Scalable Multicast in Highly-Directional 60 GHz WLANs

Sharan Naribole and Edward Knightly

Department of Electrical and Computer Engineering, Rice University, Houston, TX

Email: {nsharan, knightly}@rice.edu

Abstract—**60 GHz bands target multi-gigabit rate applications such as high definition video streaming. Unfortunately, to provide *multicast* service, the strong directionality required at 60 GHz precludes serving all clients in a multicast group with a single transmission. Instead, a multicast transmission is comprised of a sequence of beam-formed transmissions (a beam group) that together cover all multicast group members. In this paper, we design, implement, and experimentally evaluate *Scalable Directional Multicast* (SDM) as a technique to *(i)* train the access point with per-beam per-client RSSI measurements via partially traversing a codebook tree. The training balances the objectives of limiting overhead with collecting sufficient data to form efficient beam groups. *(ii)* Using the available training information, we design a scalable beam grouping algorithm that approximates the minimum multicast group data transmission time. We implement the key components of SDM and evaluate with a combination of over-the-air experiments and trace-driven simulations. Our results show that the gains provided by SDM increase with group size and provide near-optimal group selection with significantly reduced training time, yielding up to 1.8x throughput gains over exhaustive-search training and grouping.**

I. INTRODUCTION

Unlicensed access in the 7-14 GHz wide band available at 60 GHz has the potential to enable high-rate multimedia applications via directional transmission and reception [1], [2]. A *multicast* service provides multiple clients (a multicast group) with the same data from the Access Point (AP). Because using the widest possible beam at 60 GHz severely limits the data rate and range, the AP needs to partition the multicast group into multiple subsets and select an appropriate beam and data rate to serve each subset. Moreover, current 60 GHz systems employ a single RF chain per antenna array (unlike 2.4/5 GHz MIMO) such that the AP acts as a switched-beam system and generates a single beam at a time [3], [2].

In this paper, our objective is to maximize the throughput delivered to multicast groups incorporating the overhead in beam training and the subsequent selection of the *beam group*, or group of beams covering all of the group's clients for data transmission. Each beam is defined via a multi-level codebook in which the codebook level corresponds to beamwidth and the codes within a level span different directions [4], [5]. In particular, we propose **S**calable **D**irectional **M**ulticast (SDM), the first 60 GHz multicast protocol to incorporate overhead in training and beam grouping, and make the following contributions:

Scalable Training. Beam training enables the AP to obtain per-client per-beam RSSI measurements for the multicast group members. To ensure that beam training only occurs when necessary, SDM precedes each multicast transmission with a multicast group announcement and a short packet exchange with each client. Training is only invoked if a group member fails to respond. To limit overhead, we utilize a tree-based codebook structure that links the beams of different levels based on their spatial similarity.

In an idealized propagation environment with line of sight (LOS) to the AP for all codebook entries, one could simply find the strongest beam at each level from client feedback and use only its children for the next level training. For a general scenario, SDM's key strategy is to first perform training at the finest beam level, thus ensuring every client is reachable and has at least one beam with high directivity gain. Then, SDM performs a pruned tree traversal up the tree in wider beam levels. For the pruned set of beams to be used for each level's training, SDM selects the parents of the strongest beams of the previous level. In this way, SDM obtains sufficient, but not exhaustive, training that we will show enables near-optimal beam grouping.

Scalable Beam Grouping. Using the training information, SDM next selects the beam group. First, we formulate an optimization problem of minimizing the data sweep time, i.e., the time taken to transmit a fixed number of bits via sequential generation of the beams in the beam group using the Modulation and Coding Scheme (MCS) for each beam as determined by the beam training. We show that performing exhaustive search over all possible combinations of beams and clients incurs overhead of order $O(c^{K-1}N^{\frac{N}{2}+1})$, where N is the multicast group size, K is the number of levels and c is the average ratio of beamwidth between two neighboring levels. Second, we present SDM's beam grouping algorithm. The key strategy is to begin with an initial solution consisting of only the finest beams that provide high directivity gain. Then, when beneficial, SDM iteratively replaces the finest beams with wider beams in descending order of each wide beam's improvement ratio over the initial solution. By considering only the reachable client subset for each codebook, SDM searches over a reduced space of order $O(KN^3)$ which our experiments show closely track exhaustive search.

SDM Implementation & Experimental Evaluation. We implement the key components of SDM in software and use a mechanically steerable 60 GHz RF-frontend combined with the software-defined radio platform WARP [6] for transmissions and training. As a baseline for comparison, we consider

Fig. 1: (a) An indoor conference room scenario where multicast group 1 comprising of clients A, B and C wins the contention and the AP selects two beams for the data transmission. (b) The different stages of our multicast timeline model for the scenario considered in (a).

sequential unicast. Namely, because the IEEE 802.11ad standard [7], [8] does not define a multicast protocol, providing a multicast service could be realized via sequential unicast transmissions, i.e., generation of beams directed to individual clients of the multicast group. While such an approach can provide high signal strength at the clients, the total transmission time increases linearly with group size. To assess SDM, we collect training information in a typical indoor conference room setting for different client locations and codebook trees. Our results show that SDM consistently outperforms sequential unicast. Moreover, SDM provides over 1.8x throughput gains with up to 45% reduction in training overhead and 12x reduction in beam grouping overhead vs. exhaustive search and grouping.

The remainder of this paper is organized as follows. In Section II, we provide a design overview of SDM. In Section III, we present SDM's scalable training protocol. In Section IV, we present SDM's beam grouping algorithm. In Section V, we describe SDM's implementation and the data collection. In Section VI, we discuss the key results from our evaluation of SDM. In Section VII, we review related works and the paper is concluded in Section VIII.

II. DESIGN OVERVIEW

In this section, we present an overview of SDM's design. First, we discuss the network model considered in this paper. Second, we describe the beam group quality test conducted by SDM to test whether training is necessary. Third, we describe the training period conducted by SDM to update clients' signal strength information if the beam group quality test fails. Lastly, we describe the beam grouping for data transmission using the training information.

A. Network Model

We consider a highly directional environment in which both the AP and the clients are equipped with antenna arrays capable of generating a fixed set of transmission and reception beams of different beamwidths. This fixed set of beams is defined by a codebook in which each beam corresponds to a particular entry in the codebook including a combination of weights assigned to the antenna elements. We can adjust beamwidth and steer the beam using discrete phase shifts of the antenna weights [5], [4]. We consider that each antenna array utilizes a single RF chain and can generate only a single beam at a time.

We consider a network with unicast clients and multiple multicast groups, with each group comprising of multiple clients. For example, a multicast group could represent a particular TV content. When the clients request multicast service either during the association phase or as a separate request, we place the client in the corresponding multicast group and inform the client about its group number. Otherwise, we list it as a unicast client. In Fig.1(a), clients A, B and C are in multicast group 1, clients D and E are in multicast group 2 and clients F and G are unicast clients. Like 802.11, a multicast transmission begins with a group announcement when the AP wins contention to serve a multicast group (group 1 in Fig. 1(b)). Next, we describe the SDM's functioning in the different stages of the timeline.

B. Beam Group Quality Test

Except for the first transmission, a beam group will have previously been established for the prior transmission. If there was negligible client and environmental mobility since the last transmission, the same beam group can be used again for the current data transmission without performing any beam training nor a new beam grouping. Because the AP is oblivious to such mobility, to learn if the existing beam group can be used or not, SDM tests the existing beam group via transmitting a short data packet on each beam using its corresponding data rate as shown in Figure 1(b). In these packets, SDM includes information about the multicast group selected for data transmission and a pre-assigned order for clients to send ACKs.

If the AP receives an ACK from every client of the multicast group, then SDM uses the existing beam group. However, if

this test fails, SDM will find a new beam group. In Figure 1(b), the AP fails to receive an ACK from client C and consequently SDM invokes beam training.

C. Training Period

If the beam group quality test fails, SDM conducts training that provides it with the clients' latest signal strength information for the AP's different beams. In order to provide the AP with latest signal strength information for finding the best beam group, we consider every client in the multicast group takes part in the training even if a client successfully receives the beam group quality test packet. Alternatively, a network controller can invoke only the clients that fail the beam group quality test to participate in training. The key concept of training is that the AP transmits a beacon at the base rate (MCS 0 in [9]) using a particular beam from the codebook followed by a feedback packet from every client consisting of the received power measure of the transmitted beacon. Because different beams correspond to different levels, SDM conducts the training of each level separately. For simplicity, two-level training is illustrated by the wide beam training and fine beam training in Fig. 1(b). The training beacons include information about the multicast group selected for data transmission and the time the clients outside the multicast group should defer.

To limit feedback overhead, the AP transmits beacons with all selected beams of a particular level before receiving feedback from each client. A client's feedback includes the received power measures for the different beams. Although a beacon might be detected at the client as it is transmitted at the base rate, the power measure might be lower than the minimum required for a data transmission (MCS 1 in [9]). To minimize collisions, SDM pre-assigns the feedback order and this information is included in the training beacons. We consider the AP to be in quasi-omni reception mode during the feedback period.

D. Beam Group Selection and Data Transmission

Given the training information, SDM next finds the beam group to be used for data transmission. Each beam is defined by its codebook entry, the clients that it serves and the data rate used for transmission. This leads to a sequential generation of beams one after the other which we define as a *data sweep*. As the AP sends the same data for all beams in the beam group, a client receiving the same packet via more than one beam doesn't increase its throughput. We consider that the AP can sweep multiple times during the transmission opportunity (TXOP) period analogous to frame aggregation in unicast communication. Fig. 1(b), depicts two data sweeps during the data transmission period. As many clients might be served by the multicast data transmission, we consider the TXOP to begin after the beam grouping selection by the AP. Alternatively, a network controller might include the beam group quality test, beam training and beam grouping computation within the TXOP duration. However, that might lead to a significantly reduced airtime for the data transmission.

III. SCALABLE MULTICAST TRAINING

In this section, we firstly introduce the concept of multi-level codebook-based beamforming and the codebook tree architecture as a useful means to reduce the training overhead. Secondly, we describe the training strategy that minimizes overhead in ideal indoor environments followed by its challenges in a general setting. To address these challenges, we present SDM's training protocol.

A. Multi-level codebook-based beamforming

As discussed in Section II, we consider the AP and the clients are equipped with antenna arrays capable of generating a fixed set of beams of discrete beamwidths. In Fig. 1(b), the AP uses two levels of transmission beamwidth indicated by wide beam and fine beam. In general, we consider a multi-level codebook at the AP of K levels of beamwidth such that at each level, the beams are uniformly spread out $360°$ around the AP. For multicast, multiple beamwidth levels provide flexibility in selection of a beam used to serve multiple clients simultaneously in order to reduce the total transmission time. Beamwidth decreases with increasing codebook level with the 1st level representing the widest beams. If $\phi(k)$ represents the beamwidth in radians of the beam in the kth level codebook, the number of beams $M(k)$ at kth level is given by $\lceil \frac{2\pi}{\phi(k)} \rceil$.

Exhaustive training that would require every beam in the entire codebook for sending the training beacons has overhead $O(KN + c^K)$, where c represents the average ratio of the number of beams of two neighboring beamwidth levels. This overhead would have a significant impact on multicast throughput scalability. Next, we show how the codebook tree architecture can be used to reduce training overhead.

B. Codebook Trees for Partial Traversal

To scale group size with limited training overhead, we leverage the clients' feedback information after each codebook level training to select only a partial set of beams to be used in the next level training. We need to establish a relationship between the beams or codebook entries of different levels. As the number of codes increases with codebook level, we establish an edge between beam p of level k to the set of beams in level $k + 1$, each of which has the highest spatial correlation with p in comparison to any other beam of level k. The formation of such a graph results in a tree structure termed a *codebook tree* [4], [5]. Fig. 2(a) shows an example codebook tree construction in which beam 3 of an idealized widest beam level (Level 1) has two children in Level 2 and each of them have 2 children in the finest beam level (Level 3).

Basic Traversal. Firstly, we define a client to be *reachable* at level k if there exists at least one beam used for training in that level such that its received power measure is greater than or equal to the minimum required for data transmission (MCS 1 in [9]). We define this beam as the *primary beam* at level k for this client. A basic traversal of the codebook tree represents the network state in which every client is reachable
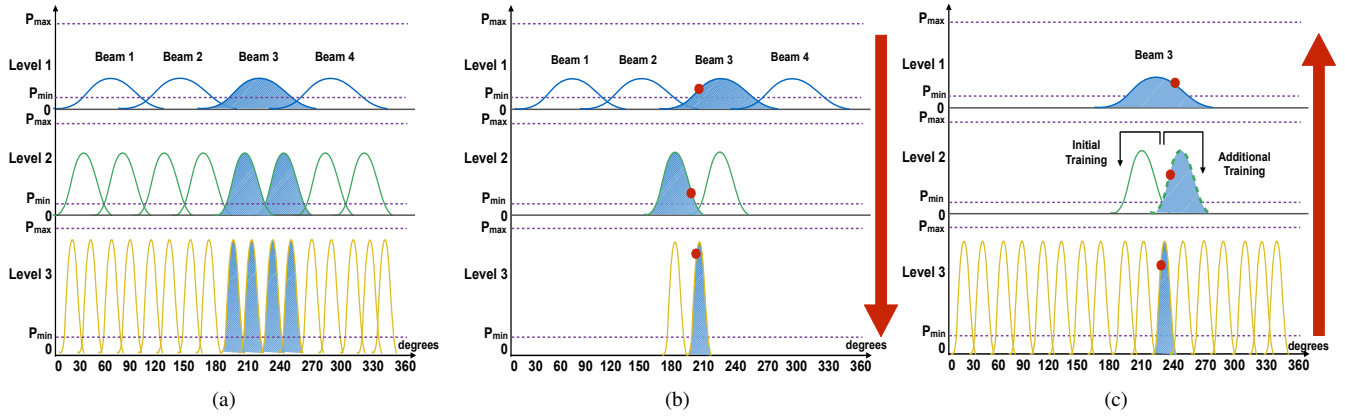
Fig. 2: (a) An example codebook tree construction (b) A basic traversal (c) SDM's traversal. In (a), (b) and (c), the x-axis indicates a pre-defined orientation angle of the AP's transmit array and the y-axis indicates the power distribution of beams, where P_{min} is the minimum required for the lowest data rate while P_{max} corresponds to the minimum for the highest data rate in the 802.11ad MCS table. In (b) and (c), the filled beams indicate the primary beams and the non-filled beams are used during the training but are not the primary beam. The red dot shows the client's power measure for the primary beam.

at all levels and the primary beam of any level is a child of the primary beam of the previous wider level.

In the basic traversal, the key strategy is to, at each level, find the union set of beams that provided the strongest beacon to the clients. Then, we use only their children in the codebook tree for the next finer level training. An example of basic traversal is illustrated for a single client in Fig. 2(b) in which training begins with the widest beams all of which are used for sending beacons. From the next level onwards, we select a partial set of beams based on the client's feedback information.

Challenges in Real Environments. There are two main challenges that exist in realistic indoor environments that make the basic traversal fall back to exhaustive training.

(i) Unreachability: A client's distance from the AP might be such that it is unreachable at a wider beamwidth level training due to the reduced directivity gain. In this case, none of the beams of this level can be used for serving data to this client. As there is no primary beam obtained for a client in a wider codebook level, the AP can't select a pruned set of beams for finer level training in order to reduce the training overhead.

(ii) Non-monotonicity: The codebook tree might be fixed for the AP's antenna array and is independent of the environment the AP is deployed. Due to presence of reflectors in the environment, if a client that was reachable at a wider level through a non line-of–sight (NLOS) path might be unreachable at a finer level due to factors such as blockage. Even in this case, the AP cannot select a pruned set of beams for the next finer level training.

In the worst case, both of the above challenges lead to the AP falling back to the exhaustive training which has significant overhead. To address the scalability of training with the presence of the above challenges in realistic indoor settings, we next present SDM's training protocol.

C. SDM's Training Protocol

The key concepts of SDM's training protocol are as follows:

(i) Descending order Traversal: Due to the high directivity gain provided by the beams of the finest beam level, every client is reachable by at least one of those beams. Otherwise, the client wouldn't be able to associate with the AP. Transmission via only finest level beams represents a *sequential unicast* beam grouping. If we performed an ascending order traversal, as in the basic traversal, then only a partial set of beams might be used in the finest level training leading to incomplete information in comparison to an exhaustive approach or a strategy in which only the finest level is trained. Thus, an only-finest-beams solution using ascending order traversal information might be worse than alternative approaches.

In contrast, with descending order traversal, as we begin with training for all finest beams, we ensure any beam grouping algorithm would generate at least the sequential unicast solution. Therefore, SDM's key strategy is to perform descending order traversal. SDM selects the parents of the primary beams found for the clients in a codebook level training as the beams to be used for the next lower level and wider beam level training.

(ii) Non-monotonicity training: If any client reachable in the previous level training is found to be unreachable in the current level training, SDM performs additional training in this level. Ideally, for each client that was reachable in the previous level, the parent of its primary beam in the previous level should be the primary beam in the current codebook level. However, if it is not, we include in the additional training the sibling of the expected ideal beam for each unreachable client if this beam wasn't already used in the initial training. Similar to the initial training, the AP sends a beacon using each selected beam for the additional training except that the feedback period has only

the unreachable clients provide the feedback.

If a client is not reachable even after this additional training, we do not consider this client in selecting the set of beams for the next level training. In the worst case, this client might not be reachable in the training of any of the remaining levels. However, as this client was reachable in the finest beam level in which all beams were used for training, there exists at least one beam that can be used for data transmission to serve each client.

An example of SDM's traversal is illustrated for a single client in Fig 2(c) in which training begins with the finest beams, all of which are used for sending beacons. In Level 2, the client is found to be unreachable during the initial training. Then, additional training is performed before proceeding to Level 1 (widest beam). At the end of training period, for each client we have a 2-dimensional training vector of all beams used in the training and their corresponding power measures. In the next section, we describe how the beam group is selected using this training information.

IV. SCALABLE MULTICAST BEAM GROUPING

Using the training information, we next select the beam group for data transmission as shown in Fig. 1(b). First, we formulate an optimization problem of minimizing the data sweep time. Secondly, we present SDM's beam grouping algorithm. Table I provides a comprehensive list of notations used in this section.

A. Problem Formulation

The training information consists of each client c's training vector that maps a beam $\psi(i,j)$ to its corresponding power measure $P(i,j,c)$. For beams not used in the training or not reachable at the client, the power measure is zero watts. As discussed in Section II, the data transmission occurs in a sweep of the selected beams with each beam transmitting the same data and the total time of a sweep is called the data sweep time. SDM's objective of beam group selection is to minimize this data sweep time.

As we send the same data from each selected beam, a client receiving the same packet from more than one beam doesn't increase its throughput. Therefore, we need to judiciously find $S(i,j)$, the set of clients to be served by a beam $\psi(i,j)$ in the final beam group so that none of the clients in this set are also assigned to another beam in the beam group. As each client is assigned to a single beam, the number of beams in the optimal beam group ranges from one beam to at most N beams where N is the number of clients. The client assignment determines the data rate $R(i,j)$ that can be used by the beam for successful reception at its serving clients. Mathematically, we select the data rate by

$$R(i,j) = \text{MCS}\left(\min_{c \in S(i,j)} P(i,j,c) \right), \qquad (1)$$

where MCS() outputs the highest data rate that can be used for transmission given the power measure. For determining which client is best served by which beam, for each $\psi(i,j)$,

we find the set of clients $C_{th}(i,j) \in \mathbb{U}$ that have a power measure greater than P_{min}, the minimum required for data transmission (MCS 1).

$\mathbb{U}$	set of multicast group clients
N	number of clients $\in \mathbb{U}$
K	number of levels
$\psi(i,j)$	ith beam in jth level
$P(i,j,c)$	power measure of client c for $\psi(i,j)$
$C_{th}(i,j)$	$\{c \mid P(i,j,c) \geq P_{min}\}$
$S(i,j)$	clients $\in \mathbb{U}$ assigned to be served by $\psi(i,j)$
$R(i,j)$	data rate selected from 802.11ad MCS table for $\psi(i,j)$ for successful reception at clients $\in S(i,j)$
$\mathbb{B}$	Beam group for data transmission
B	number of beams in a beam group
$\mathbb{I}$	initial solution beam group solution using only finest beams
$T(G)$	Data sweep time of beam group G
$\text{WIR}(G)$	WIR of a set of wide beams G when finest beams serve clients not served by beams $\in G$
C_W	clients served by wide beams in a beam group selection

TABLE I: Description of notations used in the problem formulation and algorithms description

Let $\mathbb{B} = \left\{ \big(\psi(i_1,j_1), S(i_1,j_1)\big), ..., \big(\psi(i_B,j_B), S(i_B,j_B)\big) \right\}$ be a beam group composed of B beams. We express the optimization problem as follows:

$$\min_{B,i_1,...,i_B,j_1,...,j_B,S(i_1,j_1),...,S(i_B,j_B)} \sum_{b=1}^{B} \frac{1}{R(i_b,j_b)} \qquad (2a)$$

$$\text{s.t.} \bigcup_{b=1}^{B} S(i_b,j_b) = \mathbb{U} \qquad (2b)$$

$$S(i_b,j_b) \subseteq C_{th}(i_b,j_b), \ 1 \leq b \leq B. \qquad (2c)$$

Equation 2(a) represents the cost function of the optimization proportional to the data sweep time with the search space being the beams in the codebook and corresponding client assignment to each beam. Equation 2(b) ensures each client in the multicast group is assigned to at least one beam that serves it. Equation 2(c) ensures that each client assigned to a beam received a power measure $> P_{min}$ from this beam during beam training.

B. SDM's Beam Grouping Algorithm

Here, we describe the key steps of SDM's beam grouping algorithm.

1. Initial Solution: Using the training information, we begin with an initial solution composed of only finest level beams representing a sequential unicast solution. For simplicity of explanation, we assume each client is served by a distinct finest beam although our analysis is applicable to a more general setting. Let this initial solution be denoted by $\mathbb{I} = \left\{ \big(\psi(i_1,K), c_1\big), ..., \big(\psi(i_N,K), c_N\big) \right\}$ consisting of N beams of the finest level K and $c_1, ..., c_N$ represent the clients.

We observe that $\mathbb{I}$ will not be the best solution if there exists at least one beam $\psi(i^*,j^*)$ with a client assignment $S(i^*,j^*)$ such that

$$\frac{1}{R(i^*,j^*)} < \sum_{b=1}^{N^*} \frac{1}{R(i_b,K)}, \qquad (3)$$

where for simplicity we consider $S(i^*,j^*)$ corresponds to the clients served by the first N^* beams $\in \mathbb{I}$. Equation (3) means that the transmission time using $\psi(i^*,j^*)$ for clients $\in S(i^*,j^*)$ is smaller than serving them with the finest beams. In a general scenario, Equation (3) could be satisfied for multiple such client assignments which are a subset of $C_{th}(i^*,j^*)$ for the same beam and multiple of such beams could exist.

2. Wide Beam Improvement Ratio and Hashmap: To obtain the best solution, we need to exhaustively traverse every combination of a *wide beam* (any beam not belonging to the finest codebook level) and its client assignment. Unfortunately, this exhaustive search has overhead of order $O\big(c^{K-1}N^{\frac{N}{2}+1}\big)$, where c is the mean beamwidth ratio of two neighboring levels ([10]). To overcome this infeasible overhead, SDM's key strategy is to have a unique client assignment for each wide beam. SDM utilizes only the client set $C_{th}(i,j)$ of a beam $\psi(i,j)$ as its client assignment. By selecting $C_{th}(i,j)$ for the client assignment, we are allowing this pattern to serve every client that it reached in training thereby reducing the total number of beams in the beam group. This modification along with SDM's partial training information reduces the wide beam search overhead to order $O(KN^3)$ ([10]).

We identify every beam $\psi(i^*,j^*)$ that improves upon $\mathbb{I}$ when this is the only beam added to $\mathbb{I}$ along with removal of finest beams that were serving clients $\in C_{th}(i^*,j^*)$. Let this modified beam group be denoted by $\mathbb{B}^*$. To rank all such beams in order of their improvement over $\mathbb{I}$, we define the metric *wide beam improvement ratio* (WIR) expressed mathematically as

$$\mathrm{WIR}\big(\{\psi(i^*,j^*)\}\big) = \frac{T(\mathbb{I})}{T(\mathbb{B}^*)} \qquad (4)$$

where $T(x)$ is the data sweep time of beam group x. SDM stores the information in an initially empty hashmap that takes $\mathrm{WIR}\big(\{\psi(i^*,j^*)\}\big)$ as the value and $\big(\psi(i^*,j^*)\big)$ as its key. SDM utilizes separate chaining technique [11] to store multiple values having the same key. After complete traversal of the codebook tree using the training information, SDM obtains a hashmap of wide beams that can improve the data sweep time.

If the hashmap is empty at the end of this step, then there exists no wide beam that can improve upon the initial solution of only finest beams. In that case, the sequential unicast is the best solution based on the training information provided and SDM terminates the algorithm.

3. Beam Group Selection: In this step, SDM finds the final beam group solution in an iterative manner. The initial solution is the only finest beams solution. SDM initializes an empty set C_W, that represents the clients served by wider beams. In each iteration, SDM's key strategy is to select the key from the hashmap with the largest WIR as it corresponds to the

maximum improvement possible over the initial solution. SDM adds the corresponding beam $\psi(i,j)$ to the final beam group solution and the clients $\in C_{th}(i,j)$ to C_W.

As the clients newly added to C_W need not be served by any other beam, we delete every key from the hashmap that has any client $\in C_W$ as a part of the corresponding beam's client assignment. Also, we remove the finest beams serving clients $\in C_W$ from our beam group solution. Thus, every beam added to the final beam group solution is serving a different client subset. The iterative mechanism terminates when every client of the multicast group is part of C_W or if the hashmap becomes empty due to the key deletion after each iteration. If any client is absent from C_W after this iterative procedure, then we serve such clients using the finest beams still present in the solution since the first iteration. In this manner, SDM finds the beam group for data transmission.

SDM's final beam group might be composed of a mixture of wide beams and finest beams based on the training information provided. Let SDM's beam group solution be composed of a set of Z wide beams defined by $G = \{\big(\psi(i_1,j_1)\big),...,\big(\psi(i_Z,j_Z)\big)\}$ along with finest beams serving the clients not served by the wide beams. Then, we derive ([10]) the resultant WIR of this beam group to be

$$\mathrm{WIR}(G^*) = \frac{1}{\left(\sum_{a=1}^{Z} \frac{1}{\mathrm{WIR}\big(\{\psi(i_a,j_a)\}\big)}\right) - (Z-1)} \qquad (5)$$

In this manner, SDM provides an efficient beam group based on the training information. Next, we describe our implementation of SDM and the data collection from a typical 60 GHz indoor scenario.

V. SDM IMPLEMENTATION AND DATA COLLECTION

We implement SDM and perform over-the-air data collection to evaluate its key components. In this section, we describe the implementation of the 60 GHz system followed by our methodology of data collection.

A. SDM Implementation

Our SDM implementation consists of one set of transmitter and receiver modules that are capable of communicating in the 57-64 GHz unlicensed band with up to 1.8 GHz modulation bandwidth via the VubIQ platform [12]. These modules accept and output I/Q baseband signals. For this paper, we use the transmitter module as the AP and the receiver module as the client. In order to streamline the measurement process, we integrate these modules with two WARP v1 boards according to the flow outlined in Figure 3(a).

One computer running MATLAB, WARP-Lab [13], and the VubIQ control panels controls the entire system. Using WARPLab, we generate a random set of binary data and modulate it using BPSK with a modulation bandwidth of 10 MHz (WARP v1 is capable of a transmission bandwidth of up to 20 MHz with a sampling rate of 40 MSps). WARPLab then sends the digital samples to the AP, where the WARP analog
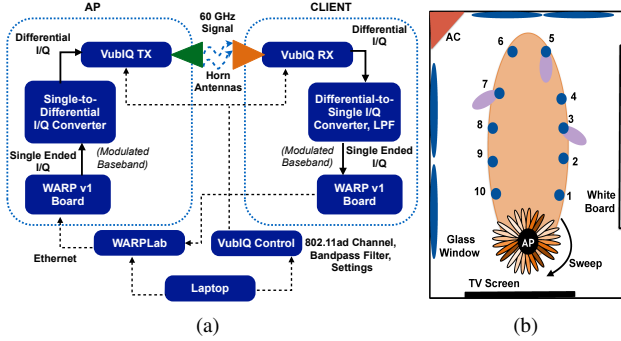
Fig. 3: (a) 60 GHz system signal and control flow. (b) Map of the room used for data collection.

daughtercard converts these samples into single-ended analog I/Q signals. These signals are passed to an evaluation board with the ADL5565 differential amplifier [14], which removes the analog daughtercard's DC offset and converts the single-ended signals into differential signals. This differential I/Q is then passed to the AP's VubIQ module where it is upconverted to 60 GHz for over-the-air transmission. The client's VubIQ module then receives this transmission and downconverts it back to analog I/Q baseband. We pass the differential signal to an off-the-shelf 15 MHz low pass filter (LT6600-15) to clean up the baseband signal. These signals are then sampled to by the client's WARP board and processed/demodulated in WARPLab.

Directional transmission and reception is achieved by using MI-WAVE's WR-15 60 GHz gain horns. To emulate the different beamwidth levels in a codebook tree, we use $7°$, $20°$, and $80°$ gain (antenna) horns. To collect received power measures at different client locations and for different receive antenna orientations, we use a mechanical motor, DC microstep driver and a commercial motion control setup [15] to steer the beams with sub-degree accuracy.

We implement an 60 GHz WLAN trace-driven emulator that is fed the over-the-air signal strength traces as inputs. Parameters and frame times are incorporated from the 802.11ad standard. We use the Single Carrier (SC)-PHY (MCS 1-12) defined in 802.11ad MCS table which is the only modulation in the first generation of chip sets.

B. Data Collection

Depending upon the client's location and the objects in the environment, the client might not be reachable from the AP for a particular codebook level. Even if the client is reachable, then its primary beam for this codebook level will vary with it's location and the reception path could either be a LOS or NLOS path. Using our 60 GHz system, we collect signal measures for a rich topology of client distributions in an indoor conference room setting illustrated in Figure 3(b). The room is composed of different reflectors including a white board, large TV screen and glass windows.

We fix the AP location at one end of the conference table. We place the client's location in 10 different positions. To emulate blockage, for each client position, we use 3 different orientations uniformly spaced in an angular range of $60°$. One of the orientations provides a LOS path to the AP from each client position whereas the other two represent client's receive beams for forced NLOS paths. We select the $20°$ horn for the client's receive antenna as it provides an efficient trade-off between receiver sensitivity of $7°$ and receive capture area of $80°$. For each client position and orientation, we perform a $360°$ sweep of the AP in steps of $5°$. To emulate the multi-level codebook structure, we conduct the AP's sweep using $7°$, $20°$, and $80°$ horns. At each point of AP's sweep, we take RMS baseband measurements to estimate the received signal strength. We normalize the signal strength measurements based on the maximum observed in the entire data set.

For the protocol evaluation, we construct over 72 5-level codebook trees using the correlation technique presented in [4] with beamwidth levels of $80°, 40°, 20°, 10°,$ and $5°$. We estimate the signal strength measurements at the clients for $40°, 10°$ and $5°$ by weighted translation of the collected measures for $80°, 20°$ and $7°$. We use an inverse relationship between signal strength and beamwidth for the translation. We convert the baseband RMS measurements to lie within the received sensitivity range provided in the 802.11ad MCS table [9] for SC-PHY modulation. To achieve this, we map the maximum value in our RMS baseband measurements to the received power of -53 dBm required for the highest data rate of 4.62 Gbps. Accordingly, we select the data rate for a given received power measure using the 802.11ad MCS table.

VI. EXPERIMENTAL EVALUATION

In this section, we evaluate the performance of SDM's training protocol and beam group algorithm with the help of our collected measurements. Moreover, we compare its performance against the following baseline models: *(i) Exhaustive approach:* This approach performs exhaustive training followed by beam grouping using an exhaustive wide beam search as discussed in Section IV; *(ii) Sequential Unicast:* This approach performs training only in the finest level followed by beam grouping consisting of only the finest level beams representing a sequential 802.11ad unicast beam generation.

Firstly, we analyze the performance in different stages of the multicast timeline (Fig. 1(b)) individually. Then, we analyze the throughput performance incorporating the overhead in training and beam group computation. For every experiment analyzed in this section, the x-axis of its corresponding figure is a given number of clients. We collect over a thousand snapshots for every x-axis point. The y-axis in each figure reports the mean and standard deviation of the metric under consideration over all the snapshots. Each snapshot is a combination of: *(i) Client Location:* a random client location selection from the 10 locations used in our collected data (Section V); *(ii) Client Orientation:* For each client, a random orientation out of the 3 receiver antenna orientations used
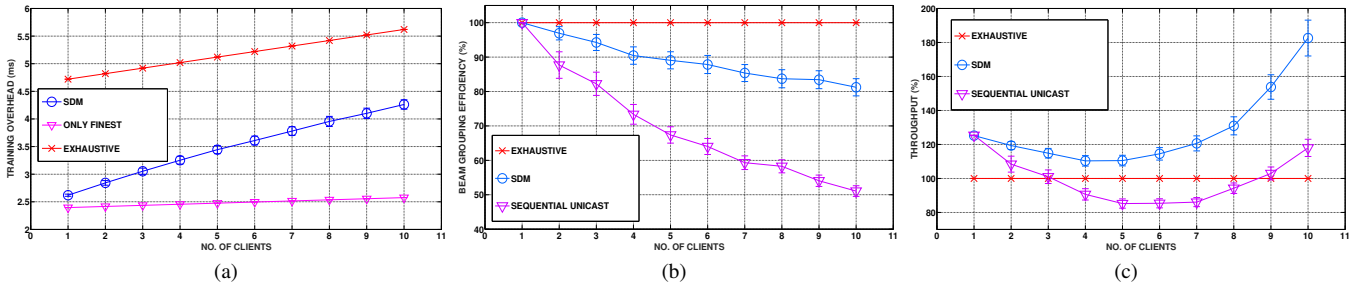
Fig. 4: Performance evaluation of SDM in different stages of the multicast timeline using real data collected with our 60 GHz system. (a) Training overhead (b) Beam Grouping Efficiency (c) Throughput.

in our collected data to emulate forced NLOS paths due to blockage; *(iii) Codebook Tree:* A random codebook tree out of the 72 codebook trees constructed by our 60 GHz WLAN trace-driven emulator.

We use the same snapshots for the evaluation of our designs in the different stages of the mutlicast timeline (Section II) including the training, beam grouping and data transmission. Figure 4 presents our results that we analyze next.

Scalable Training. For each snapshot, we conduct training independently using SDM, exhaustive training and only-finest-level training corresponding to sequential unicast. In Fig. 4(a), expectedly, exhaustive training has the highest overhead and the only-finest-level training has the lowest overhead independent of the multicast group size. Secondly, SDM has a higher slope than the other approaches because exhaustive training and only-finest-level training each have a fixed number of beacons used for training independent of the group size and only the number of feedback packets the AP receives increases with the group size. In contrast, in SDM, not only is the number of feedback packets increasing but the number of beam patterns used increases with the group size resulting in a higher slope. Thirdly, although the gain using SDM in relation to exhaustive training decreases with client size, this represents the scenario when the AP conducts training for all the clients even if a single client failed the beam group quality test. If only the clients that fail the beam group quality test take part in training then the gains would be mainly that of a small client size in Figure 4(a). *Finding: SDM consistently provides a reduced overhead with up to 44.5% reduction over exhaustive training through its feedback-controlled pruned codebook tree traversal.*

Scalable Beam Grouping. Performing an exhaustive beam and client assignment search using exhaustive training information leads to the best beam grouping solution. Therefore, we utilize exhaustive beam grouping as the baseline and compare the performance of SDM and sequential unicast. To analyze only the performance of beam grouping algorithms with appropriate training inputs, we focus on group throughput during the data transmission period of Figure 1(b) and denote this metric as the *beam grouping efficiency*. For each snapshot, we consider an 8 kB aggregated frame transmitted by each beam of the beam group during a data sweep. We consider

the data transmission period to be the maximum limit of 8.192 ms for transmit opportunity as defined in IEEE 802.11. Therefore, there may be multiple data sweeps during a single data transmission period.

In Fig. 4(b), when there is a single client, all approaches provide the same performance as all use the same finest beam and data rate to serve the client. Secondly, sequential unicast consistently has the worst performance and the difference from other algorithms increases with the group size as the number of wide beam patterns that can provide an improved solution over the sequential unicast increases. Also, the scenario of this figure represents the throughput performance if the beam group quality test (Section II) is a success such that the AP begins data transmission without performing any training and a new beam grouping. *Finding: SDM consistently provides a better performance in comparison to the sequential unicast because of its diverse beam search. The search space although limited in comparison to the exhaustive search yet has a performance within 80% of exhaustive search and grouping solution.*

Throughput. Now, we analyze the throughput performance incorporating time overhead for training and beam grouping computation. We analyze the gains provided by SDM if the time saved in training and beam grouping computation was utilized for data transmission. Once again we utilize the exhaustive approach as the baseline. Note that the sequential unicast approach has the lowest training overhead as well as the lowest beam grouping computation time.

In Fig. 4(c), first, the results indicate that when there is a single client, the exhaustive approach is even worse than sequential unicast. This is because of the larger training and beam grouping computation time although the beam group solution is the same for all as shown in Figure 4(b). Second, as the client size increases, we initially observe a performance drop for SDM and sequantial unicast. This is because of the best beam group solution provided by exhaustive approach in comparison to SDM and sequential unicast. Third, with larger group sizes, both SDM and sequential unicast performance increases. This is because the increased training time and beam group computation time of the exhaustive approach is better utilized by the other strategies for data transmission. *Finding: SDM consistently performs better than the baseline strategies*

and provides over 80% throughput gains over the exhaustive approach using its scalable codebook tree traversal during the training and beam grouping.

VII. RELATED WORK

To the best of our knowledge, this paper presents the first 60 GHz multicast protocol incorporating the overhead in training and beam grouping.

Multicast Communication. Few works have presented algorithms for optimal beam scheduling [16] and beam grouping [17]. With multiple RF chains, users can be localized in distance and angle [17] and beams can be shaped non-symmetrically [16]. In contrast, we focus on multicasting with a single-lobe pattern generation, as it requires only a single RF chain as all state-of-the-art commercial 802.11ad chipsets employ.

Unicast Beamforming. A few recent works present solutions to reducing 60 GHz beam training overhead with the objective of establishing a fine beam unicast link. The protocol in [18] optimizes the codewords used in the wider beam levels using signal strength gradient change techniques. In our work, as the training is conducted for all clients at the same time, the gradient changes in the beacon signal strengths could be highly uncorrelated across the different clients thereby preventing gradient-change based optimization. Beamforming techniques are presented in [19] to find a strong unicast link inspite of imperfect quasi-omni patterns. The wider beam training is altogether skipped in [20] by training in legacy Wi-Fi band instead. In our work, the 60 GHz channel gain information even for a wide beam is important in finding an efficient beam group.

VIII. CONCLUSION

In this paper, we addressed the challenges imposed by directional communication for a scalable multicast service at 60 GHz. We presented SDM, a novel design that includes a scalable training protocol and scalable beam grouping algorithm. Using over-the-air measurements and trace driven simulations, we showed that SDM provides the best performance in comparison to alternative strategies independent of the group size. Our future work includes extending SDM to incorporate reliability for the 60 GHz multicast transmissions.

IX. ACKNOWLEDGEMENTS

REFERENCES

[1] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, "Demystifying 60GHz outdoor picocells," in *Proc. of ACM MobiCom*, 2014.

[2] T. S. Rappaport, R. W. Heath Jr, R. C. Daniels, and J. N. Murdock, *Millimeter Wave Wireless Communications*. Pearson Education, 2014.

[3] Dell, "Dell Wireless Dw1601 Qca9005 8v256 WiGig 802.11ad Wireless Card."

[4] H.-H. Lee and Y.-C. Ko, "Low Complexity Codebook-Based Beam-forming for MIMO-OFDM Systems in Millimeter-Wave WPAN," *IEEE Transactions on Wireless Communications*, pp. 3607–3612, November 2011.

[5] S. Hur, T. Kim, D. Love, J. Krogmeier, T. Thomas, and A. Ghosh, "Multilevel millimeter wave beamforming for wireless backhaul," in *Proc. of IEEE GLOBECOM*, 2011.

[6] "Warp project." [Online]. Available: http://warpproject.org

[7] IEEE 802.11 Working Group, "IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," 2012.

[8] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and J. C. Widmer, "IEEE 802.11 ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi," *IEEE Communications Magazine*, pp. 132–141, Dec 2014.

[9] "IEEE Draft Standard for Local and Metropolitan Area Networks - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," *IEEE P802.11ad/D8.0*, June 2012.

[10] S. Naribole and E. Knightly, "Scalable Multicast in Highly-Directional 60 GHz WLANs," http://nsharan.blogs.rice.edu/files/2016/03/SDM_TR-11ssvgy.pdf, Technical Report.

[11] D. Karabeg and I. Yu, "INF2220: algorithms and data structures Series 9."

[12] Vubiq, "V60WGD03 60 GHz Waveguide Development System."

[13] N. Anand, E. Aryafar, and E. W. Knightly, "WARPlab: a flexible framework for rapid physical layer design," in *Proc. of ACM workshop on Wireless of the students, by the students, for the students*, 2010.

[14] "Analog Devices ADL5565-EVALZ."

[15] Cinetics, "Axis360 Pro modular motion control system."

[16] K. Sundaresan, K. Ramachandran, and S. Rangarajan, "Optimal Beam Scheduling for Multicasting in Wireless Networks," in *Proc. of ACM MobiCom*, 2009.

[17] H. Park, S. Park, T. Song, and S. Pack, "An Incremental Multicast Grouping Scheme for mmWave Networks with Directional Antennas," *IEEE Communications Letters*, pp. 616–619, March 2013.

[18] B. Li, Z. Zhou, W. Zou, X. Sun, and G. Du, "On the Efficient Beam-Forming Training for 60GHz Wireless Personal Area Networks," *IEEE Transactions on Wireless Communications*, pp. 504–515, February 2013.

[19] K. Hosoya, N. Prasad, K. Ramachandran, N. Orihashi, S. Kishimoto, S. Rangarajan, and K. Maruhashi, "Multiple sector id capture (midc): A novel beamforming technique for 60-ghz band multi-gbps wlan/pan systems," *IEEE Transactions on Antennas and Propagation*, pp. 81–96, Jan 2015.

[20] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement," in *Proc. of IEEE INFOCOM*, 2015.

[21] S. Sen, J. Xiong, R. Ghosh, and R. Choudhury, "Link layer multicasting with smart antennas: No client left behind," in *Proc. of IEEE ICNP*, 2008.

[22] B. Gao, Z. Xiao, C. Zhang, L. Su, D. Jin, and L. Zeng, "Double-link beam tracking against human blockage and device mobility for 60-GHz WLAN," in *Proc. of IEEE WCNC*, 2014.

[23] E. Aryafar, M. Khojastepour, K. Sundaresan, S. Rangarajan, and E. Knightly, "ADAM: An Adaptive Beamforming System for Multicasting in Wireless LANs," *IEEE/ACM Trans. on Networking*, pp. 1595–1608, Oct 2013.

[24] Y. Niu, Y. Li, D. Jin, L. Su, and D. Wu, "Blockage Robust and Efficient Scheduling for Directional mmWave WPANs," *IEEE Transactions on Vehicular Technology*, pp. 728–742, Feb 2015.

[25] S. Sur, V. Venkateswaran, and P. Zhang, X.and Ramanathan, "60 GHz Indoor Networking Through Flexible Beams: A Link-Level Profiling," in *Proc. of ACM SIGMETRICS*, 2015.

[26] K. Ramachandran, N. Prasad, K. Hosoya, K. Maruhashi, and S. Rangarajan, "Adaptive Beamforming for 60 GHz Radios: Challenges and Preliminary Solutions," in *Proc. of ACM mmCom*, 2010.

[27] J. Wang, Z. Lan, C.-S. Sum, C.-W. Pyo, J. Gao, T. Baykas, A. Rahman, R. Funada, F. Kojima, I. Lakkis, H. Harada, and S. Kato, "Beamforming Codebook Design and Performance Evaluation for 60GHz Wideband WPANs," in *Proc. of IEEE VTC*, 2009.

[28] X. An, C. Sum, R. Prasad, J. Wang, Z. Lan, J. Wang, R. Hekmat, H. Harada, and I. Niemegeers, "Beam switching support to resolve link-blockage problem in 60 GHz WPANs," in *Proc. of IEEE PIMRC*, 2009.