

Networked Beamforming in Dense mmWave WLANs

Ding Zhang
George Mason University
dzhang13@gmu.edu

Panneer Selvam
Santhalingam
George Mason University
psanthal@gmu.edu

Parth Pathak
George Mason University
pppathak@gmu.edu

Zizhan Zheng
Tulane University
zzheng3@tulane.edu

ABSTRACT

Densely deployed base stations/access points (APs) are becoming increasingly common in mmWave networks with the need to provide high capacity and reliability to clients. However, frequent beamforming between clients and APs incurs an unacceptable overhead in densely deployed mmWave WLANs with many APs. This paper presents a novel approach of “networked beamforming” where only a small subset of APs are selected for beamforming in each beacon interval in mmWave WLANs. By building a prediction model based on the concept of uncertainty, our networked beamforming scheme predicts the APs whose beamforming information is likely outdated and needs updating.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. With experimentation, we show that our scheme can significantly reduce beamforming overhead and improve network capacity for dense mmWave WLANs.

CCS CONCEPTS

- Networks → Network experimentation; Wireless local area networks;
- Hardware → Beamforming.

ACM Reference Format:

Ding Zhang, Panneer Selvam Santhalingam, Parth Pathak, and Zizhan Zheng. 2022. Networked Beamforming in Dense mmWave WLANs. In *The 23rd International Workshop on Mobile Computing Systems and Applications (HotMobile '22)*, March 9–10, 2022, Tempe, AZ, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3508396.3512871>

1 INTRODUCTION

60 GHz millimeter-wave (mmWave) WLANs using IEEE 802.11ad/ay [7, 8] standards can provide multi-gigabit per second data rates to support bandwidth-intensive applications such as multi-user volumetric video streaming, AR/VR, robotic manufacturing, etc. In mmWave WLANs, dense deployment of APs is becoming increasingly common for capacity scaling and coverage [9, 24, 26]. 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 [13, 18, 21]. The dense deployments can also greatly benefit dynamic AP-user associations to combat interference [9, 11, 21, 26, 27]. In the case of 802.11ay WLANs, densely deployed APs combined with MIMO, multi-patch antenna

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HotMobile '22, March 9–10, 2022, Tempe, AZ, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9218-1/22/03...\$15.00

<https://doi.org/10.1145/3508396.3512871>

arrays, and flexible channelization can truly realize the spatial and frequency diversity gains [3, 7]. However, dense deployment of APs and clients in mmWave WLANs poses a critical challenge where the beamforming between a large number of clients and APs incurs a formidable overhead. For example, it takes approximately 5ms to train the downlink Tx and Rx sectors of one AP to its clients in mmWave WLANs. With 10 APs, the overhead could be approximately 50ms which would consume half of the 100ms beacon interval (BI) time just for the beamforming.

We try to address the following question in this work: How can we reduce the beamforming overhead when there are a large number of densely deployed APs in a mmWave WLAN? Existing solutions proposed in the literature primarily focus on reducing the beamforming overhead on a per-link basis. These link-level solutions [4–6, 10, 15, 17–19] try and reduce the overhead by intelligently searching fewer Tx and Rx sectors for a single AP. However, the beamforming overhead of the network still increases with the number of APs. Another type of solution leverages the quasi-optical properties of mmWave channels. Here, a client’s location and orientation are estimated based on channel state information (CSI) [13] or motion sensors [21]. Then the beam sectors calculated from one AP can be used to derive the beam of remaining APs to the client without actual beamforming using techniques like triangulation [18]. These location-based methods incur additional overhead to frequently localize mobile clients. Furthermore, the reliance on quasi-optical properties makes the location-based methods less suitable for deriving NLoS reflected paths without also extensively localizing the ambient reflectors [22, 25].

In this paper, we present NetBF, a networked 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 whose beamforming information has changed to perform beamforming, we can reduce the beamforming overhead at the network level while ensuring the beamforming information is up-to-date. *However, the real challenge is how to predict when an AP’s beamforming information to its clients has changed without actually performing beamforming from that AP.*

To motivate our design, we first capture and investigate detailed beamforming (BF) information in a mmWave WLAN to find that: (i) not all APs need to do BF in each BI, and (ii) the SNR change of different APs have different levels of correlations among each other. Based on these observations, we design a scheme to identify the necessary subset of APs to do BF. We introduce a concept of uncertainty which is the probability that an AP’s SNR to its clients has changed significantly since its last BF. We then develop a prediction model where APs which performed BF recently can predict the uncertainty of other APs based on the observed correlations between APs. So, in each BI, only a subset of APs that have high uncertainty is selected for BF.




Figure 1: 802.11ay enhanced beacon frame & beamforming.

To the best of our knowledge, our solution is the first to reduce the BF overhead from a network-level perspective. A salient advantage of the proposed networked BF 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 further reduce the BF overhead. Furthermore, our scheme is *location-agnostic* and can be readily integrated into the off-the-shelf mmWave WLAN devices without MAC or PHY protocol modifications. We summarize the key contributions of NetBF as follow:

(1) We present a networked beamforming scheme (called NetBF) that can reduce the BF overhead in dense mmWave WLANs by judiciously selecting a subset of APs to do BF. We define “uncertainty” to quantify the BF information change of mmWave APs. We also propose a prediction model that uses a small amount of war-walking data to predict the AP uncertainty based on other AP’s BF observations. Finally, the model is used to select a subset of APs with high uncertainty to perform BF.

(2) We use commodity 802.11ad devices enhanced with our receive beamforming implementation to evaluate the uncertainty prediction model and the effectiveness of networked BF. Our evaluation results show that: (i) The proposed prediction model can effectively capture the uncertainty and SNR change of APs, and it is robust to small environmental changes. (ii) NetBF improves the network throughput by effectively reducing the beamforming overhead within our 802.11ad testbed.

2 MOTIVATION AND OUR APPROACH

2.1 BF overhead in dense mmWave WLANs

In 802.11ay mmWave WLAN, a beacon interval (BI) (Fig. 1a) includes beacon header interval (BHI) where EDMG (Enhanced Directional Multi-Gigabit) APs transmit beacons and control frames to train the downlink sectors, and data transmission interval (DTI) where APs and STAs exchange data frames. To train both downlink Tx and Rx sectors, the standard [7] proposes to use enhanced beacon frames where each of the beacon frames has training fields (TRN-R) appended to it (Fig. 1b). An STA receives the original part of the beacon frame using the quasi-omni sector for Tx beam training, then sweeps through all its receive sectors when receiving the appended TRN-R frames for Rx sector training.

With the enhanced beacon frames, 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 beam 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-




Figure 2: Measurement setup in a university atrium (AT) frame spacing between the consecutive beacon frames. Based on the specifications [7, 8], $t_{beacon} = 14.5\mu s$ per sector (with beacon frame size of 50 bytes and transmission rate (MCS0) of 27.5 Mbps), $t_{TRN-R} = 2.2\mu s/8.7\mu s$ 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 beamforming 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 beamforming out of the 100ms of beacon interval, significantly reducing the useful time for data communication. *Even though the current link-level beamforming schemes can reduce the per-link overhead, the total overhead still increases with the number of APs.*

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 networked BF scheme.

We consider a densely deployed WLAN with many APs and a client moving around on different routes as shown in Fig. 2. 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 [23] to implement the Rx beamforming and to extract the BF information to user space.

In Fig. 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 changes frequently, so not all APs need to perform BF frequently. As shown in Figs. 3a and 3c, when the client moves towards the AP 5 on Route 1 (first 9 steps) and towards AP 0 in Route 2 (first 10 steps), their best Tx sector does not change until the client turns around. This, albeit, is not true for all APs due to their relative position and 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 be changed significantly. However, our first question is how to predict when an AP’s BF information has changed significantly to a client without actually performing BF.




Figure 3: (a-d) Sector change and SNR change for selected route 1&2. We found that 1) Not all APs need to do BF in every BI due to beam coherence. 2) SNR change of different APs have correlations which are used in our prediction model. (e) Correlation of SNR change between APs by accumulating all instances.

(2) APs can be correlated in terms of the change in their BF information. In Figs. 3, we also 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 locations or makes particular movements. For instance, the SNR of AP 5 (Fig. 3a) and AP 6 (Fig. 3b) changed by more than 4 dB in Route 1 at step 11 when the client turns around. Similarly, the SNR of AP 0 (Fig. 3c) 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.3e 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. We note the following points about the correlations between different APs' BF information change: (1) It depends on various factors, including the location and orientation of APs, the underlying layout of the room with blockages and reflections, and the mobility of clients. (2) Pairwise correlation is asymmetric. To further verify, we conduct channel simulations with Remcom [16] for different types of indoor environments and observe that the correlations between APs widely exist for environments with varying levels of blockages and reflections due to channel correlation [18–21] in mmWave band. The correlation is a coarse representation of if two APs' BF information for a client changes similarly or not. *However, 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.*

In this paper, our proposed NetBF 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 collected war-walking data for different APs’ BF information change to build a prediction model. Our prediction model can predict the uncertainty of the APs that have not done the BF by using a small number of APs that have done BF recently. In addition, our model is location agnostic and does not require the location of APs, clients, reflectors, or blockages. Lastly, only a subset of APs with high uncertainty is selected to perform BF to reduce the BF overhead.




Figure 4: (a) AP selection overview (b) Example showing uncertainty and AP selection.

3 NETBF DESIGN

3.1 AP Uncertainty

Intuitively, the uncertainty of an AP’s BF information can be regarded as the probability that its last known BF information (SNR of Tx and Rx sectors) to its clients has 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 have observed that APs in mmWave WLANs exhibit correlation which can enable us to estimate the uncertainty of one AP based on the BF information change observed by another AP. Fig. 4 (b) 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 from BI $t - n$ has changed significantly. The uncertainty of an AP is the cumulative link uncertainty for all its clients.

3.2 Uncertainty Prediction Model

We first define uncertainty formally and use conditional probability to capture the correlations. Then we describe how to build the uncertainty prediction model.

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 $\Delta(AP_i, c_k)_{t-l}^t = \sum_{s_p=1}^{|S_T|} \sum_{s_q=1}^{|S_R|} \delta[s_p, s_q]$ 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

$\Delta(AP_i, c_k)_{t-l}^t > \kappa_{TH}$ where κ_{TH} is a predetermined threshold indicating a significant change. Furthermore, 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 $t-l$ due to the underlying correlation. Specifically, given an observed change in SNR $\delta(AP_i, c_k)_{t-l}^t$ 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 $\Delta(AP_j, c_k)_{t-n}^{t+1} > \kappa_{TH}$ (conditioned on the observed SNR change for AP_i).

Pairwise prediction model. Our uncertainty prediction model is a pairwise model 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. The pairwise models are developed by the following steps.

(1) Warwalking data collection. First, during an offline data collection phase, a client walks around in the WLAN area, collecting SNR values for all Tx and Rx sectors with each AP in its range at different locations. The warwalking can be performed on more frequently used indoor routes/areas. This offline step needs to be completed only once.

(2) Data curation. The collected data is traversed to calculate all observed pairs of input and output APs' BF information changes using a sliding time window W . We refer to W as the *beam relevance window* as shown in Fig. 4 (b). 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 BF instances for AP_i (say at BI t and $t-l$) within the last W BIs, we calculate the BF information change $\delta(AP_i, c_k)_{t-l}^t$. Similarly, for every pair of BF instances for AP_j at BI t and $t-n$ within the last W BIs, we calculate the total BF information change $\Delta(AP_j, c_k)_{t-n}^t$. The tuples $e(AP_i, AP_j) = \{\delta(AP_i, c_k)_{t-l}^t, \Delta(AP_j, c_k)_{t-n}^t\}$ are added to $E(AP_i, AP_j)$ to create a pairwise set using all warwalking clients c_k . We normalize the change using the time difference between BIs. The process is repeated while traversing the entire collected warwalking 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 BF 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 (ϵ_{TH}) 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 BF 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.




Figure 5: NetBF uncertainty prediction model (a) can predict the BF change effectively, and (b) is robust to small environment changes (Different furniture layout between set1 and set2).

Here, we make three 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 during warwalking and variations in multipath at different locations, two different locations away from each other might end up having similar BF information change and are clustered together. Second, a relatively small amount of warwalking 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). Lastly, the clustering model could be extended to an online version to consider the real-time data, which is part of our ongoing work.

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. To estimate the link uncertainty $p_{AP_j}^{c_k}$ between AP_j and each client c_k , we find 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 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 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.

After calculating the λ for all APs, the top K APs are selected to do BF in the next BI by considering $\lambda > \gamma_{TH}$, so that if the predicted uncertainty is high for a large group of APs, more APs are selected for BF. On the other hand, when all APs observed low BF information change such as in more stationary scenarios, the predicted uncertainty is small and fewer APs are selected to do BF.

4 EVALUATION

We developed a 60 GHz WLAN testbed to evaluate the uncertainty prediction model and the feasibility of networked beamforming. The testbed consists of 8 802.11ad APs deployed in a university atrium with a lobby (referred to as AT with size $18m \times 11m$ as shown in Fig. 2). We use Airfide [2] 802.11ad radios as the APs (64 sectors




Figure 6: (a,b,c) NetBF can maintain good SNR, improve the network throughput, and reduce the BF overhead. (d) The prediction model needs reasonable war-walking data.

in the codebook and 8 phased array patches) and Acer TravelMate-648 laptops [1] as the stations (32 sectors in the codebook), both of which use the Qualcomm 802.11ad chipset and the open-source wil6210 driver [23]. We modify the wil6210 driver to implement downlink receive beamforming (RxBF is not implemented by default in 802.11ad devices) with the same codebook of TxBF, set a specific Tx and Rx sector for a link, and extract channel information including per sector SNR and MCS.

Uncertainty prediction model. To build the prediction model and learn the uncertainty correlations between APs, we move a client to different locations and collect the Tx/Rx BF data. At each location, the client collects the SNR for AP Tx sectors and client Rx sectors (64×36) for all APs in the range at that location. We collected BF information data for 250 different locations in AT and then stitched them together offline to build the prediction model as shown in section 3. We empirically selected κ_{TH} as 70th percentile of all Δ to define the significant BF information change and selected $\epsilon_{TH} = 10$ (Euclidean distance between two centroids of clusters) as the similarity threshold to form clusters.

Fig. 5a shows how well our uncertainty prediction model captures the BF SNR change. Here, we include all samples of predicted uncertainty in our experimental data and compare that with the 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. 5b shows the CDF of normalized BF information change among two pairs of correlated APs (AP5-AP6 and AP4-AP7) for two different sets of warwalking data collected for the same route but with different furniture layout (desk and chair positions). The CDF difference for the two sets of warwalking data is observed to be 8.5% on average. It is because of (1) the sparsity of mmWave channel, where only the LoS and a few strong NLoS contribute to the received signal, and (2) the robustness of our prediction model, which focuses on learning uncertainty relationships through clustering (coarser than precise channel change prediction), making it tolerant to small changes in the environment.

Networked Beamforming. To evaluate the feasibility and effectiveness of networked BF, we conducted a controlled experiment with 4 APs (AP4, AP5, AP6, and AP7) in AT. A mobile client moves step by step in the room and we run the BF AP selection process in real-time to select which APs to do BF. After triggering the APs to do BF, the client connects with the AP with the highest SNR and

runs a 10s downlink Iperf. Here, the client also acts as a central controller connecting to all 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. We note that NetBF can be implemented on a central controller that connects with all APs over a wired or wireless backhaul. The BF AP selection process is triggered by a BF information change (SNR drop) when the client moves between different locations, activating a set of APs to perform BF.

We evaluate our scheme (NetBF) with 12 different client walking traces in AT and compare it with three other AP selection schemes: (1) *FixedOneAP*: A client is always connected with one fixed AP and performs BF with only that AP (default 802.11ad implementation without any central controller); (2) *AllAP*: All APs perform BF and the client receives data from the best AP with the highest Tx-Rx sector SNR. (3) *OracleBF*: We assume that the highest SNR AP is known to the client in advance and just that AP conducts BF if needed (Upper bound). In contrast, our NetBF scheme selects only a subset of APs to perform BF based on the uncertainty model developed on the war-walking data.

SNR difference. We first try to understand the following: given that NetBF scheme limits the number of APs doing BF, how well can it maintain the link SNR? Fig. 6a shows the selected sector's SNR difference between NetBF and *AllAP* (where all APs do BF) for each AP. We find that the median SNR difference is below 1 dB and the 75 percentile is below 1.5 dB. This shows that our NetBF scheme can achieve comparable SNR as *AllAP* as our BF AP selection process can always select the APs that have high potential SNR change (high uncertainty) to probe.

Throughput and BF overhead. The Iperf throughput of different schemes is shown in Fig. 6b. We find that NetBF scheme achieves higher throughput than *FixedOneAP* and *AllAP* schemes because *FixedOneAP* scheme cannot exploit the gains achievable through connecting with other potentially higher SNR APs while *AllAP* has much higher BF overhead. As shown in Fig. 6c, NetBF utilizes the uncertainty model 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 69% compared to *AllAP*, resulting in 30.8% higher throughput. *OracleBF* achieves the highest throughput as it can select the best SNR AP while incurring the BF overhead of at most 1 AP. Besides, there are times when NetBF uses all 4 APs to do BF (high uncertainty for all APs in high mobility scenario) or none of the APs to do BF (no SNR change in a static scenario), which shows that NetBF can capture the mobility and adapt to the

environment change by utilizing the uncertainty model. So, NetBF not only reduces the number of APs doing BF but it carefully selects the BF APs such that the client can consistently achieve high SNR. We expect even better performance when the APs in the mmWave WLAN become denser.

Different amount of warwalking data. We also evaluate how much warwalking data is sufficient to build the uncertainty prediction model. Fig. 6d shows the percentage throughput increment as more and more warwalking instances are added, starting with no warwalking data at all (i.e., random select an AP). Fig. 6d shows that 30 steps warwalking in AT achieves 23.7% throughput increment and convergence is observed after 120 steps with 32.4% throughput increment. Therefore, building NetBF uncertainty model requires a reasonable amount of warwalking data which can be easily collected in practice. Note that carefully choosing the locations and mobility routes for warwalking could further reduce the warwalking time.

5 RELATED WORK

mmWave beamforming. The problem of reducing beamforming overhead has been studied extensively in recent years. Hierarchical searching [19] by iteratively probing from wider to narrower beamwidth sectors can reduce the overhead to $\log(N)$ but needs additional overhead of feedback after each round. Compressive searching [10, 15, 17] can reduce the BF overhead to $O(\log(N))$ to find the LoS path without additional feedback overhead. Authors in [6] designed a multi-armed probing framework with the complexity of $O(K \log(N))$ where K is the underlying number of paths of the mmWave link. Authors in [18] reduced the BF overhead to $4K$ using the signal's power delay profile. Authors in [14] use $4N$ probing to get the CSI information on COTS devices and use that for adaptive beamforming. Our networked beamforming is orthogonal to the link-level BF strategies and could be augmented by them to further reduce the overhead.

mmWave WLANs. Given the BF information of a subset of APs, mmWave localization through WLAN APs [12, 13], triangulation between APs and a client [18] or sensor-based prediction of client's mobility [21] can be used to calculate beams for other APs. This could also reduce the BF overhead. The localization and triangulation approach can work in calculating the LoS path but is likely to perform poorly in predicting reflected paths without extensive profiling of reflectors in indoor space [22]. Additionally, frequent changes due to rotation and mobility require continuous localization and path discovery which themselves can incur overhead. Compared to this, NetBF takes a different approach where the number of APs that perform BF is intelligently reduced to reduce the BF overhead. BounceNet[9] focuses on addressing the problem of interference along with link scheduling in dense mmWave WLANs. Their algorithms utilize complete BF information at each BI for scheduling to optimize spatial reuse in dense deployments. On the other hand, our work focuses on reducing the BF overhead in the presence of multiple APs. In this sense, our work can complement BounceNet to further enhance its performance.

6 DISCUSSION

Fine-tuning the prediction model. Currently, in the prediction model, we empirically tune the parameters such as SNR change

threshold κ_{TH} , and the clustering distance ϵ_{TH} to improve the performance of the model. We are conducting studies to investigate how these parameters will affect the system's performance. For example, reducing the clustering distance will increase the number of clusters with more fine-grained clustering able to match users' observed BF change better. It improves the uncertainty prediction but increases the lookup time due to more clusters. In addition, our uncertainty prediction model is built using the warwalking data. We plan to improve the model by using *online learning* to include new users' BF information to create a more customized model.

Combining with link-level BF schemes. NetBF reduces the BF overhead by reducing the number of APs that beamform in each BI. However, when an AP performs the BF, it searches all Tx/Rx sectors with the enhanced beacons in compliance with the current standards [7, 8]. NetBF can be combined with other link-level BF schemes to further reduce the overhead within each AP's BF process. For example, compressive searching [17] could reduce the complexity to find the best Tx/Rx sector to logarithmic order by randomly probing a limited number of BF sectors while UbiG [18] reduces the complexity to a constant order using Power Delay Profile (PDP) to estimate all paths.

Resource allocation with limited BF information. NetBF can effectively reduce the BF overhead at the network level but the gap between our scheme and the oracle comes from some of the suboptimally chosen sectors in our limited BF information. Under this limited channel information, the potential interference between multiple users will affect the scheduling process. To further improve the performance of NetBF, we plan to introduce the uncertainty in the scheduling process, such as prioritizing the APs with low uncertainty to serve clients, utilizing multi-channel and channel bonding, etc.

System overhead. Our model is currently designed primarily for downlink traffic. However, if we consider the channel reciprocity, both Tx and Rx sectors selected in our model can be directly used for uplink traffic without increasing uplink overhead. In addition, we note that after building the prediction model with the war-walking process, the lookup time in the running process is negligible. Lastly, NetBF utilizes seamless handoff between dense deployed APs in the enterprise-level mmWave WLAN. The control overhead of these handoffs should be further investigated to deploy NetBF in the mmWave WLAN.

7 CONCLUSIONS

In this work, we presented a novel approach of networked beamforming (NetBF) to reduce the BF overhead in dense mmWave WLANs. NetBF judiciously selects a small subset of APs based on AP uncertainty to perform BF. Our experiments and preliminary results showed that NetBF can substantially reduce BF overhead and improve network throughput. We also show that the uncertainty prediction model can be built using a reasonably small amount of warwalking data in practice.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers and our shepherd Dr. Yao Liu for their valuable comments and feedback. This research is supported by NSF grants CNS-1815945 and CNS-1730083.

REFERENCES

- [1] Acer TravelMate P648. 2016. <https://www.acer.com/ac/en/US/press/2016/175243>.
- [2] Airfide AFN2200. 2020. <https://airfidenet.com/>.
- [3] 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.
- [4] Yasaman Ghasempour, Muhammad K. Haider, Carlos Cordeiro, Dimitrios Koutsikolas, 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>
- [5] 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>
- [6] 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>
- [7] IEEE 802.11ay: Enhanced Throughput for Operation in Bands above 45 GHz. 2021. <https://standards.ieee.org/ieee/802.11ay/6142/>.
- [8] IEEE P802.11adTM/D4.0. 2012. Part 11: WLAN MAC/PHY Enhancements for Very High Throughput in the 60 GHz Band.
- [9] 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>
- [10] 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>
- [11] 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>
- [12] 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.
- [13] 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.
- [14] 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>
- [15] Maryam Esfandi 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* (Soma, CA, USA) (*HotMobile '17*). Association for Computing Machinery, New York, NY, USA, 13–18. <https://doi.org/10.1145/3032970.3032974>
- [16] Remcom Wireless InSite 3D Wireless Prediction Software. 2021. <https://www.remcom.com/wireless-insite-em-propagation-software>.
- [17] 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>
- [18] 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>
- [19] 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>
- [20] 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.
- [21] 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>
- [22] 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.
- [23] wil6210. 2021. <https://wireless.wiki.kernel.org/en/users/drivers/wil6210>.
- [24] 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.
- [25] 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>
- [26] 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>
- [27] 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.