

Beam-forecast: Facilitating Mobile 60 GHz Networks via Model-driven Beam Steering

Anfu Zhou* Xinyu Zhang† Huadong Ma*

* Beijing Key Lab of Intell. Telecomm. Software and Multimedia, Beijing University of Posts and Telecomm.

† University of Wisconsin-Madison

Email: zhouanfu@bupt.edu.cn, xyzhang@ece.wisc.edu, mhd@bupt.edu.cn

Abstract—Low robustness under mobility is the Achilles’ heel of the emerging 60 GHz networking technology. Instead of using omni-directional antennas as in existing Wi-Fi/cellular networks, 60 GHz radios communicate via highly-directional links formed by phased-array beam-forming, so as to upgrade wireless link throughput to multi-Gbps. However, user motion causes misalignment between the Tx’s and Rx’s beam directions, and often leads to link outage. Legacy 60 GHz protocols realign the beams by scanning alternative Tx/Rx beams. But unfortunately this tedious process can easily overwhelm the useful channel time, leaving the Tx/Rx in misalignment most of the time during mobility. In this paper, we propose *Beam-forecast*, a novel model-driven beam steering approach that can sustain high performance for mobile 60 GHz links. *Beam-forecast* is built on the observation that 60 GHz channel profiles at nearby locations are highly-correlated. By exploiting this correlation, *Beam-forecast* can reconstruct the channel profile as the Tx/Rx moves, without explicit channel scanning. In this way, it can predict new optimal beams and realign links for mobile users with minimal overhead. We evaluate *Beam-forecast* using a reconfigurable 60 GHz testbed along with a trace-driven simulator. Our experiments demonstrate multi-fold throughput gain compared with state-of-the-art under various practical scenarios.

I. INTRODUCTION

60 GHz millimeter-wave (mm-Wave) radio, featured with a vast amount of spectrum resource and highly-directional links, is promised to overcome the spectrum crunch, and bring wireless networks to multi-Gbps era. With up to 7 Gbps data rate as specified in recent IEEE standards, mm-Wave technology will enable numerous ultra-high-speed wireless applications. Conventional usage of this technology mainly focuses on static or quasi-static scenarios including cellular back-haul [1], [2], wireless data center [3] and cable-replacement like wireless HDMI [4].

Beyond static networks, 60 GHz technology is anticipated to enable Gbps mobile access. *First*, both research and industry associations have endorsed the huge potential of 60 GHz picocells in providing outdoor mobile cellular access [5]–[7]. Similarly, the state-of-the-art IEEE 60 GHz MAC/PHY standard 802.11ad [8] advocates WiFi-like usage cases to provide indoor wireless access, where human mobility is pervasive. *Second*, new mobile applications of 60 GHz radios are rapidly emerging. For instance, 60 GHz radios are expected to eliminate a critical pain point in virtual reality (VR) applications [9], where massive amount of data should be transferred (traditionally through a cumbersome cable) in real-time from a media-rendering server to VR helmets.

However, to date, these usage scenarios remain in blueprint, because the robustness of 60 GHz mobile connections remains an open issue. The problem roots in two unique characteristics

of 60 GHz radios: (i) *High directionality*. 60 GHz radios face much stronger attenuation than existing Wi-Fi (26dB and 21.6dB worse when compared with 2.4 GHz and 5 GHz WiFi, respectively [10]). To remedy such strong attenuation, 60 GHz radios employ phased-array antennas to form highly directional beams. But high directionality results in high vulnerability to mobility — whenever the Tx/Rx moves, their beam directions may become misaligned, causing high risk of link outage. (ii) *Channel sparsity*. As reported widely by existing work [11]–[14], the 60 GHz spatial channel profile (SCP) is extremely sparse, *i.e.*, the Tx can reach the Rx only through a small number of line-of-sight or reflection paths. Note that the final channel quality is determined jointly by the beam in use and the underlying SCP. To achieve high link quality, the Tx/Rx need to find the beam that best matches SCP, *i.e.*, the beam direction that aligns with the strongest SCP paths. Searching for the optimal beam can entail huge overhead, especially considering the constantly changing SCP during mobility [7], [10], [15]–[17], and the large number of beam directions a 60 GHz radio can choose from.

Amid the lack of SCP knowledge, conventional solutions resort to extensive searching (as in 802.11ad [8]) or heuristic shortcut (*e.g.*, prioritizing pre-computed backup beam directions [17]). The former incurs formidable overhead and may leave the Tx/Rx in beam searching rather than data transmission most of the time [15]. The latter is far from optimum, as we will demonstrate through experiments (Sec. IV), because of the wild variation of SCP during mobility.

In this paper, we propose a new principled solution, called *Beam-forecast*, to sustain high performance for mobile 60 GHz links. *In contrast to the ‘blind’ trail-and-error beam searching, we design a model-driven method to directly predict the optimal beams*. The key insight is that SCPs at nearby locations are tightly correlated. We can leverage this correlation to reconstruct SCP at a new location, and then predict the optimal beam with minimal overhead.

To realize *Beam-forecast*, we propose two key techniques. *First*, we design a geometry model to quantitatively capture the SCP correlation. In particular, given the SCP at an initial location (also referred to as an *anchor point*), we can derive all SCP candidates for any nearby locations using the model. We emphasize that the spatial correlation does not necessarily imply the same SCPs among close-by locations. Thus, a simple spatial average may not work well. Instead, our model leverages the sparsity and pseudo-optical profile of 60 GHz channels to predict the angular channel response, and hence the location-specific channel quality. *Second*, we develop a reverse-




Fig. 1. WiMi 60 GHz software radio platform [18].




Fig. 2. Importance of beam alignment.




Fig. 3. Floor map of measurement environment.

engineering optimization to predict the optimal beam. Previous beam searching approaches [8] try all beams over unknown SCPs and use control/feedback frames to convey searching results — a tedious overhead that may overwhelm useful data transmission time. In contrast, Beam-forecast ‘virtually’ tries each candidate beam and predicts its quality without actual probing. Specifically, leveraging the correlation model, Beam-forecast first reconstructs an entire SCP based on the SCP measured by the current Tx/Rx beams (which is a partial observation of the SCP across all angles). It then predicts the link quality for different candidate beam directions, and selects the best one accordingly. Such a model-driven prediction is a pure computational process, which totally avoids the time-consuming beam scanning.

Beam-forecast can be seamlessly integrated into 802.11ad-like 60 GHz protocols, and substitute their original beam scanning operations so as to efficiently handle mobility. When a user moves too far away from the anchor point, the geometry model will lose accuracy since the multipath environment may have totally changed. In this case, the Tx/Rx need to reinitiate the 802.11ad beam scanning to get an accurate anchoring SCP, and deem this location as a new anchor point that benefits a new round of model-driven beam prediction. In other words, Beam-forecast progressively handles continuous mobility with trivial overhead by initiating beam scanning only when necessary.

We implement Beam-forecast on a custom-built 60 GHz platform, along with a trace-driven simulator. We evaluate Beam-forecast’s performance against the IEEE 802.11ad standard beam steering mechanism and a state-of-the-art PFR protocol [17]. Experimental results show that Beam-forecast can accurately predict the optimal beam with less than 5% of the 802.11ad beam scanning overhead. Overall, the optimal beam prediction and the saved scanning time translate to significant throughput gain, specifically, $8.18\times$ over 802.11ad and $4.27\times$ over PFR averaging over various networking scenarios. The advantage of Beam-forecast is most significant under crucial scenes with fast mobility or massive-element phased arrays. The multi-fold gain also reflects how mobility causes huge channel underutilization in existing 60 GHz protocols, despite their nominal Gbps bit-rates in static scenarios.

The key contribution of Beam-forecast lies in a novel model-driven approach to sustain high performance for mobile 60 GHz links. Our contribution breaks down into the following three aspects:

(i) We use testbed experiments (Sec. II) to characterize the SCP correlation of 60 GHz networks. More importantly, we develop a geometry model (Sec. III-B) to quantitatively capture

the correlation.

(ii) We propose a novel optimization framework, which employs a reverse-engineering approach to reconstruct SCPs at new locations, by simply inspecting the channels of a small fraction of beams. Furthermore, we can predict the optimal beams directly as the Tx/Rx moves to new locations, without painful beam scanning (Sec. III-C).

(iii) We experiment with Beam-forecast based on a custom-built 60 GHz software-radio platform, and demonstrate its remarkable performance gain over existing work (Sec. IV).

II. MEASUREMENTS OF SCP CORRELATION

A. Background and Motivation

60 GHz radios employ phased-array beamforming to concentrate most signal power along certain angular direction. As a result, the channel quality of a 60 GHz link is a joint effect of the directional beamforming gain and the channel response. Whereas the former can be controlled by the Tx/Rx, the latter depends on the radio environment, which we profile below.

1) *Measurement Platform*: We use a custom-built 60 GHz software-radio platform as shown in Fig. 1. The baseband signal is generated in a connected PC, and then is processed on the WARP FPGA. Afterwards, the signal passes through a high-speed 245.76 Msps DAC module, and then feeds into the RF frontend transmitter, which operates on 57-64 GHz with 10 dBm output power. The received signal goes through the path in reverse and then is analyzed on the PC. For channel measurement, we use an omni-directional antenna on the transmitter, and adopt a mechanically steerable horn antenna with 34 dBi gain and 3.4° beam-width, which is equivalent to a phased-array with 50×50 antenna elements¹.

2) *Spatial Channel Profile*: As reported widely by existing works [11]–[15], 60 GHz channels are extremely sparse. Our testbed experiments corroborate this feature. We use the 60 GHz platform, and rotate the horn antenna with 1° steps, and measure the spatial channel profile (SCP), *i.e.*, channel gains along each angular direction, represented by a vector. The black curve in Fig. 2 shows the RSS pattern along different Angle-of-Arrival (AoA). We see that the signal energy concentrates around two directions constituting one Line-of-Sight (LoS) path and one None-Line-of-Sight (NLoS) path from a reflector with strong reflectivity. The case is similar for angle-of-departure (AoD) pattern, omitted due to space limit.

¹To our knowledge, there exists no programmable phased-array that allows low-level control/access. But the highly directional horn antennas can accurately measure the spatial channel profile and emulate a phased-array (Sec. IV-A).




Fig. 4. SCP correlation: an example. From left to right: AoA_1 , AoA_2 , AoA_3 .

3) *Phased-array Beamforming*: A phased-array generates directional beams by applying precoding to its antenna elements. In particular, a beam pattern is formed by using a vector of specifically designed beamforming weights (complex numbers, one for each antenna element). Signal out of each antenna element will be coded by a corresponding beamforming weight before the signal is emitted. The resulting signal from superposing all antenna elements forms a particular directional beam: signal power at some directions is greatly strengthened while other directions weakened.

We employ a commonly-used trace-driven approach to emulate the joint effect of phased-array beamforming and channel response (see more details in Sec. IV-A). For this example, we generate 64 beams following 802.11ad codebook of a simulated 32-element linear phased-array, and then convolve the beams with the above collected AoA to get the channel quality. We plot two example beams in Fig. 2, and it is clear that *beam 1* is much better than *beam 2*, since its main beam direction closely aligns with the strongest spike of the AoA. In particular, we compute the RSS under the two beams, and *beam 1* achieves an RSS 14.7 dBm higher than *beam 2*, which translates to a bit-rate gap of 2 Gbps, according to IEEE 802.11ad coding/modulation specification [8].

Clearly, the objective of beam steering is to find the optimal beam that is best-aligned with the underlying AoA pattern, so as to produce the maximum channel quality. Without knowledge of the AoA pattern, existing beam steering resorts to trial-and-error scanning: a phased-array keeps probing all beams one by one until it finds the beam with the best channel. Such beam scanning may be needed only once for a static link. However, for a mobile link, according to measurements in [10], [15], frequent beam-scanning is required to maintain quality connections. The scanning time can easily overwhelm the channel coherence time, causing link outage most of the time during mobility.

B. SCP Correlation

In this paper, we find that brute-force beam-scanning can actually be averted, by leveraging an unexploited characteristic of 60 GHz channels, *i.e.*, the SCP correlation. Generally speaking, we find that SCPs (including AoA and AoD patterns) at nearby locations are highly-correlated with each other, and we can estimate the SCP at one location given the SCP of another. In consequence, we can predict the optimal beams directly without requiring time-consuming beam-scanning. Before we propose the detailed design, we first profile the SCP correlation via measurements. For simplicity, here we show the AoA correlation, but the case is the same for AoD.




Fig. 5. SCP correlation v.s location distance.

1) *A Showcase of SCP Correlation*: Using the 60 GHz platform, we measure the AoA patterns of three links in an office environment in Fig. 3, where we mark the transmitting location and three receiving locations, *i.e.*, l_1 , l_2 and l_3 . We intentionally set l_1 and l_2 close (*i.e.*, their distance is 0.5 meter), and l_3 far away, *i.e.*, 4 meters away from l_1 . At each location, we rotate the horn antenna in steps of 3° and record the RSS along each direction. The resulting AoA patterns (referred to as AoA_1 , AoA_2 , AoA_3) are plotted in Fig. 4. We make two observations: (i) AoA_2 is quite similar to AoA_1 , irrespective of an overall clockwise shift. In particular, the distribution of angular clusters of AoA_2 is analogous to that of AoA_1 , and the signal strength along each direction is almost the same. (ii) AoA_3 is far from AoA_1 and AoA_2 , in terms of both distribution of angular clusters and signal strength.

The reason behind these observations is channel sparsity [10]–[15] and the pseudo-optical propagation characteristics of 60 GHz signals. Specifically, the AoA pattern is dominated by the LoS path (if any), plus a few NLoS paths reflecting from a few main reflectors (with strong reflectivity, like metal furniture). When a user moves to a nearby location, signals are still coming from the same reflectors, but from different reflecting points. As a result, the signal's incident angle changes slightly, but the overall AoA pattern preserves much similarity. However, when a user moves too far away, the signal will come from a different set of reflectors with random incident angles, and the AoA pattern becomes less correlated.

2) *A Quantitative Analysis of SCP Correlation*: We conduct a quantitative study of SCP correlation by examining the cross correlation of AoA patterns from different locations. Specifically, we collect the AoAs at 30 locations in the office as shown in Fig. 10, and the distance between any two locations ranges from 0.43m to 5.82m. We then compute the cross correlation between any two locations' AoA patterns and plot the result as distance increases in Fig. 5. The result shows a clear trend: the cross correlation is high when location distance is small, while becoming weaker with distance. Overall, the trend indicates that SCP correlation at nearby locations is ready to be exploited, but also calls for solution when the user moves continuously and the SCP becomes less correlated.

III. DESIGN

A. Design Overview

The above measurement insights provide the foundation for our Beam-forecast design. Beam-forecast comprises three major modules, as illustrated in Fig. 6. The first is a geometry model, which quantitatively captures the SCP correlation. More

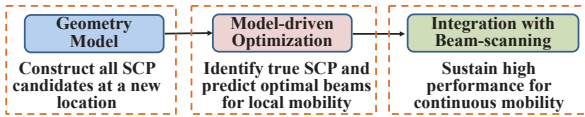


Fig. 6. Design framework.




Fig. 7. Geometry model illus.




Fig. 8. Environment mapping illus.

specifically, the model can estimate the SCP at an arbitrary nearby location, given the SCP at an anchor point. Based on the geometry model, the second component, *i.e.*, model-driven optimization, predicts optimal Tx/Rx beam-pairs as the user moves nearby. The prediction incurs minimal overhead, because it is model-driven and only needs parameterization based on one-shot measurement of the currently-used beam (or a small fraction of extra random beams), instead of using the exhaustive beam-scanning.

Under continuous mobility, a user (client device) may move too far away from the anchor point, and in consequence the SCP correlation weakens and invalidates the prediction. To meet this challenge, we leverage the IEEE 802.11ad beam-scanning to refresh the SCP measurement. Afterwards, we deem this location as a new anchor point that benefits later predictions. In this way, we progressively handle continuous mobility with trivial overhead. In what follows, we proceed to detail Beam-forecast's three components.

B. Geometry Model

1) *Modeling SCP*: The objective of the geometry model is: given the SCP at an initial anchor point, it models the SCPs of nearby locations. As illustrated in Fig.7, we assume a user moves to a new location denoted as l_1 , with distance d along a direction θ from the anchor point l_0 with coordinate of $\{x_0, y_0\}$. Then the coordinate of l_1 is as follows,

$$\begin{aligned} x_1 &= x_0 + d \cos(\theta) \\ y_1 &= y_0 + d \sin(\theta) \end{aligned} \quad (1)$$

Note that an SCP is comprised of multiple signal paths, and each path p can be represented with a quadruple as $\{\alpha_p, \beta_p, s_p, \phi_p\}$ where α_p is the path's angle of arrival (AoA), β_p the angle of departure (AoD), s_p the signal strength and ϕ_p the phase change after traversing this path. We proceed to show how to get such path information one by one.

Deriving AoA: When the user is at the anchor point l_0 , it performs a beam scanning, and then it can distill surrounding information via environment mapping (see more details in next subsection). The surrounding information includes the AP's coordinate $\{x_{AP}, y_{AP}\}$, and the layout/reflectivity of main reflectors. For simplicity, we only plot one reflector R_p in Fig. 7 and demonstrate how to get information of the LoS path and the NLoS path reflected by R_p . But the approach applies to all paths and all reflectors.

Based on the AP and the reflector R_p 's locations, we can derive the location of a *virtual AP* \tilde{AP}_p , *i.e.*, the mirror-reflecting point of the AP against reflector R_p . Denote the virtual AP's coordinate as $\{x_{\tilde{AP}_p}, y_{\tilde{AP}_p}\}$. Following basic optical principles, the path 'bent' by R_p can be treated as coming from \tilde{AP}_p . Therefore, the incident angle α_p satisfies

$$\tan(\alpha_p) = \frac{y_{\tilde{AP}_p} - y_1}{x_{\tilde{AP}_p} - x_1} \quad (2)$$

For the LoS path, without any reflector's 'bending' effect, AoA can be derived directly from the AP's and user's coordinate. Specifically, the AoA α'_p can be characterized as:

$$\tan(\alpha'_p) = \frac{y_{AP} - y_1}{x_{AP} - x_1} \quad (3)$$

From Eq. (2) and Eq. (3) we can immediately have AoA.

Deriving AoD: From the line stringing \tilde{AP}_p and l_1 , we derive the exact reflecting point on R_p referred to as P_{ref} with coordinate $\{x_{ref}, y_{ref}\}$. Then the AoD, denoted as β_p , satisfies:

$$\tan(\beta_p) = \frac{y_{ref} - y_{AP}}{x_{ref} - x_{AP}} \quad (4)$$

When path p is an LoS, the relationship between AoD β'_p and AoA α'_p is more straightforward as shown in Fig. 7(b), and we can derive β'_p as follows,

$$\beta'_p = \alpha'_p + \pi \quad (5)$$

Deriving Signal Strength: The path length denoted with Δ_p is the Euclidean distance between \tilde{AP}_p and l_1 for reflection path, *i.e.*,

$$\Delta_p = \|\tilde{AP}_p - l_1\| \quad (6)$$

and length for LoS path is as follows,

$$\Delta'_p = \|AP - l_1\| \quad (7)$$

In consequence, we can derive signal strength (in dB scale) along this path based on Friis model [10] as follows,

$$s_p = 20 \log_{10} \left(\frac{4\pi \Delta_p}{\lambda} \right) - \eta_p \quad (8)$$

where λ is the carrier wavelength and η_p the reflectivity loss.

Deriving Phase: the phase change after traversing path p can be computed from the path length as follows,

$$\phi_p = \frac{2\pi \Delta_p}{\lambda} \quad (9)$$

2) *Environment Mapping*: To bootstrap the geometry model, we propose a straightforward method to distill environment information from mmWave signals at an anchor point. We first find that only part of the environment information is necessary for handling local movement. For instance, consider a radio environment as shown in Fig. 8, which includes one AP, one user marked with 'Rx' and 4 reflectors, *i.e.*, from R_1 to R_4 . We can find that the reflecting signal from R_1 and R_2 cannot reach the user at l_0 due to the pseudo-optical propagation characteristic of 60 GHz radios. In other words, the effective surrounding environment only includes R_3 and R_4 for the user. Moreover, when the user moves nearby, it is very likely that major reflectors (*e.g.*, walls and metal furniture) in the environment remain static. Whereas human obstacles can move, they mostly block/absorb instead of reflecting the signals [10].

Based on the insight, the environment mapping works as follow: at the anchor point l_0 , we perform a beam scanning,

from which we obtain the sparse set of strong beam directions from both the AP and Rx's prospective. In consequence, we can identify the effective reflecting points as ones that lie at the intersections of strong beam directions from the AP and Rx, as illustrated in Fig. 8. Using basic triangulation, we can pinpoint the locations of the virtual AP. Finally, we compute the reflectivity by analyzing the signal strength collected during beam scanning using the Friss model [10]. Based on such information, the geometry model is able to estimate SCPs for locations near l_0 as described above. During the process, we assume that the locations of AP and l_0 are known, and in practice we can drive them with the single-AP mmWave localization method [19].

Though the straightforward method cannot achieve as good of an 'imaging' effect as advanced mmWave imaging radars [20], [21], it can construct a rough environment including major reflectors, which already supports geometry model and finally brings significant performance gain (see Sec. IV).

When a user moves too far away from the anchor point, the environment information may become stale. For example, the effective reflector will change to R_1 if the user moves to l_1 . We can handle this problem by progressively updating the anchor point, as we will present in Sec. III-D.

To make the environment mapping and the geometry model tractable, we ignore higher order reflection by assuming that strong signals all come from zero or first order reflection paths. We find that this does not impact accuracy, which is consistent with previous studies [7], [10], since signal from higher-order paths is usually much weaker.

C. Model Driven Optimization

Ideally, given the user's moving distance d and direction θ from the anchor point, the geometry model can construct the SCP at the new location, and then predict the new optimal beam directly. However, the exact values of d and θ are unavailable. One may tackle this problem using GPS/WiFi or sensor-assisted localization [22], [23]. But GPS/WiFi or sensor-assisted localization have error of meter or sub-meter level, which is far from the required accuracy for SCP reconstruction.

In Beam-forecast, we adopt a reverse-engineering approach to predict the true SCP at the new location, which is fast and accurate. In a nutshell, we enumerate all possible combinations of moving distance d and moving direction θ . For each combination, we compute one candidate SCP using the geometry model, multiply it angle-wise with the current beam's angular response, thus obtaining a candidate channel impulse response (CIR). We evaluate the modeled candidate CIRs and the actual measured CIR of the current beam, and select the SCP that produces the minimum gap as the true underlying SCP. We present a formal description below.

Suppose we have a set of L nearby locations, created by discretizing moving distance d and moving direction θ . For each location $l \in L$, its SCP consists of a sparse set of P paths, including LoS and NLoS paths caused by main reflectors. Each path $p \in P$ can be represented with a quadruple as $SCP_p^l = \{\alpha_p, \beta_p, s_p, \phi_p\}$, and we can derive such path information using the geometry model.

We then perform an element-wise multiplication between the SCP and the current beam's known angular gain pattern, so as to get the final channel CIR. More specifically, suppose a receiver has a phased-array with N_R antenna elements, and it is currently using the k_{th} beam, then the beamforming gain at an arbitrary spatial direction α is:

$$A_R^k(\alpha) = \sum_{n=1}^{N_R} w(n, k) \exp(j2\pi\mu\cos\alpha/\lambda) \quad (10)$$

where $w(n, k)$ is the beamforming weight applied to the n_{th} element when using this k_{th} beam. μ is the space between antenna elements, and is usually set to be half wavelength $\frac{\lambda}{2}$.

Similarly, for a transmitter with N_T elements and using the i_{th} beam, the beamforming gain at a spatial direction β is

$$A_T^i(\beta) = \sum_{n=1}^{N_T} w(n, i) \exp(j2\pi\mu\cos\beta/\lambda) \quad (11)$$

Then we can compute the expected CIR at location l , denoted with $h_{i,k}^l$ as follows,

$$h_{i,k}^l = \sum_{p=1}^P A_R^k(\alpha_p) A_T^i(\beta_p) s_p \exp(j\phi_p) \quad (12)$$

To sum up, the final CIR $h_{i,k}^l$ is the aggregated effect of signals from all paths, which are reshaped by beamforming at both the transmitter and receiver side, and then summed together. Meanwhile, we can derive the actual CIR $h_{i,k}^{ms}$ after the user moves to a new location. Specifically, each frame includes the built-in channel-training preambles in order to demodulate the packet payload, according to the PHY layer standard of 802.11ad [8]. The receiver extracts $h_{i,k}^{ms}$ after it receives the frame.

Then we can identify the most possible new location, l^* , by selecting the one that minimizes the gap between modeled CIRs and measured CIR,

$$l^* = \arg \min_{l \in L} \|h_{i,k}^l - h_{i,k}^{ms}\|^2 \quad (13)$$

Given the estimated l^* and the corresponding SCP^{l^*} , we can compute the final channel CIRs under all beam pairs, according to Eq. (12). Finally, we can identify the optimal beam pair $\{i^*, k^*\}$ at this new location as follows,

$$\{i^*, k^*\} = \arg \max_{i, k} \|h_{i,k}^{l^*}\|^2 \quad (14)$$

Combating Channel Noise: Ideally, inspecting the performance under the beam in use is good enough to estimate true SCP and the optimal beam pair. However, due to channel noise, the optimization in Eq. (13) may over-fit the SCP. As a result, the estimated location l^* and the predicted beam may not be accurate. To handle this problem, we can use more measurements to alleviate the channel noise impact. In particular, we use Q random beam pairs (denote with $\{(i_1, k_1), (i_2, k_2), \dots, (i_Q, k_Q)\}$ instead of just the only one in use, and measure the CIRs as $H^{ms} = \{h_{i_1, k_1}^{ms}, h_{i_2, k_2}^{ms}, \dots, h_{i_Q, k_Q}^{ms}\}$. Accordingly, we compute a set of theoretical channel for each l as $H^l = \{h_{i_1, k_1}^l, h_{i_2, k_2}^l, \dots, h_{i_Q, k_Q}^l\}$. Finally, we use H instead of h in Eq. (13) and Eq. (14) to derive the optimal beam pair. We show in Sec. IV that we can successfully derive the optimal beam by using less than 5% of the total beam pairs.

Computational Complexity: Essentially, Beam-forecast saves channel time by doing model-driven computation, *i.e.*, decreasing channel probing time while increasing computational time. From the models in previous two sections, we can see



Fig. 9. Experiment environment




Fig. 10. AP and client locations.

that the computational time of Beam-forecast is proportional to the product of the total number of used beam pairs Q and the total number of SCP candidates L (in turn depends on the granularity of the discretized distance d and direction θ), both of which are limited to the order of several dozens. Moreover, Beam-forecast can run in low-layer firmware in practice, so the computation time cost of the finite scale multiply-and-add operations is negligible.

D. Protocol Operations

To handle continuous mobility, we incorporate beam prediction with the existing beam-searching protocols in 802.11ad. Basically, when a user moves far away from the anchor point, the model-driven prediction may lose effect, *i.e.*, the channel with the predicted beam cannot support the lowest Data PHY bit-rate (referred to as δ) defined in 802.11ad specification [8]. In this case, we will trigger the beam-scanning, which can refresh the ground-truth SCP and identify the best beam². Then, we deem this location as a new anchor point, based on which a new round of model-driven prediction begins.

Workflow. We present the protocol procedure of Beam-forecast as follows.

(i): At the initial anchor point, Beam-forecast performs a beam-scanning to find the best Tx/Rx beams and bootstraps the geometry model. While a user is moving, Tx/Rx beams lose alignment gradually, hence the PHY-layer bit-rate keeps dropping.

(ii): To prevent link outage, Beam-forecast keeps doing model-driven beam steering. In particular, at the beginning of each Beacon Interval, Beam-forecast estimates the new SCP using the geometry model, and then predicts the optimal beam pair. We denote corresponding bit-rate under the predicted beam pair as r_l . We let step (ii) iterate as long as Beam-forecast prevents link outage, *i.e.*,

$$r_l \geq \delta \quad (15)$$

(iii): Otherwise, Beam-forecast cannot recover link since the user is too far from the anchor point. In this case, we set the current location as the new anchor point, and return to step (i).

Intuitively, the procedure is a progressive approach for mobile 60 GHz links: instead of using time-consuming beam-scanning each time, we always prioritize Beam-forecast, and rollback to beam-scanning only when necessary.

IV. EVALUATION

A. Implementation and Evaluation Methodology

Hardware: We build WiMi, a 60 GHz software radio platform as described in Sec. II. Using WiMi, we carry out over-

²Link outage may also attribute to blockage, besides user mobility. Recent work like BeamSpy [15] has proposed efficient mechanisms to enhance link robustness to blockage, which is orthogonal and complementary to our work.




Fig. 14. An example of SCP gap between geometry model and groundtruth. Left: AoA gap = 0.040. Right: AoD gap = 0.041.

the-air measurements and collect SCPs in an office environment that incorporate signal propagation, attenuation and reflection effect. These SCPs serve as the ground-truth in the micro-benchmark evaluation.

Trace-driven Emulation of Phased-array Beamforming:

Due to the lack of reconfigurable phased-array antennas, we use a classical trace-driven approach to emulate phased-array beamforming, which has been verified and commonly employed in previous studies [7], [10], [15]. Specifically, we generate the phased-array's angular gain pattern using a standard 802.11ad codebook [15], and then multiply it angle-wise with the SCP patterns measured at a specific Tx/Rx location, so as to form the channel response. Using different codebook configurations, we can emulate the beamforming effect for various phased-arrays.

Ray-tracing Simulation: Evaluating system-level performance of mobile 60 GHz networks requires SCP measurement with micro-second granularity as a user moves, which is infeasible for existing 60 GHz platforms with mechanically rotated horn antennas. To overcome this limitation, we adopt the classical ray-tracing approach [24], which extensively enumerates the possible LoS/NLoS propagation paths, traces their path-loss, reflection and refraction effects. Ray-tracing has been verified to accurately reproduce mmWave multipath environment in comparison with ground truth [25].

On top of the ray-tracing simulator, we implement the control/signaling/timing primitives of 802.11ad's TDMA MAC. A virtual clock is used to time each protocol procedure, including beam scanning, inter-frame spacing, *etc.*, with 802.11ad default timing parameters [8]. We also implement all design components of Beam-forecast, including the geometry model and reverse-engineering beam prediction, which augment 802.11ad's beam scanning.

Performance Benchmarks: For performance comparison, we have also implemented the default beam-scanning in IEEE 802.11ad, and state-of-the-art preemptive fast recovery (PFR) protocol proposed recently in [17].

(i) 802.11ad baseline [8]. We implement an 802.11ad baseline protocol that uses a heuristic to handle mobility: once link quality drops below the lowest bit-rate threshold, it performs an exhaustive beam scanning and then steers to the beam pair with the best link quality.

(ii) PFR [17]. Instead of immediately calling for an exhaustive beam-scanning, PFR preferentially probes two opportunistic beam directions upon detecting link outage. First, PFR dilates the current beam into a wider one, in the hope that the wider beam can cover the new location. Subsequently, if beam dilation fails, PFR will probe on a pre-computed backup




Fig. 11. Distribution of SCP gaps.




Fig. 12. Beam prediction accuracy vs. # beams.




Fig. 13. Beam prediction overhead vs. phased-array.

beam-pair (*i.e.*, the second strongest beam pair found in the last round of beam scanning), in the hope that the backup beam-pair is not blocked or misaligned at the new location. Note that a complete PFR [17] also includes a rate adaptation mechanism, but here we focus on beam steering.

For simplicity, we adopt exhaustive beam-scanning for 802.11ad. Both PFR and Beam-forecast augment exhaustive scanning and need to invoke it occasionally, so it is fair to compare how much gain each mechanism has atop the exhaustive baseline. We have also implemented more advanced beam-scanning mechanisms like hierarchical beam searching [8], [10]. Though these schemes have lower absolute overhead than exhaustive scanning, the overhead still grows quadratically [10] with phased-array size, similar to the exhaustive scanning. We thus omit the detailed results due to space constraint.

B. Micro-benchmark Evaluation

1) *Accuracy of Geometry Model*: We examine the accuracy of the proposed geometry model by comparing the estimated SCP against the measured ground-truth. We carry out experiments in an office environment (Fig. 9). We set one AP near the center and 30 arbitrary client locations around the AP, whose coordinates are plotted in Fig. 10. The distance between any two client locations lies in [0.43m, 5.82m] with an average of 3.09m. We first collect the ground-truth SCP at each location using the 60 GHz testbed. Then we feed the SCPs into the simulator, in which we alternately set one client location as the anchor point, and estimate the SCPs for remaining locations using the geometry model. We compute the gap between the estimated SCPs and the ground-truth (SCP gap, which is the difference of two vectors representing the estimated and the ground-truth SCP, respectively), resulting in 870 values in total. We plot the distribution of these gaps in Fig. 11, which shows small gaps of less than 0.04 for about 80% of the estimations.

Note that the SCP, or channel angular gain patterns, are scaled by our 60 GHz radio hardware. To have a more intuitive understanding of how small the gap is, we present a specific example in Fig. 14. From the figure, the estimated AoA/AoD pattern and the ground-truth match closely in terms of both angular directions and signal strength, despite minor deviations. We note that the absolute AoA/AoD gap for this example is about 0.04, which means that about 80% of the estimations in Fig. 11 can achieve accuracy above this level.

2) *Accuracy of Beam Prediction*: Based on the geometry model, Beam-forecast can predict new optimal beams after a user moves. In contrast to brute-forcing beam scanning, Beam-forecast achieves the prediction by only inspecting CIR of the beam in use or a few extra randomly-selected beams. The number of extra beams determines Beam-forecast's overhead,

and now we investigate the relationship between the number of beams and prediction accuracy.

Similar to the above experimental scenario, we alternately choose one client location as the anchor point, and then predict the beam pairs after a user moves to remaining locations. We compare the resulting RSS at the new location under three beam configurations: predicted beam from Beam-forecast, the optimal beam (upperbound from an oracle perspective), and the old optimal beam at the anchor point (*i.e.*, representing the case without beam steering). For Beam-forecast, we gradually increase the number of beams that it uses. Due to space limit, we show the RSS when using a medium-size 16-antenna phased-array, but the trend is similar for arrays of different size.

We plot the average RSS in Fig. 12 for the total of 870 links under three kinds of beam configurations, from which we observe two insights: (i) Beam-forecast's model-driven beam prediction quickly converges towards the optimal bound. In particular, when the number of used beam pairs is 40, the resulting RSS is only 1.9 dBm less than the upperbound, which is far better than the 9.4 dBm gap between no beam-steering and the upperbound. Note that 40 beam pairs are only a very small fraction, *i.e.*, about 4% among all beam pairs. (ii) The model-driven optimization cannot bring further gain when the number of used beams goes beyond a certain amount. The reason is that some locations are far from its anchor points, and thus the geometry model loses effect. This result testifies the necessity to occasionally use beam scanning to recalibrate the model-driven prediction.

3) *Beam-steering Overhead*: Continuing with the above experiment, we now examine how much channel time that Beam-forecast can save. In particular, we stop Beam-forecast when the resulting RSS is beyond 80% of the optimal RSS, and then we count how many beam pairs have been tried. We compare time spent on trying these beams with the 802.11ad beam-scanning time, and plot the result in Fig. 13.

802.11ad's beam scanning time grows quadratically with the number of antenna elements in a phased-array, which corroborates with existing measurements [10], [15]. For instance, the scanning time is about 28 ms for a 32-antenna phased-array, which is a large overhead especially in mobile scenarios where beam-scanning runs frequently. However, Beam-forecast is able to keep the steering overhead lower than 1 ms. In particular, the overhead is only 1.56%, 5.08%, 2.54% and 2.83% of that under IEEE 802.11ad, for the 4 phased array configurations, respectively. Therefore, we set Beam-forecast to use only 5% of the total beams in later system-level evaluation.




Fig. 15. Throughput vs. moving speed.




Fig. 16. Throughput vs. phased array size.




Fig. 17. Throughput vs. # extra reflectors.

C. System-level Evaluation

Now we investigate the overall performance of Beam-forecast under various scenarios.

1) *Mobility Speed*: We set 10 random trajectories with length within the range of [10m, 20m] in a bigger office environment (Fig. 10). We simulate a client user with saturated traffic demand who moves along these trajectories with speed of 1Km/h to 5 Km/h with steps of 1Km/h, which is the common indoor walking speed.

From the result in Fig. 15, we make three observations. (i) Mobility has an adverse effect on 802.11ad and Beam-forecast: the larger the speed, the more frequent beam scanning, and in consequence the smaller the throughput. Despite the same trend, Beam-forecast shows significant gain over 802.11ad especially under high moving speeds, due to Beam-forecast's ability to avoid most beam scanning. Specifically, the gain increases from $2.13\times$ to $22.10\times$ as mobility speed increasing from 1Km/h to 5Km/h. (ii) PFR is much less sensitive to mobility speed, since its wider beams or backup beams are able to prevent sudden link outage under high mobility. However, wider beams/non-optimal backup beams imply much less beam-forming gain and in consequence result in smaller bitrate³. In particular, the gain of Beam-forecast over PFR is from $10.08\times$ to $3.06\times$. (iii) PFR is worse than 802.11ad under slow mobility, and better under fast mobility. This demonstrates the trade-off between beam optimality and the overhead of finding the beam. More specifically, 802.11ad's beam scanning ensures the optimal beams and the largest bit-rates, but the overhead of the beam scanning is large. However, the frequency of beam scanning is low for slow-moving users, thus achieving higher performance.

2) *Phased-array Size*: We repeat the above experiments using different phased-arrays with the number of antenna elements of {4, 8, 16, 32, 64}, while fixing the moving speed to 3Km/h. We plot the average throughput under the three approaches in Fig. 16.

We have two insights: (i) The gain of Beam-forecast is marginal for small-size phased-array, whereas it becomes remarkable with phased-array size. In particular, the gain of Beam-forecast over 802.11ad and PFR increases from $1.05\times$ to $6.50\times$, with average gains of $3.37\times$ and $2.76\times$, respectively. This attributes to two factors. First, for small-size phased-array, the beam is wide and the beam-forming gain is low. In consequence, the bit-rate discrepancy between a random beam and the optimal beam predicted by Beam-forecast is small. Second, the total number of beams is limited for small-size phased-arrays, therefore the overhead of brute-forcing beam

scanning is marginal. As a result, the advantage of avoiding beam scanning is not obvious for small phased-arrays. (ii) The throughput of all methods starts decreasing when phased array size is beyond a certain value (*i.e.*, 32). The reason is that beams become narrower and more directional for large-size phased-array, which incur more frequent beam re-alignment. In consequence, the beam steering overhead becomes larger, which counteracts the increasing beam-forming gain of higher directionality. However, the overhead of model-driven Beam-forecast increases very slowly, and thus its performance gain over the other two methods becomes more significant, though the absolute value of its throughput also drops.

3) *Reflecting Environment*: As is well-known, performance of 60 GHz radios is sensitive to surrounding environment. We model different environments by gradually adding more reflectors upon the basic topology of Fig. 9 in the ray-tracing setup. In this experiment, we use 32-antenna phased-arrays and 3Km/h mobility speed, and plot the result in Fig. 17. We see that the throughput of all three methods decreases with the number of reflectors, since more reflectors worsen the variation of SCP during mobility. Irrespective of the decreasing trend, Beam-forecast still has significant performance gain over 802.11ad and PFR, about $9.82\times$ and $3.20\times$, respectively, averaging over all environment settings.

To summarize, Beam-forecast shows consistent and significant performance gain over 802.11ad and PFR under various practical scenarios. The gain is multi-fold on average and can be up to $22\times$, which is attributed to Beam-forecast's ability of optimal beam prediction with minimal beam steering overhead.

V. RELATED WORKS

60 GHz Networking: Earlier effort on 60 GHz networks mainly focused on measuring/modeling physical characteristics like propagation, oxygen absorption *etc.* (see [12]–[14] and references therein). These works discovered unique features of 60 GHz communications such as channel sparsity and severe sensitivity to reflecting environment, which are fundamentally different from conventional Wi-Fi/cellular networks operating at low frequencies.

As 60 GHz hardware and standard become mature, system-level studies have attracted increasing attention recently. Zhu *et al.* [7] proposed a 60 GHz picocell architecture and showed the huge potential of 60 GHz radio for providing outdoor access. Sur *et al.* [10] focused on 60 GHz indoor networking and investigated network coverage, impact of beam steering and multi-link spatial reuse issues. These works pioneered 60 GHz networking in practical environment where mobility is ubiquitous, and also highlighted the inefficiency of existing protocols under mobility, which inspired our work.

³The result also indicates that incorporating the speed-robustness property of PFR with Beam-forecast can further enhance performance under fast-moving scenarios, which is an interesting topic but out of the scope of this paper.

Mobility in 60 GHz Networking: How to track user mobility and sustain high performance is largely an unsolved problem for mobile 60 GHz networks. Earlier works designed complicated digital beamforming patterns [26]–[28] or beam steering protocols [29] to narrow down the search space and thus reduce overhead. Despite the valuable theoretical insights, they are not applicable to practical 60 GHz nodes, whose phased-arrays only support codebook-based analog beam-forming [10].

Alternative solutions [16], [30] leveraged legacy 2.4/5 GHz Wi-Fi bands to predict the 60 GHz beam quality and guide the beam steering. However, the channel profiles of 2.4/5 GHz and 60 GHz differ drastically: 2.4/5 GHz signals have much lower penetration loss and stronger diffraction effects. So the prediction can only provide coarse guidelines. In addition, requirement of additional 2.4/5 GHz antenna array infrastructure is *not* feasible on low-cost/small-size client devices. In addition, much effort has been devoted to the parallel mobility-support problem in directional networks operating at 2.4/5 GHz frequencies [31]–[33], but the solutions are ineffective for 60 GHz networks due to fundamentally different channel [7], [10], [16], [17].

Beam-forecast is partly inspired by BeamSpy [15], which can efficiently recover a static 60 GHz link from human blockage. Unlike BeamSpy, however, Beam-forecast exploits a different feature of 60 GHz channels, *i.e.*, the SCP correlation, to handle the more challenging mobility problem.

VI. CONCLUSION

In this paper, we have demonstrated that the 60 GHz spatial channel profiles are highly correlated with locations. More importantly, the correlation can be captured by modeling the 60 GHz signals' sparse propagation paths and their interactions with reflectors. Leveraging this observation, we propose Beam-forecast as a model-driven solution to facilitate mobile 60 GHz networking. Beam-forecast is able to predict optimal beams with minimal overhead, and hence sustains high performance as the 60 GHz Tx/Rx change their locations. Experiments show that Beam-forecast achieves multi-fold throughput gain compared with state-of-the-art under various network scenarios. Beyond efficient beam steering, we believe Beam-forecast can augment with other 60 GHz networking primitives, such as multi-AP handoff and interference management.

VII. ACKNOWLEDGMENT

The work reported was supported in part by the National Natural Science Foundation of China (No. 61572476), the Funds for Creative Research Groups of China under Grant No.61421061, and the National Natural Science Foundation of China (No. 61532012, No. 61332005) and Beijing Training Project For The Leading Talents in S&T (ljrc201502). The work was also supported in part by US NSF funding CNS-1343363, CNS-1350039, CNS-1506657, CNS-1518728.

REFERENCES

- [1] S. Singh, R. Mudumbai, and U. Madhow, "Distributed coordination with deaf neighbors: Efficient medium access for 60 ghz mesh network," in *IEEE INFOCOM*, 2010.
- [2] R. Singh, R. Mudumbai, and U. Madhow, "Interference analysis for highly directional 60-ghz mesh networks: The case for rethinking medium access control," *IEEE/ACM Transactions on Networking*, vol. 19, no. 5, 2011.
- [3] X. Zhou and etc, "Mirror mirror on the ceiling: Flexible wireless links for data centers," in *ACM SIGCOMM*, 2012.
- [4] R. Daniels and R. H. Jr., "60 ghz wireless communications: Emerging requirements and design recommendations," *IEEE Vehicular Technology Magazine*, vol. 2, no. 3, 2007.
- [5] "The 5g infrastructure public private partnership," <https://5g-ppp.eu/>, 2014.
- [6] T. Rappaport *et al.*, "Millimeter Wave Mobile Communications for 5G Cellular: It Will Work!" *IEEE Access*, vol. 1, 2013.
- [7] Y. Zhu *et al.*, "Demystifying 60ghz outdoor picocells," in *ACM MobiCom*, 2014.
- [8] IEEE Standards Association, "IEEE Standards 802.11ad-2012: Enhancements for Very High Throughput in the 60 GHz Band," 2012.
- [9] "Pai eyes 50-60 ghz spectrum for virtual reality," <http://www.broadcastingcable.com/news/washington/pai-eyes-50-60-ghz-spectrum-virtual-reality/146980>, 2016.
- [10] S. Sur *et al.*, "60 ghz indoor networking through flexible beams: A link-level profiling," in *ACM SIGMETRICS*, 2015.
- [11] H. Xu, V. Kukshya, and T. Rappaport, "Spatial and temporal characteristics of 60-ghz indoor channels," *IEEE JSAC*, vol. 20, no. 3, 2002.
- [12] C. Anderson and T. Rappaport, "In-building wideband partition loss measurements at 2.5 and 60 ghz," *IEEE Transactions on Wireless Communications*, vol. 3, no. 3, 2004.
- [13] P. F. M. Smulders, "Statistical characterization of 60-ghz indoor radio channels," *IEEE Transactions on Ant. and Prop.*, vol. 57, no. 10, 2009.
- [14] T. Rappaport and etc, "Broadband millimeter-wave propagation measurements and models using adaptive-beam antennas for outdoor urban cellular communications," *IEEE Trans. on Ant. and Prop.*, vol. 61, no. 4, 2013.
- [15] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, "Beamspy: Enabling robust 60 ghz links under blockage," in *USENIX NSDI*, 2016.
- [16] N. T *et al.*, "Steering with eyes closed: mm-wave beam steering without in-band measurement," in *IEEE INFOCOM*, 2015.
- [17] M. K. Haider and E. W. Knightly, "Mobility resilience and overhead constrained adaptation in directional 60 ghz wlans: Protocol design and system implementation," in *ACM MobiHoc*, 2016.
- [18] "WiMi," <http://xyzhang.ece.wisc.edu/wimi/index.html>.
- [19] T. Wei, A. Zhou, and X. Zhang, "Facilitating robust 60 ghz network deployment by sensing ambient reflectors," in *USENIX NSDI*, 2017.
- [20] F. Guidi, A. Guerra, and D. Dardari, "Millimeter-wave massive arrays for indoor slam," in *IEEE ICC*, 2014.
- [21] Y. Zhu, Y. Zhu, B. Y. Zhao, and H. Zheng, "Reusing 60ghz radios for mobile radar imaging," in *Proc. of ACM MobiCom*, 2015.
- [22] N. K. Santhapuri *et al.*, "Sensor assisted wireless communication," in *IEEE LANMAN*, 2010.
- [23] L. Ravindranath, C. Newport, H. Balakrishnan, and S. Madden, "Improving wireless network performance using sensor hints," in *Proc. of USENIX NSDI*, 2011.
- [24] E. De Groot, T. Bose, C. Cooper, and M. Kruse, "Remote transmitter tracking with raytraced fingerprint database," in *IEEE MILCOM*, 2014.
- [25] D. Esposti *et al.*, "Ray-tracing-based mm-wave beamforming assessment," *IEEE Access*, 2014.
- [26] Y. Tsang, S. Lin, and A. Poon, "Fast beam training for mmwave communication system: From algorithm to circuits," in *ACM Int. Workshop on mmWave Communications*, 2010.
- [27] Y. M. Tsang, A. S. Y. Poon, and S. Addepalli, "Coding the beams: Improving beamforming training in mmwave communication system," in *IEEE GLOBECOM*, 2011.
- [28] J. Wang *et al.*, "Beam codebook based beamforming protocol for multi-gbps millimeter-wave wpan systems," *IEEE JSAC*, vol. 27, 2009.
- [29] B. Li, Z. Zhou, H. Zhang, and A. Nallanathan, "Efficient beamforming training for 60-ghz millimeter-wave communications: A novel numerical optimization framework," *IEEE Transactions on Vehicular Technology*, vol. 63, no. 2, 2014.
- [30] H. Park, K. Yongsun, S. Taewon, and S. Pack, "Multiband directional neighbor discovery in self-organized mmwave ad hoc networks," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 3, 2015.
- [31] A. Amiri Sani *et al.*, "Directional antenna diversity for mobile devices: characterizations and solutions," in *ACM MobiCom*, 2010.
- [32] V. Navda *et al.*, "Mobisteer: using steerable beam directional antenna for vehicular network access," in *ACM MobiSys*, 2007.
- [33] H. Ma, Z. Xi, and A. Ming, "A coverage-enhancing method for 3d directional sensor networks," in *IEEE INFOCOM*, 2009.