FastND: Accelerating Directional Neighbor Discovery for 60-GHz Millimeter-Wave Wireless Networks

Anfu Zhou, Teng Wei, Xinyu Zhang, and Huadong Ma

***Abstract*— Neighbor discovery (ND) is a critical primitive for 60-GHz wireless networks with highly directional radios. Prior work has attempted to improve the ND efficiency but overlooks the unique properties of 60-GHz phased-array antennas and spatial channel profile. In this paper, we first conduct a systematic study of the ND problem using a reconfigurable 60-GHz radio. Combined with an analytical model, we find that environmental characteristics and client mobility substantially affect 60-GHz ND latency, and due to inherent spatial channel sparsity of 60-GHz channels, even short-distance links can experience intolerable latency. To solve these new challenges, we propose a mechanism called FastND that accelerates ND by actively learning the spatial channel profile. FastND leverages steerability of 60-GHz phased-array antennas and accumulates channel information by overhearing beacon preambles along different beam directions. Using a compressive sensing framework, together with a strategical beam selection mechanism, FastND can infer the strongest spatial angle to listen to, thereby increasing the likelihood to quickly decode beacons and achieve ND. Our testbed experiments and ray-tracing tests demonstrate that FastND can reduce 802.11ad ND latency to 1/10–1/2, with different levels of mobility, human blockage, environmental sparsity, and non-line-of-sight links.**

***Index Terms*— 60 GHz wireless networks, directional neighbor discovery, beam steering.**

I. Introduction

DRIVEN by the unprecedented growth of mobile devices and bandwidth-hungry applications, mobile network traffic will increase dramatically, by 1000× from 2012 to 2020 as predicted by market research [1]. Millimeter-wave (mmWave) communication is considered as the killer technology to tackle the 1000× challenge [2], [3]. Two IEEE network standards, 802.11ad [4] and 802.15.3c [5], both targeting the 60 GHz unlicensed spectrum, have recently been ratified to deliver multi-Gbps wireless bitrate, and thus to enable a new class of device-to-device applications such as uncompressed video streaming, wireless docking and instant file synchronization among nodes without fixed infrastructure. Different from the Independent Basic Service Set (IBSS infrastructure) in the legacy low-frequency Wi-Fi networks such as 802.11n/ac, a new network architecture referred to as Personal Basic Service Set (PBSS) is adopted in the mmWave standards. In PBSS, one of the stations is chosen to play the role of PBSS Central Point (PCP), which coordinates channel access and perform network management for multiple clients.

However, due to unique properties of the 60 GHz signals, realizing the mmWave vision entails a reinvestigation of many protocol primitives, the first and foremost being *neighbor discovery* (ND). ND is needed in both infrastructure networks and device-to-device connections, where a server node (*e.g.*, PCP in PBSS) periodically broadcasts beacons so as to be identified by clients requesting association. ND is a straightforward procedure in legacy WiFi or cellular networks: as long as a client falls in the server's range, the client can decode the beacon and discover the server within one beaconing period. However, ND becomes far more challenging for 60 GHz mmWave networks, especially for the PBSS without preinstalled infrastructure, due to the following characteristics.

(i) High directionality. To counteract inherently strong attenuation at high frequencies, 60 GHz radios use phased-array antennas to concentrate signals within ultra-narrow beams—as narrow as a few degrees [6]. A client can sense server's beacon only if their beams are properly aligned. This procedure can incur high latency, considering the enormous number of directions each node can point to, and the fact that the client might be mobile and varying its beam direction. Existing work proposed to leverage the omni-directional mode to speed up ND. While the approach is effective in infrastructure Wi-Fi, the effect is opposite in PBSS. Concurrent beacons from nearby neighbors in PBSS may collide and prolong the ND latency [7]. Moreover, the resulting neighbor set can be quite different from that obtained via directional ND [8], due to the disparate antenna gains and spatial coverage. *(ii) Spatial channel sparsity*. Measurement studies [6] have shown that the 60 GHz spatial channel response is dominated by a few paths: even if a server is omni-directional, the client can hear it only from a few angular directions. To identify one of these few paths, a directional client needs to scan across all spatial angles, which leads to formidable ND latency. *(iii) Vulnerability to blockage*. Unlike in low frequencies, mmWave signals can hardly diffract around human body and other obstacles due to small wavelengths. The high directionality exacerbates this limitation. Thus, a mmWave ND protocol

Manuscript received December 31, 2017; revised July 2, 2018; accepted August 18, 2018; approved by IEEE/ACM Transactions on Networking Editor T. He. Date of publication September 10, 2018; date of current version October 15, 2018. This work was supported in part by the National Natural Science Foundation of China under Grants 61772084, 61720106007, 61332005, and 61532012, and in part by the Fundamental Research Funds for the Central Universities under Grant 2017RC43. The work of T. Wei and X. Zhang was supported by the U.S. NSF Funding under Grants CNS-1350039, CNS-1506657, CNS-1518728, and CNS-1617321. *(Corresponding author: Anfu Zhou.)*

A. Zhou and H. Ma are with the Beijing Key Laboratory of Intelligent Telecommunications Software and Multimedia, Beijing University of Posts and Telecommunications, Beijing 100876, China (e-mail: zhouanfu@bupt.edu.cn; mhd@bupt.edu.cn).

T. Wei is with the Department of Computer Engineering, University of Wisconsin–Madison, Madison, WI 53706 USA (e-mail: twei7@wisc.edu).

X. Zhang is with the Department of Electrical and Computer Engineering, University of California, San Diego, CA 92093 USA (e-mail: xyzhang@ucsd.edu).

Digital Object Identifier 10.1109/TNET.2018.2867044

often needs to leverage non-line-of-sight (NLOS) reflection paths, which can be created owing to the phased-array's ability to electronically steer beam directions.

These challenges also distinguish 60 GHz ND from conventional directional neighbor discovery problems. Most prior work in directional ND [9]–[12] assumes a line-of-sight (LOS) between PCP and client and tries to maximize the geometrical overlapping between their beams. Moreover, it is often assumed that the transmitter/receiver beam has a well-defined conical shape. The assumption no longer holds for 60 GHz phased-array radios, which rely on analog beamforming and generate one or more main beams along with many sidelobes. In addition, device motion and blockage is less of a problem at low frequencies, because the beamwidths are much wider (due to inherent antenna size constraint) and penetration loss much smaller.

In this paper, we systematically investigate the 60 GHz ND problem and propose a practical solution to address the aforementioned challenges. We target the PBSS scenario, where a PCP broadcasts identity information along several quasi-omni directions in the beginning of each beacon period. A client (*i.e.*, any non-PCP node in the PBSS) chooses one of its own listening beam directions within one beacon period. The *ND problem is to strategically steer the listening beam direction, so as to quickly reach the first direction that can decode any of the beacons from the PCP*.

Using a reconfigurable 60 GHz testbed that we custom-built, we first conduct a measurement study to thoroughly understand the 60 GHz channel characteristics pertinent to 802.11ad ND. Our measurement testifies the high channel sparsity: even at a distance of $2m$, a client may only hear the PCP along 2% of all angular directions.[1] The sparsity worsens as link distance increases. Based on the measurement, we build a theoretical model to understand the ND latency. Our analysis provides closed-form average delay results for static links and, for the first time to our knowledge, models the average ND delay for mobile 60 GHz clients under different degrees of channel sparsity. Our analysis differs fundamentally from previous theoretical/simulation studies [15], [16], in the sense that we model the impact of environment/link dynamics on ND performance following practical measurement observations.

With the measurement and analytical insights, we develop a mechanism called FastND to uniquely address the directional ND problem in 60 GHz networks. In contrast to existing 60 GHz ND protocols [4] that rely on passive trial-and-error scanning, FastND adopts an *active learning* principle[2]: it allows the client to continuously accumulate spatial channel information from previous undecodeable beacons (whose preambles can still be useful), and then strategically steer the listening beam direction to the next most promising spatial angle, so as to quickly reach the first effective direction that can decode any of the beacons from the PCP.

To realize the principle, FastND incorporates two key techniques. *First*, FastND harnesses the inherent channel sparsity, and formulates a compressive sensing framework to predict the most promising listening beam direction. *Second*, we find that traditional compressive sensing is prone to over-fitting and often traps the client in the same locally-optimal beam prediction. FastND overcomes this problem by designing a maximum distance based beam prediction, which can choose beams with the most orthogonal channel information, thus circumventing the over-fitting issue.

[1]The observation is different from previous experiments using COTS phased array, in which the beams are much wider due to limited number of antenna elements and manufacture imperfection [13], [14].

[2] Active learning is a general principle in statistical sampling. The idea is to dynamically decide on next sampling operations based on observation from previous operations, instead of following any prescribed deterministic or random operation schedule [17].

TABLE I
KEY PARAMETERS

BI	Beacon interval	100 ms
$TH_{decoding}$	Decoding threshold	-78dBm
TH_{CSI}	Carrier sense threshold	-90.55dBm

We validate the performance of FastND using our 60 GHz testbed along with ray-tracing simulation. We use the exhaustive neighbor scanning (exND), as defined in the 802.11ad beam training procedure [4], as a benchmark. We also compare FastND against the state-of-the-art Hunting Based Neighbor Discovery (HDND) proposed in [7]. Our experiments show that FastND can significantly speed up the ND process in practical scenarios, especially under crucial scenes when existing neighbor discovery protocols lead to very large delays, such as massive-element phased arrays, long link distances, or reflector-scarce environment. Overall, FastND reduces ND latency to as low as $\frac{1}{2}$ to $\frac{1}{10}$ of that by HDND and exND.

To our knowledge, FastND marks a first step to address practical ND challenges imposed by unique characteristics of the 60 GHz channel. Different from theoretical works [7], [15], [16], [18], [19], FastND also represents the first 60 GHz ND protocol verified based on a real 60 GHz platform. The specific contributions of FastND can be summarized as follows:

(i) We use testbed experiments (Sec. III) combined with rigorous analysis (Sec. IV) to thoroughly understand how the unique characteristics of 60 GHz channel and client mobility substantially affect 60 GHz ND latency.

(ii) We propose a novel scheme FastND, which adopts an active learning principle to cumulatively approximate the spatial channel profile, and identify the best listening beam direction to minimize ND latency (Sec. V).

(iii) We experiment with FastND on a custom-built 60 GHz software-radio platform, and demonstrate its ability to tame ND latency under practical settings (Sec. VI).

II. SYSTEM MODEL

A. A Primer of Directional Neighbor Discovery in 60 GHz Networks

60 GHz signals suffer from 28 dB worse attenuation than 2.4 GHz Wi-Fi signals, due to very short wavelength, Oxygen absorption, *etc.* [20]. To remedy high path loss, the IEEE 802.11ad 60 GHz radios pack dozens of antennas into a phased array, and generate highly directional beams with different directions and beamwidths.

To support directional node discovery and data transmission, an 802.11ad PCP divides time into Beacon Intervals (BIs), as shown in Fig. 1. In this paper, each BI lasts 100ms, *i.e.*, the default value in IEEE 802.11ad (Tab. I). An BI starts with a Beacon Header Interval (BHI), a short period allocated for PCP (node) discovery and association. In the beginning period of BHI (*i.e.*, BTI), an PCP sweeps a series of beacons along different quasi-omni beams, in the hope that a client can receive at least one. To decode a beacon, a client needs to steer

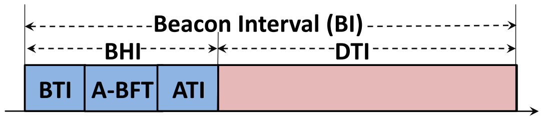

Fig. 1. Beacon interval.

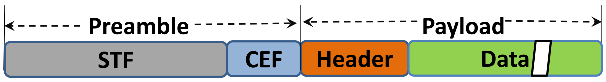

Fig. 2. Frame structure of a beacon.

itself to a listening beam direction, and it can decode a beacon if the RSS (received signal strength) exceeds the receiver sensitivity threshold (*i.e.*, $TH_{decoding} = -78dBm$) corresponding to the lowest modulation level passes a threshold, given the default noise level.[3]

However, due to the sparsity of 60 GHz channels [6], majority of the listening beams' RSS may fall below $TH_{decoding}$. To find a useful listening beam, an exhaustive beam scanning procedure in the PBSS mode works as follows: within each BI, the client needs to hold on to the same listening beam direction, so as to traverse all transmitting directions from the PCP. If none of the beacons is heard within an BI, the client will wait until the next BI and switch to a new listening beam direction. The procedure differs from the AP discovery procedure in the infrastructure mode, in which a client listens in omni-directional mode. As noted above, omni-directional mode is not suitable for PBSS due to neighbor interference and operational range difference [7], [8].

After a successful ND (*i.e.*, being able to decode a beacon), the client performs an association beamform training (within the A-BFT field), and then exchanges management information (within the ATI). Only afterwards can the client exchange data frames with the PCP (within the Data Transmission Interval, DTI).

B. PBSS Coordination

The PCP plays a key role in a PBSS and acts like the AP in infrastructure networks. Only after a client establishes association with the PCP, it can join the PBSS and start data transmission thereafter. In practice, a PCP can be selected using randomization, polling or more complex node synchronization operations [8]. In this work, we assume that the PCP has been chosen, and focus on the ND, *i.e.*, how a client can find the PCP with minimum latency.

III. MEASUREMENT

Measurement Platform

To understand the practical aspects of the 60 GHz directional ND problem, we have custom-built a 60 GHz software-radio testbed (Fig. 3). The testbed allows programmable waveform generation and received-signal processing on a PC host. It comprises Mango WARP SDR board [23] for baseband signal processing, and a high-speed ADC/DAC module connected to a 60 GHz RF front-end [24]. The RF module operates on 57–64 GHz with 10 dBm output power and 245.76 Msps baseband sampling rate. For a preliminary measurement, we plug fixed-beam horn antennas and omni-directional antennas, and steer antenna beams using an automatic motion control system [25]. We emphasize that the FastND design is applicable to real electronically steerable phased-array antennas.

[3] Strictly speaking, the sensitivity threshold is decided by SNR. We adopt a methodology (developed and commonly used in [6], [21], and [22]) to map the measured RSS to the sensitivity.

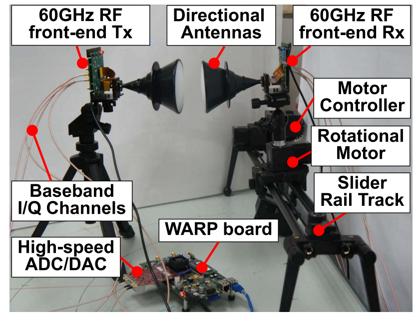

Fig. 3. Custom-built 60 GHz SDR testbed.

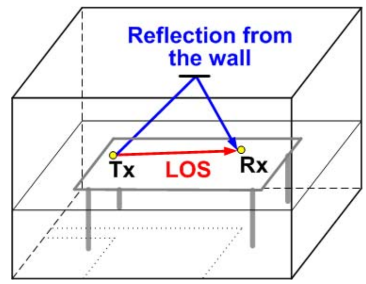

Fig. 4. Intuition behind channel sparsity.

We deploy the testbed in a typical $7{\times}8m^2$ office environment, where the Tx and Rx are mounted 1.5m above the floor. Then we conduct a set of measurements to examine the spatial channel sparsity of 60 GHz channel that are pertinent to ND performance. Note that spatial sparsity has been discovered and discussed adequately in previous studies, including measurements on a variety of mmWave band such as 28 GHz, 38 GHz, 60 GHz and 73 GHz (see [6], [21], [26]–[30], and references therein). In this work, we mainly focus on the impacts of channel sparsity on ND performance.

A Showcase of 60 GHz Channel Sparsity

Recent measurements have shown that an exhaustive search based ND leads to intolerable latency for 60 GHz networks [6], [31]. Specifically, *the spatial channel response is dominated by a few paths from a few angular directions, due to the fact that mmWave signal energy tends to concentrate on the LOS path (if it exists) and a few NLOS paths with strong reflections*. Fig. 4 illustrates an example case with one LOS path and a NLOS one from a strong reflector (*e.g.*, a metal object) on the wall. Fig. 5 testifies this problem in an example testbed link setup. The Tx is omni-directional, and client steers a 3.4° horn antenna across the azimuth plane to measure the AoA. We observe that the AoA pattern is extremely sparse: although it is omni-directional, the client can only receive strong signals from two densely concentrated directions. This implies that even for an optimistic setup with omni-directional receiving mode, it may take a long time for the client to encounter a direction with strong AoA.

2-D Angular RSS Distribution

We now set both the PCP and client to directional mode, with 2m physical distance. We emulate an 802.11ad ND procedure: within each BI, the PCP rotates across 120 beams

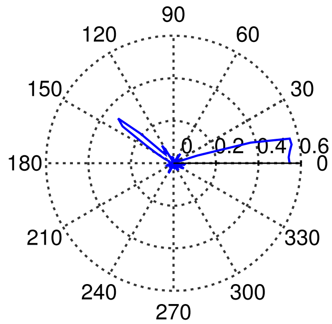

Fig. 5. An AoA pattern example.

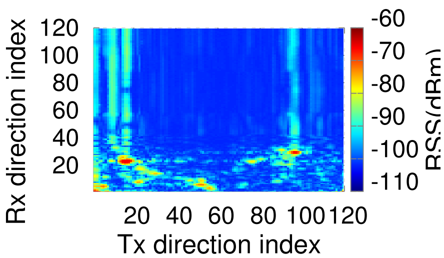

Fig. 6. RSS across direction pairs.

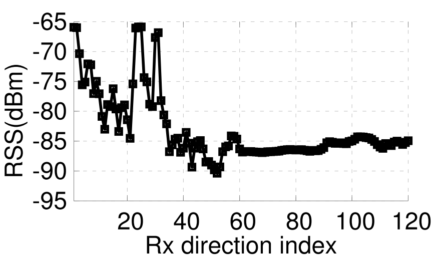

Fig. 7. Maximum RSS for each Rx direction.

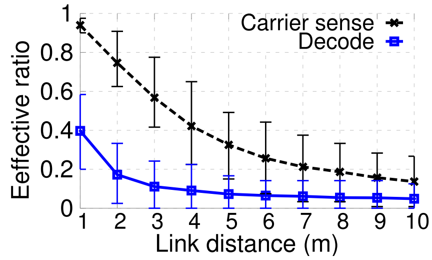

Fig. 8. Effective direction ratio v.s. link distance.

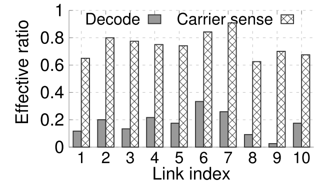

Fig. 9. Effective direction ratio variation for 2-m links.

(3° each) and sends beacons on each beam, while the client chooses its own listening beam. After an BI, the client moves to a new listening beam and the procedure repeats. The heatmap in Fig. 6 plots the RSS on each pair of Tx/Rx beams (120×120 in total). For 99.3% of the beam-pairs, the RSS is lower than -78dBm, which means that the client cannot decode the beacon most of the time. We reorganize the result and plot the maximum RSS for each listening direction in Fig. 7. We find that the client has a chance to decode a beacon for only 19 out of 120 directions.

RSS Distribution V.s. Link Distance

Now we demonstrate that the problem exists pervasively for 60 GHz radios. We increase the PCP-client distance from 1m to 10m with steps of 1m, and sample 10 distinct client locations for each distance. For simplicity, we use an omni-direction antenna on the (optimistically emulating the sweeping behavior of during an BI), and the client rotates for 120 distinct directions.

We characterize the ND latency indirectly using a metric called *ratio of effective directions*, or η-factor, which is defined as follows,

$$\eta = \frac{M_{eff}}{M} \tag{1}$$

where M_{eff} is the number of directions with RSS larger than the beacon decoding threshold $TH_{decoding}$, and M is the number of all directions. A higher η implies a client has a higher chance to quickly encounter one direction where ND can succeed. From the results in Fig. 8, we make two observations: *(i)* The effective ratio, η factor, decreases exponentially as the link distance increases. For example, the average η drops from 40% to 5% as distance increases from 1m to 10m. The reason lies in that the signal strengths of more paths fall below the beacon decoding threshold, along with link distance. This implies that ND will become much longer for longer PCP-client distances. We note that the phenomenon does not exist in omni-directional WiFi ND, or directional ND which commonly assume fan-shaped beam patterns without considering channel sparsity [9]–[12]—in such cases, ND latency is a constant as long as the client is within the PCP's "range". *(ii)* η shows large variation, meaning that same-distance links can have quite different ND latency. As an example, we depict the η of all link at 2m distance in Fig. 9. We see that *even at a very short range of 2m, some links have a very low η value of 2%, which indicates a long ND delay*. On the other hand, Fig. 9 plots the effective ratio of directions where the client can sense the PCP (but not necessarily decode beacons). This ratio is much higher than η and will be leveraged by our FastND design to speed up the ND procedure.

IV. Analysis

Now we establish an analytical framework to examine the implication of the foregoing measurement study of ND latency. The analysis is fundamentally different from [15] and [16], in the sense that it is grounded on our testbed measurement and takes into account the *channel sparsity characteristics*. Moreover, ours is the first analysis to model the average ND delay for *60 GHz mobile clients*.

A. ND Delay for Static Links

Suppose a client has M listening beams, and the η-factor is K/M. The ND delay equals the earliest time when the client's listening beam falls in one of the K effective directions that support beacon decoding. Suppose ND succeeds some time before the m_{th} BI. For simplicity, we assume that the distribution of effective beams is uniform over spatial angles (The analysis can be extended to more complicated distribution model, but it is out of the scope of this work). The probability that the non of K effective direction appears before the m_{th}

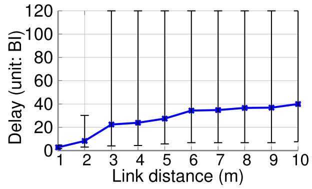

Fig. 10. Average delays for static links.

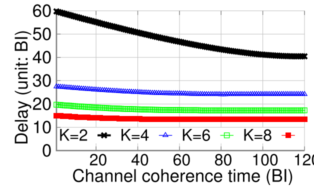

Fig. 11. Average delays for mobile client under homogeneous links.

BI is as follows,

$$G(m) = \begin{cases} \frac{C^m_{M-K}}{C^m_M} & \text{if } m \le M-K \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

where $C^b_a = \binom{a}{b}$. In consequence, the probability that ND succeeds within m BIs, denoted with $F(m)$, is as follows

$$F(m) = 1 - G(m) \tag{3}$$

Then we can derive the probability distribution function $f(m)$ (*i.e.*, the probability that ND succeeds at exactly the m_{th} as follows) is:

$$f(m) = F(m) - F(m-1) = G(m-1) - G(m) \tag{4}$$

Then the ND latency, denoted with $\overline{m}$, is

$$\overline{m} = E(m) = \sum_{m=1}^{M-K} mf(m)dm \tag{5}$$

We connect the analysis with practical channels by feeding the measured AoA patterns across 100 links (Fig. 8) into the model, and compute the average latency. Fig. 10 plots the average delay, and also the max and min delay. We make two observations: *(i)* The average delay increases with the link distance, due to the decreasing K shown in Fig. 8. *(ii)* The delay has very large variance — Even for some links with small distance (say 3m), the delay could exceed 100 BIs, implying intolerable worst-case experience for users who need association. The reason lies in the high location sensitivity of 60G channel. Due to short wavelength and high directionality, a minor displacement may shift a client out of the coverage of PCP's beam.

B. ND Delay for Mobile Links

When a client moves, the spatial channel profile (*i.e.*, AoA pattern) changes frequently, and the ND procedure needs to keep starting over. Note that the channel coherence time is determined by user's mobility model. For simplicity, we assume a continuously mobility model [32], in which a client moves at a constant speed v following an arbitrary direction.[4] The channel coherence time $c \approx \frac{1}{v}$. *i.e.*, the channel condition remains roughly consistent for every c BIs.

Suppose an ND succeeds at the m_{th} BI, we can divide this event into two sub-events. The first is that the user visits $u(m)$ locations but has not succeeded in ND. Since the user spends c BIs in each location, we have

$$u(m) = \lfloor \frac{m}{c} \rfloor \tag{6}$$

We denote the probability of this sub-event as $P_1(m)$. The other sub-event is that ND succeeds at the $u(m)+1$ location, and specifically at the $s(m)_{th}$ BI after the user enters the location, where $s(m)$ is computed as: $s(m) = mod(m,c)$. We denote the probability of this sub-event as $P_2(m)$. Clearly, the probability that ND succeeds at the m_{th} BI is

$$P(m) = P_1(m)P_2(m) \tag{7}$$

We proceed to derive $P_1(m)$. For any location among the $u(m)$ unsuccessful locations, say the i_{th} ($i \in [1, u(m)]$) location with $\vec{K}(i)$ effective beams, we can re-write Eq. (2) as follows,

$$G^i(m) = \begin{cases} \dfrac{C^m_{M-\vec{K}(i)}}{C^m_M} & \text{if } m \le M - \vec{K}(i) \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

Then probability for *not* having a successful ND in this location is $G^i(c)$, according to the analysis of static scenarios. Then we have

$$P_1(m) = \prod_{i=1}^{u(m)} G^i(c) \tag{9}$$

Then for the second sub-event, we can compute the probability that ND succeeds exactly in the $s(m)$ BI slot at the $u(m)+1$ location as follows,

$$P_2(m) = F^{u(m)+1}(s(m)) - F^{u(m)+1}(s(m)-1) \tag{10}$$

where the function $F^{u(m)+1}$ is defined like Eq. (3), *i.e.*, $F^{u(m)+1} = 1 - G^{u(m)+1}$. Putting Eq. (9) and Eq. (10) into Eq. (7), we have the probability that a mobile user achieves ND in the m_{th} BI. Based on this, we compute the expectation time for ND as follows,

$$\overline{m} = \sum\nolimits_{m=1}^{\infty} mP(m) \tag{11}$$

Using Eq. (11), we can compute the expected ND latency as client's velocity v changes under both homogeneous network scenarios (*i.e.*, different locations have the same degree of channel sparsity, likely to happen in outdoor environment) and heterogeneous scenarios (which may happen in complicated indoor multipath environment with sporadic reflectors). Interestingly, we find that *the impact of mobility is opposite under the two scenarios*, as detailed below.

1) Homogeneous Environment: Fig. 11 plots the average ND delay of a mobile client, which moves from a random location to another. We observe that: *(i) Mobility increases the average ND latency in a homogeneous environment*. Taking the case $K = 2$ as an example, we can see that as velocity decreases (the channel coherence time increases), the delay drops steadily. When the velocity becomes 0, the delay is 40 BIs, which matches the static case computed from Eq. (5). The reason can be understood intuitively: suppose a static client fails ND in one BI. Conditioned on this event, the probability that it finds an effective direction becomes higher in the next BI. However, a mobile user forfeits this higher probability and starts over again in the new location, resulting in larger ND latency. *(ii) The impact of mobility weakens as the channel*

[4]Here we set the time unit of v to be BI, *i.e.*, v is measured as meters per BI.

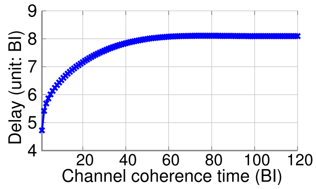

Fig. 12. Average delays for mobile client under heterogeneous links.

sparsity decreases. The latency gap between a static user and a fastest-moving user becomes smaller as K increases. The reason is that when K is larger, the ND ends faster, which leaves less space for the mobility to cause much difference.

2) Heterogeneous Environment: We now analyze ND latency in our lab environment which shows highly heterogeneous AoA patterns across locations (Fig. 9): The η-factor of 10 AoAs collected from 10 spots vary from 0.02 to 0.33, and the number of effective directions (*i.e.*, $\vec{K}$) varies between 3 and 40. We feed this $\vec{K}$ into Eq. (11) and get the average latency under different moving speeds as shown in Fig. 12. Contrary to the previous case, *mobility decreases the ND delay in such heterogeneous environment*. We find the underlying reason is as follows. In heterogeneous environment, there exist certain good locations with many effective directions (*e.g.*, the link with $K = 40$). A faster moving client can meet one of such good locations more quickly. In the extreme case, when a static user happens to be in a bad location where no effective direction exists, it will experience an infinitely long ND latency.

V. FastND Design

A. Design Overview

High-Level Idea: In contrast to the 802.11ad ND protocol that passively tries each beam direction, FastND employs *an algorithm to accumulate spatial channel information, and strategically steer the listening beam direction, so as to quickly reach the first effective direction that can decode any of the beacons from the PCP*. Specifically, FastND entails a closed-loop procedure: in each loop, the client accumulates the channel state information, more specifically the RSS, along its current listening direction. Such RSS sensing is feasible even when it cannot decode the beacon. Based on current and prior channel sensing results, the client predicts the next most promising beam direction, and steers its phased-array beamforming accordingly. ND succeeds if the client can decode a beacon along this listening direction. Otherwise, it repeats the procedure.

The key challenge in FastND lies in how to derive the next most promising beam direction during each loop. We design two complementary mechanisms to meet this challenge: the Compressive Sensing based Beam Prediction (CSBP) module and the Maximum Distance based Beam Prediction (MDBP) module. CSBP takes advantage of the sparsity characteristic of 60 GHz channel, and utilizes compressive sensing to recover the spatial channel profile (*i.e.*, the AoA pattern). Afterwards, it steers to the listening beam whose direction is aligned with the strongest spike of AoA.

Ideally, CSBP can generate an accurate channel profile after accumulating sufficient channel measurements, which is usually assumed in previous compressive sensing based channel recovery works [33], [34]. However, For the ND process, especially at the beginning phase, very few channel measurements are available. We find that with these limited measurements, compressive sensing is prone to over-fitting, *i.e.*, it results in very similar AoA peaks as in previous rounds, which traps ND in a long loop.

The MDBP is designed to handle this problem. MDBP incorporates a metric to guide the selection of next beam that is likely to provide the most orthogonal channel information, thus preventing compressive sensing from over-fitting. It is noteworthy that the active learning principle plays a key role in the entire procedure, i.e, helping to decide the most promising beam direction in CSBP, and also to decide the most complimentary direction in MDBP.

Workflow: FastND preserves the 802.11ad PHY layer and MAC layer beaconing mechanism, except that it guides the client to decide which beam to use in the next beacon period. Fig. 13 illustrates the basic workflow of FastND.

(i): An unassociated client keeps listening with a selected beam (at the very beginning, it picks a beam randomly) for a whole beacon interval (BI). If it can decode any beacon, the ND ends successfully; otherwise it proceeds to the next step.

(ii): The client senses the RSS of undecodeable beacons, and runs CSBP to compute the next listening beam direction (Sec. V-C) with this sensing result and all previous information.

(iii): Check the usability of the beam: if it is a new beam, FastND jumps to step *(v)*, else to step *(iv)*.

(iv): The MDBP module generates a new beam that is farthest to previously used beams, according to Sec. V-D.

(v): The client starts a new round of listening with the generated beam direction.

Steps *(i)*-*(v)* are repeated for each beacon interval, until one beacon is decoded by the client, which then starts the user association procedure as defined in 802.11ad. In what follows, we first conduct a feasibility analysis of accumulating channel information from undecodable 802.11ad beacons, and then proceed to detail the CSBP and MDBP.

B. Accumulating Spatial Channel Information

FastND builds on a key observation that the channel state information (CSI) of undecodeable beacons can be utilized to predict the most promising beam direction a client should listen to. The reason why CSI can be extracted from undecodable beacons lies in 802.11ad's high redundancy protection over beacon preambles (Fig. 2). A beacon preamble comprises 40 Golay-128 sequences for short training field (STF), and 9 Golay-128 sequences for channel estimation field (CEF), both modulated with a low-order but robust $\frac{\pi}{2}$-BPSK. The STF is used for detecting the presence of a beacon and performing time/frequency synchronization; whereas the CEF is used to estimate the CSI. Payload (*i.e.*, header and data) is modulated using higher-order constellation (up to 16-QAM), compounded with either direct sequence spread spectrum (DSSS) or OFDM. The DSSS modulated payload has a $32\times$ spreading factor protected by $\frac{1}{2}$-LDPC coding. It is more robust than OFDM, and mainly used to carry low-rate control data.

A simple back-of-the-envelop analysis can verify that the known preamble has much more redundancy than payload and thus is more resilient under high noise or weak signal strength. In particular, we have $9 \times 128 = 1152$ symbols for estimating CSI, in contrast to only $32 \times 2 = 64$ symbols for decoding each bit of payload data. *The SNR gap of resilience is about* $10log10(\frac{1152}{64}) = 12.55dB$. In other words, given that the RSS threshold for decoding data bits is -78 dBm [4], the threshold for extracting RSS can be as low as $TH_{CSI} = -78 - 12.55 = -90.55$ dBm. Moreover, the 12.55 dB higher sensitivity of preamble decoding translates into approximately $4.2\times$ *wider CSI sensing range than data decoding range*, assuming signals attenuate following the free-space model (which is reasonable for highly directional LOS links [6]).

The above analysis is also corroborated by our measurement. We compare the ratio of directions with RSS larger than TH_{CSI} against the ratio of directions with RSS larger than $TH_{decoding}$ in Fig. 8, as the link distance increases. We can observe that there is a large gap persisting irrespective of link distance. On average, the number of directions supporting carrier sense is 350% of that supporting decoding. This demonstrates that useful CSI can still be extracted from the preamble even when the data bits in the beacon cannot be decoded.

We make three additional notes about beacon sensing: *(i)* The client needs to distinguish the beacons from other packets. This could be realized by sensing the beacon frame duration and its unique periodicity pattern [35]. *(ii)* Just as in the original 802.11ad, FastND focuses on discovering a single PCP, which suits the 802.11ad's anticipated short-range deployment/usage scenarios [4]. In case when multiple PCPs are closely located, the PCPs need to avoid interference by using different channels, or unique Golay code preambles. Detailed exploration of such cases are beyond the scope of FastND. *(iii)* One concern is that in mobile scenarios, the accumulated beacons may become staled gradually. However, prior work showed that consecutive locations exhibit much consistency on AoA patterns [36], and thus the beacons are still useful under quasi-static or slow-movement (e.g., user walking) scenarios (which is the targeted scenarios of highly-directional 60 GHz networks). We will further validate this in Sec. VI.

C. Compressive Sensing Based Beam Prediction

Suppose the client's phased-array can generate U different beam patterns. When the PCP is sending beacons, the spatial channel response (*i.e.*, AoA pattern between PCP and client) can be recovered if the client can sequentially scan all the U beams. The resulting AoA pattern can be represented by a series of variables $\boldsymbol{x} = \{x(1), x(2), \ldots, x(N)\}$, and each variable $x(n)$ (a complex number) is the spatial channel response on the n^{th} direction. However, this straightforward scanning is highly costly because the 802.11ad client can only listen using one beam in each beacon interval (Sec. II-A). *The objective of CSBP* is equivalent to estimating $\boldsymbol{x}$, *i.e.*, to recover the AoA pattern, by listening to only a few beams. Below we detail CSBP. Note that a client can recover one AoA pattern for each quasi-omni transmitting beam direction from the PCP, but without loss of generality, we focus on how to recover one of them.

When receiving signals, an 802.11ad client uses phased-array antenna to perform a codebook based beamforming. The resulting *signal along each spatial direction is a multiplication between the spatial channel response and antenna gain along that direction*. To generate the u_{th} beam ($u \in [1, U]$, where U is the number of total beams), the phased-array applies a series of beam weights to each of its antenna elements. Consider a linear phase-array with N_R elements, and $w(u,i)$ is the weight vector (one entry from a predefined codebook) applied to the i_{th} antenna element, then the beamforming gain at an arbitrary spatial direction

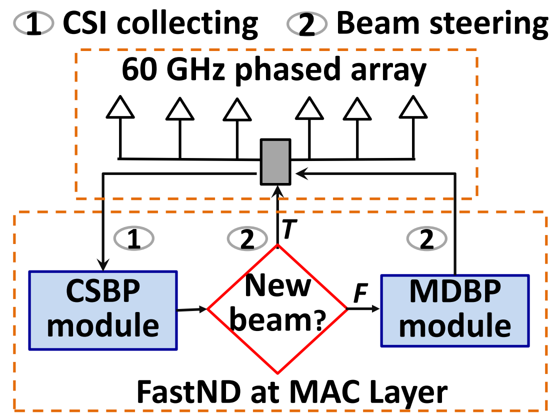

Fig. 13. Workflow of FastND.

$\theta(n)$ is as follows:

$$A(u,n) = \sum\nolimits_{i=1}^{N_R} w(u,i) exp(j2\pi i d cos(\theta(n))/\lambda) \quad (12)$$

where d is the distance between antenna elements, and λ is the carrier wavelength. Then the beam's CSI is:

$$\sum\nolimits_{n=1}^{N} A(u,n)x(n) \quad (13)$$

Meanwhile, the client can measure a channel coefficient (*i.e.*, CSI, denoted with $h^{ms}(u)$) when using the u_{th} beam. Then, we have a constraint equation as follows,

$$\sum\nolimits_{n=1}^{N} A(u,n)x(n) = h^{ms}(u) \quad (14)$$

After the client obtains a total of P CSI measurements (denoted with a vector $H^{ms} = \{h^{ms}(1), \ldots, h^{ms}(P)\}$) along a set of beams denoted as Φ, we can concatenate the measurement results as

$$A_{\Phi,N}\boldsymbol{x} = H^{ms} \quad (15)$$

Theoretically, this system of equations is highly underdetermined, because $\boldsymbol{x}$ is the continuous spatial channel response spanning all angles. However, due to the inherent sparsity of practical 60 GHz channels, we can approximate $\boldsymbol{x}$ by recovering a few dominating spatial angles from the measurement matrix H^{ms} and beamforming matrix $A_{\Phi,N}$. More specifically, we can apply compressive sensing theory [37] and formulate the recovery problem as an optimization problem:

$$\min \|\boldsymbol{x}\|_1, \text{subject to constraint (15)} \quad (16)$$

The objective is convex and constraint is affine, so the problem can be easily solved using any convex optimization algorithm. Resulting solution vector, denoted by $\tilde{\boldsymbol{x}}$, provides an estimation of the AoA pattern. Then the client can find the index of the best beam to listen to, *i.e.*, the beam that produces the strongest channel when combining the antenna gain and channel gain:

$$\tilde{u} = \arg \max_{u=1,\ldots,U} \sum\nolimits_{n=1}^{N} A(u,n)\tilde{x}(n) \quad (17)$$

It is noteworthy that unlike prior work in directional ND [9]–[12], CSBP does not assume a specific geometry for the beams. It directly models the codebook-based beam generation from a phased-arrays, which results in irregular-shaped beams.

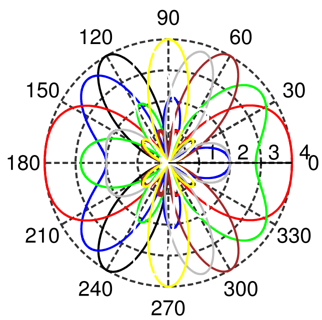

Fig. 14. A practical phased-array creates irregular beam patterns.

D. Maximum Distance Based Beam Prediction

The compressive sensing optimization may over-fit the spatial channel profile $\boldsymbol{x}$ when the number of measurements is insufficient. As a result, the best beam $\tilde{u}$ predicted by CSBP may fall in the set of beams Φ that the client already tried. In such cases, FastND invokes its MDBP module, which reselects a beam instead—a beam most orthogonal to those in Φ. The intuition is that such beams add the richest information about the channel. By accumulating such complementary information, MDBP helps the CSBP to quickly converge to an accurate channel profile estimation, and hence a small ND delay.

To determine the "orthogonality" between beams, MDBP needs to introduce a distance metric. Different from the ideal homogeneous fan-shape beams assumed in previous works, practical beams generated by a phased-array are heterogeneous in terms of spanning-area and beam strength. A representative example is shown in Fig. 14, which includes 8 beam patterns (Each beam has a different color) from a simulated 60 GHz phased array with 4 antenna elements (linearly spaced with half wavelength) following the standard practice of codebook design [4], [38]. As a result, we cannot define the distance between two beams simply as index difference or angular distance.

Instead, in MDBP, we define $\Delta_{i,j}$, the distance between beam i and beam j, as the l_1 norm of the difference between the corresponding beamforming gain vectors.[5] Formally,

$$\Delta_{i,j} = \|A(i,:) - A(j,:)\|_1 \tag{18}$$

Combined with Eq. (12), this definition takes into account both the amplitude and angular discrepancy, and thus can accurately reflect the difference between irregular-shaped beams. Based on this pairwise distance definition, we further define the distance between a beam u and the *set of previously-used beams* Φ as,

$$\Delta_{i,\Phi} = \min_{j\in\Phi} \|A(i,:) - A(j,:)\|_1 \tag{19}$$

Then, if a beam $\tilde{u}$ has the largest distance with Φ,

$$\tilde{u} = \arg \max_{u=1,\ldots,U} \Delta(u,\Phi) \tag{20}$$

we say that it is most orthogonal to already used beams in Φ. So, MDBP will dictate the client to listen to $\tilde{u}$.

[5]We have also tried the l_2 norm, but its performance (*i.e.*, the final ND delay) is worse. Intuitively, l_1 grows at the same rate in all directions for a beam pattern, but l_2 amplifies the impact of directions with larger discrepancy. In the beam distance computation, all directions should be of equal importance, thus l_1 shall be a more appropriate norm. A more rigorous and theoretical analysis is interesting, but it is out the scope of the work.

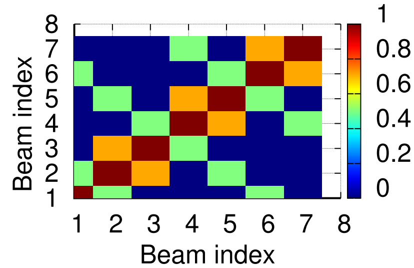

Fig. 15. Cross correlation of beams.

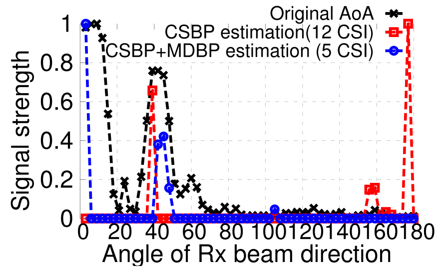

Fig. 16. Comparison of estimated AoA with original AoA.

E. Summary and Showcase

Algorithm 1 summarizes FastND's detailed operations involving both CSBP and MDBP. As a case study to testify the algorithm, we set up one 60 GHz link using our testbed in an office environment. The PCP uses a 180° quasi-omni antenna, and the client emulates a phased array with 64 antenna elements (Sec. VI details the emulation). We first measure the AoA pattern (*i.e.*, $\boldsymbol{x}$) by rotating a ultra-narrow 3.4° antenna at client's position. Then, the client runs FastND using CSBP alone, and CSBP plus MDBP, respectively. We make a few observations from the results (Fig. 16): *(i)* FastND, powered with the active learning, can indeed recover AoA shape with very few number of measurements, *i.e.*, 5 measurements in this case, while the passive error-and-trail methods require $\frac{180}{3.4} \approx 53$ measurements to scan all directions. Although the recovered AoA shape differs from the original one, the *angular direction* where the AoA peaks is very close when using CSBP&MDBP—this is exactly needed by FastND to guide the client to listen to the direction with best signal strength. *(ii)* CSBP alone requires a large number of measurement to recover the AoA, but the recovery is still not as accurate as CSBP&MDBP, which verifies the necessity to integrate both modules.

F. Computational Overhead

The computational overhead of FastND is negligible due to the following reasons. *(i)* Although a plain convex optimization problem may require $O(N^3)$ to solve, a compressive sensing optimization convex problem like (16) inside CSBP has a much lower computational complexity of $O(Nlog(K))$, by leveraging the sparsity feature of compressive sensing problems [37]. Here N is the dimension of variable vector $\boldsymbol{x}$, and K is the number of non-zero entries in the vector. *(ii)* The most time consuming part within MDBP lies in the beam distance computation (Eq. (18)), which has a complexity of $O(U^2N)$ (U is the number of beams). But the computation can be done offline, because the beam weights are fixed and thus the beam distances won't change at run time. The runtime MDBP only involves basic arithmetic operations (19) and (20). *(iii)*: FastND runs once every Beacon Interval (BI), and an BI typically lasts 100ms, which provides sufficient time for the computation.

G. Handling Practical Issues

Phase Non-Coherence: In CSBP, specifically in Eq. (15), we assumed that perfect phase of the measured channel coefficient H^{ms} is available. However, the channel measurements are incoherent and the phase gets corrupted due to carrier frequency offset (CFO), on practical COTS mmWave radios [39], [40]. To handle the problem, we adapt the original FastND and design a non-coherent FastND, which uses only amplitudes of the measured channels. In particular, we adapt Eq. (16) as follows,

$$\min \|\boldsymbol{x}\|_1 + \|(|A_{\Phi,N}\boldsymbol{x}| - |H^{ms}|)\|_2 \tag{21}$$

We solve the non-linear optimization problem using the Nelder-Mead Simplex method [41] to derive an optimal $\tilde{\boldsymbol{x}}$. Except such process of deriving $\tilde{\boldsymbol{x}}$, non-coherent FastND operates exactly the same with FastND following Alg. 1.

Algorithm 1 FastND Algorithm Running on a Client

INPUT: The set of all beams $A_{U,N}$
OUTPUT: Decoded Beacon
Φ = [] /*The set of used beams*/
H = [] /*The set of channel measurements */
Select a random beam u from A
/* Iterative loop until a beacon is decoded*/
while TRUE
    Perform receive beamforming with beam u
    if can decode a beacon
        return the beacon
    endforif
    Distill the resulting channel $h^{ms}(u)$
    $\Phi = \Phi \cup \{u\}$ /*Update the used beam set*/
    $H = H \cup \{h^{ms}(u)\}$ /*Update the channel set*/
    /*Compute the next promising beam using CSBP */
    $\tilde{u}$ = CSBP(A, Φ, H)
    if $\tilde{u} \in \Phi$
    /*Compute the next promising beam using MDBP */
        $\tilde{u}$ = MDBP(A, Φ)
    endforif
    $u = \tilde{u}$ /*Loop with the new beam*/
endforwhile

We illustrate the impact of phase non-coherence by comparing the performance of FastND and its non-coherent version using an experiment in Fig. 17. For the case, we use a 32-element linear phased-array and each antenna element is controlled by a 2-bit phase shifter. We plot the number of beams required by exhaustive ND (exND), FastND and non-coherent FastND for 30 links (setup details in Sec. VI-B.1). We have two observations: *(i)* The lack of phase information indeed has an adverse effect on ND. In particular, FastND uses only 2.5 beams in average for successful ND, while its non-coherent counterpart requires 4.5 beams. *(ii)* Despite of the adverse effect, non-coherent FastND is still much faster than exND. Specifically, non-coherent FastND only costs 25.8% beams compared with exND. Moreover, plenty of experiments under various scenarios in Sec. VI validate that FastND, after taking into account the effect of phase non-coherence, still significantly outperforms the state-of-the-aft.

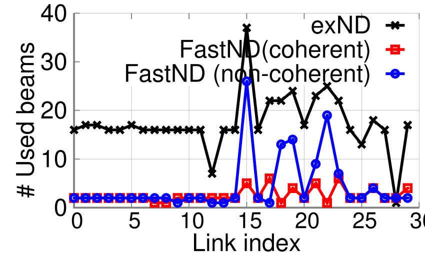

Fig. 17. Impact of phase non-coherence.

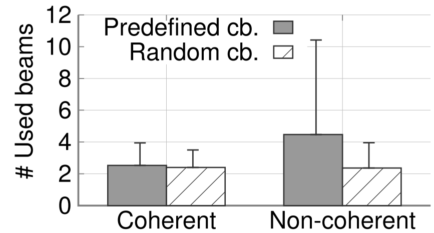

Fig. 18. Impact of random codebook.

Beam Pattern Randomness: FastND re-uses the existing codebook, *i.e.*, it samples the channel using pre-defined beam patterns defined in the codebook. The beam patters are designed with specific objectives (*e.g.*, high-directionality, coverage maximization), rather than being generated randomly [4], [38]. On the other hand, random sampling is required before applying compressive sensing, or restricted isometry property (RIP) should be satisfied, *i.e.*, the sampling matrices should be nearly orthonormal, at least when operating on sparse vectors [42].

We find that the used codebook in FastND, though not fully random, still satisfies the RIP to a certain extent. In particular, the beam weight for the n_{th} antenna element of the u_{th} beam, $w(n,u)$, is defined as follows [38],

$$w(n,u) = j^{\left\lfloor \frac{(n-1)\times \mod((u-1)+U/4,\,U)}{\frac{U}{4}} \right\rfloor} \tag{22}$$

where $j \triangleq \sqrt{(-1)}$, U is the total number of beam patterns. From Eq. (22), the legitimate choice for each $w(n,u)$ are limited to {1, j, -1, -j}, which can be realized by a 2-bit phase shifter. Said differently, $w(n,u)$ has been quantized by the coarse-grained phase shifter, which introduces randomness. Moreover, the generated beam patterns from Eq. (22) are nearly orthogonal. For instance, we generate 8 beams for a 4-element phased array, and plot the cross-correlation of the beams in Fig. 15. We can observe that the correlation matrix is an approximated block-diagonal matrix, which indicates good orthogonality.

We then compare the ND delays under our codebook with that under a purely random codebook. To generate random codebook, the beam weight of each antenna element are chosen randomly from {1, j, -1, -j}. We plot the average and std. of the number of used beams over 30 static links (setup details in Sec. VI-B.1) in Fig. 18 and have two insights: *(i)* For coherent FastND, random codebook doesn't have substantial effect, which implies that the predefined codebook already suffices for coherent compressive sensing. *(ii)* In contrast, random codebook saves about 44.4% beams in non-coherent case, *i.e.*, the impact of randomness becomes much more significant when the phase information is missing. Note that

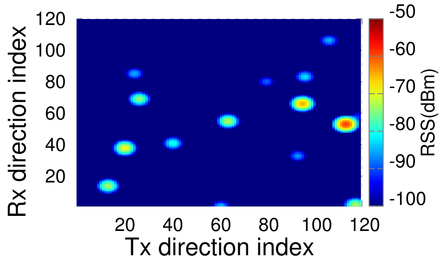

Fig. 19. Example of simulated RSS pattern.

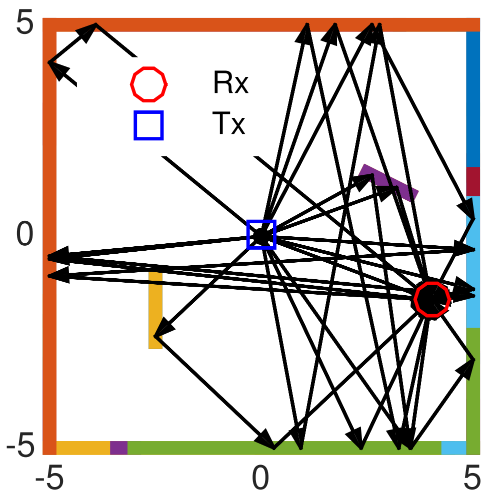

Fig. 20. Ray-tracing environment setup.

through random codebook helps to speed up ND, it's wider beam patters bear the disadvantages of low-directionality and more interference [39]. In this work, we still use the predefined codebook generated by Eq. (22), and leave the integration of random codebook for future exploration.

VI. Evaluation

A. Evaluation Methodology

We have implemented FastND on the baseband processing unit of our 60 GHz software radio testbed. To evaluate FastND, we would ideally need a reconfigurable phased-array antenna, but COTS 60 GHz phased-arrays are usually built on chip and do not allow access to low-level information. So we follow a classical approach to emulate phased-array beamforming [43] on the testbed. Specifically, we use linear phased-arrays with six kinds of dimensions, *i.e.*, the number of antenna elements increases from 16 to 96 with steps of 16. Each antenna element is controlled by a 2-bit phase shifter, *i.e.*, the phase shift values are $\{0, \pi/2, \pi, 3\pi/2\}$. We generate antenna gain pattern of a linear phased-array using 802.11ad codebook, and then convolve the antenna gain pattern of the phased-array antenna, with the spatial channel response (*i.e.*, AOA pattern) between the Tx and Rx. The AoA pattern is measured by sweeping a ultra-narrow horn antenna ($3.4°$) at both side, using a programmable motion control system.

Our 60 GHz platform plus the phased-array emulator can directly run ND algorithms for static links. Yet, emulating a mobile link requires us to collect the AoA trace at every timestamp, which is infeasible for the motion control system. Thus, for dynamic links, we develop a ray-tracing simulator to generate the fine-grained AoA trace. Ray-tracing [44] has shown to be an accurate way to simulate millimeter wave propagation in both indoor and outdoor environment. In particular, Neekzad *et al.* [45] demonstrate that AoA patterns from ray-tracing is reasonably close to real measurement in both LOS and NLOS scenarios. We develop our 60 GHz ray-tracing simulator following [46], which models the radio propagation pathloss, reflection and refraction effects in a 2D plane using image-based methods. Fig. 20 depicts the ray-tracing setup of a living room environment. It shows an example of propagation paths from an PCP (located at the origin) to a client (located at coordinate [4, -1.5]), which are reflected from two plastic reflectors and wooden walls. The corresponding RSS distribution is plotted in Fig. 19, and the channel sparsity is similar as in our measurement.

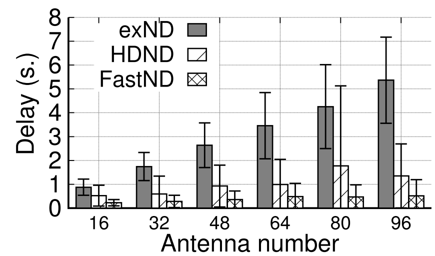

Fig. 21. Static links: average delay.

We have also implemented the 802.11ad's ND related TDMA-based MAC components on the 60 GHz testbed, including packetization with preambles, beacon sweeping, inter-frame spacing, *etc.* We use a virtual clock to time all protocol actions, with 802.11ad's default timing parameters. With this setup, we investigate FastND's performance under various practical factors, including mobility, blockage, link distance and surrounding environment. Almost all experiments run on our 60 GHz platform plus the phased-array emulator, except the mobility experiment that uses ray-tracing simulation.

We evaluate the performance of FastND against the exhaustive ND in 802.11ad, and the state-of-the-art Hunting based Directional Neighbor Discovery (HDND) proposed in [7]. In HDND, Tx and Rx node continuously sweep their directional beams towards opposite directions (*e.g.*, clockwise for Tx and counter-clockwise for Rx), until a beacon from the Tx is received and decoded at the Rx side. Using specifically-designed rotating strategy (under constraint of rotating speed and beamwidth), HDND is proved to achieve successful ND within a delay bound. It is noteworthy that HDND [7] requires to differentiate the beam steering speed of Tx and Rx, which incurs more protocol complexity and does not follow the 802.11ad standards. In addition, in spite of many alternative 60 GHz ND protocols in literature (see [15], [16], [18], and references therein), they assume simplified beam patterns that are inconsistent with those generated by phased-arrays, or require major protocol modifications [7], [8] that are incompatible with the 802.11ad standard. In contrast, FastND operates in full compatibility with 802.11ad and takes into account practical factors such as code-book beamforming of phased-array and signal propagation characteristics derived from testbed measurements.

B. Evaluation Results

1) Static Links: We first compare FastND against exND and HDND using 30 static links randomly deployed in an 10×10m office environment, as illustrated in Fig. 20. The average ND latency results are plotted in Fig. 21

We have two major observations: *(i)* exND latency scales dramatically with the number of antenna elements (*i.e.*, N_R) on the phased-array. This poses significant challenge to the

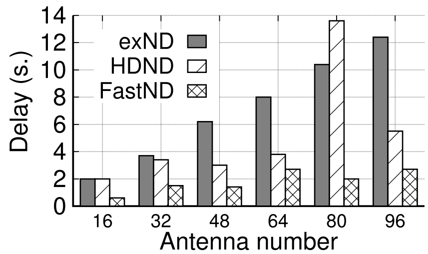

Fig. 22. Static links: worst delay.

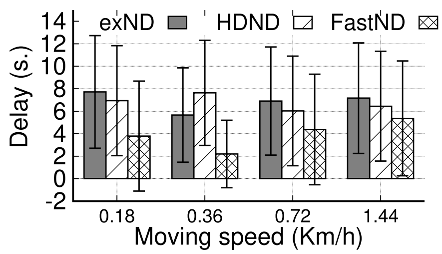

Fig. 23. Average delay of mobile links.

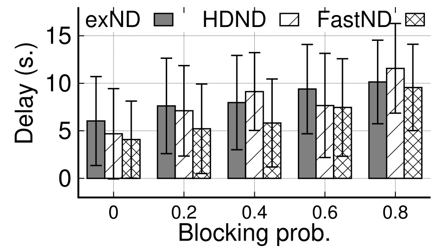

Fig. 24. Ave. delay with human blockage.

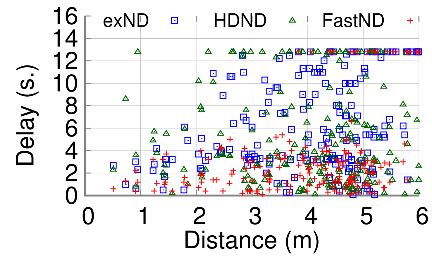

Fig. 25. Delay distribution of 200 instances.

practical use of 802.11ad, since hundreds of antenna elements are expected to be packed in a small-sized phased array to support multi-Gbps data rate [3]. In contrast, FastND's delay grows very slowly and keeps under 1 second on average, so the latency gap with 802.11ad increases as N_R grows. When $N_R = 96$, the mean latency under FastND is only about $\frac{1}{10}$ of that under IEEE 802.11ad. The result clearly validates the advantage of active learning of FastND, while accumulating more information, FastND can predict and converge to the optimal beam direction quickly, in contrast to the slow "blind" trail-and-error exND. *(ii)* Both HDND and FastND significantly reduce the ND latency. In particular, the average delay of HDND is 37.42% of exND, while that of FastND is only 15.02% of exND. Moreover, we find that HDND exhibits much more variance than FastND. From the worst-case delay given in Fig. 22, HDND's delay can exceed exND and reaches as high as 13.6s (*e.g.*, when $N_R = 96$), while that of FastND is bounded within 3s. The result corroborates the findings in [7], *i.e.*, due to the deterministic nature of HDND and 60 GHz spatial channel sparsity, it may take very long time for the clockwise-rotating beams of Tx and the counter-clockwise-rotating beams of Rx rendezvous at a path with RSS strong enough.

The worst-case latency shows a similar trend, so we omit the discussion here.

2) Mobile Links: We set up mobile links by allowing the client to walk through 100 different trajectories within a 10m×10m office. It is expected that the highly directional 60 GHz links is meant to be used in quasi-static scenarios such as hand movements or postural changes when a static person is holding a client device, so here we mainly focus on speed from 0.18 to 1.44Km/h [6]. A user starts beam scanning at a random orientation at each location along its moving trajectory.

We plot the average and standard deviation of mobile links' delay in Fig. 23, from which we have two major observations. *(i)* Different from static scenarios, mobile link delays are much more dynamic, for all three methods. In particular, there is *not* a monotonic increasing or decreasing relationship between ND delay and moving speed. The reason, by analyzing the mobility traces, is that mobile trajectories in practice is usually consisted of both homogeneous and heterogeneous locations, as the reflecting environment keep evolving during mobility [36]. Since moving speeds have opposite effect on these two cases (as analyzed in Sec. IV), the mixed effect shows non-monotonic trend. *(ii)* Despite of large fluctuation, FastND is able to control the delay inside a small range, compared with exND and HDND. On average, FastND's delay is about $\frac{1}{2}$ of that of exND and HDND.

3) Human Blockage: Human blockage is a major hindrance to the highly directional 60 GHz links, especially for indoor deployments that often have heavy human activity like people sitting or moving in the conference room. We model blockage by assigning a blockage probability for each beam direction, following [47]. Fig. 24 shows the ND delays under different levels of blockage. We observe that human blockage prolongs the discovery latency for all the three methods, since more effective directions are blocked and less CSI can be collected to recover the channel profiles. Though FastND's gain becomes marginalized with increasing blockage severity, it still outperforms exND and HDND in the common cases.

4) Link Distance: We next evaluate the impact of link distance. 200 client locations are randomly generated within the ray-tracing environment (Fig. 20), with link distance distributed within $[0, 6]$ meters. Fig. 25 plots the scatter diagram over their ND delays. We observe that: *(i)* both exND and HDND delay show a drastic variation beyond a short distance of 2 meters. Half of the clients experience large delays within $[6, 13]$ seconds, implying very low quality of experience. *(ii)* In contrast, the delays of most clients under FastND are less than 5 seconds, except for a small number of special locations where the channel condition is so bad that no useful CSI can be collected.

5) Environment Sparsity: In practice, 60 GHz radios can be deployed in scatter-rich environment like conference rooms, or open space like corridors or even outdoor. We model such diverse surrounding environment by varying the number of reflectors in the ray-tracing setup, and then we derive the ND delays of the three comparing methods. From Fig. 26, we have two major observations: *(i)* all methods cost more time in sparse environment (with few reflectors), since the η-factor is very small there. For example, the delay of exND decreases from 8.8s to 2.6s when the number of AoA clusters increases from 1 to 7. *(ii)* Despite of the same trend,

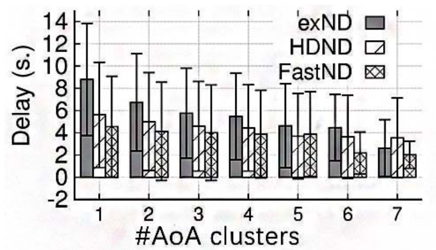

Fig. 26. Ave. delay under various multipaths.

FastND outperforms exND and HDND under each setting. In particular, the advantage of FastND over HDND becomes more significant under scatter-rich environment, since FastND can quickly converge to one of many effective directions, while HDND still requires many full rounds of clockwise/counter-clockwise rotation to achieve effective Tx/Rx rendezvous.

To sum up, FastND outperforms exND and HDND across a wide range of practical scenarios. More importantly, the advantage of FastND is more significant under crucial scenes, such as massive-element phased arrays, long link distances and open environment with a limited number of reflectors. In these scenarios, FastND achieves ND using only $\frac{1}{10}$ to $\frac{1}{2}$ of the time compared with the other two methods on average.

VII. Discussion and Future Work

Effect of AP/PCP Deployment

In this work, we examine the ND delay under random AP/PCP deployment. An in-depth research is to optimize the AP/PCP's position to facilitate ND. Our recent work [48] shows that the performance of 60 GHz networks is highly sensitive to environment structure and reflectivity, then it places AP/PCP in an optimal way so as to maximize network coverage and link robustness. Similarly, joint optimization of AP/PCP deployment and adaptive beam steering shall further speed up neighbor discovery, which is left for future work.

Element Gain

For simplicity, we ignore the element gain of phased array, but assume each element to be omni-directional. From Fig. 14, we can see that the beam patterns together cover a full circle of $360°$. However, element gains exist on practical phased arrays. For instance, front-to-back gain ratio is actually very large on COTS phased arrays such as Qualcomm [49] and Intel radios [50]. Fortunately, FastND can easily incorporate the element gain. In particular, suppose the element gain is $e(i, \theta)$ for the i_{th} element on direction $\theta(n)$, we can modify the original Eq. (12) to be

$$A(u,n) = \sum_{i=1}^{N_R} w(u,i)e(i,\theta(n))exp(j2\pi idcos(\theta(n))/\lambda) \tag{23}$$

Despite such modification, other designs of FastND keep the same, and in principle FastND will work with element gain. However, the problem is that element gain values proprietary to manufacturer are not open to researchers, as far as we know. Element gain measurement/estimation should be an interesting problem for future research.

ND in 3-D Space

In this work we focus on ND in 2-D space. ND shall become much more complicated in 3-D space, because beam scanning along elevation and azimuth dimensions are coupled together, which leads to a high scanning overhead at $O(N^2)$ order (here N is the total number of beam directions along each dimension). However, a recent work [51] shows that the two dimensions can be decoupled by exploiting a unique interaction between 3-D spatial channel profile and beam patterns. In consequence, the 3-D beam scanning can be degraded as a limited number of 2-D beam scanning, which is of $O(N)$ scanning overhead. The work in [51] focuses on fast beam tracking (*i.e.*, after successful neighbor discovery) for emerging mobile application including untethered virtual Reality and miracst, in which 3-D movement or rotation are heavily involved. The principle there can also be applied to speed up neighbor discovery in 3-D space, which is left for future exploration.

VIII. Related Works

60 GHz Networking

Amid the commercialization of 60 GHz devices [49], [50] and maturity of protocol standardization [4], [5], 60 GHz networking has been attracting more and more attention. Based on prior theoretical works (see [29] and references therein), recently much research effort has been devoted to solve practical challenges when implementing both outdoor [3], [21] and indoor 60 GHz multi-Gbps wireless networks. In particular, many solutions are proposed to combat adverse effect of blockage or mobility [36], [48], [52]–[56]. In addition, 60 GHz radios are widely considered as the enabling technology for the emerging untethered VR [57], 4K/8K Miracast [58] and wireless data centers [22]. Recent works [13], [22], [59]–[62] start to design system-level solutions to guarantee application performance. While these works focus on improving 60 GHz performance after link establishment (*i.e.*, after ND procedures finishes), our work aims to speed up the initial ND process.

ND in 60 GHz Networks

Previous study [7], [8], [15], [16], [18] has analyzed or simulated ND performance in 60 GHz networks and has attempted to improve the ND efficiency, but they assume simplified antenna beams, and ignore unique properties of 60 GHz spatial channel profile and practical phased-array antenna patterns.

It is possible to conduct ND via the 2.4 GHz band using omni-directional antennas, while transmitting data through the 60 GHz directional phased-array [19], [53]. Besides requiring extra radio hardware, a more serious concern is that channel characteristic of 2.4 GHz and 60 GHz differ drastically: two radios that see each other at 2.4 GHz may not be able to establish connections at 60 GHz.

Compressive sensing has been used in 60 GHz beam searching, or in a more general setting of multipath channel recovery [14], [33], [34], [39]. While FastND is inspired by these works, it differs in several aspects. *First*, previous works focus on the compressive channel recovery problem exclusively, and assume enough channel state information has been aggregated at once, which is not feasible in ND state [39]. *Second*, they do not address the practical impacts of channel sparsity and mobility on the latency of ND protocols. *Third*, performance of channel recovery is evaluated theoretically [34] or in an 8 GHz anechoic chamber [33], whose channel profiles differ significantly from 60 GHz multipath environment.

Hierarchical approach is typically used to speed up channel training in 60 GHz networks, where sector-level sweeping is firstly used to derive a coarse-grained beam-forming sector, and then a fine-grained beam searching is used to refine the sector [4], [6]. Note that such hierarchical approach requires Tx/Rx to exchange control messages so as to progressively narrow down searching space, which is not feasible in ND stage where Tx/Rx link has not been established yet. Moreover, hierarchical approach is of low efficiency as the population of users grows [39].

Active Learning

As a general principle, active learning has been applied in machine learning, statistical inference, mobile sensing, and *etc.* (see [17], [63], [64]). FastND is inspired by these works, but solves unique challenges proprietary to 60 GHz mmWave radios. *First*, to bootstrap the learning process, FastND exploits the coding/modulation redundancy in the preamble of 802.11ad frames to distill hidden channel information. *Second*, to derive next sampling operations, *i.e.*, the next most promising beam direction, FastND leverage both inherent spatial channel sparsity and practical ir-regular beam shape generated by phased array, to design two complimentary beam prediction mechanisms, in order to converge to the effective beam direction with limited information in ND stage.

ND in Omni-Directional Wireless Networks

In legacy Wi-Fi and cellular networks with persistently active PCP (or BS), ND is a straightforward procedure. The problem becomes complicated in self-organized wireless sensor networks, where sensor nodes are activated intermittently to save power. Two sensor nodes can discover each other only when their wake up periods overlap. An ideal ND protocol as such should make an optimal trade-off between energy efficiency and rendezvous latency. To this end, many ND protocols have been proposed (see a survey [65]). These protocols can be considered as ways to intentionally *sparsify* the temporal schedule to save energy. 60 GHz ND faces a completely different challenge, *i.e.*, spatial channel sparsity, which depends on the radio environment and cannot be manipulated by time-domain scheduling.

ND in Low-Frequency Directional Wireless Networks

In directional wireless networks, communicating parties must agree on where and when to point their directional beams. Extensive work has developed algorithmic solutions to speed up ND for low-frequency directional WiFi and cellular networks (see [9]–[12], and the references therein). These algorithms typically assume a fan-shaped antenna beam pattern, and design geometry-based beam switching methods to allow the transmitter and receiver beam to fast rendezvous. The underlying assumption are: the nodes fall in LOS, beam patterns are regular, channel profile (AoA pattern) overlaps exactly with beam pattern, *etc.* These assumptions are no longer valid in 60 GHz networks (Sec. I), and hence a new and principled design is needed.

IX. Conclusion

We have presented FastND, a practical mechanism that enables fast neighbor discovery for 60 GHz wireless networking. FastND is built on our insights from test-bed measurements and theoretical analysis, which capture the impact of the inherent sparsity characteristic of 60 GHz channels. FastND incorporates a compressive channel recovery and an adaptive beam selection technique to speed up the neighbor discovery process. Experiment results show that FastND can achieve delay as low as $\frac{1}{2}$ to $\frac{1}{10}$ of that by existing de-facto protocol under adverse network scenarios.

References

[1] Qualcomm Inc. (2014). *The 1000x Data Challenge*. [Online]. Available: https://www.qualcomm.com/1000x

[2] (2014). *The 5G Infrastructure Public Private Partnership*. https://5g-ppp.eu/

[3] T. Rappaport *et al.*, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.

[4] *Enhancements for Very High Throughput in the 60 GHz Band*, IEEE Standards 802.11ad-2012, IEEE Standards Association, 2012.

[5] *Millimeter-Wave-Based Alternate Physical Layer Extension*, IEEE Standards 802.15.3c-2009, IEEE Standards Association, 2009.

[6] S. Sanjib, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 GHz indoor networking through flexible beams: A link-level profiling," in *Proc. ACM SIGMETRICS*, 2015, pp. 71–84.

[7] Y. Wang, S. Mao, and T. S. Rappaport, "On directional neighbor discovery in mmWave networks," in *Proc. IEEE ICDCS*, Jun. 2017, pp. 1704–1713.

[8] L. Chen, Y. Li, and A. V. Vasilakos, "Oblivious neighbor discovery for wireless devices with directional antennas," in *Proc. IEEE INFOCOM*, Apr. 2016, pp. 1–9.

[9] S. Vasudevan, J. Kurose, and D. Towsley, "On neighbor discovery in wireless networks with directional antennas," in *Proc. INFOCOM*, 2005, pp. 2502–2512.

[10] Z. Zhang and B. Li, "Neighbor discovery in mobile ad hoc self-configuring networks with directional antennas: Algorithms and comparisons," *IEEE Trans. Wireless Commun.*, vol. 7, no. 5, pp. 1540–1549, May 2008.

[11] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit, "Ad hoc networking with directional antennas: A complete system solution," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 3, pp. 496–506, Mar. 2005.

[12] F. Tian, B. Liu, H. Cai, H. Zhou, and L. Gui, "Practical asynchronous neighbor discovery in ad hoc networks with directional antennas," *IEEE Trans. Veh. Technol.*, vol. 65, no. 5, pp. 3614–3627, May 2015.

[13] T. Wei and X. Zhang, "Pose information assisted 60 GHz networks: Towards seamless coverage and mobility support," in *Proc. ACM MobiCom*, 2017, pp. 42–55.

[14] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive millimeter-wave sector selection in off-the-shelf IEEE 802.1 and devices," in *Proc. ACM CoNEXT*, 2017, pp. 414–425.

[15] X. An, R. V. Prasad, and I. Niemegeers, "Impact of antenna pattern and link model on directional neighbor discovery in 60 GHz networks," *IEEE Trans. Wireless Commun.*, vol. 10, no. 5, pp. 1435–1447, May 2011.

[16] J. Ning, T.-S. Kim, S. V. Krishnamurthy, and C. Cordeiro, "Directional neighbor discovery in 60 GHz indoor wireless networks," in *ACM MSWiM*, 2009, pp. 365–373.

[17] B. Settles, "Active learning literature survey," Dept. Comput. Sci., Univ. Wisconsin-Madison, Madison, WI, USA, Tech. Rep. 1648, 2010.

[18] G. M. Ölçer, Z. Genç, and E. Onur, "Smart neighbor scanning with directional antennas in 60 GHz indoor networks," in *Proc. IEEE PIMRC*, Sep. 2010, pp. 2393–2398.

[19] H. Park, Y. Kim, T. Song, and S. Pack, "Multiband directional neighbor discovery in self-organized mmWave ad hoc networks," *IEEE Trans. Veh. Technol.*, vol. 64, no. 3, pp. 1143–1155, Mar. 2015.

[20] T. S. Rappaport, R. W. Heath, Jr., R. C. Daniels, and J. N. Murdock, *Millimeter Wave Wireless Communications*. Englewood Cliffs, NJ, USA: Prentice-Hall, 2014.

[21] Y. Zhu *et al.*, "Demystifying 60GHz outdoor picocells," in *Proc. ACM MobiCom*, 2014, pp. 5–16.

[22] X. Zhou *et al.*, "Mirror mirror on the ceiling: Flexible wireless links for data centers," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 4, pp. 443–454, Oct. 2012.

[23] Rice University. (2013). *Wireless Open-Access Research Platform*. [Online]. Available: http://warp.rice.edu/trac/wiki

[24] Pasternack. (2014). *60 GHz Transmit/Receive Development Systems*. [Online]. Available: http://www.pasternack.com

[25] (2014). *Axis360 System*. [Online]. Available: http://cinetics.com/two-axis360/

[26] T. S. Rappaport *et al.*, "Overview of millimeter wave communications for fifth-generation (5G) wireless networks—with a focus on propagation models," *IEEE Trans. Antennas Propag.*, vol. 65, no. 12, pp. 6213–6230, Dec. 2017.

[27] A. I. Sulyman, A. Alwarafy, G. R. MacCartney, T. S. Rappaport, and A. Alsanie, "Directional radio propagation path loss models for millimeter-wave wireless networks in the 28-, 60-, and 73-GHz bands," *IEEE Trans. Wireless Commun.*, vol. 15, no. 10, pp. 6939–6947, Oct. 2016.

[28] C. Gustafson, K. Haneda, S. Wyne, and F. Tufvesson, "On mm-wave multipath clustering and channel modeling," *IEEE Trans. Antennas Propag.*, vol. 62, no. 3, pp. 1445–1455, Mar. 2014.
[29] H. Xu, V. Kukshya, and T. S. Rappaport, "Spatial and temporal characteristics of 60-GHz indoor channels," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 3, pp. 620–630, Apr. 2002.
[30] T. Nitsche, G. Bielsa, I. Tejado, A. Loch, and J. Widmer, "Boon and bane of 60 GHz networks: Practical insights into beamforming, interference, and frame level operation," in *Proc. ACM CoNEXT*, 2015, Art. no. 17.
[31] P. F. M. Smulders, "Statistical characterization of 60-GHz indoor radio channels," *IEEE Trans. Antennas Propag.*, vol. 57, no. 10, pp. 2820–2829, Oct. 2009.
[32] T. Camp, J. Boleng, and V. Davies, "A survey of mobility models for ad hoc network research," *Wireless Commun. Mobile Comput.*, vol. 2, no. 5, pp. 483–502, 2002.
[33] D. E. Berraki, S. M. D. Armour, and A. R. Nix, "Application of compressive sensing in sparse spatial channel recovery for beamforming in mmWave outdoor systems," in *Proc. IEEE WCNC*, Apr. 2014, pp. 887–892.
[34] W. U. Bajwa, J. Haupt, A. M. Sayeed, and R. Nowak, "Compressed channel sensing: A new approach to estimating sparse multipath channels," *Proc. IEEE*, vol. 98, no. 6, pp. 1058–1076, Jun. 2010.
[35] R. Zhou, Y. Xiong, G. Xing, L. Sun, and J. Ma, "ZiFi: Wireless LAN discovery via ZigBee interference signatures," in *Proc. ACM MobiCom*, 2010, pp. 49–60.
[36] A. Zhou, X. Zhang, and H. Ma, "Beam-forecast: Facilitating mobile 60 GHz networks via model-driven beam steering," in *Proc. IEEE INFOCOM*, May 2017, pp. 1–9.
[37] Y. C. Eldar and G. Kutyniok, Eds., *Compressed Sensing: Theory and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2012.
[38] J. Wang *et al.*, "Beamforming codebook design and performance evaluation for 60 GHz wideband WPANs," in *Proc. IEEE Veh. Technol. Conf. Fall*, Sep. 2009, pp. 1–6.
[39] M. E. Rasekh, Z. Marzi, Y. Zhu, U. Madhow, and H. Zheng, "Noncoherent mmWave path tracking," in *Proc. ACM HotMobile*, 2017, pp. 13–18.
[40] Y. Zhu, Y. Zhu, B. Y. Zhao, and H. Zheng, "Reusing 60 GHz radios for mobile radar imaging," in *Proc. ACM MobiCom*, 2015, pp. 103–116.
[41] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the Nelder-Mead simplex method in low dimensions," *SIAM J. Optim.*, vol. 9, no. 1, pp. 112–147, 1998.
[42] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Comp. Rendus Math.*, vol. 346, nos. 9–10, pp. 589–592, May 2008.
[43] M. Park and P. Gopalakrishnan, "Analysis on spatial reuse and interference in 60-GHz wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 27, no. 8, pp. 1443–1452, Oct. 2009.
[44] V. Degli-Esposti *et al.*, "Ray-tracing-based mm-wave beamforming assessment," *IEEE Access*, vol. 2, pp. 1314–1325, 2014.
[45] B. Neekzad, K. Sayrafian-Pour, J. Perez, and J. S. Baras, "Comparison of ray tracing simulations and millimeter wave channel sounding measurements," in *Proc. IEEE PIMRC*, Sep. 2007, pp. 1–5.
[46] E. De Groot, T. Bose, C. Cooper, and M. Kruse, "Remote transmitter tracking with raytraced fingerprint database," in *Proc. IEEE MILCOM*, Oct. 2014, pp. 325–328.
[47] A. Maltsev *et al.*, *Channel Models for 60 GHz WLAN Systems*, IEEE Standard P802.11-09/0334r8, 2010.
[48] T. Wei, A. Zhou, and X. Zhang, "Facilitating robust 60 ghz network deployment by sensing ambient reflectors," in *Proc. USENIX NSDI*, 2017, pp. 213–226.
[49] Qualcomm Inc. (2015). *Qualcomm VIVE 802.11ad*. [Online]. Available: https://www.qualcomm.com/products/vive/11ad
[50] Intel. (2015). *Intel Tri-Band Wireless-AC 18260*. [Online]. Available: http://www.intel.com/content/dam/www/public/us/en/-documents/product-briefs/tri-band-wireless-ac17265-brief.pdf
[51] Z. Anfu *et al.*, "Following the shadow: Agile 3-D beam-steering for 60 GHz wireless networks," in *Proc. IEEE INFOCOM*, 2018, pp. 1–9.
[52] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, "BeamSpy: Enabling robust 60 GHz links under blockage," in *Proc. USENIX NSDI*, 2016, pp. 193–206.
[53] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: Mm-wave beam steering without in-band measurement," in *Proc. IEEE INFOCOM*, Apr./May 2015, pp. 2416–2424.
[54] M. K. Haider and E. W. Knightly, "Mobility resilience and overhead constrained adaptation in directional 60 GHz WLANs: Protocol design and system implementation," in *Proc. ACM MobiHoc*, 2016, pp. 61–70.
[55] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "WiFi-assisted 60 GHz wireless networks," in *Proc. ACM MobiCom*, 2017, pp. 28–41.
[56] A. Loch. H. Assasa, J. Palacios, J. Widmer, H. Suys, and B. Debaillie, "Zero overhead device tracking in 60 GHz wireless networks using multi-lobe beam patterns," in *Proc. ACM CoNEXT*, 2017, pp. 224–237.
[57] (2017). *Desperately Seeking Wireless: VR's Aiming to Cut the Cord*. [Online]. Available: https://www.cnet.com/news/vr-desperately-seeking-wireless/
[58] *Wi-Fi Display Technical Specification V1.2n*, Wi-Fi Alliance, Austin, TX, USA, 2011.
[59] O. Abari, D. Bharadia, A. Duffield, and D. Katabi, "Enabling high-quality untethered virtual reality," in *Proc. USENIX NSDI*, 2017, p. 49.
[60] W. Zhang *et al.*, "3D beamforming for wireless data centers," in *Proc. ACM HotNets*, 2011, Art. no. 4.
[61] Y. Zhu *et al.*, "Cutting the cord: A robust wireless facilities network for data centers," in *Proc. ACM MobiCom*, 2014, pp. 581–592.
[62] Y. Cui *et al.*, "Diamond: Nesting the data center network with wireless rings in 3D space," in *Proc. USENIX NSDI*, 2016, pp. 657–669.
[63] A. Singh, R. Nowak, and P. Ramanathan, "Active learning for adaptive mobile sensing networks," in *Proc. ACM IPSN*, Apr. 2006, pp. 60–68.
[64] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, "Active learning with statistical models," *J. Artif. Int. Res.*, vol. 4, no. 1, pp. 129–145, Mar. 1996.
[65] W. Sun, Z. Yang, X. Zhang, and Y. Liu, "Energy-efficient neighbor discovery in mobile ad hoc and wireless sensor networks: A survey," *IEEE Commun. Surveys Tuts.*, vol. 16, no. 3, pp. 1448–1459, 2014.

Anfu Zhou received the B.S. degree from the Renmin University of China in 2004 and the Ph.D. degree in computer science from the Institute of Computing Technology, Chinese Academy Sciences, in 2012. He is currently an Associate Professor with the School of Computer Science, Beijing University of Posts and Telecommunications. His research interests lie in mobile computing, wireless networking, and Internet of Things systems.

Teng Wei received the B.S. degree from Shanghai Jiao Tong University in 2013. He is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, University of Wisconsin–Madison. His research interests lie in wireless networks and mobile systems, with emphasis on ubiquitous wireless sensing for enabling new Internet-of-Things applications.

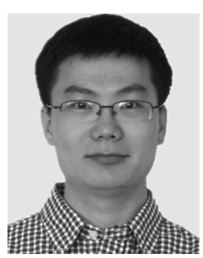

Xinyu Zhang received the Ph.D. degree from the University of Michigan in 2012. He was an Assistant Professor with the University of Wisconsin–Madison from 2012 to 2017. He is currently an Associate Professor with the Department of Electrical and Computer Engineering, University of California, San Diego. His research interests lie in wireless systems and ubiquitous computing. He was a recipient of the ACM MobiCom Best Paper Award in 2011, the NSF CAREER Award in 2014, and the Google Research Award in 2017 and 2018.

Huadong Ma received the Ph.D. degree in computer science from the Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, in 1995. He is currently a Chang Jiang Scholar Professor with the Beijing University of Posts and Telecommunications, Beijing, where he is also the Director of the Beijing Key Laboratory of Intelligent Telecommunications Software and Multimedia and the Executive Dean of the School of Computer Science. He has authored over 200 papers and four books. His research interests include multimedia system and networking, sensor networks, and Internet of Things. He was a recipient of the National Funds for Distinguished Young Scientists in 2009.