, and hence the resulting beamformed SNR at the clients is reduced. Hence, it becomes important to intelligently choose composite beams that tradeoff coverage and beamforming gain [7].

In adaptive beamforming, the channel to each of the clients is estimated and fed back to the access point (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

Omni Antennas and Multicast: Link-layer multicast solutions with omnidirectional antennas have been proposed in [11] and [12]. While these solutions are restricted to theory, recently [13] 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 omnidirectional 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 [14]–[16] and multicast [7], [8], [17]–[19] applications. For unicast applications, these include both theoretical [20] and practical [14]–[16] systems that leverage switched beam antennas. Practical unicast systems that leverage switched beam antennas were considered in [14]–[16]. The joint problem of multicasting and (adaptive) beamforming has received significant attention in the physical-layer community [17]–[19] from a theoretical perspective. While these works target the continuous (power, rate) version of the problem without addressing the scheduling aspect, we consider both in this paper, which makes the problem different. More importantly, we 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 [7], [8]. In addition to these solutions being less effective in practical indoor environments (shown experimentally later), the problem formulation and hence solutions are significantly different when it comes to adaptive beamforming.

MU-MIMO Protocols: Multiuser multiple-input–multiple-output (MU-MIMO) has been recently explored in [5] and [21]–[23] 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

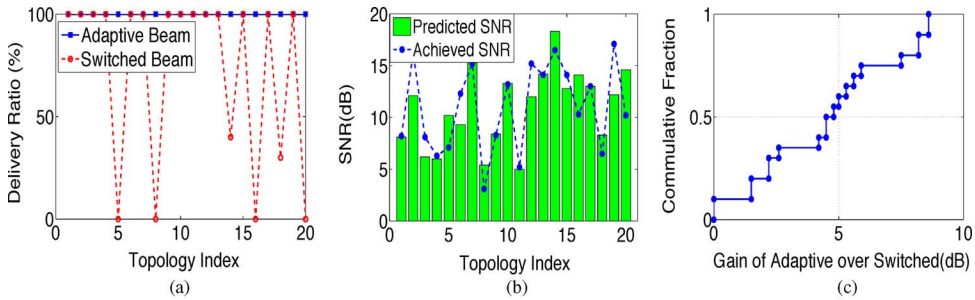

Fig. 1. Adaptive versus switched beamforming performance comparison. (a) PDR. (b) Modeling accuracy in switched beam. Switched beam can achieve low PDR due to SNR modeling inaccuracy in indoor multipath rich environments. (c) SNR gain. Adaptive beam can provide large SNR gains in indoor environments.

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.

III. Motivation

Current beamforming solutions for improving the multicast performance (e.g., [7]–[9]) advocate the use of switched beamforming. Hence, in order to motivate the need for adaptive beamforming, we address the following questions.

1) 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. We conduct an experiment in an indoor environment (for detailed topological information, please refer to Section VII) by considering three clients in a multicast group. A circular array of four antennas with four predetermined 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 [7]. 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 the 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. 1(a) are quite the contrary, where the PDR is less than 50% in 30% of the topologies, 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. 1(b). It is clear that the modeling assumption behind the predicted SNR, which may hold in line-of-sight environments, does not hold in many of our indoor topologies. Our results in Fig. 1(b) indicate that in 60% of the topologies the modeling assumption have an average error of 3.2 dB, where it either underestimates or overestimates the actual SNR. This, in turn, can be attributed to the multipath nature of indoor environments, which makes it hard to predict the effect of composite beams needed for multicasting. Note that the selected packet transmission rate (modulation and coding scheme) depends on the predicted SNR. However, if the achieved SNR is even 1 dB less than the predicted one, the corresponding PDR can drop significantly (for detailed discussion on the correspondence between PDR, transmission rate, and SNR, please refer to Section VI-C). This is also verified in the experimental results of Fig. 1(a).

2) Given That Switched Beamforming Cannot Address Multicasting Efficiently in Indoor Wireless Environments, the Next Question to Understand Is Whether Adaptive Beamforming (Designed to Handle Multipath) can Address the Deficiencies of Switched Beamforming for Multicasting. Toward 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 V). The resulting PDR for each of the topologies is compared against switched beamforming in Fig. 1(a). It can be clearly seen that adaptive beamforming is capable of delivering the predicted performance in contrast to switched beamforming. Furthermore, the commulative fraction (CDF) of the SNR gain (in dB) of adaptive over switched beamforming over all the topologies, depicted in Fig. 1(c), clearly indicates the large potential of adaptive beamforming for improving multicast performance.

IV. Design Challenges

In this section, we describe the system model and the challenges in realizing a practical adaptive beamforming 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: 1) how to group (schedule) the clients that belong to a multicast session, into one or multiple transmissions; 2) how to calculate the adaptive beamformer for each of the transmissions; and 3) the transmission rate for each of the groups.

We consider a narrowband system model, where the received baseband signal y_k of the kth user is given by

$$y_k = \mathbf{h}_k x + z_k, \qquad k = 1, \ldots, K \tag{1}$$

where $\mathbf{x}$ is the transmitted symbol from the base station antennas, $\mathbf{h}_k = [h_{1k} h_{2k} \ldots h_{Nk}]$ is the channel gain for the

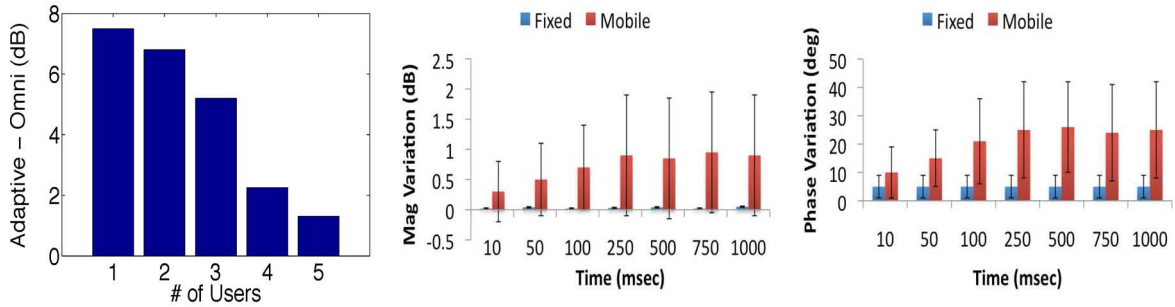

(a) (b) (c)

Fig. 2. (a) Impact of multicast group size on adaptive beam. As the size of the multicast group increases, gains of adaptive beamforming start to diminish. This advocates the partitioning of a multicast group into smaller groups. (b) Channel magnitude variation. (c) Channel phase variation. (b), (c) Channel dynamics (phase and amplitude) investigation for fixed and mobile clients. The large channel variations with mobile clients indicates the need for frequent channel feedback in such scenarios.

kth 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. 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 rate of the kth user can be written as

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

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 nonconvex problem, and its discrete version is NP-hard as well [18]. This makes it challenging to design an efficient algorithm to compute an adaptive multicast beamformer.

C. 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. We perform an experiment, where we increase the number of users in the multicast group from one to five in the topology of Fig. 5(a). The adaptive beamformer is determined for each group, and the gain of the resulting minimum SNR of the beamformed transmission over omnidirectional transmission is plotted in Fig. 2(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 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 omnidirectional transmission so as to cater to all the users. This, in turn, advocates the partitioning of users in a large multicast group into subgroups of smaller size and enabling beamforming to improve transmissions in each of the subgroups. The need for such partitioning (scheduling) is exacerbated in the presence of discrete rate tables. For example, consider two users that each achieve a 5-dB SNR when jointly beamformed to. With 802.11 rate table of (for detailed SNR-rate mapping of 802.11, please refer to Section VI), the transmission rate would be 1 Mb/s. 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 Mb/s. 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 subgroups, there is a (time) multiplexing loss with different subgroups 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 subgroup but incur the multiplexing loss.

D. 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 the 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, where the AP transmits 100 packets/s 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. 2. 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.

V. Design of ADAM

In this section, we describe the design of ADAM, our adaptive beamforming-based multicast system that addresses the identified challenges. We 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, we propose efficient algorithms that are implemented in ADAM and are suitable for a practical system design. We address the impact of channel dynamics and ADAM's solutions to increase robustness in Section VIII.

A. 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 the multicast data.

The main algorithmic component of ADAM is to design efficient user partitioning and multicast beamformer for Step 4. To evaluate this, we use the notion of schedule length (delay) required for multicast data transfer to the entire group as our metric of optimization. We assume a PDR requirement of 100% for all of the clients. If some of the clients can tolerate a lower PDR, it can be incorporated in our 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.

B. 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

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

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 $R(\mathrm{SNR})$ 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 our work is on discrete rates. J, S_j, and $\mathbf{w}_j$ are the outputs of the problem.

C. Algorithm Overview

As described in Section IV-B, finding a multicast beamformer is nonconvex, and its discrete version is also NP-Hard. Thus, it is not feasible to find an optimal beamformer with general channel vectors, even for a small-size single multicast group. The problem formulation in $\mathcal{P}_2$ 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, we adopt a decomposition approach that divides the problem into two subproblems in a manner that allows the two subproblems to reinforce each other. For a given number of groups, we 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. We then employ a greedy, one-shot algorithm to provide a near-optimal multicast beamformer within each group.

By combining the above two subproblems, we 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 subproblems: 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.

D. 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. We use the notion of chordal distance [24] between two vectors as our 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

$$d_{\mathrm{c}}(\mathbf{h}_i, \mathbf{h}_j) = \sqrt{1 - \frac{|\mathbf{h}_i \mathbf{h}_j^*|^2}{|\mathbf{h}_i|^2 |\mathbf{h}_j|^2}}. \tag{3}$$

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

$$|d_{\mathrm{c}}(\mathbf{h}_i, \mathbf{w}) - d_{\mathrm{c}}(\mathbf{h}_j, \mathbf{w})| \leq d_{\mathrm{c}}(\mathbf{h}_i, \mathbf{h}_j). \tag{4}$$

Second, based on (3), minimizing $d_{\mathrm{c}}(\mathbf{h}_j, \mathbf{w})$ is equivalent to maximazing $\tilde{\mathbf{w}} \tilde{\mathbf{h}}_j \tilde{\mathbf{h}}_j^* \tilde{\mathbf{w}}^*$ (SNR) where $\tilde{\mathbf{w}}$, and $\tilde{\mathbf{h}}_j$ are the normalized beamforming vector and normalized channel vector.

Hence, when we 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: 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 $(S_1, \ldots, S_n)$
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: $\forall j = 1, \ldots, n$: Let $S_j^t = \{k : d_{\mathrm{c}}(\mathbf{h}_k, m_j^{(t-1)}) \leq d_{\mathrm{c}}(\mathbf{h}_k, m_i^{(t-1)}), \forall k = 1, \ldots, K, \forall i = 1, \ldots, n, j \neq i\}$
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: $S_i = S_i^Q \ \forall i \in \{1, \ldots n\}$

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: (Line 11) Partitioning: During this, users are assigned to partitions that 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.

E. 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:

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

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 necessary Karush–Kuhn–Tucker (KKT) conditions for the optimality of $\mathcal{P}_3$ can be written as

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

$$\nabla_{\mathbf{w}} L(\mathbf{w}, \underline{\lambda}) = 2\mathbf{w} - 2 \sum_{k=1}^{K} \lambda_k \mathbf{h}_k^* \mathbf{h}_k \mathbf{w} = 0 \tag{5}$$

$$\lambda_k (\alpha - \mathbf{w}^* \mathbf{h}_k^* \mathbf{h}_k \mathbf{w}) = 0 \tag{6}$$

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

Based on the optimality conditions in (5) and (6), 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 $\mathbf{h}_k^* s$.

This can be inferred from (5). The reason is that (5) can be written as

$$\mathbf{w} = \sum_{k=1}^{K} \lambda_k \mathbf{h}_k^* \mathbf{h}_k \mathbf{w} = \sum_{k=1}^{K} \lambda_k \mathbf{h}_k^* a_k = \sum_{k=1}^{K} \beta_k \mathbf{h}_k^* \tag{7}$$

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 (6). Suppose that M out of K values of λ_k are nonzero 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 $i = 1, \ldots, k-1$. We can rewrite (7) as

$$\mathbf{w} = \sum_{k=1}^{K} \beta_{\pi,k} \mathbf{h}_{\pi,\pi(k)}^*. \tag{8}$$

We 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 (8) and orthogonal construction of $\mathbf{h}_{\pi,\pi(k)}^*$, we have the following for $j = 1, \ldots, M$:

$$\begin{aligned} \alpha &= \left(\sum_{k=1}^{K} \beta_{\pi,k}^* \mathbf{h}_{\pi,\pi(k)} \right) \mathbf{h}_{\pi(j)}^* \mathbf{h}_{\pi(j)} \left(\sum_{l=1}^{K} h_{\pi,\pi(l)}^* \beta_{\pi,l} \right) \\ &= \left(\sum_{k=1}^{j} \beta_{\pi,k}^* \mathbf{h}_{\pi,\pi(k)} \right) \mathbf{h}_{\pi(j)}^* \mathbf{h}_{\pi(j)} \left(\sum_{l=1}^{j} h_{\pi,\pi(l)}^* \beta_{\pi,l} \right) \end{aligned} \tag{9}$$

The expression (9) has the following interpretation, which can be used to build a greedy solution. When $j = 1$, we have

$$\alpha = \beta^*_{\pi,1}\mathbf{h}_{\pi,\pi(1)}\mathbf{h}^*_{\pi(1)}\mathbf{h}_{\pi(1)}\mathbf{h}^*_{\pi,\pi(1)}\beta_{\pi,1}. \quad (10)$$

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

$$\begin{aligned}\alpha = {} & \beta^*_{\pi,1}\mathbf{h}_{\pi,\pi(1)}\mathbf{h}^*_{\pi(2)}\mathbf{h}_{\pi(2)}\mathbf{h}^*_{\pi,\pi(1)}\beta_{\pi,1} \\ & + \beta^*_{\pi,2}\mathbf{h}_{\pi,\pi(2)}\mathbf{h}^*_{\pi(2)}\mathbf{h}_{\pi(2)}\mathbf{h}^*_{\pi,\pi(2)}\beta_{\pi,2} \\ & + 2\mathfrak{Re}\left\{\beta^*_{\pi,1}\mathbf{h}_{\pi,\pi(1)}\mathbf{h}^*_{\pi(2)}\mathbf{h}_{\pi(2)}\mathbf{h}^*_{\pi,\pi(2)}\beta_{\pi,2}\right\}. \quad (11)\end{aligned}$$

Now, $\beta_{\pi,2}$ can be found to satisfy this condition given that $\beta_{\pi,1}$ from the previous step is used. We 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 we find the values for the next $\beta_{\pi,k}$. However, at each step, the value for $\beta_{\pi,k}$ that satisfies condition (9) 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. With the aim of minimizing the norm of $\mathbf{w}$, at each step we find the value of $\beta_{\pi,k}$ that minimizes the partial norm of $\mathbf{w}$, which in turn is defined as

$$\|\mathbf{w}\|^2_{1\to k} = \left\|\sum_{j=1}^{k}\beta_{\pi,j}\mathbf{h}^*_{\pi,\pi(j)}\right\|^2_2. \quad (12)$$

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

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

TABLE I
WARPLAB PHYSICAL-LAYER PARAMETERS

Frequency	2.4 GHz - Channel 14
Bandwidth	625 KHz
Payload size	100 bits
Modulation	BPSK, QPSK 16-QAM, 64 QAM
Convolutional Codes	Generator polynomials: $g0(133)$ $g1(171)$, Rate = 1/2, 2/3, 3/4
Base data rate	156 Kbps at 1/2 rate

Algorithm 2 summarizes our proposed algorithm for multicast beamformer design. The key steps of our 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 (lines 9–11).

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 $\mathbf{w}$ (lines 12–22).

Step 3: Steps 1 and 2 are repeated for every permutation π to obtain the corresponding beamforming vector $\mathbf{w}_\pi$. The final beamforming vector is obtained as the one that has the minimum norm over all of the permutations (line 25).

The key advantage of the proposed algorithm is that there is no need for an iterative approach as in prior works [17]; such iterative approaches require fine adjustments to the solution parameters to obtain fast convergence and avoid divergence and are not amenable to practical implementations.

Time Complexity Analysis: For a given permutation of users, Algorithm 2 takes $O(N \times K^2)$ time to compute a beamformer (lines 9–23). Here, N is the number of antennas and is equal to the size of the channel vectors. Therefore, by considering all possible permutations, the total complexity of Algorithm 2 is $O(N \times K^2 \times K!)$. This indicates 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, we can consider a small number of randomly selected permutations such that the overall algorithm is computationally tractable.

VI. EXPERIMENTAL SETUP

In this section, we describe the implementation of ADAM as well as switched beamforming solutions for multicasting.

A. Hardware and Software

Our implementation is based on the WARPLab framework [25]. 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 I specifies the PHY parameters used in our evaluation. Our APs use four radio boards that 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).

Our implementation uses a channel bandwidth of 652 kHz. This channel bandwidth is smaller than the 20-MHz channel bandwidth used in 802.11a/b/g. We emphasize that similar experimental results would be obtained with a higher channel

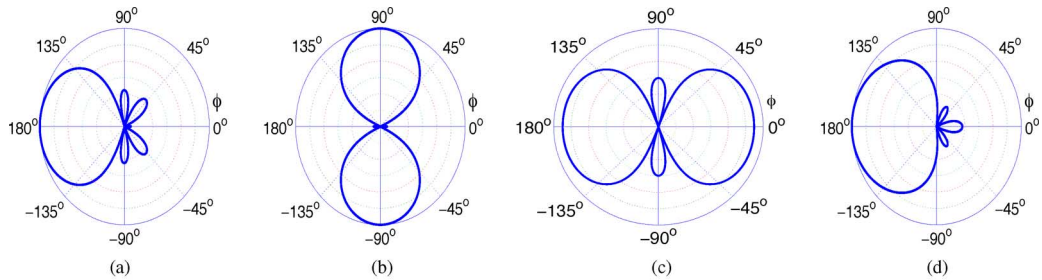

Fig. 3. Linear array antenna structure allows for construction of three orthogonal beams. Two of such beam patterns are depicted in (a) linear beam pattern 1 and (b) linear beam pattern 2. The third beam is similar to (a), however, it points toward 0°. A composite beam constructed from linear beams pointing toward 0° and 180° is shown in (c). Circular array structure allows for construction of four orthogonal beams pointing toward 0°, 90°, 180°, and −90°. The beam pointing toward 180° is shown in (d) circular beam pattern 1.

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., 802.11a/g) where the channel is divided into many subcarriers, per-subcarrier (or a group of subcarrier) channel information provides accurate channel information.

B. Multicasting Framework

We implemented three multicast mechanisms on our 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: We have considered Linear and Circular arrays for switched beamforming. In a linear array with antenna separation distance of $\frac{15*\lambda}{100}$, three orthogonal beams can be created [26]. Fig. 3(a), and (b) depicts 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. 3(a), which however will point toward the 0° direction. Fig. 3(c) depicts a composite beam that is composed of the two linear beams pointing toward the 0° and 180° directions. In circular arrays, antenna elements are placed in a circle with equal distance between each two neighbor antennas. Fig. 3(d) depicts one of the resulting beam patterns for a separation distance of $\frac{\lambda}{2}$ [26]. With appropriate shifting of the phase across the antennas, the beam pattern of Fig. 3(d) can be rotated to point toward the $-90°$, 0°, and 90° directions, thus, providing four orthogonal beams.

We have implemented switched multicast beamforming according to [7], whose solutions search over beam patterns that are a superset of those considered in [8] and show considerable gains compared to [8]. 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 SNR value. 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.

ADAM: We have implemented the components of ADAM based on our discussion in Section V.

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 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 our implemntation is approximately 50 ms.

Modulation and Coding Scheme (MCS) Selection: All of the studied protocols in this paper select an 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 400 WB channel emulator [27] 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 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 profile is composed of 14 Rayleigh fading channels with multipath RMS delay spread of 30 ns and maximum delay of 200 ns. Fig. 4(a) shows the 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 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.

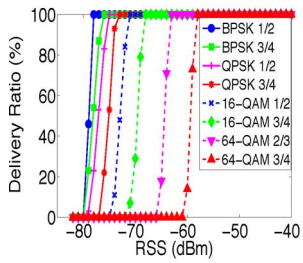

(a)

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

(b)

Fig. 4. (a) WARP board rate table and (b) 802.11 rate table as a function of SNR.

D. Performance Metrics

All of our indoor experiments are conducted during the 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 US. As observed in Fig. 2, 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 our indoor over-the-air experiments is an average of 50 samples. Due to the coherent channel conditions, the observed variation across each data point is less than 5% of the averaged value. Hence, for ease of presentation we only plot the average values. 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 for comparison of different schemes studied in this paper. 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 the multicast throughput as the average of PDRs and throughputs over all of the clients.

VII. Gains of Adaptive Beamforming

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

Scenario: Fig. 5(a) depicts our experimental setup in which we deployed six nodes in an office environment. Nodes 1 and 2 each have four antennas and can thus be used as transmitters or single-antenna receivers. We first consider node one as our transmitter, and among 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 the multicasting systems considered in this paper.

A. Impact of Discrete Rates

Performance Gains: Fig. 5(b) shows the schedule length of ADAM when the rate is selected according to the WARP SNR-rate relation of Fig. 4(a). 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. 5(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.

Suboptimality of Partitioning: Fig. 5(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. 5(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 beamformer selection components (Algorithms 1 and 2) depend only on the client channel vectors and are not affected by the specific SNR-rate mapping of the hardware. However, the joint partitioning and beamfomer selection algorithms (JPB-A and JPB-S) select the partition that results in the minimum schedule length by taking into account the specific SNR-rate mapping of the implementation. Hence, 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. 4(b). Fig. 5(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 the 802.11 rate table of Fig. 4(b) and calculate the resulting schedule length for each of the schemes.

Fig. 5(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 6–54 Mb/s. It also supports basic rates of 1 and 2 Mb/s 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.

B. Impact of the Number of Antennas

Fig. 5(d) shows the measured schedule length of ADAM as a function of the number of antennas and the number of clients across all the topologies. Fig. 5(d) shows that with two antennas at the transmitter and with five clients, ADAM slightly decreases the schedule length compared to an omnidirectional

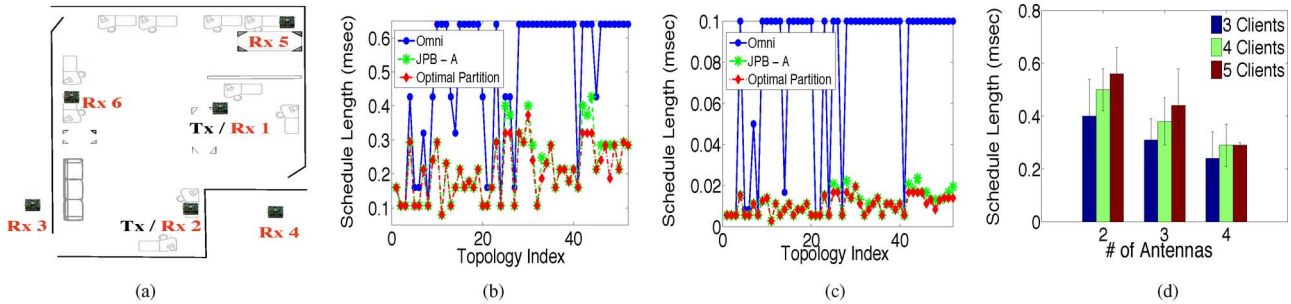

Fig. 5. ADAM's performance evaluation in an indoor environment. (a) Map of the environment. (b) Schedule length with WARP. ADAM provides an average gain of 2.8 compared to omni with WARP board's rate table. Furthermore, ADAM's greedy user partitioning achieves a performance close to the optimal partitioning. (c) Gains with 802.11 rates. When the dynamic range of the rate table is high, ADAM can provide even higher gains. (d) Gains with varying number of antennas. ADAM's gains as a function of the number of antennas.

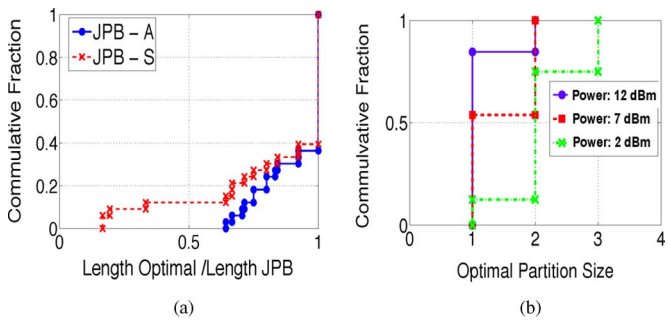

Fig. 6. (a) Algorithm comparison. JPB-S can be trapped in a local minimum, whereas JPB-A considers all possible user partitions and thus has a better performance. (b) Optimal partition size for different transmit powers.

transmission as depicted in Fig. 5(b). This is because as the size of the clients with respect to the number of antennas increases, the randomness of the channel vectors of different clients coupled with the high number of users makes the beamformer vector to tend to that of an omnidirectional transmission. On the other hand, Fig. 5(d) shows that increasing the number of antennas for a fixed number of clients can significantly reduce the schedule length.

C. Algorithm Evaluation

We now evaluate the algorithmic components of ADAM. We 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 discussed in Section V-C).

Performance Versus Complexity: Fig. 6(a) depicts the CDF of the ratio between the schedule length of the optimal user partitioning to that of the proposed partitioning algorithms. We 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.

In all our experiments we observed that the user partitioning component of algorithms JPB-A and JPB-S partitions the users into any given number of groups in less than 20 iterations. However, as discussed in Section V-C, selecting an appropriate beamformer for a given group of users is proportional to the number of user permutations and has a factorial time complexity. Thus, in order to reduce the time complexity with a large number of users, one can consider a small number of randomly selected permutations such that the overall algorithm is computationally tractable.

Optimal Partition Size: Fig. 6(b) shows the CDF of the optimal partition sizes for three different transmission powers. For high transmission powers (Tx Power $=$ 12 dBm), up to 85% of topologies do not require partitioning. As we reduce the Tx power, the need for partitioning increases. Fig. 6(b) shows that with 10 dB reduction in transmission power (Tx Power $=$ 2 dBm), 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 Tx power, it may not be feasible to serve all of the clients in the same group. Second, with low Tx 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, our experimental results show that the typical number of optimal partitions is low. Thus, in order to reduce the computational complexity, we can limit the number of partitions to a small constant, independent of K.

D. Adaptive Versus 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. 5(a). For each topology, we first perform adaptive beamforming. Next, without changing the antenna array, we perform switched multicast beamforming by using the predetermined 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. 5(b).

Relative Gains: We now compare the schedule length of switched beamforming to that of adaptive beamforming. Fig. 7(a) shows that ADAM provides an average gain of 1.8 and 2.1 over switched beamforming with circular and linear

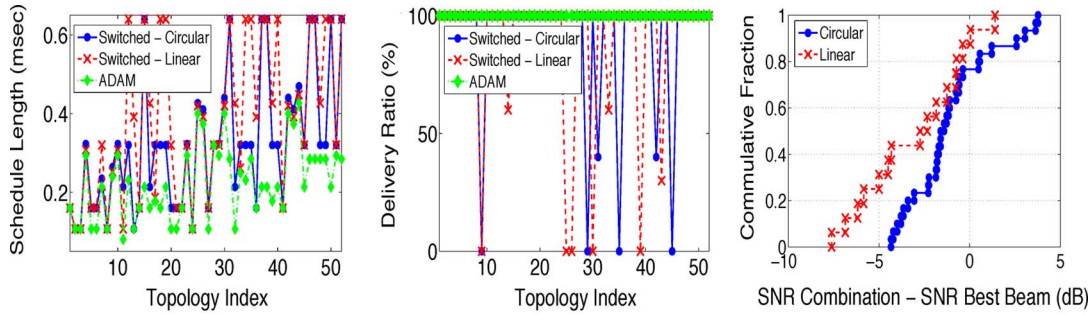

Fig. 7. Comparison between ADAM and switched beam in indoor environments. (a) Predicted schedule length of switched-beam. Unlike switched beam, ADAM benefits from the multipath and thus provides significant gains in terms of schedule length. (b) PDR. Furthermore, composite beam SNR prediction inaccuracy that is used in switched beamforming, results in low PDR in many scenarios. (c) Impact of beam combining. The composite beam's real SNR can be significantly lower/higher from the predicted SNR in indoor environments.

arrays, respectively. Furthermore, ADAM consistently outperforms switched beamforming in every topology. This can be attributed to the fact that switched beam uses only a finite set of predetermined beams, which might even have a lower gain compared to an omni transmission in the presence of multipath. Indeed, by comparing Figs. 5(b) and 7(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. 7(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 an 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. 5(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. 7(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.

Overhead Comparison: In switched beamforming, the index of the best beam and the resulting SNR is fed back by each client. This results in 2-bit overhead for beam selection and 6-bit overhead for SNR (out of 64 levels), resulting in 8-bit total overhead.

In our current implementation of ADAM, 8 bits are fed back for each antenna, resulting in a total of 32 bits overhead per client. Note that our implementation does not use any codebook for channel estimation, which can be used for significant overhead reduction. Recent implementations of adaptive beamforming [28] have shown that a codebook size of 64 (and hence 6 bits) provides similar performance to infinite codebook for a four-antenna transmitter. Note that for both schemes, the overall impact of feedback is small compared to a multicast packet size. Also, as a channel estimate can be used for multiple packet transmissions, the impact of overhead can be further reduced.

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.

VIII. 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 timescale 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 channel 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 VI. 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 beamfoming in changing channel conditions as a function of feedback rate. Specifically,

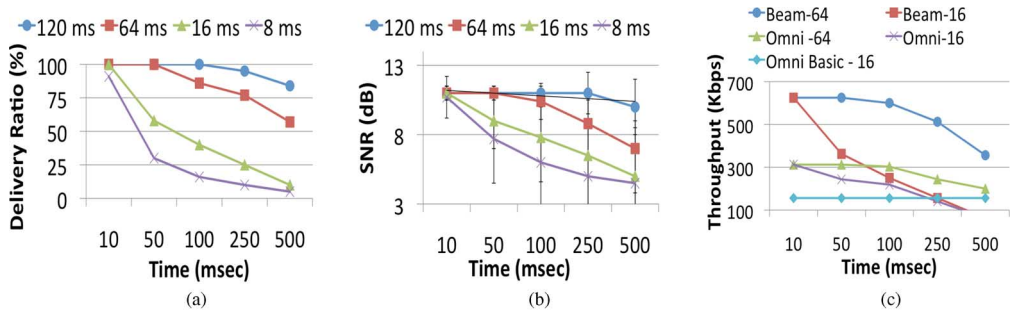

Fig. 8. Impact of coherence time and feedback rate on ADAM's performance. As the rate of channel information feedback decreases, (a) multicast PDR, (b) average SNR, and (c) throughput decrease.

we vary the timescale of channel information feedback (t_{f}) that is available at the transmitter. Once the transmitter obtains the channel information, it jointly beamforms toward 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. 8(a) shows the average PDR as a function of channel feedback timescale for different coherence times. We observe that the PDR of multicast beamforming drops as the timescale of channel feedback increases for a given coherence time, or as the coherence time decreases for a fixed feedback timescale. 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 timescale 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. 8(b) as a function of channel feedback timescale. 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 timescale 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 $s = \frac{t_{\mathrm{f}}}{t_{\mathrm{c}}}$) 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 $s = \frac{t_{\mathrm{f}}}{t_{\mathrm{c}}}$

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. 8(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 as well as omni transmission with base rate. However, at extremely reduced feedback rate ($t_{\mathrm{f}} = 500$ ms) and small coherence time ($t_{\mathrm{c}} = 16$ 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.

B. Reduced Feedback and Mobility

In any multicast system, the required PDR is dependent on the application. As seen in Fig. 8(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. 8(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 $s = \frac{t_{\mathrm{f}}}{t_{\mathrm{c}}}$, 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 $s = \frac{t_{\mathrm{f}}}{t_{\mathrm{c}}}$, 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

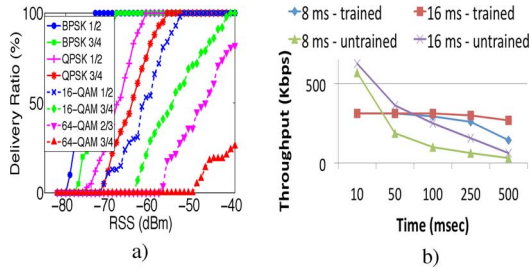

Fig. 9. (a) WARP SNR-rate for $s = \frac{50}{8}$. Training WARP boards according to $s = \frac{50}{8}$, and (b) the resulting impact on throughput.

duration of $t_{]rmf}$, measures the PDR, and repeats this experiment for a thousand trials. The emulator uses the same configuration parameters of Section VI. However, instead of using a static channel ($t_c = \infty$), its t_c value is based on the s parameter.

Fig. 9(a) 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(a) to Fig. 9(a), 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. 8. However, we obtain our rate table according to Fig. 9(a) for $s = \frac{50}{8}$. Fig. 9(b) 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 timescale 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 timescales smaller than t_f, and higher for the timescales larger than t_f. Note that apart from throughput, PDR is another metric that should be considered in selecting between a trained versus 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) 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.

IX. Discussion and Future Work

ADAM's protocol design is similar to the IEEE 802.11ac standard [1], in which a base station broadcasts a channel probing message for clients to estimate and feedback the channel information. As an alternative approach for channel estimation, a base station can obtain channel information based on the preexisting uplink traffic (due to TDD channel reciprocity). This can potentially help with reducing overhead, provided that timely estimates of the channel information are available.

ADAM's design requires channel information feedback from the clients in the multicast group. If some of these clients do not support channel estimation and feedback capability, ADAM similar to 802.11ac [1] should use a default fixed rate for multicasting.

Our proposed solution for handling infrequent channel feedback (or high user mobility) assumes knowledge of the coherence time of the corresponding clients. A client can estimate the coherence time based on transmission of known waveform signals (e.g., pilots) by the base station and analysis of the variation of the received signal samples over time [29], [30]. This information can then be reported to the base station along with the channel information. This approach improves the robustness of the beamforming solution by requiring additional feedback bits to denote the coherence time. Joint design of channel information and coherence time estimation mechanisms or other solutions to add robustness when employing adaptive beamforming is a topic of our ongoing work.

X. Conclusion

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

References

[1] *IEEE Draft Standard for IT—Telecommunications and Information Exchange Between Systems—LAN/MAN—Specific Requirements—Part 11: Wireless LAN Medium Access Control and Physical Layer Specifications—Amd 4: Enhancements for Very High Throughput for Operation in Bands Below 6 GHz*, IEEE P802.11ac/D3.0, Jun. 2012, pp. 1–385.

[2] D. Gesbert, F. Tosato, C. V. Rensburg, and F. Kaltenberger, *UMTS Long Term Evolution: From Theory to Practice*. Hoboken, NJ: Wiley, 2009.

[3] J. Andrews, A. Ghosh, and R. Muhamed, *Fundamentals of WiMAX: Understanding Broadband Wireless Networking*. Upper Saddle River, NJ: Prentice-Hall, 2007.

[4] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannak, S. Seshan, and P. Steenkiste, "DIRC: Increasing indoor wireless capacity using directional antennas," in *Proc. ACM SIGCOMM*, Aug. 2009, pp. 171–182.

[5] E. Aryafar, N. Anand, T. Salonidis, and E. Knightly, "Design and experimental evaluation of multi-user beamforming in wireless LANs," in *Proc. ACM MobiCom*, Sep. 2010, pp. 197–208.

[6] S. Lakshmanan, K. Sundaresan, R. Kokku, A. Khojestepour, and S. Rangarajan, "Towards adaptive beamforming in indoor wireless networks: An experimental approach," in *Proc. IEEE INFOCOM MiniConf.*, Apr. 2009, pp. 2621–2625.

[7] K. Sundaresan, K. Ramachandran, and S. Rangarajan, "Optimal beam scheduling for multicasting in wireless networks," in *Proc. ACM MobiCom*, Sep. 2009, pp. 205–216.

[8] S. Sen, J. Xiong, R. Ghosh, and R. R. Choudhury, "Link layer multicasting with smart antennas: No client left behind," in *Proc. IEEE ICNP*, Nov. 2008, pp. 53–62.

[9] H. Zhang, Y. Jiang, S. Rangarajan, and B. Zhao, "Wireless data multicasting with switched beamforming antennas," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 526–530.

[10] A. Prabhu, H. Lundgren, and T. Salonidis, "Experimental characterization of sectorized antennas in dense 802.11 wireless mesh networks," in *Proc. ACM MobiHoc*, May 2009, pp. 259–268.
[11] S. Jain and S. R. Das, "Mac layer multicast in wireless multihop networks," in *Proc. IEEE COMSWARE*, Jan. 2006, pp. 1–10.
[12] P. Chaporkar, A. Bhat, and S. Sarkar, "An adaptive strategy for maximizing throughput in MAC layer wireless multicast," in *Proc. ACM MobiHoc*, May 2004, pp. 256–267.
[13] R. Chandra, S. Karanth, T. Moscibroda, V. Navda, J. Padhye, R. Ramjee, and L. Ravindranath, "DirCast: A practical and efficient Wi-Fi multicast system," in *Proc. IEEE ICNP*, Oct. 2009, pp. 161–170.
[14] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan, "On the effectiveness of switched beam antennas in indoor environments," in *Proc. ACM PAM*, Apr. 2008, pp. 122–131.
[15] V. Navda, A. P. Subramanian, K. Dhansekaran, A. Timm-Giel, and S. Das, "MobiSteer: Using steerable beam directional antenna for vehicular network access," in *Proc. ACM MobiSys*, June 2007, pp. 192–205.
[16] A. Subramanian, P. Deshpande, J. Gao, and S. Das, "Drive-by localization of roadside WiFi networks," in *Proc. IEEE INFOCOM*, Apr. 2008, pp. 718–725.
[17] A. Lozano, "Long-term transmit beamforming for wireless multicasting," in *Proc. IEEE ICASSP*, Apr. 2007, vol. 3, pp. III-417–III-420.
[18] N. Sidiropoulos, T. Davidson, and Z. Luo, "Transmit beamforming for physical-layer multicasting," *IEEE Trans. Signal Process.*, vol. 54, no. 6, pp. 2239–2251, Jun. 2006.
[19] E. Matskani, N. Sidiropoulos, and L. Tassiulas, "On multicast beamforming and admission control for UMTS-LTE," in *Proc. IEEE ICASSP*, Mar. 2008, pp. 2361–2364.
[20] R. Ramanathan, "On the performance of ad hoc networks with beamforming antennas," in *Proc. ACM MobiHoc*, Oct. 2001, pp. 95–105.
[21] 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 *Proc. ACM MobiCom*, Sep. 2009, pp. 49–60.
[22] S. Gollakota, S. D. Perli, and D. Katabi, "Interference alignment and cancellation," in *Proc. ACM SIGCOMM*, Aug. 2009, pp. 159–170.
[23] S. Barghi, H. Jafarkhani, and H. Yousefi'zadeh, "MIMO-assisted MPRaware MAC design for asynchronous WLANs," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 1652–1665, Dec. 2011.
[24] J. H. Conway, R. H. Hardin, and N. J. Sloane, "Packing lines, planes, etc.: Packings in grassmannian spaces," *Exp. Math.*, vol. 5, no. 2, pp. 139–159, 1996.
[25] Rice University, Houston, TX, "Rice University WARP project," [Online]. Available: http://warp.rice.edu/
[26] C. A. Balanis, *Antenna Theory: Analysis and Design*. Hoboken, NJ: Wiley, 2005.
[27] Azimuth Systems, Acton, MA, "Azimuth Systems," [Online]. Available: http://www.azimuthsystems.com/
[28] M. Duarte, A. Sabharwal, C. Dick, and R. Rao, "Beamfoming in MISO systems: Empirical results and EVM-based analysis," *IEEE Trans. Wireless Commun.*, vol. 9, no. 10, Oct. 2010.
[29] T. S. Rappaport, *Wireless Communications: Principles and Practice*. Upper Saddle River, NJ: Prentice-Hall, 1996.
[30] J. Camp and E. Knightly, "Modulation rate adaptation in urban and vehicular environments: Cross-layer implementation and experimental evaluation," *IEEE/ACM Trans. Netw.*, vol. 18, no. 6, pp. 1949–1962, Dec. 2010.

Ehsan Aryafar (S'05–M'11) received the B.S. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 2005, and the M.S. and Ph.D. degrees in electrical and computer engineering from Rice University, Houston, TX, in 2007 and 2011, respectively.

He is a Post-Doctoral Research Associate with Princeton University, Princeton, NJ. His research interests are in the areas of wireless networks, high-performance MAC protocol design, network deployment and resource provisioning, and network measurements.

Mohammad Ali (Amir) Khojastepour (S'02–M'05) received the B.Sc. and M.Sc. degrees from Shiraz University, Shiraz, Iran, in 1993 and 1996, respectively, and the Ph.D. degree from Rice University, Houston, TX, in 2004, all in electrical and computer engineering.

Since 2004, he has been a Member of Technical Staff with the Mobile Communications and Networking Research Department, NEC Laboratories America, Princeton, NJ. His research interests are in the areas of information theory and coding, communication theory and signal processing with emphasis on multiuser communications, and wireless networks.

Karthik Sundaresan (M'05–SM'12) received the Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanata, in 2006.

He is a Research Staff Member with the Mobile Communications and Networking Research Department, NEC Laboratories America, Princeton, NJ. His research interests are in the areas of wireless networks and mobile computing and span both algorithm design as well as system prototyping.

Dr. Sundaresan has been the recipient of several best paper awards at ACM MobiHoc 2008, IEEE ICNP 2005, and IEEE SECON 2005.

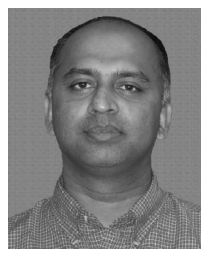

Sampath Rangarajan (M'91–SM'05) received the Ph.D. degree in computer science from the University of Texas at Austin in 1990.

He heads the Mobile Communications and Networking Research Department, NEC Laboratories America, Princeton, NJ. Previously, he was with the Networking Research Center, Bell Laboratories, Holmdel, NJ. Prior to that, he was co-founder and Vice President of Technology with Ranch Networks, Morganville, NJ, a venture-funded startup in the IP networking space. Earlier, he was a Researcher with the Systems and Software Research Center, Bell Laboratories, Murray Hill, NJ. Before joining Bell Laboratories, he was an Assistant Professor with the Electrical and Computer Engineering Department, Northeastern University, Boston, MA. His research interests span the areas of mobile communications, mobile networks, and distributed systems.

Dr. Rangarajan has been on the editorial boards of IEEE TRANSACTIONS ON COMPUTERS, IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, and *Mobile Computing and Communications Review*.

Edward Knightly (S'91–M'96–SM'04–F'09) received the B.S. degree from Auburn University, Auburn, AL, in 1991, and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 1992 and 1996, respectively, all in electrical engineering.

He is a Professor of electrical and computer engineering with Rice University, Houston, TX. His research interests are in the areas of mobile and wireless networks and high-performance and denial-of-service resilient protocol design.

Prof. Knightly is a Sloan Fellow. He served as Associate Editor of numerous journals and special issues including the IEEE/ACM TRANSACTIONS ON NETWORKING and the IEEE JOURNAL ON SELECTED AREAS OF COMMUNICATIONS Special Issue on Multi-Hop Wireless Mesh Networks. He served as Technical Co-Chair of IEEE INFOCOM 2005 and General Chair of ACM MobiHoc 2009 and ACM MobiSys 2007, and served on the program committee for numerous networking conferences including ICNP, INFOCOM, MobiCom, and SIGMETRICS. He is a recipient of National Science Foundation CAREER Award. He received the Best Paper Award from ACM MobiCom 2008.