CoBF: Coordinated Beamforming in Dense mmWave Networks

DING ZHANG, George Mason University, USA
PANNEER SELVAM SANTHALINGAM, George Mason University, USA
PARTH PATHAK, George Mason University, USA
ZIZHAN ZHENG, Tulane University, USA

With MIMO and enhanced beamforming features, IEEE 802.11ay is poised to create the next generation of mmWave WLANs that can provide over 100 Gbps data rate. However, beamforming between densely deployed APs and clients incurs unacceptable overhead. On the other hand, the absence of up-to-date beamforming information restricts the diversity gains available through MIMO and multi-users, reducing the overall network capacity. This paper presents a novel approach of "coordinated beamforming" (called CoBF) where only a small subset of APs are selected for beamforming in the 802.11ay mmWave WLANs. Based on the concept of uncertainty, CoBF predicts the APs whose beamforming information is likely outdated and needs updating. The proposed approach complements the existing per-link beamforming solutions and extends their effectiveness from link-level to network-level. Furthermore, CoBF leverages the AP uncertainty to create MU-MIMO groups through interference-aware scheduling in 802.11ay WLANs. With extensive experimentation and simulations, we show that CoBF can significantly reduce beamforming overhead and improve network capacity for 802.11ay WLANs.

CCS Concepts: • **Networks** → **Network experimentation**; **Wireless local area networks**; • **Hardware** → **Beamforming**.

Additional Key Words and Phrases: mmWave; Beamforming; WLAN

ACM Reference Format:
Ding Zhang, Panneer Selvam Santhalingam, Parth Pathak, and Zizhan Zheng. 2023. CoBF: Coordinated Beamforming in Dense mmWave Networks. *Proc. ACM Meas. Anal. Comput. Syst.* 7, 2, Article 31 (June 2023), 26 pages. https://doi.org/10.1145/3589975

1 INTRODUCTION

Building upon the current 802.11ad standard [14, 21], the upcoming IEEE 802.11ay [9, 15, 54] refines the PHY and MAC specifications to enable the next generation of 60 GHz millimeter-wave (mmWave) WLANs. With the newly introduced support for multiple-input multiple-output (MIMO) and flexible channelization, 802.11ay can provide up to 100 Gbps of data rate. Such high data rates will make it possible to support hundreds of densely deployed devices in WLANs and their bandwidth-hungry applications such as immersive volumetric video streaming, augmented and virtual reality (AR/VR), robotic manufacturing, etc.

Authors' addresses: Ding Zhang, George Mason University, Fairfax, VA, USA, 22030, dzhang13@gmu.edu; Panneer Selvam Santhalingam, George Mason University, Fairfax, VA, USA, 22030, psanthal@gmu.edu; Parth Pathak, George Mason University, Fairfax, VA, USA, 22030, phpathak@gmu.edu; Zizhan Zheng, Tulane University, New Orleans, LA, USA, 70118, zzheng3@tulane.edu.

2476-1249/2023/6-ART31
https://doi.org/10.1145/3589975

Scaling the WLAN infrastructure to hundreds of IoT and edge computing devices will require the dense deployment of access points (APs) in future mmWave WLANs. Apart from the density necessary for capacity scaling and coverage, such dense deployment of APs can increase the robustness of mmWave WLANs, where fast AP hand-offs can be used to protect against link blockages and high attenuation [23, 38, 43]. The dense deployments can also greatly benefit dynamic AP-user associations to combat interference [17, 46, 48]. In the case of 802.11ay WLANs, densely deployed APs combined with MIMO, multi-patch antenna arrays, and flexible channelization can truly realize the spatial and frequency diversity gains [9, 15].

However, dense deployment of APs and clients in mmWave 802.11ay WLANs poses a critical challenge where the beamforming (BF) between a large number of clients and APs incurs a formidable overhead. Frequent beamforming between APs and clients is necessary to adjust the beams reacting to blockages and mobility. If up-to-date beamforming information is available, it is possible to dynamically determine the AP-client association and allocate network resources in an efficient way. But frequent BF can take a significant amount of time. For example, in 802.11ay, it takes approximately 5ms to train the downlink transmitter (Tx) and receiver (Rx) sectors of one AP to all its clients (detailed calculation in Sec. 2). With ten collocated APs, the overhead could be approximately 50ms which would consume half of the 100ms beacon interval (BI) just for the BF.

Existing solutions proposed in the literature focus on reducing the beamforming overhead per-link basis. Solutions such as [10–12, 19, 29, 34, 37–39] reduce the overhead by intelligently searching fewer Tx and Rx sectors for a single AP. However, the BF overhead of the network still increases with the number of APs [17]. Another type of solution leverages the quasi-optical properties of mmWave channels [20, 40, 53]. Here, a client's location and/or orientation is estimated based on channel state information (CSI) [23] or motion sensors [43]. Then beams calculated from one AP can be used to derive the beam of the remaining APs to the client without actual BF using techniques like triangulation [38]. These location-based methods incur additional overhead to localize mobile clients frequently. Furthermore, the reliance on quasi-optical properties makes the location-based methods less suitable for deriving Non-line-of-sight (NLoS) reflected paths without extensively localizing the ambient reflectors [44, 47].

In this paper, we present CoBF, a coordinated beamforming system that addresses the problem of high beamforming overhead in dense mmWave WLANs. Our key idea is that if we can identify only a small subset of APs that need to perform BF in each BI, we can reduce the BF overhead at the network level while ensuring that the BF information is up-to-date. To identify the subset of APs, we introduce a concept of uncertainty which is the probability that an AP's beams to its clients have changed significantly since its last BF. We develop a prediction model where APs which performed BF recently can predict the uncertainty of other APs. Then the APs with high uncertainty can be selected for BF. In addition, we use the network-level uncertainty to dynamically choose the number of APs that need to perform BF in each BI, allowing CoBF to adapt to the network dynamics seamlessly. Furthermore, the uncertainty value is leveraged for interference management, multi-user multiple input, multiple output (MU-MIMO) grouping, and scheduling.

To the best of our knowledge, our solution is the first to reduce the beamforming overhead from a network-level perspective. A salient advantage of the proposed coordinated beamforming is that it is conceptually *orthogonal to the link-level BF schemes* that are widely studied and can be used in conjunction with them to reduce the BF overhead further. Besides, our scheme is *location-agnostic* and can be readily integrated into off-the-shelf devices without any MAC or PHY protocol modifications. It makes our coordinated BF scheme applicable to different types of mmWave networks, such as 802.11ad/ay WLANs and 5G New Radio (NR) cellular networks. However, developing the coordinated BF scheme requires us to address two critical challenges:

(1) How to model uncertainty and use it to select APs for BF? We address this by developing an uncertainty prediction model. The model enables us to capture the BF information change observed (across different sectors and patches) for one AP which recently performed BF and use it to predict the potential BF information change for other APs. The prediction model can identify the probability of BF information change (i.e., whether an AP's last known beams will still hold true or not). The uncertainty prediction models are developed using a small amount of observed BF data, which typically incorporate the underlying indoor multi-path, stationary blockages, and user mobility. We propose a solution where these models for predicting uncertainty can be trained using an offline (perform warwalking and train the models offline) or online (collecting training data while predicting the APs to BF) approach. Using the predicted uncertainty, the AP selection procedure is performed in run-time to find a small subset of APs that need to update their BF information for the network efficiently.

(2) How to manage interference and schedule MU-MIMO with limited BF information? Since not all APs perform BF in each BI, we exploit the AP uncertainty to develop an SINR (Signal to Interference and Noise Ratio) model by considering both intra- and inter-group interference between APs and clients. The BF information and the uncertainty are utilized to form MU-MIMO groups with the objective of improving network throughput. Lastly, the observed network performance after scheduling is leveraged to determine the number of BF APs in the next BI, creating a closed-loop system of BF AP selection and scheduling.

In this work, we make the following contributions: (1) We present a first-of-its-kind coordinated beamforming scheme (called CoBF) that can reduce the BF overhead in dense 802.11ay WLANs by judiciously selecting a subset of APs to do BF. Through modeling the uncertainty of BF information, we present a framework that can be used for AP selection and for determining the number of APs to beamform based on network-level uncertainty. (2) We develop an uncertainty-aware scheduler that considers the uncertainty of APs' BF information to select APs and links during scheduling and resource allocation in 802.11ay WLANs. We formulate the uncertainty-based SINR model and propose a greedy scheduler that performs MU-MIMO grouping. (3) We implement and evaluate CoBF using extensive real-world experiments and large-scale trace-driven simulations. First, we use commodity 802.11ad devices enhanced with Rx beamforming implementation to evaluate the uncertainty prediction model and the effectiveness of coordinated BF. Second, we use the Remcom InSite [32] mmWave channel simulator to evaluate CoBF in two large, densely deployed WLAN scenarios. Our evaluation shows that:

- The coordinated beamforming improves the network throughput by 30.8% compared to 802.11ad BF with a 71% reduction in BF overhead in our 802.11ad testbed with 4 APs. Furthermore, the mean difference in AP-client SNR (Signal to Noise Ratio) is observed to be less than 1 dB for the sector selected by our coordinated BF scheme and the scheme where all APs perform BF.
- CoBF with MU-MIMO can achieve up to 5.5 times higher average throughput and 4.2 times lesser BF overhead with limited control latency compared to the default 802.11ay BF scheme where all APs perform beamforming in each beacon interval.
- The warwalking efforts required to build the uncertainty model are reasonably small. For example, approximately 120 warwalking steps are required in a room of 198 m^2 area in the experiment scenario. Our results also show that carefully updating the training data using an online approach is feasible, eliminating the need for the time-consuming warwalking process. The online model achieves a comparable performance in terms of uncertainty prediction.
- We also evaluate the control latency of CoBF and find that our uncertainty based beamforming and centralized scheduling incur much smaller latency overhead compared to the time taken by default exhaustive beamforming of 802.11ad/ay.

- CoBF can be combined with existing link-level beamforming schemes to achieve 7.9%, 9.3%, and 15.9% network throughput increment with CoBF + compressive sensing [34], CoBF + UbiG [38] and CoBF + MUTE [10] respectively compared to CoBF + 802.11ay link-level BF.

2 MOTIVATION AND SYSTEM OVERVIEW

2.1 802.11ay Overview

In 802.11ay, enhanced directional multi-gigabit (EDMG) stations (STAs) get access to the medium within beacon intervals (BI) as shown in Fig. 1(a). Similar to 802.11ad, a BI includes the beacon header interval (BHI) and the data transmission interval (DTI). The BHI is further divided into three sub-intervals: beacon transmission interval (BTI), where APs transmit beacons to train the downlink sectors; association beamforming training (A-BFT), where STAs perform uplink sector training in a slotted CSMA manner; announcement transmission interval (ATI) in which the AP exchanges management frames with STAs. In addition, the APs and STAs exchange data frames in DTI based on beamforming information (e.g., Tx and Rx antenna sectors from BHI) either using contention-based access periods (CBAPs) or time-divided service periods (SPs).

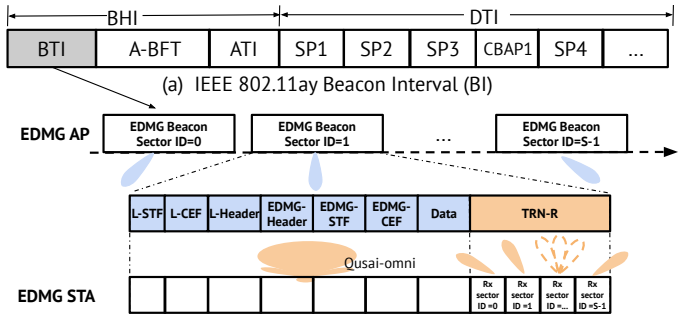

Fig. 1. 802.11ay enhanced beacon frame & beamforming

Enhanced beacon frames. 802.11ay [9, 15, 54] proposes an enhancement to the beacon frames (which we refer to as enhanced beacon frames) where the beacon frames can be used not only for downlink Tx beamforming (I-TXSS) but also downlink Rx beamforming (I-RXSS). Here, the AP transmits a beacon frame in all of its sectors one by one, where each of the beacon frames has training fields (TRN-R) appended to it (see Fig. 1(b)). An STA receives the original part of the beacon frame using the quasi-omni sector while sweeping through all its receive sectors when receiving the appended TRN-R frames. The STAs then provide feedback about observed sectors and their SNR to the AP during the following A-BFT. In CoBF, we use these enhanced beacon frames to train downlink Tx and Rx sectors from APs to clients.

MU-MIMO. Downlink MU-MIMO is another promising feature of 802.11ay, in which up to eight concurrent spatial streams could be enabled to realize up to 100 Gbps data rate. 802.11ay MU-MIMO BF process includes two phases. The SISO phase is similar to the BF training process but extended for sectors of all antennas on the AP. In the MIMO phase, multiple streams are simultaneously probed over different sectors for a subset of selected sectors and users based on SISO phase results.

2.2 Problem, Motivation and Our Approach

2.2.1 BF overhead in dense mmWave WLANs. With the enhanced beacon frames in 802.11ay WLANs, the time taken to complete the downlink BF for one AP is $T_{AP} = |S_T|(t_{beacon} + |S_R|t_{TRN-R} + t_{BIFS})$, where $|S_T|$ and $|S_R|$ are the number of sectors on AP and client, respectively. t_{beacon} is the time to transmit one beacon frame (without the following TRN-R), t_{TRN-R} is the time to transmit one TRN-R training sequence, and t_{BIFS} is the inter-frame spacing between the consecutive beacon

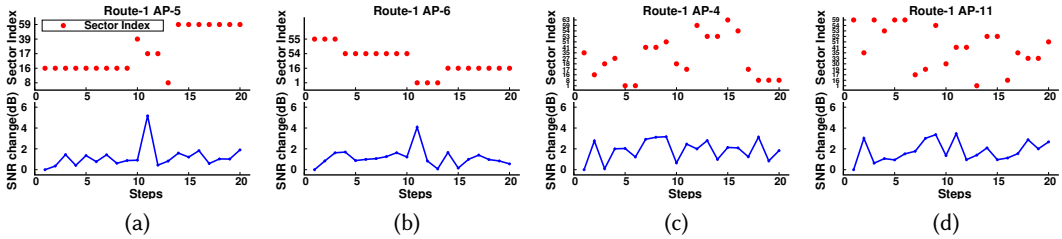

Fig. 2. Sector change and SNR change for Route 1.

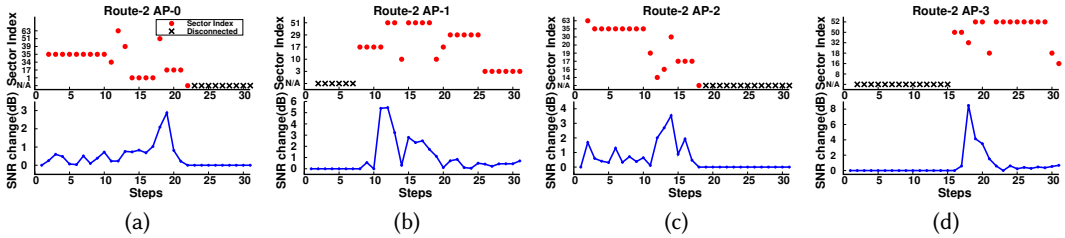

Fig. 3. Sector change and SNR change for Route 2.

frames. Based on 802.11ay specifications [15], $t_{beacon} = 14.5\mu s$ per sector (with beacon frame size of 50 bytes and transmission rate of 27.5 Mbps(Modulation and Coding Scheme(MCS) 0)), $t_{TRN-R} = 2.2\mu s/8.7\mu s$ with 64/256 length Golay sequence for one TRN unit and t_{BIFS} varies from $1\mu s$ to $18\mu s$. As an example, the BF overhead with 64 Tx sectors and 16 Rx sectors could be between $3.2ms$ to $9.9ms$ for one AP. Furthermore, the overhead linearly increases with the number of APs. In a dense network of 10 APs, $32ms$ to $99ms$ could be wasted in BF out of the $100ms$ of beacon interval, significantly reducing the useful time for data communication. The link-level BF schemes can reduce the per-link overhead, but the total overhead still increases with the number of APs.

2.2.2 Observations in Multi-AP BF. We first conduct a measurement study to understand the BF patterns among different APs in a mmWave WLAN. The observations then form the basis of our coordinated BF scheme. We consider a densely deployed WLAN with many APs and a client moving around on different routes, as shown in Fig. 4a. We measure the BF information, including the Tx sector index and their corresponding SNR for different APs for different routes. These routes are selected based on typical moving patterns of users. For instance, in Route 1, the client enters the room, goes to a desk at the corner, and goes back to the door, while Route 2 represents a typical route taken by the client when she leaves the building. In our experiments, the APs [2] and clients [1] both are equipped with 802.11ad radios. We modified the 802.11ad driver [45] to implement the Rx BF and to extract the BF information to user space. In Figs. 2 and 3, we show the downlink APs' highest SNR Tx sector index (selected by the AP) and their corresponding SNR change for each subsequent step as the client moves on Routes 1 and 2. Here, each step is approximately 1 m. We make two important observations that motivate our system design:

(1) Not all APs Tx sector frequently changes, so not all APs need to perform BF frequently. As shown in Figs. 2a and 3a, when the client moves towards AP 6 on Route 1 (first nine steps) and towards AP 0 in Route 2 (first ten steps), their best Tx sector does not change until the client turns

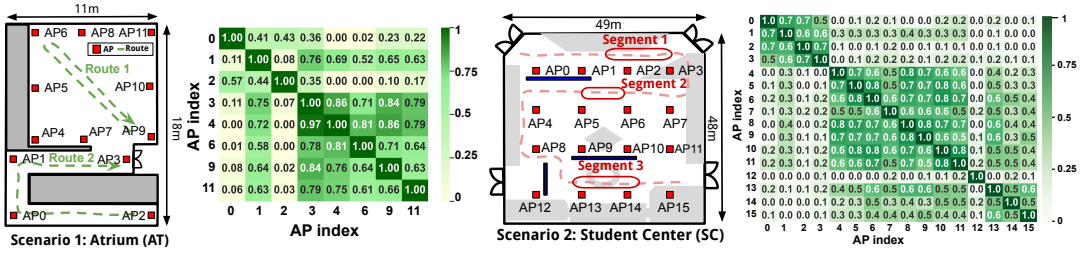

Fig. 4. (a) Network layout and client routes used in testbed experiments. (b) Correlation of SNR change between APs by accumulating all instances. (c) Room layout and client route used in channel simulation. (d) Correlation of AoA between different pairs of APs by accumulating all instances.

around. This, albeit, is not true for all APs. The best Tx sector for AP 4 and AP 11 on Route 1 changes frequently, as shown in Fig. 2c and 2d due to their relative position and the user mobility. This longer beam coherence for some APs means that not all APs need to perform BF in every beacon interval. They can be triggered to do BF only when their BF information is expected to change significantly. However, the question is how to predict when an AP's BF information has changed significantly to a client without actually performing BF.

(2) APs can be correlated in terms of the change in their BF information. In Figs. 2 and 3, we plot the (absolute) SNR change of every subsequent measurement step and observe similar patterns of SNR change between different APs when the client moves at various different locations or makes particular movements. For instance, the SNR of AP 6 (Fig. 2a) and AP 9 (Fig. 2b) changed by more than 4 dB in Route 1 at step 11 when the client turns around. Similarly, the SNR of AP 0 (Fig. 3a) and AP 3 (Fig. 3d) also changed together in Route 2 at step 18. These similar patterns of SNR change show that the APs can have correlations in terms of their BF information to clients. Fig.4b shows the correlation matrix of SNR change between APs by accumulating all instances among 12 different routes. When two APs' SNR change pattern has high similarities with each other, they have a high correlation value.

Before developing a model that can leverage this correlation, we first try to understand the fundamental underlying reason behind the existence of the correlation. The pairwise SNR correlation (which we later use to formally define uncertainty) emerges from the correlation between the angle of arrival/departure (AoA/AoD) of paths from the APs to a client. As the client moves, the AoA/AoD of LoS and NLoS paths from APs to the client changes accordingly. However, depending on the mobility of the client, location and orientation of APs, and most importantly blockages and reflections in the room, different pairs of APs demonstrate different AoA/AoD correlations. This in turn results in different SNR correlation and uncertainty dependencies between the pair of APs.

To validate this, we carry out channel propagation simulations in the Remcom Wireless Insite simulator [32]. We simulate the 3D environment and AP positions shown in Fig. 4c with a client walking on the shown route. At each step, we calculate the AoA of the path with the highest signal strength between the client and all APs. We then calculate the Pearson's correlation of the AoA between different AP pairs for the collected data as shown in Fig. 4d. We observe that when a pair of APs can provide a direct LoS path to the client, their AoA correlation is high. To better understand this, we divide the route into three segments and analyze the segment-specific correlations. When the client is walking on Segment 1, both AP 0 and AP 1 can provide an LoS path to it. Since this path's AoA changes with client mobility, the AoA correlation is observed to be high

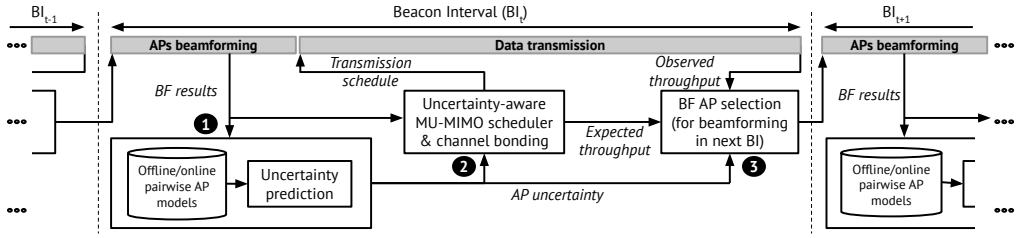

Fig. 5. CoBF system overview.

(0.81 for Segment 1 and 0.7 as in Fig. 4d for the entire route). For Segment 2, AP 10 can provide an LoS path to the client while AP 14 can provide a non-LoS path only due to a wall blockage. In this case, their AoA correlation is still relatively high (0.53 for Segment 2 and 0.5 as in Fig. 4d for the entire route) as the LoS path depends on client mobility only while the NLoS path also depends on the position of reflecting objects and blockages. Lastly, for Segment 3, both AP 9 and AP 10 can only provide a non-LoS path to the client, resulting in a lower correlation of 0.22. Since the reflected path depends on available reflecting surfaces and the position of blockages along with client mobility, the correlation reduces. The overall correlation between APs 9 and 10 is still high (0.6 in Fig. 4d.) given that they have LoS paths to many parts of the entire route shown in Fig. 4c. This way, the correlation in how different AoA/AoD changes for different pairs of APs translate to different SNR change and SNR correlation.

The SNR correlation is a coarse representation of if two APs' BF information for a client changes similarly or not. *We need a fine-grained model that can predict the significant BF information change for one AP given the BF information change for another AP for various instances of client location and mobility.* Our proposed CoBF scheme reduces the BF overhead by judiciously selecting a small subset of APs based on how important it is for those APs to perform the BF. This importance is formally modeled as "uncertainty" which is the probability that an AP's BF information (SNR of different sectors) to its clients has changed significantly since its last BF. We mine a small amount of observed data for different APs' BF information change to build a prediction model. Furthermore, we use the uncertainty of BF information from the prediction model for AP grouping and sector selection during MU-MIMO scheduling to further improve the network throughput.

2.3 CoBF **Overview**

Fig. 5 shows a high-level overview of CoBF. We assume a centrally-managed enterprise WLAN scenario where APs are connected to a controller through a high-speed backhaul. CoBF consists of three important modules: (i) the uncertainty prediction module, (ii) the BF AP selection module, and (iii) uncertainty-aware scheduler with MU-MIMO grouping. CoBF executes the following steps:

(1) At the start of a BI, a small subset of selected APs perform BF using enhanced beacon frames to train their downlink Tx and Rx sectors of all antennas on the AP and all clients. The updated BF information is input to the uncertainty prediction module. The module consists of AP uncertainty prediction models which can be trained online or offline and then looked up during the run-time. Based on the BF information of the APs that performed the BF, the module predicts the uncertainty (i.e., the probability of significant change in BF information) for other APs.

(2) The current BF information and the predicted AP uncertainty values are input to the uncertainty-aware scheduler. The scheduler uses this information to calculate SINR with intra-/inter- interference in a probabilistic manner to perform Tx and Rx sector selection, MU-MIMO

grouping, and scheduling. The calculated schedule is then executed by the APs during the data transmission interval.

(3) The AP uncertainty is also input to the BF AP selection module, which ranks the APs based on predicted uncertainty. Based on the feedback ratio of observed throughput to expected throughput and the predicted uncertainty, a subset of APs (more if the ratio is small, less otherwise) is selected from the ranked list for BF in the next BI. The process is repeated every BI creating a closed-loop, feedback-based coordinated BF, and scheduling system for 802.11ay WLANs.

3 BEAMFORMING AP SELECTION

3.1 AP Uncertainty

Intuitively, the uncertainty of an AP is the probability that its last known channel and BF information to its clients have changed significantly. If an AP's uncertainty has increased substantially, it could be a potential candidate for BF in the next BI. A key challenge here is how we can predict the uncertainty of an AP. We claim that APs in mmWave WLANs exhibit correlation that can enable us to estimate the uncertainty of one AP based on the BF information change observed by another AP. Fig. 6(a) demonstrates this concept. Here, AP_a performed BF in BI t and $t-l$. Similarly, AP_e performed BF in BI $t-n$. Given the change observed by AP_a at BI t compared to BI $t-l$ for a client c_k, the link uncertainty of AP_e to c_k is the probability that its last known BF information (SNR of Tx and Rx sectors) from BI $t-n$ has *changed significantly. Put more simply, given an AP's observed channel change (BF information change) to a client, the link uncertainty of a different AP to that client is the probability that its last known channel has changed significantly. The uncertainty of an AP is the cumulative link uncertainty for all its clients.* In mmWave WLANs, APs can exhibit varying levels of correlation depending on their relative positions, distance, ambient reflectors, and blockages in the environment. As we show in the next subsection, the correlation information can be captured in the pairwise uncertainty prediction model.

Uncertainty vs. channel estimation. The uncertainty does not predict the actual change of channel (and the channel itself by that means) but instead tries to capture the likelihood that the channel has changed significantly or not. Predicting the change in the channel through observations of different antennae or frequencies is a different problem that has been studied recently in [5, 41].

3.2 Uncertainty Prediction Model

We first define uncertainty formally before describing how we build the uncertainty prediction model. We then discuss how the model is used for BF AP selection.

The channel quality observed when an $AP_i \in M$ performs BF with a client $c_k \in N$ in a BI t can be represented using a $|S_T| \times |S_R|$ matrix $D(AP_i, c_k)^t$ where $d_i^k \in D$ is the observed SNR value, S_T is the set of Tx sectors on AP_i, S_R is the set of Rx sectors on client c_k and $s_p \in S_T$, $s_q \in S_R$. If AP_i performed BF in BI $t-l$ and is also selected for BF in BI t, let $\delta(AP_i, c_k)_{t-l}^{t}$ be the matrix of element-wise absolute difference between the two matrices $D(AP_i, c_k)^t$ and $D(AP_i, c_k)^{t-l}$, and $\Delta(AP_i, c_k)_{t-l}^{t} = \Sigma_{s_p=1}^{|S_T|}\Sigma_{s_q=1}^{|S_R|}\delta[s_p, s_q]$ is the total observed change in SNR between the BF at two BIs. The link uncertainty between an AP_i and c_k at BI t is the probability that

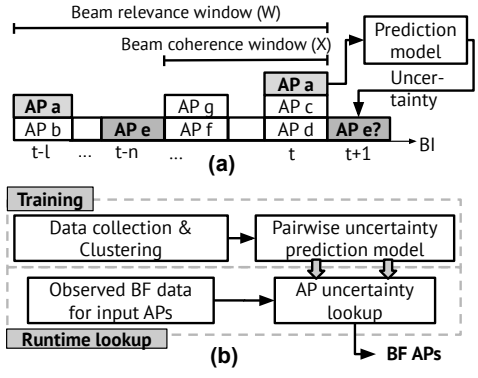

Fig. 6. (a) Example showing uncertainty and AP selection. (b) BF AP selection overview.

$\Delta(AP_i, c_k)^t_{t-l} > \kappa_{TH}$ where κ_{TH} is a predetermined threshold indicating a significant change. Furthermore, based on the pairwise prediction models that we build, the link uncertainty between AP_j (which did BF at $t-n$) and c_k in BI $t+1$ can be predicted by AP_i which performed BF at t and $t-l$. Specifically, given an observed change in SNR $\delta(AP_i, c_k)^t_{t-l}$ for input AP_i, the link uncertainty for output AP_j and c_k in the next BI $t+1$ can be predicted as the probability that $\Delta(AP_j, c_k)^{t+1}_{t-n} > \kappa_{TH}$ (conditioned on the observed SNR change for AP_i).

Our uncertainty prediction model is constructed in a pairwise manner, where the *input AP* is the one where channel change is observed and the *output AP* is the one for which the uncertainty is being predicted. We evaluate the performance of other types of models (multiple AP correlation scheme) in Section 5. Our models are developed as shown in Fig. 6(b).

(1) Data collection. First, during the data collection phase, a client walks around in the WLAN area collecting SNR values for Tx and Rx sectors with different APs in their range. CoBF does not require the client to follow any specific mobility pattern. As we demonstrate in Section 5, a moderate amount of warwalking effort is adequate for gathering all APs' BF information offline in practical scenarios. Alternatively, the uncertainty prediction models can also be trained in an online manner where the BF data collected by clients from carefully chosen APs is used for training while predicting the next subset of APs to perform BF. As we show in Section 5, such online learning can not only reduce/eliminate the warwalking efforts but can achieve a comparable network performance.

(2) Data curation. The collected data is traversed to calculate observed pairs of input and output AP changes using a sliding time window W. We refer to W as the *beam relevance window* as shown in Fig.6(a). It is used to specify the last W BIs when calculating the BF information change for APs. BF performed in BIs older than W are considered outdated and are disregarded.

For every pair of beamforming instances for AP_i (say at BI t and $t-l$) within the last W BIs, we calculate the BF information change $\delta(AP_i, c_k)^t_{t-l}$. Similarly, for every pair of beamforming instances for AP_j at $BI\ t$ and $t-n$ within the last W BIs, we calculate the total beamforming information change $\Delta(AP_j, c_k)^t_{t-n}$. The tuples $e(AP_i, AP_j) = \{\delta(AP_i, c_k)^t_{t-l}, \Delta(AP_j, c_k)^t_{t-n}\}$ are added to $E(AP_i, AP_j)$ to create a pairwise set using moving clients c_k. We normalize the change using the time difference between BIs. The process is repeated while traversing the entire collected data to create pairwise sets $E(AP_i, AP_j)$ consisting of tuples of $e(AP_i, AP_j)$ calculated as above.

(3) Clustering and uncertainty estimation. Our next step is to cluster the pairwise AP set $E(AP_i, AP_j)$ based on the similarity in input AP_i's observed beamforming change (i.e., $\delta(AP_i, c_k)$) to calculate the probability of significant change for the output AP_j (i.e., uncertainty). We use clustering where an instance $e(AP_i, AP_j) = \{\delta(AP_i, c_k), \Delta(AP_j, c_k)\}$ is randomly selected from set $E(AP_i, AP_j)$ and all instances in $E(AP_i, AP_j)$ that are within a predetermined Euclidean distance from e's input $\delta(AP_i, c_k)$ are grouped to form a cluster. Note that $\delta(AP_i, c_k)$ is a matrix of size $|S_T| \times |S_R|$ which can be reduced to $|S_T|$ by selecting the best Rx sector for each Tx sector. The process is repeated until all instances in $E(AP_i, AP_j)$ are clustered. *For each cluster, we find the uncertainty of the output AP_j by calculating the fraction of instances that have total beamforming change higher than a predetermined threshold (i.e., $\Delta(AP_j, c_k) > \kappa_{TH}$).* Each cluster is represented by $f(v_n, p_n) \in F(AP_i, AP_j)$ where v_n is the centroid of the cluster and p_n is the calculated uncertainty. The process is repeated for all pairs of APs to create their clusters (i.e., set F) and uncertainty.

Here, we make two important remarks. First, we note that the clusters formed after the clustering process do not simply depend on locations. This is because due to different orientations and variations in multipath at different locations, it is possible that two different locations away from each other end up having similar BF information changes and are clustered together. Second, a relatively small amount of training data is sufficient because the model focuses on learning the

uncertainty probability through clustering which is coarser than predicting the actual channel change (e.g., predicting how much the SNR of each sector between an AP and client has changed). We also compare the offline model with online models by dynamically increasing the training data in the evaluation.

3.3 AP selection Process

During the runtime, the uncertainty prediction model developed in the last section is used for selecting a subset of APs to do BF. The objective here is to identify the APs that potentially have high (predicted) uncertainty. This objective is captured by accumulating link uncertainty among all clients for that AP with $\lambda_{AP_j} = (\sum_{k=1}^{|N|} p_{AP_j}^{c_k})/|N|$ where p is the link uncertainty. After calculating the λ for all APs, the top K APs are selected to do BF in the next BI. Here, $K \propto h_o/h_e$ depends on the throughput feedback. When the observed network performance h_o is close to the expected throughput h_e (calculated from the scheduler introduced in Sec. 4) in the current BI, it can be inferred that the BF information is relatively up-to-date, and fewer APs are selected for BF in the next BI. On the other hand, when the gap between the observed and expected network throughput increases, the AP selection module aggressively selects more APs for BF, reducing the overall uncertainty in the network. The process of this beamforming AP selection is described below:

(1) After BF in the current BI t, find the APs that did not do BF in last $\mathcal{X}$ BIs (i.e., $AP_j \notin A_{t-\mathcal{X}}^t$), and mark them as output APs. Here, $\mathcal{X}$ is referred as the *beam coherence window* (as shown in Fig. 6a) where APs that performed BF within $\mathcal{X}$ do not have to repeat their BF again in the next BI as their observed BF information is likely to be unchanged (i.e., beam coherence). $A_{t-\mathcal{X}}^t$ is the set of APs that did BF during BI $t-\mathcal{X}$ to BI t. Then for each of these output APs, find the input AP_i(s) in set $A_{t-\mathcal{X}}^t$ using the maximum correlation value as shown in Fig.4b from the training model.
(2) Calculate link uncertainty $p_{AP_j}^{c_k}$ between AP_j and each client c_k by finding the nearest cluster $f(v_n, p_n) \in F(AP_i, AP_j)$ from the observed change $\delta(AP_i, c_k)$ (smallest distance between centroid v_n and $\delta(AP_i, c_k)$) and use its uncertainty p_n or weighted uncertainty from all input APs as the link uncertainty, i.e., $p_{AP_j}^{c_k} = p_n$. *The idea here is to match the current observed BF change of the input AP to a client with the warwalking clusters, find the cluster with similar change, and use the probability of significant change for that cluster as the uncertainty of the output AP to the client.*
(3) Rank all output APs based on the score λ. Select top $K = (|M| - |A_{t-x}^t|) \times h_o/h_e$ APs to do BF for next BI. Goto step 1 for the next BI.

It is worth noting that we select the input AP that has the highest correlation with the output AP and use it to predict the uncertainty for the output AP in (1) above. The intuition behind this pairwise model is that APs with high correlation can better predict the uncertainty values for each other. An alternative model could consider all input APs to predict the uncertainty of the output AP. However, this requires a large number of models (compared to the pairwise models) where every combination of subsets of input and output APs should be modeled, resulting in very high computational and lookup costs. In Section 5, we compare our pairwise approach with a multiple AP correlation scheme (Fig. 11d) where instead of selecting the highest correlation AP for uncertainty prediction, the predicted uncertainty value of all input APs are weighted with their correlation with the output AP and the weighted mean of uncertainty is used.

4 UNCERTAINTY-AWARE SCHEDULING

As shown in Fig. 5, the APs selected in the previous section will conduct BF at the beginning of the BI. After the BF, the uncertainty values are updated before the scheduler can schedule links and initiate data transmission. In CoBF, we build an SINR model to capture interference based on the BF information and their uncertainty. In this section, we first introduce our uncertainty-based SINR

modeling and then focus on the centralized scheduling with the MU-MIMO scheduler. Specifically, CoBF uses the uncertainties to guide the MU-MIMO scheduler. Here, instead of selecting "uncertain" APs as in the BF, the scheduler selects "certain" APs in scheduling to improve network throughput.

4.1 Interference with Uncertainty

Scheduling in 802.11ay WLANs requires accounting for three types of interference: (i) Path/sector diversity: The use of different Tx and Rx sectors can change link interference relationship. (ii) User diversity: With the use of MU-MIMO, different groups of clients can be formed for each AP, resulting in varying levels of *intra-group interference*. Similarly, in a WLAN, there can be *inter-group interference* between the client groups of different APs. (iii) Channel diversity: The use of different channels changes the adjacent and non-adjacent channel interference (ACI and NACI). The above-mentioned interference can be captured using an SINR model, however, a critical challenge here is that such a model assumes that complete BF information is available at the scheduler before the link interference modeling and scheduling can be carried out. As stated earlier, getting the complete BF information for all APs and clients incurs unacceptable measurement overhead. To address this problem, *we integrate the uncertainty of BF information directly into the interference model and scheduling problem.*

Expanding our previous notations, set D represents the BF information at the controller. Each entry in D is d_i^k representing the SNR between AP_i and client c_k on channel b with Tx sector s_{ik}^p and Rx sector s_{ik}^q. Here, $b \in B$ where B is the set of available channels with different channel index (8 channels in 802.11ay). Because the current BF information might not be up-to-date, we use the uncertainty captured by our prediction model described in last section to estimate the SNR value. Therefore, the SNR between AP_i and c_k can be calculated as

$$S(AP_i, c_k) = d_i^k (1 - p_{AP_i}^{c_k}) \tag{1}$$

, where $(1 - p_{AP_i}^{c_k})$ is the probability that the last known beams from AP_i to c_k have not changed significantly (i.e., certainty) and can still provide the last known SNR.

The intra-group interference from the concurrent transmissions $d_i^{k'}$ using another RF chain from the same AP_i to a different client $c_{k'} \neq c_k$ can be calculated as the summation of signals from the transmit sector $s_{ik'}^{p'}$ of $d_i^{k'}$ but for the same receive sector s_{ik}^q in d_i^k. The intra-group interference can be represented as

$$I_{intra}(AP_i, c_k) = \sum_{k' \neq k} d_i^{k'} (1 - p_{AP_i}^{c_{k'}}) \tag{2}$$

On the other hand, in dense mmWave WLAN, the APs forming different groups can interfere with each other. So the inter-group interference from transmissions of other APs for the same receive sector of the same link at c_k can be calculated as

$$I_{inter}(AP_i, c_k) = \sum_{i' \neq i} \sum_{k'' \neq k} d_{i'}^{k''} (1 - p_{AP_{i'}}^{c_{k''}}) \mathcal{I}_{b,b'} \tag{3}$$

The parameter $\mathcal{I}_{b,b'}$ represents the interference factor between channels b and b'. We calculate the factor for 802.11ay channels [15] using the spectrum mask model presented in [30]. Given different APs could be dynamically assigned different channels, the factor affects inter-group interference. The SINR of the link can be calculated as

$$SINR(AP_i, c_k) = \frac{S(AP_i, c_k)}{N_0 + I_{intra}(AP_i, c_k) + I_{inter}(AP_i, c_k)} \tag{4}$$

The advantage of our SINR model is that it directly operates on sector-level SNR where a sector can cover the LoS path, a NLoS path or both (multi-path). So, interference between different links can be calculated regardless of whether they operate on LoS and/or reflected paths.

4.2 Uncertainty-aware Scheduling

In CoBF MU-MIMO scheduling, our objective is to maximize the total sum-rate of all potential links in the network by finding an optimal MU-MIMO groups set with suitable Tx sectors, Rx sectors and channels for each link. The complexity of this problem is $O(|M|^{|N|}\mathbb{C}^{|M|})$ where $\mathbb{C} = (|S_T||S_R|)^{N_{RF}}|B|$ is the search space within the group of one AP. $|M|, |N|, |S_T|, |S_R|, N_{RF}, |B|$ are the total number of APs, clients, Tx/Rx sectors, per AP RF chains, and number of channels respectively. Due to this high complexity, we present next a heuristic based approach that leverages the AP uncertainty.

AP searching order. Our heuristic-based scheduler greedily picks APs with low link uncertainty and high signal strength (together they yield high expected throughput) to clients. Firstly, CoBF creates an AP searching order for scheduling. The scheduler creates a sorted list U where the APs are sorted based on $\sum_{k=1}^{|N|} \alpha_{AP_i}^{c_k} \times (1 - p_{AP_i}^{c_k})$ where $\alpha_{AP_i}^{c_k}$ is the link quality factor calculated based on updated D and $(1 - p_{AP_i}^{c_k})$ is the updated certainty values after APs complete their BF.

MU-MIMO grouping and scheduling. For each AP in the sorted list U, CoBF first forms its MU-MIMO group. The grouping process greedily adds clients with high demand to a group. For each newly added client, it iteratively searches for alternate Tx and Rx sectors for all clients in the group until the sum rate of SINR of links in the group cannot be improved. The client is added to the group only if the SINR of all links in the group are higher than the demand after its joining. As stated earlier, our intra-group interference calculation accounts for link uncertainty. Since the scheduler uses the SINR interference model, the interference between links in a MU-MIMO group is directly calculated based on the SISO BF information (SNR of different Tx and Rx sectors) and uncertainty. The scheduler then goes over the 802.11ay spectrum to allocate channels to each group and tries to pack multiple groups in the same scheduling slot while considering inter-group interference. In this process, the scheduler incrementally adds links of a newly formed group to the existing schedule. A link is added if adding it does not violate the SINR requirement of existing schedule. If no link in the new group can be scheduled, next adjacent channel is used to reduce the inter-group interference. The process is repeated for every AP's group in order. The APs carry out the transmission based on the schedule and get the observed throughput h_o. The expected throughput h_e is calculated as the summation of scheduled links' demand. The ratio h_o/h_e then becomes input to the uncertainty-based AP selection for the next BI.

The complexity of the scheduling procedure is $O(|M|(\mathbb{G} + 8(|M|N_{RF})^2))$ where 8 is the number of channels in 802.11ay, and $(|M|N_{RF})$ is the maximum available number of links in the MU-MIMO WLAN. Here, $\mathbb{G} = |N||S_T||S_R|\mu$ is the complexity of the grouping algorithm where μ is the number of iterations (less than 10 in our case). Therefore, our greedy scheduler reduces the complexity from exponential to polynomial.

5 EVALUATION

We evaluate CoBF using (i) an 802.11ad WLAN testbed with a modified driver and (ii) large-scale trace-driven 802.11ay simulations where traces are collected through real channel measurements and Remcom mmWave channel simulator. The idea of coordinated BF presented in this work is general and can be applied to 802.11ay as well as 802.11ad WLANs. Since no 802.11ay devices are currently commercially available, we first use our 802.11ad testbed to evaluate CoBF. We then use measurement and Remcom channel traces to simulate 802.11ay WLANs for a large-scale evaluation with MU-MIMO.

Testbed setup and implementation. We develop a mmWave WLAN testbed to evaluate the feasibility of coordinated BF and collect measurement channel traces for 802.11ay WLAN simulations. The testbed consists of 12 802.11ad APs deployed in a university atrium with a lobby (referred as AT with size $18m \times 11m$ as shown in Fig. 4a and 7). We use Airfide [2] 802.11ad radios as the APs. The Airfide devices use Qualcomm QCA9500 802.11ad chip with QCA6335 baseband chip and QCA6310 RF front-end chip. Each AP is equipped with 8 phased array antennas (Fig. 7) where each array is made of 32 antenna elements. We use Acer TravelMate-648 laptops [1] as the stations. The laptops use the same Qualcomm chipset but are equipped with only one phased array antenna (32 elements as shown in [25]). Both devices use the open-source wil6210 driver [45] to implement 802.11ad MAC and PHY. We use a 64-sector codebook on the Airfide APs and a 36-sector codebook on the laptops. The 3dB beamwidth of the Airfide codebook sectors is approximately 40° with sidelobes (shown in [52]). The sectors of Acer laptop antenna also have non-trivial sidelobes [25, 34].

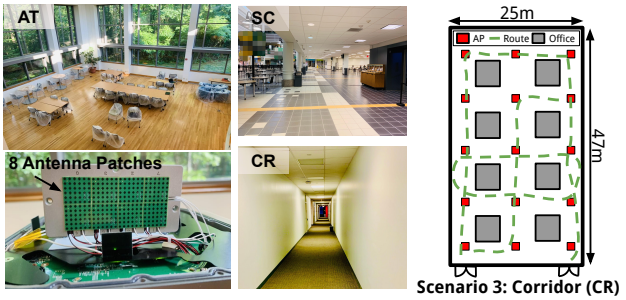

Fig. 7. mmWave testbed experiment and trace-driven simulation scenarios.

We modify the wil6210 driver to achieve the following functionality: 1) implement receive BF (not implemented by default in 802.11ad devices) with the same codebook of transmit BF; 2) set a specific Tx and Rx sector of the codebook for a link; 3) set the number, order, and index of sectors used in the Tx/Rx BF; 4) extract channel information including per sector SNR, MCS for any combination of Tx and Rx sector. With these functionalities, we utilize both Tx and Rx BF and implement our network beamforming scheme on the testbed.

Prediction model implementation. To build the prediction model and learn the uncertainty between APs, we place a client at different locations and collect the enhanced Tx/Rx BF data. At each location, the client collects the SNR for all AP Tx sectors and all client Rx sectors (64×36) for all APs in the range at that location. The data is then stitched together to be used by different walking traces in training. We run both offline and online approaches based on these data to evaluate the CoBF. The data is collected for 250 different locations in AT but as we show later, a relatively small number of locations are sufficient to achieve an accurate uncertainty prediction model.

5.1 Micro-benchmarking Uncertainty Prediction Model

To evaluate the effectiveness of uncertainty prediction between different APs, we first calculate the pairwise BF information change Δ with our experimental data for different APs, followed by the validation of how well our uncertainty prediction model can predict these BF information change.

In Figs. 8a and 8b, we show the CDF of normalized total BF information change (Δ) and uncertainty of clusters ($\Delta > \kappa_{TH}$, κ_{TH} is set to 70th percentile of all Δ) for 5 different output APs based on the same input AP in Scenario AT (shown in Fig. 4a and Fig. 7). When AP-6 is the input AP, Fig. 8a shows that AP-2 does not have any BF information change. That is expected given the room layout which results in large physical separation and blockage between the two APs. Other APs observe a range of BF information change values given the change of AP-6. AP-1 has close to 20% instances where no change is observed in relation to AP-6's change. This is also evident from the uncertainty

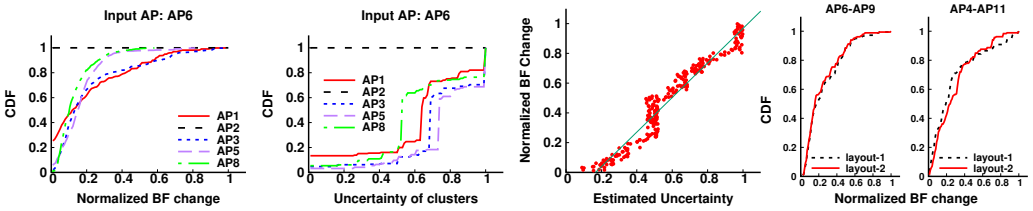

(a) (b) (c) (d)

Fig. 8. Micro-benchmarking uncertainty prediction model on 802.11ad testbed.

of clusters. Here, given a change in observed BF information in AP-6, there is a higher probability that AP-3's BF changes significantly compared to AP-1. This is in agreement with the AT layout shown in Fig. 4a which shows that our model can capture the effects of the underlying layout of the room without knowing the exact room map. In addition, in the same large open space and with the same distance to AP-6, AP-5 observes much higher BF information change and uncertainty of cluster compared to AP-8. It is because the uncertainty model does not just depend on the relative location of APs, but also depends on the room layout and the user mobility in the room, which can also be captured in our prediction model through the training process.

Fig. 8c shows the validity of our prediction model. It shows how well our uncertainty estimation model captures the BF SNR change. Here, we include all samples of estimated uncertainty in our experimental data and compare that with BF SNR change observed by the APs. We find that our estimated uncertainty accurately represents the BF information change and we can use it to guide the BF AP selection process.

Next, we further show that our uncertainty prediction model is not sensitive to small changes in the environment. Fig.8d shows the CDF of normalized BF information change among two pairs of correlated APs (AP6-AP9 and AP4-AP11) for two different sets of warwalking data collected for the same route but with different furniture layouts (desk and chair positions). The CDF difference for the two sets of warwalking data is observed to be 8.5% on average. This can be attributed to the fact that our model focuses on learning uncertainty relationships through clustering (coarser than precise channel change prediction), making it tolerant to small changes in the environment.

5.2 Micro-benchmarking Coordinated Beamforming

To evaluate the feasibility and effectiveness of coordinated BF, we conduct experiments with 4 APs (AP6, AP9 AP4 and AP11) in AT. We first run the warwalking process to develop the AP uncertainty models offline. We then run the BF AP selection process in real time on the mobile client. Here, the client also acts as a central controller connecting to all 4 APs over a 5 GHz 802.11ac control channel. The control channel is used to communicate which AP(s) should do BF and which AP should send downlink data to the client. Then the selected AP sends Iperf data to the client. Here, AP utilizes the default 802.11ad CSMA MAC for channel access. We note that CoBF can be implemented on a central controller that connects with all APs over a wired backhaul. The BF AP selection process is triggered by a drop in SNR at the client, activating a set of APs to perform BF depending on the BF AP selection scheme. We note that this micro-benchmarking aims to show the effectiveness of our coordinated BF model without considering the uncertainty-aware scheduling and MU-MIMO which are evaluated in the next subsection.

Comparison. We evaluate our coordinated beamforming scheme with different client walking traces in AT and compare it with three other schemes- **(1) 11ad-FixedAP:** A client is always

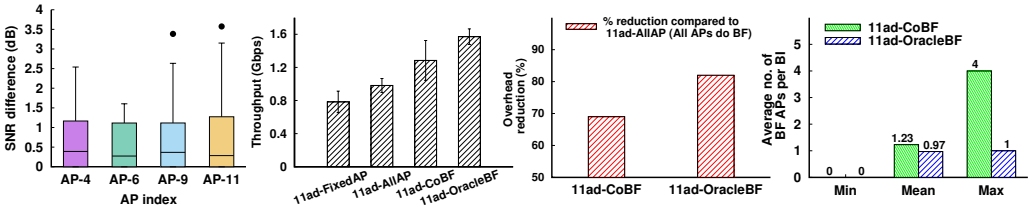

Fig. 9. Micro-benchmarking coordinated beamforming on 802.11ad testbed.

connected with one fixed AP and performs BF with only that AP. This is the default 802.11ad implementation without any central controller. **(2) 11ad-AllAP:** Here, all APs perform BF and the client receives data from the best AP with the highest Tx-Rx sector SNR. The choice of the AP and Tx-Rx sectors change as the client moves around in the room corresponding to the channel changes. **(3) 11ad-OracleBF:** The client receives data from the highest SNR AP at any time similar to the 11ad-AllAP scheme. However, we assume that the highest SNR AP is known to the client in advance. This means that the client does not perform BF with all APs, but instead only performs BF with the highest SNR AP to determine Tx and Rx sectors. Since knowing the best AP in advance is difficult in practice, the 11ad-OracleBF scheme throughput is calculated based on the 11ad-AllAP scheme implementation with BF overhead of only the best AP. In contrast, our **11ad-CoBF** scheme selects only a subset of APs to perform BF based on the uncertainty model. Then the AP with the highest SNR based on available BF information is selected to send Iperf downlink data to the client.

SNR difference. We first try to understand the following: given that our CoBF scheme limits the number of APs doing BF, how well can it maintain the link SNR? Fig. 9a shows the SNR difference between CoBF scheme and 11ad-AllAP SNR (where all APs do BF) for all APs. We find that the median SNR difference for all 4 APs is below 0.5 dB, and the 75 percentile is below 1.2 dB. This shows that our CoBF can achieve comparable SNR as 11ad-AllAP as our BF AP selection process can always select the APs with significant SNR change to do the BF.

Throughput and BF overhead. The throughput of different schemes is shown in Fig. 9b. To evaluate the packet level BF overhead of different schemes, we use an 802.11ad laptop in monitor/sniffer mode to capture the packets in the air and calculate the overhead of BF packets. The overhead is shown in Fig. 9c and 9d. We find that CoBF scheme achieves higher throughput than the 11ad-FixedAP and 11ad-AllAP schemes. 11ad-FixedAP performs the worst in terms of throughput as it cannot exploit the gains achievable through connecting with other potentially higher SNR APs. 11ad-AllAP can always select the highest SNR AP but at the cost of very high BF overhead (4 BF APs at all times). In comparison, 11ad-CoBF utilizes the uncertainty model in conjunction with throughput feedback to trigger only a small number of APs to do BF (1.23 on average). This reduces the BF overhead (amount of airtime used for BF) by 71% compared to 11ad-AllAP as shown in Fig. 9c, resulting in 30.8% higher throughput. 11ad-OracleBF achieves the highest throughput as it can select the highest SNR AP while incurring the BF overhead of only 1 AP at any point in time. As shown in Fig. 9d, there are times when 11ad-CoBF uses all 4 APs to do BF or none of the APs to do BF depending on the current state of BF uncertainty. So, CoBF not only reduces the number of APs doing BF but also carefully selects the BF APs such that the client can consistently achieve high SNR while ensuring low BF overhead.

5.3 Large-scale Evaluation

5.3.1 Trace-driven simulations. To evaluate CoBF at scale, we collect two types of channel traces:
(1) Channel traces using testbed. We use our modified 802.11ad driver to collect SNR of all Tx and Rx sectors for 12 Airfide APs in AT scenario (Fig. 4a and Fig. 7). The channel traces are collected for 250 different locations at a granularity of 0.8m.
(2) Channel traces using Remcom InSite. We create two more enterprise WLAN scenarios (Fig. 4c and Fig. 7): (i) SC: a large university student center (49m × 48m) with few stationary blockages, and (ii) CR: a corridor (25m × 47m) with a substantial number of blockages (e.g., offices, labs, etc.) in the middle (Fig. 7). We create 3D models of the SC and CR rooms and import them into Remcom InSite [32]. Here, 16 APs are deployed in SC, and 15 APs are deployed in CR. Each AP is deployed at 2m height and is equipped with eight antenna array patches (groups of 2 pointing in 4 different directions 90° apart). Each patch has 8 × 8 element arrangement. The eight patches are also used to enable 8 RF chains and spatial streams in 802.11ay MU-MIMO. Each client is equipped with 1 4 × 4 patches, pointing in one of the four different directions, and has one RF chain. The antenna gain pattern of the Tx antenna array has a 3dB beamwidth of approximately 12° to 15°.

The Remcom Wireless Insite utilizes an X3D ray model [31] which can provide exact path calculations and detailed multi-path estimation for MIMO simulations. We collect the channel matrix H for every point at 0.1m distance granularity in the SC and CR spaces. We use AWS servers with 4 GPUs, 32 CPUs, and 244 GB of memory to run the Remcom simulations for over 29 days to collect accurate channel traces. We then use the H matrices to calculate the enhanced beamforming information of APs (64 sectors) and clients (32 sectors). The received signal is calculated as $D = F_{kq}^H H F_{ip} F_{BB} X + F_{kq}^H N_0$ where $F_{kq}, F_{ip} \in \mathcal{F}$ represent antenna weight vectors of Tx analog sector s_{ik}^p and Rx analog sector s_{ik}^q, F_{BB} is the baseband precoder, X is the transmitted signal, and F^H is the conjugate transpose of F.

To implement realistic mobility of users, we identify typical walking routes (moving between offices, office to/from building doors/elevators, etc.). We vary the number of clients up to 25 in AT and up to 100 in SC and CR. Note that we select different routes for warwalking and evaluation. In terms of traffic, 90% of the traffic demand is uniformly generated from MCS1 to MCS8, and 10% of the traffic is uniformly generated from MCS9 to MCS21 as per the 802.11ay MCS. Following the same methodology in [10], we assume multiple virtual RF chains are co-located at the AP using all antenna patches. We also use the 802.11ay SNR-MCS mapping to estimate the network throughput as in [10, 17, 39, 48].

5.3.2 Comparison. We evaluate two variants of CoBF and compare them with other schemes.
(1) CoBF-**Uncertainty:** This is the main scheme of CoBF, which uses the uncertainty prediction model and uncertainty-aware scheduling with throughput feedback for AP selection.
(2) CoBF-**Random:** This scheme uses the throughput feedback to determine the *number* of APs to beamform in the next BI, but instead of selecting these APs using uncertainty (as in CoBF-uncertainty), the subset of APs are selected randomly. This enables us to evaluate in isolation how well uncertainty-based AP selection performs.
(3) Oracle-BF: In this scheme, we assume that the AP selection module knows the subset of APs that will be eventually scheduled in each BI. Hence, the APs that are to be scheduled but have not performed BF in the last $\mathcal{X}$ BIs are selected for BF in this scheme. Given that the Oracle scheme has complete knowledge about the (future) scheduling in each BI, it finds the exact subset of APs to beamform. We note that the Oracle-BF scheme still uses the scheduler like the other schemes, but it incurs the least BF overhead for a given schedule.
(4) 802.11ay-BF: This scheme assumes that all APs perform BF in each BI as in 802.11ay.

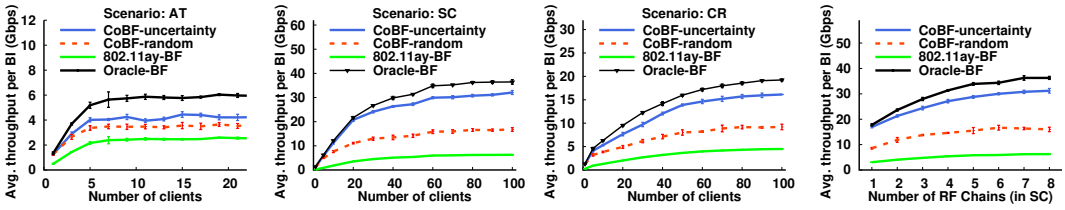

Fig. 10. Throughput comparison for different schemes with varying numbers of clients and RF chains.

5.3.3 Throughput gains due to AP selection. Figs. 10a, 10b and 10c show the average throughput per BI for different schemes as we increase the number of clients. The network throughput initially increases with the number of clients and then saturates at a point for all three scenarios. This is expected given that beyond a certain point, the spatial and frequency diversity gains saturate, and spatial reuse cannot be further increased. The saturation occurs with fewer clients in AT compared to SC and CR. This is because, in the AT experiment scenario, we use real 802.11ad devices that have larger beamwidth and non-trivial sidelobes compared to Remcom cases. This, in conjunction with the smaller area of the AT room, increases the inter and intra-group interference much faster and reduces the achievable spatial reuse.

In terms of different schemes, we observe that our scheme CoBF-uncertainty achieves 5.5 times higher throughput compared to 802.11ay and is close to 90% of the Oracle-BF scheme in SC and CR scenarios. In AT, the gain is 1.8 times higher than 802.11ay and 77% of the Oracle-BF, mostly due to the larger interference footprint of the scenario. A similar result is observed in Fig. 10d in the SC scenario with 100 clients when increasing the number of RF chains from 1 to 8. More RF chains increase the throughput for all schemes as expected, and as before CoBF schemes perform significantly better than 802.11ay.

Based on the throughput results, we make three observations. First, the gains of two CoBF schemes can be attributed to the reduction in BF overhead as far fewer APs are performing BF in CoBF schemes compared to the 802.11ay scheme, where all APs perform BF in each BI. Second, CoBF-uncertainty's better performance (average gain 70.3%) compared to CoBF-random in all three scenarios demonstrates that merely selecting fewer APs is not sufficient; the choice of APs also plays a critical role. Third, the difference between the CoBF-uncertainty and the Oracle-BF scheme can be attributed to the fact that the Oracle-BF does a better job at selecting the APs that will be eventually scheduled in that BI since it knows the future scheduling as the comparison purpose.

BF overhead. Fig. 11a compares the BF overhead of different schemes in terms of the number of APs doing BF per BI using mean, 75th, and 95th percentiles. First, as we expect, the BF overhead is high for 802.11ay (12 APs in AT), given that every AP performs BF in each BI. In comparison, CoBF-uncertainty has 4.2 times fewer APs performing BF on average. This shows that the uncertainty-based AP selection reduces the number of BF APs and their overhead. Looking at 75th and 95th percentile values, we find that there are some cases where many more APs (5 to 9 out of 12) perform BF in CoBF. Such cases are the result of the simultaneous change of orientation of many clients while walking, resulting in outdated BF information for many AP (i.e., high uncertainty). Even in such cases, CoBF-uncertainty provides valuable savings in terms of overhead compared to 802.11ay BF. In the case of SC and CR, CoBF-uncertainty also yields 5.1 times and 4.8 times lower overhead than 802.11ay-BF on average. CoBF-random achieves a bit higher overhead compared with CoBF-uncertainty, and the throughput performance loss is due to the poor choice of APs.

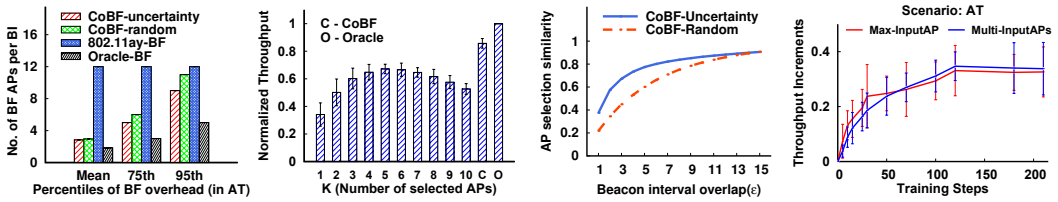

Fig. 11. (a) BF overhead, (b) Different number of selected APs, (c) AP selection similarity (d) Pairwise uncertainty correlation (Max-InputAP) vs. multiple AP correlation (Multi-InputAPs)

Impact of number of APs (K). As stated earlier, our CoBF selects top K APs to perform beamforming after the APs are ranked based on their predicted uncertainty. Since K is a crucial parameter that affects the performance of the system, we select K based on the throughput feedback of the previous beacon interval. Specifically, $K \propto h_o/h_e$ where h_o and h_e are the observed and expected network throughput in the previous BI, respectively. This means that in the current BI, fewer APs are selected if the observed network throughput closely matches the expected value. However, if the observed throughput differs significantly from the expected value, the BF information of APs is likely outdated, and hence, more APs are selected in the next BI for BF.

In Fig. 11b, we evaluate how well the throughput feedback works in terms of selecting K. To do so, we compare the CoBF with fixed values of K where a predetermined number of APs are selected in each BI without considering the dynamic throughput feedback. We also compare CoBF with the Oracle solution where the value of K that achieves the highest network throughput is assumed to be known in advance. As we observe in Fig. 11b, the fixed value of K performs poorly compared to CoBF. This is because a lower value of K reduces the BF overhead but results in outdated BF information in absence of throughput feedback. On the other hand, a higher value of K can keep the BF information relatively more up-to-date albeit at a cost of high BF overhead with unnecessary BF for many APs. In comparison, CoBF is able to better adapt to the changing network performance and use it for predicting K. This is evident from CoBF performance being close to the oracle.

Pairwise uncertainty correlation vs. multiple AP correlation. In our current AP selection process, a pairwise AP uncertainty correlation model is used where the input AP with the highest correlation is selected for the output AP. We compare this with multiple AP correlation scheme where a weighted average uncertainty value of all input APs is used instead. This scheme gives more weight to the predicted uncertainty values when the correlation is high. Fig. 11d shows the throughput increment of both schemes with different amounts of warwalking data. We find that Max-InputAP performs better than the Multi-InputAP scheme when the warwalking training data is less. Specifically, before 71 steps, the Max-InputAP achieves a 24.13% higher throughput increment compared to Multi-InputAP. This is because, with less training data, some of the calculated AP correlations are less accurate and can lead to incorrect uncertainty prediction. However, as more training data is added to the model, the Multi-InputAP scheme performs better (4.39% more throughput increment) with the calculated correlations being more representative of reality. We note that the throughput increment with multiple APs is not significantly high and the Max-InputAP can still perform reasonably well with different amounts of training data.

AP selection similarity. To take a better look at the accuracy of AP selection compared to the Oracle-BF scheme, we define a new metric SI as $|AP_O \cap AP^{\epsilon}_{CoBF}|/|AP_O|$, where $AP_O \cap AP^{\epsilon}_{CoBF}$ captures the intersection of the set of APs selected for BF by the Oracle-BF scheme and CoBF

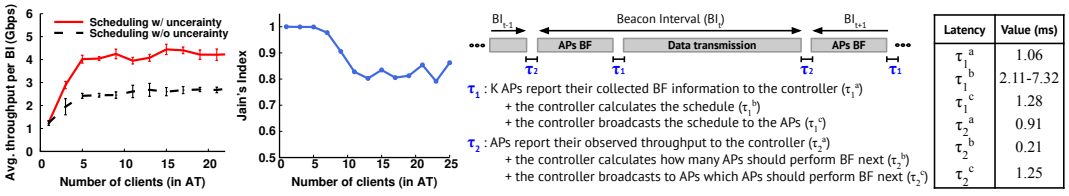

Latency	Value (ms)
τ_1^a	1.06
τ_1^b	2.11-7.32
τ_1^c	1.28
τ_2^a	0.91
τ_2^b	0.21
τ_2^c	1.25

(a) (b) (c)

Fig. 12. (a) Effectiveness of uncertainty-aware scheduling (b) Fairness evaluation for different numbers of clients, (c) Various latencies involved in CoBF and their average values.

scheme within ϵ BIs. In Fig. 11c, AP similarity *SI* represents the percentage of similarity between BF AP set AP_O for Oracle-BF and set AP_{CoBF} for CoBF scheme within ϵ BIs. The similarity for CoBF-uncertainty increases as ϵ increases. More importantly, CoBF-uncertainty consistently selects more APs that are also selected by the Oracle-BF scheme compared to the CoBF-random, signifying the effectiveness of the uncertainty-based AP selection. With 3 BIs, CoBF-uncertainty achieves 67% BF AP similarity as the Oracle-BF scheme.

5.3.4 Gains due to uncertainty-aware scheduling. Our uncertainty-aware scheduler picks APs with low uncertainty with a higher priority. We compare the uncertainty-aware scheduler with another scheduler that uses the same greedy steps but does not consider AP uncertainty while scheduling. Fig. 12a shows the network throughput with the increasing number of clients for the two schedulers in the experimental scenario AT. We observe that the uncertainty-aware scheduler provides a significant gain of 59.8% compared to the uncertainty-unaware scheduler.

We also evaluate the fairness of client throughput using Jain's fairness index. The results are shown in Fig. 12b for varying numbers of clients with CoBF-uncertainty in the AT scenario. We find that as the number of clients increases, the index decreases from 1 to 0.83 and then converges with the further increase in the number of clients. The convergence occurs around the same point when the network throughput also starts saturating, shown in Fig. 10a. It indicates that even when BF APs are selected based on their aggregate uncertainty towards clients, this selection strategy does not adversely affect the performance of a single or few clients. CoBF is capable of updating the beams towards all clients reasonably well to avoid any fairness issues or saturation.

5.3.5 Control Latency. Since our CoBF scheme relies on a central controller, periodic communication between the APs and the controller is necessary. Such a communication incurs an additional time overhead that should not be too large such that it overshadows the benefits of CoBF's intelligent BF. Fig. 12c shows a detailed description of various latencies involved in CoBF and their mean values. Here, we assume that the controller is connected to the APs over a wired (Ethernet) backhaul as commonly implemented in most enterprise WiFi networks. As shown in Fig. 12c, there are two major central controller-related latencies: scheduling latency (τ_1) and BF AP selection latency (τ_2)

Scheduling latency (τ_1):. When (K) APs complete their BF, they report their BF results to the central controller. We refer to this latency as τ_1^a. The controller then calculates the schedule based on our interference and scheduling algorithm which takes τ_1^b. Lastly, the controller sends out the schedule to APs and let's assume this communication takes τ_1^c. We find that τ_1^a is only 1.06 ms and our BF updates are 160 bytes. Furthermore, the BF process of one AP can proceed in parallel with another AP's BF update to the controller. This means that the τ_1^a is essentially sending one packet

from the last BF AP to the controller. τ_1^b takes on average 7.32 ms on a 32-core desktop based on our simulations. When the uncertainty is low (fewer updated BF, fewer changes to interference), τ_1^b can be as low as 2.11 ms. Lastly, the controller can send out the calculated schedule as a broadcast packet to all APs which takes 1.28 ms on average.

BF AP selection latency (τ_2):. After the data transmission, the APs report their observed throughput to the controller. Considering a separate packet from each AP to the controller, the latency for doing so (τ_2^a) is approximately 0.91 ms. The AP then uses this information to select the number of APs (K) that should perform BF in the next BI. We note that here only the number of APs needs to be calculated since the actual order list of APs based on their uncertainty can be calculated at the controller when APs perform data transmission. The process of doing so takes a negligible amount of time, i.e., τ_2^b = 0.21 ms. Lastly, a broadcast packet can be used by the controller to notify the APs about which APs should perform BF in the next BI where τ_2^c = 1.25 ms.

Considering the two control latencies, our CoBF scheme incurs an overhead of 7 to 12 ms. As we discussed in Sec. 2.2.1, the time overhead for BF for one AP can be up to 8-10 ms in practice. Compared with the exhaustive BF of 802.11ad/ay, CoBF significantly reduces the number of APs performing the BF. As shown earlier, in the AT scenario, 802.11ad/ay requires all 12 APs to perform BF while CoBF requires only 3 APs to perform BF on average in each BI. This results in a huge reduction (approximately 71 ms) in BF overhead for CoBF while the added control latency is much smaller in comparison. Hence, CoBF can reduce the BF overhead without incurring any considerable control overhead.

5.3.6 Different amount of warwalking data. We now evaluate how much warwalking data is sufficient to build the uncertainty prediction model. Figs. 13a, 13b, and 13c show the percentage throughput increment as more and more warwalking instances are added, starting with no warwalking data (i.e., random selection). Our warwalking scheme is referred as “Offline-All” in Figs. 13a, 13b, and 13c. Here, one warwalking instance is when a client performs beamforming with all APs in its range at a discrete point in physical space. We separately build a new uncertainty model as new warwalking data is added and rerun the throughput measurement.

Fig. 13a shows the throughput gain with the number of steps considered in warwalking for the experimental AT scenario. Here, walking for 30 steps achieves 23.7% throughput increment, and convergence is observed after 120 steps with 32.4% throughput increment. Fig. 13b shows the results for the SC scenario where the throughput increases sharply even with a small number of warwalking instances. In the SC scenario, the highest gain (87.5%) is observed with 150 discrete walking steps in the room. We note that the throughput increment after 50 steps is marginal. In the case of the CR scenario (Fig. 13c), the highest gain (52.9%) is observed with 210 steps in the room. The lower throughput gain in AT compared to SC and CR can be attributed to the coarser granularity of warwalking data collection in experiments compared to simulations. The CR scenario requires more amount of warwalking because of substantially more stationary blockages between APs. Overall, in all three scenarios, we observe that building the CoBF uncertainty model requires a reasonable amount of warwalking data which can be fairly easily collected in practice. Note that carefully choosing the locations and mobility routes could further reduce the warwalking time.

5.3.7 Offline warwalking vs. online learning. While the overhead of offline learning is small but non-negligible, we now discuss how our scheme can be adapted to online learning settings. To do so, we develop two alternate schemes and compare them with our "Offline-All" scheme. We refer to these two online learning schemes as “Online-Random” and “Online-CoBF”. We note that in our Offline-All scheme, a client performs BF with all APs at each step while moving around. This BF information is used for training the uncertainty prediction model after the offline warwalking

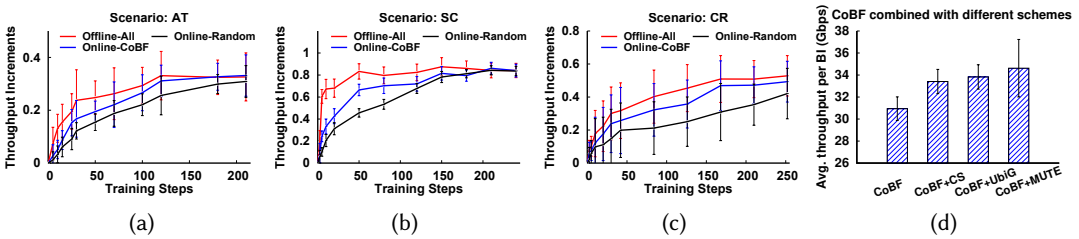

Fig. 13. (a-c) The performance of CoBF uncertainty prediction model with different amount warwalking data and its comparison with online learning alternatives (d) CoBF is orthogonal to the link-level BF schemes.

phase. In comparison, the online schemes leverage the data from a client that is moving around while simultaneously performing BF based on a given scheme and collecting the data to train the model. In the Online-Random scheme, the client only performs BF with randomly selected APs at each step. In the Online-CoBF scheme, the APs with which the client performs BF are not selected randomly but are instead selected using our CoBF scheme. In both online schemes, the uncertainty prediction model is retrained at every step with additional BF information being available.

Figs. 13a, 13b, and 13c show the performance of the online scheme as more and more training data is added compared with the offline scheme with different amount of warwalking data used for training. We find that the online models converge in a similar manner (albeit slowly) compared to adding more warwalking data to the Offline-All model. The performance gap between the Offline-All and Online-Random is significant as the client randomly selects APs to perform BF in the Online-Random model compared to all APs in the Offline-All model. In the CR scenario, the gap in performance is observed to be significant even after as many as 250 walking steps. On the other hand, the Online-CoBF model performs relatively better compared to the Online-Random model and achieves a convergence much similar to the Offline-All model. This shows that when APs for which BF information is used in training are selected carefully (as in the CoBF scheme), the resultant BF information can be very useful in training the uncertainty prediction models even with a limited amount of data. This means that it is indeed feasible to convert our offline warwalking-based model training into an alternative online learning scheme which can in fact eliminate the laborious warwalking process. However, such a model does require a careful selection of APs for BF to collect the training data as selecting them randomly cannot guarantee convergence in many practical scenarios.

5.3.8 Augmenting CoBF *with link-level beamforming schemes.* CoBF reduces the BF overhead by reducing the number of APs that beamform in each BI. However, it assumes that when an AP performs the BF, it searches all Tx/Rx sectors. CoBF can be combined with other existing link-level BF schemes to further reduce the overhead. To show how our coordinated beamforming can be augmented with link-level beamforming, we use CoBF in conjunction with three link-level BF schemes: (i) compressive searching (CS) [34] which reduces the complexity to find the best Tx/Rx sector to logarithmic order by randomly probing a limited number of BF sectors; (ii) UbiG [38] which reduces the complexity to a constant order using Power Delay Profile (PDP) to estimate all paths; (iii) MUTE [10] which firstly uses PDP to select a candidate group of sectors with maximum spatial separation and then uses zero-forcing to reduce the interference in the MIMO group.

Fig. 13d shows that the throughput can be increased by 7.9%, 9.3%, and 15.9% with 64 Tx sectors and 32 Rx sectors in our SC scenario with 8 RF chains when CoBF-uncertainty scheme is combined with CS, UbiG, and MUTE, respectively. The gains in the case of CS and UbiG can be attributed to

their lower complexity in link-level sector searching ($O(Plog(S))$ for CS and $4P$ for UbiG, where P is the number of mmWave paths and S is the number of sectors). In the case of MUTE, we combine it within each MU-MIMO group selected by CoBF to reduce the number of searched sectors, along with zero-forcing to reduce intra-group interference assuming that we can get the full H matrix for all clients at the APs. We also observe a considerable variation in MUTE in Fig. 13d because zero-forcing can reduce the transmitted signal quality below the client requirement in some cases, which can impair the total throughput. In conclusion, as per the recent trend of increasing the number of Tx and Rx antennas and sectors, the link-level beamforming can be seamlessly combined with CoBF to reduce the BF overhead and increase the achievable throughput.

6 RELATED WORK

mmWave beamforming. The problem of reducing beamforming overhead has been studied extensively in recent years (survey in [18]). Hierarchical searching [3, 13, 39] by iteratively probing from wider to narrower beamwidth sectors can reduce the overhead to $log(S)$ but needs additional overhead of feedback. Compressive searching [19, 26–29, 34] can reduce the BF overhead to $O(Plog(S))$, in which P is the underlying number of paths of the mmWave link, without additional feedback overhead. Authors in [12] designed a multi-armed probing framework with the complexity of $O(Plog(S))$. Besides, multi-armed patterns are utilized to improve reliability [16] and support multicast [51]. Authors in [38] reduced the BF overhead to $4P$ using the signal's power delay profile. Authors in [25] use $4S$ probing to get the CSI information on COTS devices and use that for adaptive beamforming. Hybrid beamforming based MIMO for mmWave network has been studied in early works [3, 4, 8, 35, 36]. In [11], authors proposed an approach that decouples hybrid beamforming from user selection. In [10], the authors propose to use the power delay profile to predict the intrinsic path skeleton first and then select a subset of candidate sectors for the group with maximum path separation. As we showed in Sec. 5, CoBF can be augmented with the above-mentioned link-level and MIMO BF schemes to further reduce the total overhead. Recently, authors in [42, 52] proposed a mmWave MIMO software radio platform that can be used for implementing hybrid beamforming in CoBF.

mmWave WLANs. In mmWave WLAN, interference is modeled in the form of conflict graphs in [17, 46, 48], and different schemes [6, 33, 49] were proposed to reduce the link interference. These prior works assume that complete BF information for all links is available before calculating link interference, while CoBF relies on only a subset of APs that performed BF and incorporates BF uncertainty directly in the interference calculation.

Given the BF information of a subset of APs, mmWave localization through WLAN APs [7, 22–24], triangulation [38] or sensor-based prediction of client's mobility [43] can be used to calculate LoS paths for other APs, which could also reduce the beamforming overhead. However, they would perform poorly in predicting reflected paths without extensive measurement-based profiling of reflectors in indoor space [44], especially in case of dynamic blockages. Compared to this, CoBF takes a different approach where the number of APs that perform BF is intelligently reduced to decrease the BF overhead. In a recent work [50], authors showed the feasibility of networked beamforming to reduce the BF overhead. In comparison, CoBF presents a holistic framework for uncertainty prediction based on both offline and online data, a closed-loop AP selection, and an uncertainty-aware MU-MIMO grouping and scheduling with limited BF information. Our work also performs extensive experimental implementation and evaluation in three different scenarios to show the effectiveness of coordinated beamforming.

7 CONCLUSIONS, LIMITATIONS, AND FUTURE WORK

In this work, we presented a novel scheme of coordinated beamforming to reduce the beamforming overhead in dense mmWave 802.11ay WLANs. We introduced CoBF, a closed-loop, feedback-based system that carefully selects a small subset of APs to perform beamforming in each BI based on BF uncertainty. Our numerical evaluation showed that CoBF can provide a substantial reduction in BF overhead and improvements in network throughput compared to default 802.11ay. The work can be extended in the future to address some of its current limitations. CoBF can be adopted in a distributed setting where the uncertainty evaluation and scheduling can be performed by the APs directly without the need for a central controller. Such a distributed scheduler can also include uplink traffic from clients to the APs. Finally, more complex machine learning models such as neural networks can be utilized to further improve the online learning of uncertainty relationships.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers and our shepherd Dr. Emiliano De Cristofaro for their valuable comments and feedback. This research is supported by NSF grants CNS-2045885, CNS-1730083, and CNS-1816943.

REFERENCES

[1] Acer TravelMate P648. 2016. https://www.acer.com/ac/en/US/press/2016/175243.
[2] Airfide AFN2200. 2020. https://airfidenet.com/.
[3] A. Alkhateeb, O. El Ayach, G. Leus, and R. W. Heath. 2014. Channel Estimation and Hybrid Precoding for Millimeter Wave Cellular Systems. *IEEE Journal of Selected Topics in Signal Processing* 8, 5 (2014), 831–846.
[4] A. Alkhateeb, G. Leus, and R. W. Heath. 2015. Limited Feedback Hybrid Precoding for Multi-User Millimeter Wave Systems. *IEEE Transactions on Wireless Communications* 14, 11 (2015), 6481–6494.
[5] Arjun Bakshi, Yifan Mao, Kannan Srinivasan, and Srinivasan Parthasarathy. 2019. Fast and Efficient Cross Band Channel Prediction Using Machine Learning. In *The 25th Annual International Conference on Mobile Computing and Networking* (Los Cabos, Mexico) *(MobiCom '19)*. Association for Computing Machinery, New York, NY, USA, Article 37, 16 pages. https://doi.org/10.1145/3300061.3345438
[6] G. Bielsa, A. Loch, and J. Widmer. 2019. Optimizing mmWave Spatial Reuse: Signal-To-Interference Aware Beamtraining. In *2019 IEEE 20th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*. 1–6.
[7] G. Bielsa, J. Palacios, A. Loch, D. Steinmetzer, P. Casari, and J. Widmer. 2018. Indoor Localization Using Commercial Off-The-Shelf 60 GHz Access Points. In *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*. 2384–2392.
[8] J. Choi. 2015. Beam Selection in mm-Wave Multiuser MIMO Systems Using Compressive Sensing. *IEEE Transactions on Communications* 63, 8 (2015), 2936–2947.
[9] Y. Ghasempour, C. R. C. M. da Silva, C. Cordeiro, and E. W. Knightly. 2017. IEEE 802.11ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine* 55, 12 (2017), 186–192.
[10] Yasaman Ghasempour, Muhammad K. Haider, Carlos Cordeiro, Dimitrios Koutsonikolas, and Edward Knightly. 2018. Multi-Stream Beam-Training for mmWave MIMO Networks. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking* (New Delhi, India) *(MobiCom '18)*. ACM, New York, NY, USA, 225–239. https://doi.org/10.1145/3241539.3241556
[11] Yasaman Ghasempour and Edward W. Knightly. 2017. Decoupling Beam Steering and User Selection for Scaling Multi-User 60 GHz WLANs. In *Proceedings of the 18th ACM International Symposium on Mobile Ad Hoc Networking and Computing* (Chennai, India) *(Mobihoc '17)*. Association for Computing Machinery, New York, NY, USA, Article 10, 10 pages. https://doi.org/10.1145/3084041.3084050
[12] Haitham Hassanieh, Omid Abari, Michael Rodriguez, Mohammed Abdelghany, Dina Katabi, and Piotr Indyk. 2018. Fast Millimeter Wave Beam Alignment. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (Budapest, Hungary) *(SIGCOMM '18)*. Association for Computing Machinery, New York, NY, USA, 432–445. https://doi.org/10.1145/3230543.3230581
[13] S. Hur, T. Kim, D. J. Love, J. V. Krogmeier, T. A. Thomas, and A. Ghosh. 2013. Millimeter Wave Beamforming for Wireless Backhaul and Access in Small Cell Networks. *IEEE Transactions on Communications* 61, 10 (2013), 4391–4403.
[14] IEEE. 2012. IEEE Standard for Information technology–Telecommunications and information exchange between systems–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 Std 802.11ad-2012 (2012). https://standards.ieee.org/ieee/802.11ad/4527/

[15] IEEE. 2021. IEEE Standard for Information Technology–Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks–Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 2: Enhanced Throughput for Operation in License-exempt Bands above 45 GHz. IEEE Std 802.11ay-2021 (2021). https://standards.ieee.org/ieee/802.11ay/6142/

[16] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. 2021. Two Beams Are Better than One: Towards Reliable and High Throughput MmWave Links. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference* (Virtual Event, USA) *(SIGCOMM '21)*. Association for Computing Machinery, New York, NY, USA, 488–502. https://doi.org/10.1145/3452296.3472924

[17] Suraj Jog, Jiaming Wang, Junfeng Guan, Thomas Moon, Haitham Hassanieh, and Romit Roy Choudhury. 2019. Many-to-Many Beam Alignment in Millimeter Wave Networks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. USENIX Association, Boston, MA, 783–800. https://www.usenix.org/conference/nsdi19/presentation/jog

[18] S. Kutty and D. Sen. 2016. Beamforming for Millimeter Wave Communications: An Inclusive Survey. *IEEE Communications Surveys Tutorials* 18, 2 (2016), 949–973.

[19] Z. Marzi, D. Ramasamy, and U. Madhow. 2016. Compressive Channel Estimation and Tracking for Large Arrays in mm-Wave Picocells. *IEEE Journal of Selected Topics in Signal Processing* 10, 3 (April 2016), 514–527. https://doi.org/10.1109/JSTSP.2016.2520899

[20] Thomas Nitsche, Guillermo Bielsa, Irene Tejado, Adrian Loch, and Joerg Widmer. 2015. Boon and Bane of 60 GHz Networks: Practical Insights into Beamforming, Interference, and Frame Level Operation. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies* (Heidelberg, Germany) *(CoNEXT '15)*. Association for Computing Machinery, New York, NY, USA, Article 17, 13 pages. https://doi.org/10.1145/2716281.2836102

[21] Thomas Nitsche, Carlos Cordeiro, Adriana B. Flores, Edward W. Knightly, Eldad Perahia, and Joerg C. Widmer. 2014. IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi [Invited Paper]. *IEEE Communications Magazine* 52, 12 (2014), 132–141. https://doi.org/10.1109/MCOM.2014.6979964

[22] J. Palacios, G. Bielsa, P. Casari, and J. Widmer. 2018. Communication-Driven Localization and Mapping for Millimeter Wave Networks. In *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*. 2402–2410.

[23] J. Palacios, P. Casari, H. Assasa, and J. Widmer. 2019. LEAP: Location Estimation and Predictive Handover with Consumer-Grade mmWave Devices. In *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*. 2377–2385.

[24] J. Palacios, P. Casari, and J. Widmer. 2017. JADE: Zero-knowledge device localization and environment mapping for millimeter wave systems. In *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*. 1–9.

[25] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. 2018. Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking* (New Delhi, India) *(MobiCom '18)*. Association for Computing Machinery, New York, NY, USA, 241–255. https://doi.org/10.1145/3241539.3241576

[26] D. Ramasamy, S. Venkateswaran, and U. Madhow. 2012. Compressive adaptation of large steerable arrays. In *2012 Information Theory and Applications Workshop*. 234–239.

[27] D. Ramasamy, S. Venkateswaran, and U. Madhow. 2012. Compressive tracking with 1000-element arrays: A framework for multi-Gbps mm wave cellular downlinks. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. 690–697.

[28] M. E. Rasekh and U. Madhow. 2018. Noncoherent compressive channel estimation for mm-wave massive MIMO. In *2018 52nd Asilomar Conference on Signals, Systems, and Computers*. 889–894. https://doi.org/10.1109/ACSSC.2018.8645127

[29] Maryam Eslami Rasekh, Zhinus Marzi, Yanzi Zhu, Upamanyu Madhow, and Haitao Zheng. 2017. Noncoherent MmWave Path Tracking. In *Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications* (Sonoma, CA, USA) *(HotMobile '17)*. Association for Computing Machinery, New York, NY, USA, 13–18. https://doi.org/10.1145/3032970.3032974

[30] Shravan Rayanchu, Vivek Shrivastava, Suman Banerjee, and Ranveer Chandra. 2011. FLUID: Improving Throughputs in Enterprise Wireless Lans Through Flexible Channelization. In *Proceedings of the 17th Annual International Conference on Mobile Computing and Networking* (Las Vegas, Nevada, USA) *(MobiCom '11)*. ACM, New York, NY, USA, 1–12. https://doi.org/10.1145/2030613.2030615

[31] Remcom InSite High Fidelity Ray Tracing. 2023. https://www.remcom.com/wireless-insite-models/high-fidelity-ray-tracing.

[32] Remcom Wireless InSite 3D Wireless Prediction Software. 2023. https://www.remcom.com/wireless-insite-em-propagation-software.

[33] Daniel Steinmetzer, Adrian Loch, Amanda García-García, Joerg Widmer, and Matthias Hollick. 2017. Mitigating Lateral Interference: Adaptive Beam Switching for Robust Millimeter-Wave Networks. In *Proceedings of the 1st ACM Workshop on Millimeter-Wave Networks and Sensing Systems 2017* (Snowbird, Utah, USA) *(mmNets '17)*. Association for

Computing Machinery, New York, NY, USA, 29–34. https://doi.org/10.1145/3130242.3130244

[34] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. 2017. Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices. In *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies* (Incheon, Republic of Korea) *(CoNEXT '17)*. Association for Computing Machinery, New York, NY, USA, 414–425. https://doi.org/10.1145/3143361.3143384

[35] R. A. Stirling-Gallacher and M. S. Rahman. 2015. Multi-user MIMO strategies for a millimeter wave communication system using hybrid beam-forming. In *2015 IEEE International Conference on Communications (ICC)*. 2437–2443.

[36] S. Sun, T. S. Rappaport, R. W. Heath, A. Nix, and S. Rangan. 2014. Mimo for millimeter-wave wireless communications: beamforming, spatial multiplexing, or both? *IEEE Communications Magazine* 52, 12 (2014), 110–121.

[37] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2017. WiFi-Assisted 60 GHz Wireless Networks. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking* (Snowbird, Utah, USA) *(MobiCom '17)*. Association for Computing Machinery, New York, NY, USA, 28–41. https://doi.org/10.1145/3117811.3117817

[38] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2018. Towards Scalable and Ubiquitous Millimeter-Wave Wireless Networks. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking* (New Delhi, India) *(MobiCom '18)*. Association for Computing Machinery, New York, NY, USA, 257–271. https://doi.org/10.1145/3241539.3241579

[39] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 2015. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems* (Portland, Oregon, USA) *(SIGMETRICS '15)*. Association for Computing Machinery, New York, NY, USA, 71–84. https://doi.org/10.1145/2745844.2745858

[40] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links under Blockage. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation* (Santa Clara, CA) *(NSDI'16)*. USENIX Association, USA, 193–206.

[41] Deepak Vasisht, Swarun Kumar, Hariharan Rahul, and Dina Katabi. 2016. Eliminating Channel Feedback in Next-Generation Cellular Networks. In *Proceedings of the 2016 ACM SIGCOMM Conference* (Florianopolis, Brazil) *(SIGCOMM '16)*. Association for Computing Machinery, New York, NY, USA, 398–411. https://doi.org/10.1145/2934872.2934895

[42] Song Wang, Jingqi Huang, Xinyu Zhang, Hyoil Kim, and Sujit Dey. 2020. X-Array: Approximating Omnidirectional Millimeter-Wave Coverage Using an Array of Phased Arrays. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking* (London, United Kingdom) *(MobiCom '20)*. Association for Computing Machinery, New York, NY, USA, Article 5, 14 pages. https://doi.org/10.1145/3372224.3380882

[43] Teng Wei and Xinyu Zhang. 2017. Pose Information Assisted 60 GHz Networks: Towards Seamless Coverage and Mobility Support. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking* (Snowbird, Utah, USA) *(MobiCom '17)*. Association for Computing Machinery, New York, NY, USA, 42–55. https://doi.org/10.1145/3117811.3117832

[44] Teng Wei, Anfu Zhou, and Xinyu Zhang. 2017. Facilitating Robust 60 GHz Network Deployment by Sensing Ambient Reflectors. In *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation* (Boston, MA, USA) *(NSDI'17)*. USENIX Association, USA, 213–226.

[45] wil6210. 2018. https://wireless.wiki.kernel.org/en/users/drivers/wil6210.

[46] Y. Yang, A. Zhou, D. Xu, S. Yang, L. Wu, H. Ma, T. Wei, and J. Liu. 2020. mmMuxing: Pushing the Limit of Spatial Reuse in Directional Millimeter-wave Wireless Networks. In *2020 17th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. 1–9.

[47] Z. Yang, P. H. Pathak, J. Pan, M. Sha, and P. Mohapatra. 2018. Sense and Deploy: Blockage-Aware Deployment of Reliable 60 GHz mmWave WLANs. In *2018 IEEE 15th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*. 397–405. https://doi.org/10.1109/MASS.2018.00063

[48] Ding Zhang, Mihir Garude, and Parth H. Pathak. 2018. MmChoir: Exploiting Joint Transmissions for Reliable 60GHz MmWave WLANs. In *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing* (Los Angeles, CA, USA) *(Mobihoc '18)*. Association for Computing Machinery, New York, NY, USA, 251–260. https://doi.org/10.1145/3209582.3209608

[49] D. Zhang, P. S. Santhalingam, P. Pathak, and Z. Zheng. 2019. Characterizing Interference Mitigation Techniques in Dense 60 GHz mmWave WLANs. In *2019 28th International Conference on Computer Communication and Networks (ICCCN)*. 1–9.

[50] Ding Zhang, Panneer Selvam Santhalingam, Parth Pathak, and Zizhan Zheng. 2022. Networked Beamforming in Dense MmWave WLANs. In *Proceedings of the 23rd Annual International Workshop on Mobile Computing Systems and Applications* (Tempe, Arizona) *(HotMobile '22)*. Association for Computing Machinery, New York, NY, USA, 102–108. https://doi.org/10.1145/3508396.3512871

[51] Ding Zhang, Puqi Zhou, Bo Han, and Parth Pathak. 2023. M5: Facilitating Multi-User Volumetric Content Delivery with Multi-Lobe Multicast over MmWave. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems* (Boston, Massachusetts) *(SenSys '22)*. Association for Computing Machinery, New York, NY, USA, 31–46. https://doi.org/10.1145/3560905.3568540

[52] Renjie Zhao, Timothy Woodford, Teng Wei, Kun Qian, and Xinyu Zhang. 2020. M-Cube: A Millimeter-Wave Massive MIMO Software Radio. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking* (London, United Kingdom) *(MobiCom '20)*. Association for Computing Machinery, New York, NY, USA, Article 15, 14 pages. https://doi.org/10.1145/3372224.3380892

[53] A. Zhou, X. Zhang, and H. Ma. 2017. Beam-forecast: Facilitating mobile 60 GHz networks via model-driven beam steering. In *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*. 1–9.

[54] P. Zhou, K. Cheng, X. Han, X. Fang, Y. Fang, R. He, Y. Long, and Y. Liu. 2018. IEEE 802.11ay-Based mmWave WLANs: Design Challenges and Solutions. *IEEE Communications Surveys Tutorials* 20, 3 (2018), 1654–1681.

Received August 2022; revised February 2023; accepted April 2023