

Towards Scalable and Ubiquitous Millimeter-Wave Wireless Networks

Sanjib Sur

University of South Carolina
sur@cse.sc.edu

Xinyu Zhang

University of California San Diego
xyzhang@ucsd.edu

Ioannis Pefkianakis

Hewlett Packard Labs
ioannis.pefkianakis@hpe.com

Kyu-Han Kim

Hewlett Packard Labs
kyu-han.kim@hpe.com

ABSTRACT

Millimeter-wave (mmWave) technology is emerging as the most promising solution to meet the multi-fold demand increase for mobile data. Very short wavelength, high directionality, together with sensitivity to rampant blockages and mobility, however, render state-of-the-art mmWave technologies unsuitable for ubiquitous wireless coverage. In this work, we design and implement *UbiG* — a mmWave wireless access network — that can deliver ubiquitous gigabits per second wireless access consistently to the commercial-off-the-shelf IEEE 802.11ad devices. *UbiG* has two key design components: (1) a fast probing based beam alignment algorithm that can identify the best beam consistently with guaranteed latency in a mmWave link, and the algorithm scales well even with a very large number of beams; and (2) an infrastructure-side predictive ranking based fast access point switching algorithm to ensure seamless gigabits per second connectivity under mobility and blockage in a dense mmWave deployment. Our IEEE 802.11ad testbed experiments show that *UbiG* performs close to an “Oracle” solution that instantaneously knows the best beam and access point for gigabits per second data transmission to users.

CCS CONCEPTS

• **Networks** → **Network protocol design; Network algorithms; Wireless access networks; Physical links; Mobile networks;**

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

MobiCom'18, October 29–November 2, 2018, New Delhi, India

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5903-0/18/10...\$15.00

<https://doi.org/10.1145/3241539.3241579>

KEYWORDS

Millimeter-Wave; 5G; 60 GHz; Beam Alignment; AP Switching

ACM Reference Format:

Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2018. Towards Scalable and Ubiquitous Millimeter-Wave Wireless Networks. In *The 24th Annual International Conference on Mobile Computing and Networking (MobiCom'18), October 29–November 2, 2018, New Delhi, India*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3241539.3241579>

1 INTRODUCTION

The skyrocketing global demand for mobile data has created a significant push to both academia and industry to pursue new wireless technologies beyond existing LTE and Wi-Fi. Millimeter-wave (mmWave) is emerging as the most promising solution [10, 33, 44] to meet the nearly $100\times$ traffic surge in the next five years [6]. mmWave technology promises multi-Gbps throughput by operating on the unlicensed multi-GHz bandwidth — several orders of magnitude wider than LTE and Wi-Fi combined. Multiple standardization efforts [4, 16, 17] and Commercial-Off-The-Shelf (COTS) products operating on the unlicensed 60 GHz mmWave spectrum already offer upto 7 Gbps of wireless bit-rate. Further, recent research advancements of new mmWave systems [2, 13, 41, 43, 53] and large-scale field-trials by popular network operators and vendors [50] led mmWave to be recommended as a key technology for the 5G mobile broadband [10, 11, 44].

mmWave signals, unfortunately, suffer from severe path-loss due to high operating frequency. So, mmWave devices need to use phased-array antennas to focus their Radio Frequency (RF) energy through narrow directional beams and compensate the attenuation loss. Aligning the beams between the Access Point (AP) and the user devices, however, remains a fundamental challenge due to the high channel dynamics under device mobility and obstacle blockage and due to the large number of possible beam directions in next-generation mmWave devices [2, 41–43, 54].




Figure 1: Our 60 GHz experimental platform: (a) A dual-band IEEE 802.11ad/ac AP with a 4×8 60 GHz phased-array antenna. (b) Same phased-array antenna mounted on a laptop. Enabling ubiquitous Gbps wireless access: (c) *UbiG* aligns beams between AP and user devices with few probes and ensures Gbps access via switching data connection between APs proactively during blockages and mobility.

IEEE 802.11ad-compliant solutions [16] employ a *multi-level scan* to align the beams. Still, this approach can take hundreds of milliseconds latency to find the best alignment [41, 54], even for a pair of nodes with a small number of $N = 64$ beams [15, 42, 54]. During continuous mobility, furthermore, the scan may be invoked persistently, exacerbating the latency. Recent compressed sensing solutions [23, 30, 31] promise to scale well by requiring only $\mathcal{O}(K \log N)$ measurements. Here, K is the number of dominant channel paths and it is typically bounded by 2 or 3 [3, 32, 34, 39, 42, 43, 49]. These solutions, unfortunately, do not work with the COTS mmWave devices because they require phase-coherence across measurements [23, 30, 31] or need customized phased-array antennas [2, 15]. Although recent non-coherent design is shown to work in commodity platforms [36], it has limitations to align the best beam under multipath mmWave channel. A fixed overhead beam alignment protocol [43] allows to infer the best beam with fast probing, but works for quasi-stationary links only. Finally, out-of-band solutions (e.g., [26, 41]) require additional MIMO Wi-Fi radios to identify the best beam; but the solutions do not scale well for systems with large number of mmWave beams and limited Wi-Fi antennas.

Even if an “Oracle” solution can quickly align the best beam, it does not guarantee ubiquitous Gbps connectivity to an AP during blockage. Identifying the best unblocked AP among multiple candidates, each with up to hundreds of beams, can take the order of seconds and easily interrupt the upper layer protocols [36].

In this work, we propose *UbiG* – a *Ubiquitous Gbps* wireless access network solution – that promises to overcome the limitations of the existing mmWave network designs. *UbiG* sets three key design goals:

(1) *Scalable beam alignment*: *UbiG* aims to align the beam between the mmWave devices at a latency that is independent of the size of the phased-array antennas. This way it can bound the link recovery latency under various channel dynamics and

even for a very large number of beam directions.

(2) *Ubiquitous Gbps connectivity*: *UbiG* seeks to maintain Gbps connectivity even under blockage or mobility. Beam alignment alone may not help when a mmWave link’s Line-Of-Sight (LOS) is blocked. Ultra-dense AP deployment and fast AP switching are required to sustain the Gbps connectivity.

(3) *Low hardware cost*: *UbiG* needs to implement the algorithms on standard-compliant COTS mmWave hardware which may not provide coherent PHY layer information or only allows for coarse-grained phase control over the phased-array antenna elements.

UbiG introduces two key design components to meet the above goals: (1) a fast beam alignment algorithm that can align the best beam between an AP and a user device with a very few beam probing; (2) a rapid and scalable AP switching algorithm which introduces a concept of predictive “AP ranking” that can steer and connect to the best alternative AP without going through the traditional slow discovery. Our design works with existing cheap COTS IEEE 802.11ad devices and scales well with future platforms that may have larger phased-array antennas. Figure 1(c) shows the design overview of *UbiG*.

UbiG’s fast beam alignment algorithm is built upon a simple observation – while the channel gain and phase change with steering different beam directions, the underlying physical paths along which the signals travel from transmitter to receiver remain the same. As long as a transmitter can extract the properties (complex gain and directionality) of the paths, it can align the best beam without scanning through the entire space. *UbiG* judiciously selects a fixed number of beams to probe and then employs a *space-time* analysis on the channel measurement of the beams to extract the properties of all dominating paths. The number of probeings under this approach is independent of the size of the phased-array antenna and thus the number of beam directions N . This space-time analysis allows *UbiG* to break away from existing sparse-signal recovery mechanisms where the measurement complexity can not




Figure 2: Radiation pattern of two example beams in our platform. Signal strength is in dB.

be reduced below $\mathcal{O}(K \log N)$ [15, 23, 30, 31].

Channel measurements from COTS mmWave platforms, however, are phase-incoherent [36, 52] due to the lack of carrier phase tracking across different packets. *UbiG* is able to overcome the limitation based on an observation — even if the phase of channel measurement across multiple probes are incoherent, the relative phase of the channel within a single probe are still coherent. Because the relative phase only depend on the distance difference among paths, it remains the same across consecutive channel measurements. *UbiG* leverages a first-arrival-path phase calibration and joint estimation of all dominating paths to overcome the limitation.

UbiG, in addition, maintains Gbps connectivity upon blockage and mobility by harnessing the cooperation among densely deployed APs. These APs are tightly coordinated through a central controller and can share a high-speed Ethernet or fixed-beam mmWave backhaul [18]. Such multi-AP architecture has been widely deployed in enterprise Wi-Fi networks, but the high directionality of mmWave brings new challenges even for a simple problem — “which AP should serve a user device?” Instead of an exhaustive probing of all beams on all APs, *UbiG* can rank the APs in the user device’s vicinity using only the current AP’s channel measurement. It achieves this by leveraging a triangulation scheme to predict the gain and direction of the dominating LOS path of these APs. Upon AP ranking, *UbiG* uses the fast probing scheme to identify the best AP and beam to seamlessly steer the user device without affecting the running applications.

We have implemented and evaluated *UbiG* using COTS IEEE 802.11ad devices, that support 64 beams, in various static and dynamic settings. Our experiments show that in 82% of the settings, the SNR loss under *UbiG*’s fast beam alignment is less than 1.5 dB that enables near-optimal beam and bit-rate selection. Under simulated settings, the algorithm scales well even for a large phased-array antenna with 1024 beams. The average SNR loss is less than 1 dB under medium to high SNR conditions. Furthermore, *UbiG*’s predictive “AP ranking”

needs to probe at most 3 APs to identify the best one even in ultra-dense mmWave network deployments; this allows for guaranteed latency and seamless Gbps connectivity.

UbiG, in summary, makes three core contributions:

- (1) We design a fast beam alignment algorithm that can align the best beam between AP and user devices with a few probes (*c.f.*, Section 3); it runs in low-cost COTS hardware and scales well even for future platforms with larger phased-array antenna.
- (2) We design a virtual cell architecture where *UbiG*’s controller can rank the best APs in user’s range without probing (*c.f.*, Section 4); this allows for a seamless AP switch without affecting running applications.
- (3) We implement and extensively evaluate our design in multiple realistic settings (*c.f.*, Section 5 and 6).

2 BACKGROUND ON MMWAVE BEAMFORMING

Millimeter-wave radios use phased-array beamforming to focus the RF energy [14, 35]. For a 1D linear phased-array antenna, the gain pattern of the m^{th} beam can be expressed as [46]:

$$A_m(\theta) = \sum_{n=1}^N \mathbf{w}(n, m) \cdot e^{(j2\pi n d \cos \theta / \lambda)} \quad (1)$$

where N is the number of antenna elements (with uniform separation d), and λ is the wavelength of the wireless signal. The weights $\mathbf{w}(n, m)$ can be tuned to generate different beams with diverse radiation patterns. COTS mmWave platforms typically use a 2D phased-array antenna (Figure 1(a, b)) and thus can generate radiation pattern in both azimuth and elevation plane denoted by $A_m(\theta^{\text{az}}, \theta^{\text{el}})$. Figure 2 shows the azimuth and elevation radiation patterns of two example beams in our IEEE 802.11ad platform. Note that, the beam patterns are far from a perfect shape due to discrete configuration weights [36, 41, 46, 52]. For each beam, however, the pattern A_m is fixed during hardware design time and known *a priori*.

Since a transmitter’s signals can reach the receiver through multiple paths, say K , the mmWave channel created by the m^{th} transmit beam is given by:

$$h_m = \sum_{k=1}^K \sum_{\theta} A_m(\theta) \cdot g_k(\theta_k) \cdot \delta(\theta - \theta_k) \quad (2)$$

where $g_k(\theta_k)$ denotes the complex channel gain of the k^{th} path towards direction θ_k . The gain g_k equals $a_k e^{j\phi_k}$ for path amplitude a_k and phase ϕ_k . The performance of a transmit beam depends on its spatial alignment with the K paths of the mmWave channel; the alignment is captured by the Dirac

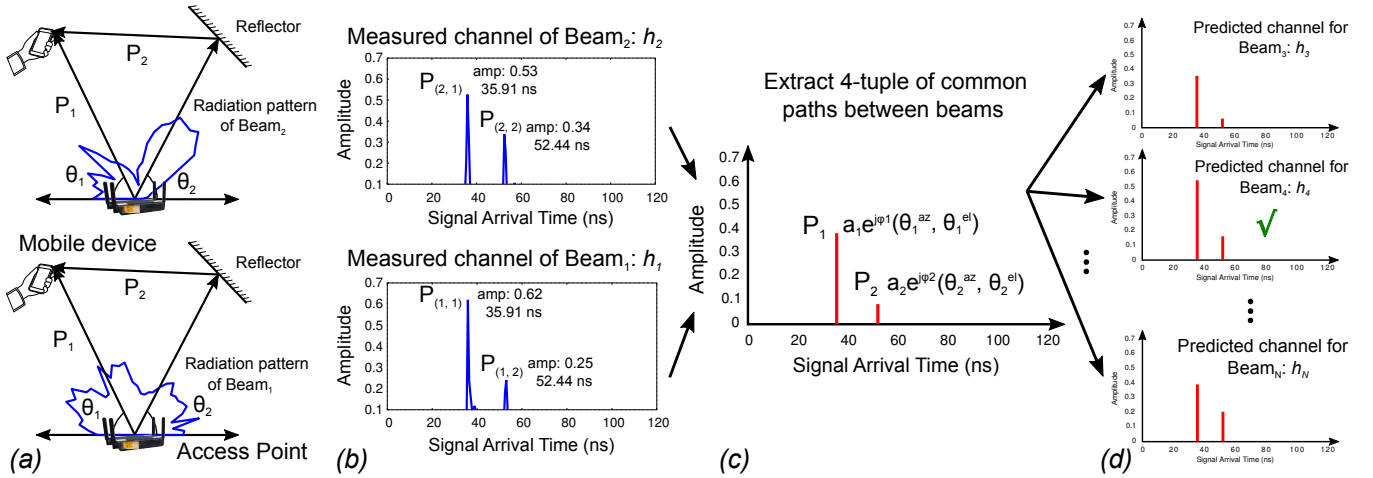


Figure 3: Intuition behind *UbiG*'s fast beam alignment: (a) Two beams with different radiation patterns share the same physical paths to reach the mobile device. (b) Measured channel responses (h_1 and h_2) show the arrival times of the two strong paths under both beams. While the arrival times for both the paths remain exactly same under both beams, the gain of the paths change due to different radiation patterns. (c) Extracting 4-tuple of common paths between the two beams. (d) Mapping the 4-tuple of the paths back to channels of N beams to predict the best.

delta function $\delta(\theta - \theta_k)$. Each path, in practice, can traverse both azimuth (θ_k^{az}) and elevation (θ_k^{el}) directions. Thus, the k^{th} path can be fully characterized by a 4-tuple in the form $(a_k, \phi_k, \theta_k^{az}, \theta_k^{el})$. Since each beam pattern A_m is known, as long as the transmitter can extract such 4-tuples for each of the K paths in the mmWave channel, it can predict the best beam without scanning through the entire space. This is the fundamental principle underlying *UbiG*'s fast beam alignment algorithm.

3 UBIG'S FAST BEAM ALIGNMENT

3.1 Intuition

If the mmWave transmitter and receiver can identify the best beams with a small fixed number of channel measurements, then they can immediately recover from beam misalignment due to blockage and mobility, with negligible response latency. To approach this ideal, *UbiG* leverages a key property of the mmWave channel and phased-array beamforming. While changing beam patterns leads to different channel measurements at the receiver, the underlying physical signal paths traversed by each of the beams remain the same and can be well-defined by the 4-tuple (Section 2). Thus, if the transmitter can extract the 4-tuple of each of the K mmWave paths — typically upper-bounded by 3 for mmWave channel [3, 32, 34, 39, 42, 43, 49] — it can accurately recreate an estimate of the channel for any transmit beam. The key goal of *UbiG*'s beam alignment algorithm is to estimate the properties of the K dominant paths using a fixed number of channel

measurements.

Figure 3 illustrates how *UbiG* approaches the goal. Suppose the mmWave wireless signals from the AP to the mobile device traverse two paths. Let the corresponding 4-tuple of the signal paths be: $(a_1, \phi_1, \theta_1^{az}, \theta_1^{el})$ and $(a_2, \phi_2, \theta_2^{az}, \theta_2^{el})$. Signals along these two paths are amplified by the two different transmit beams as shown in Figure 3(a). While the amplification effects differ due to diverse beam patterns, the arrival time of each path remains exactly the same as shown in the measured time-domain channel of the two beams in Figure 3(b). This is because these paths traverse exactly the same distance even when amplified by two different beams. If the AP can somehow extract the 4-tuple of the paths as in Figure 3(c), it can recreate the channels (Figure 3(d)) of the rest of the beams following Equation (2), since the radiation pattern A_m of each beam is known.

Without prior knowledge of $(\theta_1^{az}, \theta_1^{el})$ and $(\theta_2^{az}, \theta_2^{el})$, unfortunately, it is not possible to estimate the complex gain of the paths (a_1, ϕ_1) and (a_2, ϕ_2) because measured channel is a function of both the gain and directionality of the physical paths and the amplification by the beams towards that direction. *UbiG* computationally searches through the gain and direction of the paths to find the best 4-tuple that generates maximum matching with the measured channels. In the next section, we describe this process mathematically and formulate it as an L2-norm minimization problem.

3.2 Algorithm

3.2.1 Identifying Signal Paths. Let us assume, the mmWave signals from the AP to the mobile device arrive along K distinct paths. Each path is uniquely characterized by the 4-tuple $(a_k, \phi_k, \theta_k^{az}, \theta_k^{el})$, where $k = 1, 2, \dots, K$. Furthermore, assume that the channel measurement of M beams are available to the AP from the mobile device. Each of the channel measurement of the m^{th} beam takes the form as in Equation (2), and the multi-GHz channel measurement captures information of each individual mmWave path. Because each path potentially traverses different length and arrives at the mobile device with a slight delay; thus, the paths can be measured separately in time-domain. Figure 3(b) shows two examples of such channel measurements. Existing COTS IEEE 802.11ad devices, for example, allow 1.76 GHz channel measurement through standard packet decoding pipeline [16]. A mobile device, thus, can measure each potential path with the arrival time difference of at least 0.57 ns *w.r.t.*, another path¹.

It is non-trivial, unfortunately, to measure the complex gain *i.e.*, (a_k, ϕ_k) of each path without prior knowledge of the directions $(\theta_k^{az}, \theta_k^{el})$. This is because the channel measurement of the m^{th} beam captures information about the gain of the path convoluted with the complex gain of the beam pattern (Equation (2)). Although the radiation pattern of each beam is known and fixed *a priori*, the actual beam gain used to amplify the path depends on the unknown direction of the path, *i.e.*, $(\theta_k^{az}, \theta_k^{el})$. So, how do we extract the 4-tuple of the paths? The key observation is that the 4-tuple of the signal paths remain the same even if we measure the channel from other beams. The only thing that changes is the amplification gain of that path due to different radiation patterns. Thus, the contribution of the m^{th} beam on the k^{th} path can be expressed as a simple amplification model:

$$P_{(m,k)} = A_m(\theta_k^{az}, \theta_k^{el}) \cdot a_k e^{j\phi_k} \quad (3)$$

On the other hand, assume that the measured response of the k^{th} path from the measured channel of m^{th} beam is given by $P_{(m,k)}^c$ (*e.g.*, one single spike in the measured plot in Figure 3(b)). There could be potentially many possible directions of the path which can have similar amplitude and phase response after amplification by the beam pattern, however. *UbiG* leverages channel measurements from a few additional beams that amplify the same physical paths. Suppose the channel measurements from M beams are available at the AP-side. Then the 4-tuple of the k^{th} path can be extracted from the optimization:

$$\begin{aligned} \{a_k, \phi_k, \theta_k^{az}, \theta_k^{el}\}^* &= \underset{\{a_k, \phi_k, \theta_k^{az}, \theta_k^{el}\}}{\operatorname{argmin}} \quad ||\{P_{(m,k)} - P_{(m,k)}^c\}_{m=1}^M||^2 \\ \text{s.t., } -\pi &\leq \phi_k \leq \pi, 0 < a_k \leq \frac{\max|P_{(m,k)}^c|}{\min|A_m|} \quad \forall m \in [1, M] \end{aligned} \quad (4)$$

where the bound on the path phase and amplitude are fixed during the optimization solver time. The upper bound on the amplitude a_k is determined based on the ratio between the measured amplitude and the absolute normalized minimum strength of the m^{th} beam.

The above optimization problem, however, is non-convex and constrained over the discrete grid of directions $(\theta_k^{az}, \theta_k^{el})$. Because, in practice, $A_m(\theta_k^{az}, \theta_k^{el})$ for each beam can only be measured with finite resolution on the azimuth and elevation angle². To solve this optimization problem, we leverage a variant of the genetic algorithm [12] with mixed integer constraints. A well known problem with the non-convex optimization solver, however, is the issue of local minima convergence. We overcome this issue by applying a random variable initialization and multiple seed generation. The above optimization runs for each of the K paths. In practice, K is bounded since mmWave wireless channel is typically sparse containing only 2 or 3 dominating paths [3, 32, 34, 39, 42, 43, 49]. Once the 4-tuple of all the K paths are extracted, *UbiG* can leverage the channel model in Equation (2) to predict the performance of each individual beam.

Number of probings: The optimization formulation in Equation (4), intuitively, needs to solve for $4K$ unknown variables (4 variables for each of the K paths). Thus, the optimization needs at least $4K$ beams' channel measurement, as each of the channel may reliably measure any number of paths between 1, 2, \dots , K . With K set to 3, the $4K$ measurements really translates to only 12, independent of the number of beam N from a phased-array. Multiple issues such as measurement noises, local minima convergence, and discrete grid search in the optimization solver, however, can cause the transmitter to align onto a sub-optimal beam choice after $4K$ probes. Additional δ probings, in practice, can refine the prediction and lead it close to optimum. We will later demonstrate through experiments that a small δ suffices.

Choice of beam directions: A key question still remains: how do we select the $4K$ beams for probing? *UbiG*'s optimization module, intuitively, achieves better convergence if the measured channels are as much uncorrelated as possible. In order to reduce the correlation, *UbiG* selects to probe $4K$ beams with as much non-overlapping radiation pattern as possible.

¹Thus, as long as two dominating paths have length difference equal to the span of human palm (~ 17 cm), they can be measured separately.

²Our platform vendor only provides measurement at 5° resolution.

The *structural similarity* (SSIM) metric [47] between the 2D beam patterns (e.g., Figure 2) is used to measure the overlapping between two beam patterns. *UbiG* probes an additional $4K - 1$ beams from the set of N available beams that has the lowest SSIM with the current beam. Note that, the SSIM ranking can be computed offline and only involves a simple table lookup operation at runtime.

3.2.2 Challenges in COTS Hardware. Applying the optimization in straightforward manner on commodity platform, however, is challenging for two reasons.

Incoherent channel measurement: The phases of the channel measurement across the $4K$ probes are incoherent. This is a well-known problem in any COTS mmWave devices (e.g., the IEEE 802.11ad devices in Facebook's Terragraph project [9, 36] and our experimental platforms Figure 1(a, b), [41]). Since the COTS platforms use cheap high frequency oscillators, tracking the carrier phase across multiple channel probing packets is infeasible. Fortunately, even if the phases of the channel measurement across multiple probe packets are incoherent, the relative phases of the channel within a single probe are still coherent, as they only depend on the length difference between the mmWave paths and free from any hardware idiosyncrasies. We leverage this observation to first sanitize the phase of the entire channel of a single probe packet using the phase of the first-arrival-path (i.e., LOS path) and then modify the optimization problem in Equation (4) as follows:

$$\begin{aligned} \{a_k, \phi_k, \theta_k^{az}, \theta_k^{el}\}^*|_{\forall k} &= \operatorname{argmin} ||\{\hat{P}_{(m,k)} - \hat{P}_{(m,k)}^c\}_{m=1}^M||^2|_{\forall k} \\ \text{s.t., } \hat{P}_{(m,k)} &= P_{(m,k)} \cdot \exp(-j\angle P_{(m,1)}) \\ \& \hat{P}_{(m,k)}^c &= P_{(m,k)}^c \cdot \exp(-j\angle P_{(m,1)}^c) \quad \forall m \in [1, M] \end{aligned} \quad (5)$$

where $\hat{P}_{(m,k)}$ and $\hat{P}_{(m,k)}^c$ denote the phase calibrated modeled and measured channel respectively of each k^{th} path under m^{th} beam calibrated against the first-arrival-path ($P_{(m,1)}$ and $P_{(m,1)}^c$). Since the phase of the individual paths are now dependent on the phase of the LOS path, the optimization problem can be solved by joint extraction only (denoted by $|_{\forall k}$ in Equation (5)) of the 4-tuple of all K paths.

A side-effect of this joint extraction is that the solution space is biased towards the most dominating path (typically the open LOS path). Extracting the 4-tuple of the other multi-paths, thus, becomes challenging. We overcome this by leveraging the core intuition from existing sparse-recovery algorithm by orthogonal matching pursuit [5]. The key idea is to iteratively solve Equation (5) K times, and after each iteration, remove the most dominating path (i.e., estimated k^{th} path

Algorithm 1 *Channel Pruning*

- 1: $M = 4K$, Channel length: τ , Tap separation: $\Delta = 3$
 - 2: Input: Measured channel $\mathbf{h} : (h_1, \dots, h_M)$
 - 3: Output: Pruned channel $\hat{\mathbf{h}} : (\hat{h}_1, \dots, \hat{h}_M)$
 - 4: $\mathbf{H}_0 \leftarrow \text{FFT}(h_1, \dots, h_M)$
 - 5: **foreach** $\tau_i \in (1, \dots, \tau)$
 - 6: Assign $h'_j \leftarrow h_j$; $h'_j(\tau_i) = 0 \quad \forall j \in [1, M]$
 - 7: $\mathbf{H}' \leftarrow \text{FFT}(h'_1, \dots, h'_M)$, $\rho(\tau_i) = |\mathbf{H}' \cdot \mathbf{H}_0^H|$
 - 8: Ascend sort ρ and pick first K indices τ'_1, \dots, τ'_K ensuring $|\tau'_i - \tau'_j| \geq \Delta \quad \forall i, j \in [1, K]$
 - 9: Assign $\hat{h}_j(\tau'_i) = h_j(\tau'_i) \quad \forall j \in [1, M], \forall i \in [1, K]$
-

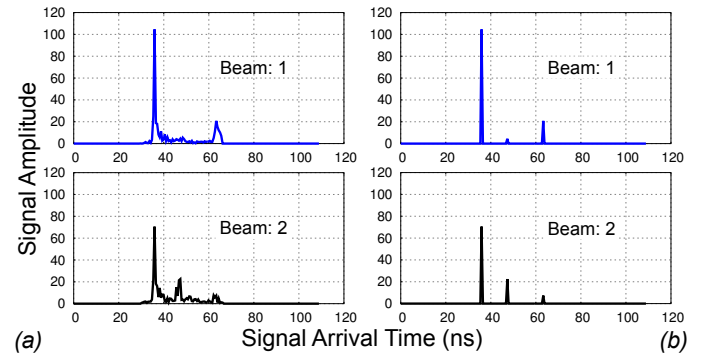


Figure 4: (a) Measured channel responses of two beams. (b) Output of the channel pruning Algorithm 1.

with strongest amplitude) from both model $\hat{P}_{(m,k)}$ and measurement $\hat{P}_{(m,k)}^c$. This allows to extract the 4-tuple of all the K paths reliably.

Extracting common paths: Still, COTS platform provides limited channel measurement resolution due to its ADC limitations. Identifying a common set of K paths from $4K$ beams' measurement, thus, becomes non-trivial. This is because, a single beam may not amplify all the K paths reliably; depending on the radiation pattern, the amplification towards certain paths may be close to noise floor. To overcome this, *UbiG* first collects the channel measurement from all $4K$ beams, and then iteratively prunes the measurements as described in the Algorithm 1. At a high level, Algorithm 1 prunes each individual path from the measured channel of all $4K$ beams, and measures the sum correlation between the original and pruned channel. Since removing the dominating paths destroys the measured channel, *UbiG* identifies K dominating paths that on pruning show the lowest sum correlation. Figure 4 also shows an example pruned result of the measured channel from two different beams in our platform.

3.3 Integrating *UbiG* with IEEE 802.11ad

UbiG's fast beam alignment algorithm can be integrated with IEEE 802.11ad COTS mmWave devices seamlessly. At a very high level, the *UbiG* transmitter and receiver follow the similar beam adaptation principle as in IEEE 802.11ad [16]. A receiver goes to quasi-omni directional mode and transmitter scans the directions with its own narrow beam directions. Instead of scanning all N beams, a *UbiG* transmitter, however, scans a fixed number of $4K$ beams and collects the feedback of 1.76 GHz channel measurement from the receiver. The provision to feed channel measurement back is already available in IEEE 802.11ad standard-compliant devices as an optional feature [16]. This allows the transmitter to rank all N beams according to their predicted performance. Finally, an additional probing on the first δ^3 beams in the ranked order can align the best beam direction. The receiver can follow the same procedure to align its own beam direction. The alignment algorithm can be triggered whenever there is a SNR change of 1.5 dB which is the average SNR difference between two MCS options in IEEE 802.11ad [16]. The entire $4K + \delta$ search allows the standard compliant devices to align the best beam within less than $300 \mu s$ under all protocol overheads even for a 1024 beams system. Existing COTS solution, in contrast, will require multiple seconds to identify such best beam alignment [15, 42, 54].

3.4 Discussions

Dominating paths colliding on same time bin: Note that, IEEE 802.11ad devices cannot measure two dominating paths with arrival time difference less than 0.57 ns separately. In practice, however, the mmWave dominating paths are typically well separated in space and arrival times [34, 37, 43]. In the worst case, if collision happens — *i.e.*, the paths arrive within 0.57 ns — the optimization solver assumes a single path which may potentially lead to convergence to local minima and a wrong beam alignment. With the additional δ refinement, however, this problem can still be alleviated, as we will show in our experiments.

Feedback collection overhead: For each of the $4K + \delta$ beams, the mobile device needs to feedback the measured channel to the AP. Current standards such as IEEE 802.11ad already incorporate such feedback mechanisms in their beam alignment protocols. *UbiG* only needs to piggy-back the complex gains of the K paths on a single feedback packet. Such a feedback packet only takes less than $1 \mu s$ in IEEE 802.11ad.

Search overhead under multi-users: While *UbiG* can align the beam with a few probes with a single user, the searching latency will potentially grow with the number of users in the network. The AP, however, can coordinate the multiple

users during probing as in [36], thus eliminating this issue. Specification for such coordination mechanisms is unavailable in existing standard COTS platforms [16], and we leave the extension and implementation for our future work.

4 ENABLING UBIQUITOUS GBPS ACCESS

Efficient beam alignment between a single AP and the user device, however, does not guarantee ubiquitous mmWave connectivity [1, 41, 43, 48]. Rampant blockage from several obstacles, including user's own body in close proximity to the device, may potentially block all the available beams. To enable seamless mmWave connectivity, *UbiG* leverages co-operation between multiple nearby APs and switch to the unblocked one before the link suffers from a catastrophic outage.

Similar to the existing virtual cell architectures in enterprise Wi-Fi networks [24, 40], *UbiG* leverages "thin APs" coordinated by a centralized controller. *UbiG*'s APs, specifically, share the same IEEE 802.11 basic service set identifier (BSSID) and appear as one AP to user devices. So, a user device needs to associate and authenticate with only one AP which will then synchronize the device's association state with all other APs in the network. This enables the users to seamlessly roam among the APs without repeating the AP scanning/authentication/association processes. Compared with the conventional Wi-Fi, however, realizing seamless mmWave connectivity introduces two key challenges.

Prohibitive AP selection latency: To select the best AP, each nearby AP needs to first identify the best beam to the device among tens to hundreds even thousands of beam options [36]. AP selection latency, thus, can easily reach the order of seconds. Even if each AP uses *UbiG*'s fast beam alignment, the latency grows with AP density significantly and may disrupt the upper layer protocols (*e.g.*, TCP [41]).

Coarse handoff: Although the virtual cell architecture allows for rapid switching among APs, it cannot guarantee seamless handoff under dynamic mmWave channels. For example, if the mmWave link between AP_i and $user_k$ rapidly breaks due to blockage, the buffered packets at AP_i will be dropped, even if $user_k$ can quickly handoff to an unblocked AP_j . WGTT [40] solves this problem by forwarding each downlink packet to all APs within user's vicinity. Such packet replication may easily overload the inter-AP backhaul, considering the Gbps per-user requirements in IEEE 802.11ad applications.

UbiG solves the above challenges by leveraging a zero-latency AP ranking and Wi-Fi traffic offloading, as we discuss next.

³ $\delta = 8$, by default.




Figure 5: (a) *UbiG* ranks group of potential APs based on the geometrical extrapolation of the dominating LOS path. (b) Extrapolating gain and directionality of LOS path from AP_2 using probe packets from only AP_1 . (c) Distribution of predicted beam strength of AP_2 and AP_3 .

4.1 AP Ranking

UbiG's AP ranking algorithm allows the controller to keep track of a small set of "good" APs, and quickly switch to the best one before the current AP fails. This is better illustrated with the example in Figure 5. Suppose that a user device currently uses AP_1 for data transmission, but has $M - 1$ alternative APs in its vicinity. The *UbiG* controller can rank the $M - 1$ APs *without additional probing* based on two key observations:

- (1) When the LOS between the AP and user device is open, all the beams' performance mostly depends on it. Thus, by simply estimating the gain and direction of the dominant path, *UbiG* can predict the potential link performance of an AP.
- (2) The LOS dominating paths between two APs and the device form a virtual triangle in space (Figure 5(b)). This allows the controller to geometrically estimate the gain and direction of the dominating LOS path from the second AP without probing it.

UbiG uses the above observations as follows. First, it leverages the channel measurements from the fast beam alignment to estimate the gain g_1 and direction θ_1 of the LOS path from the current AP_1 to the user device. Then, g_1 and θ_1 information is further sent to the controller that uses it to extrapolate the LOS path from the other $M - 1$ APs to the user device without probing. Specifically, from the LOS path gain g_1 , *UbiG* estimates the distance d_1 between AP_1 and the device. Since, the channel measurement already separates multiple paths, the distance can be directly measured following the free-space path loss model as: $d_1 = \sqrt{\alpha/|g_1|}$. Here, α is the transmit power of the AP^4 , and $|g_1|$ denotes the magnitude of the complex gain g_1 . The estimated AP-user distance d_1 , along with the relative distance d_{AP} and direction θ_{AP} between APs

— known *a priori* during network deployment time — allow *UbiG* to construct a virtual triangle among AP_1 , AP_2 , and user device (Figure 5(b)). Now, the distance and relative angle between the candidate AP_2 and the user can be estimated as:

$$d_2 = \sqrt{d_1^2 + d_{AP}^2 - 2d_1d_{AP}\cos(\theta_1')}, \quad \frac{d_2}{\sin(\theta_1')} = \frac{d_1}{\sin(\theta_2')}$$

where $\theta_1' = \theta_1 + \theta_{AP}$. Finally, the gain and direction of the LOS path from the candidate AP_2 to the user is:

$$|g_2| = \alpha/d_2^2, \quad \theta_2 = 180^\circ - \theta_2' - \theta_{AP}$$

The triangle of Figure 5(b), in practice, lies in 3D space since the signal paths have two directional components (azimuth θ^{az} and elevation θ^{el}). The extension for 3D case is trivial and hence omitted here.

The estimated gain and directionality allow the controller to predict the dominant channel of all N beams from AP_2 to user device following Equation (2). This further allows the controller to create a predictive beam strength distribution profile, denoted as $P(AP_m)$, of all APs to the user device. Figure 5(c), for example, shows the beam strength distribution of the two APs, AP_2 and AP_3 , predicted by simply reusing the channel measurements from fast beam alignment at AP_1 . The actual ranking $R(AP_m)$ of AP_m is determined based on the average of the median and maximum predicted beam strength as follows:

$$R(AP_m) = \frac{\tilde{P}(AP_m) + \max(P(AP_m))}{2}$$

where $\tilde{P}(AP_m)$ denotes the median of $P(AP_m)$. This ranking allows *UbiG* to sort the M APs in the vicinity and keep track of the best AP.

⁴ α can be updated during run-time if the AP employs transmit power adaptation.

UbiG's AP ranking, however, assumes that each AP has one open LOS path to the user device. Thus, it may not be able to achieve an absolute ranking when some LOS paths are blocked. *UbiG* alleviates the potential impact by picking the first M' best APs from the ranked list and estimating their performance using the fast beam alignment (Section 3.2). Our experiments will show that a small number of $M' = 3$ APs can ensure accurate AP selection, even in ultra-dense AP deployment (Section 6). Finally, *UbiG*'s controller selects one of the M' AP with highest SNR to serve the user device. Furthermore, *UbiG*'s AP selection can be extended to consider additional metrics such as traffic load and channel utilization.

AP switching cost: Our experiments show that AP switching latency (*i.e.*, T_{sw}) can be on the order of tens of ms in our implementation. *UbiG*'s AP switching decision takes into account such overhead. It switches from the current AP_i to the best AP_j only if the aggregate data transmission time at AP_j and the switching latency is smaller than the transmission time at the current AP_i ($T_{AP_j} + T_{sw} < T_{AP_i}$). The transmission time is calculated based on the measured or predicted bit-rate from each AP under a unit traffic load (*e.g.*, amount of data that can be delivered using the lowest MCS within the typical period that a user dwells on an AP).

4.2 Traffic Offloading

UbiG's AP ranking and fast switching, unfortunately, may not guarantee robust connectivity under certain cases where all APs in a device's vicinity are blocked (*e.g.*, the device's phased-array antenna is covered by user's hand). In addition, even if *UbiG* can identify high-throughput APs in a device's range, the current AP may host pending packets for the user device which may be dropped upon blockage. Inspired by previous work [41] and the session transfer mechanism in IEEE 802.11ad [16], *UbiG* leverages multi-band chipsets' capability, and *proactively* switches to the more robust Wi-Fi upon detecting a blockage.

Interface switching: Existing work [41] requires periodic probing at Wi-Fi to identify 60 GHz blockage. We argue that this approach has limitations to scale in ultra-dense AP deployments where an AP may serve more than 100 devices. Differently from [41], *UbiG*'s interface switching relies solely on 60 GHz channel measurements to proactively identify blockages before 60 GHz link breaks. Our extensive measurements have shown that when LOS path gets blocked, 60 GHz link is very likely to break, especially when strong reflectors are unavailable or far away. This observation echoes previous field tests [41]. *UbiG* extracts the LOS path using channel measurements (*i.e.*, LOS path is the first tap in Figure 3), and switches to Wi-Fi if its channel amplitude is below a predefined threshold. In our implementation, interface switching only takes 1.5 to

2 ms and does not affect the application (*e.g.*, running over TCP) performance.

4.3 Discussions

Network deployment density: While existing enterprise Wi-Fi deployments are already dense (*e.g.*, 1 AP per 10 m [41]), we expect future mmWave deployments will be even denser. The exact density requirement is still an open question at mmWave domain and depends on various factors (*e.g.*, network deployer's incentives and management overhead), and we leave the investigation as future work.

AP channelization: *UbiG*'s virtual cell architecture requires all APs to be configured on the same channel. Single channel configuration prevents channel scanning overhead and has been shown to allow for seamless handoffs in Wi-Fi [24]. Due to its inherently better spatial reuse [28], directional mmWave links suffer less from interference than Wi-Fi, even under dense deployments. We leave the study of interference management as future work.

5 IMPLEMENTATION

We implement and evaluate *UbiG* by collecting channel measurements from our COTS testbed (Figure 1(a, b)). The testbed includes AP and user devices each equipped with a 4×8 element phased-array antenna that can generate 64 beams in 3D space. Both the AP and user devices are IEEE 802.11ad-compliant [16], and support bit-rate upto 4.62 Gbps by operating on 1.76 GHz bandwidth at the unlicensed 60 GHz spectrum. Our AP platform runs an open-source OpenWrt [27] codebase on a dual-core 1.7 GHz CPU. The 60 GHz RF front-end is controlled via a peripheral baseband chipset that can steer the beam directions by exchanging control and coordination packets through the kernel and firmware commands. Each device can store the 1.76 GHz channel measurement from the "Channel Estimation" (CE) header field of the IEEE 802.11ad data packets [16]. The current COTS drivers and firmwares (*e.g.*, [19]), however, do not allow us to extract the channel measurement from the hardware. To overcome this issue, we use vendor-provided custom-built firmware with customized APIs to extract the channel measurement to the kernel space of the devices. We expect such API updates to be available in the COTS drivers and firmwares soon.

The 1.76 GHz bandwidth allows our devices to measure the arrival time of each mmWave path with 0.57 ns resolution. The CE field itself spans only 650 ns as per the IEEE 802.11ad standard [16]; thus, fast probing and collecting feedback from the user can finish within less than 300 μs , including all protocol overheads. Our user devices, unfortunately, do not support feeding back the channel measurement in real-time to the AP




Figure 7: Floorplan of the experimental area. Few of the user device's placements are shown with circle.

via the 60 GHz interface due to the current firmware's limitations. We feedback the measured channel using the devices' co-located Wi-Fi interface. Under this setup, however, a single measurement requires more than 600 ms due to the delays introduced in various software stacks and kernel-firmware coordinations. This limits our evaluation to only static users. While we leave the mobile user evaluation of *UbiG* as future work, we note that a software-only implementation of the genetic algorithm on the resource-constrained and cheap platform is indeed challenging. Customized and real-time hardware implementation of genetic algorithms, however, is fairly known and can converge in tens of microseconds [38, 45]. In future, vendors may choose to incorporate *UbiG*'s fast beam alignment and AP switching algorithm in customized hardware-firmware implementation to run in real-time. We leave such real-time implementation for our future work.

6 EVALUATION

In this section, we evaluate *UbiG*'s performance by conducting testbed experiments in a typical indoor office space that include cubicles, furniture, *etc.*, creating a rich multipath environment. We first conduct single-AP experiments where the user devices are randomly placed in more than 200 locations across a $10 \times 7m^2$ area (Figure 7). The AP is mounted at 10 ft height on top of a tripod stand. We further evaluate the efficacy of our AP ranking and switching algorithm by deploying 3 to 10 APs connected via an Ethernet backhaul to a central controller.

6.1 Micro-benchmarks

Effectiveness of fast beam alignment: We first empirically verify whether the model used by *UbiG*'s fast beam alignment (Section 3.2) can accurately predict the best beam. We measure the signal-to-noise ratio (SNR) loss of the model in comparison with an oracle beam alignment scheme that runs *exhaustive search*. The *exhaustive search* is performed by steering all 64 beams from the AP and user side and measuring the “ground-truth” channel information. The pair that yields the best SNR

performance is the optimal choice. We evaluate *UbiG*'s SNR loss for both *single* and *multi path* scenarios. We identified 84 *single path* settings which show only one spike in the measured channel of all beams after applying our channel pruning algorithm (Section 1)⁵. The rest of the settings, which show multiple spikes in the channel measurement, are considered *multi path*. Figure 6(a) shows that in the *single path* case, the median SNR loss is only 1.43 dB with $4K$ ($K = 3$ by default) probings. While the worse case SNR loss can go up 7.7 dB, the loss quickly drops with additional δ refinement over the $4K$ probes. Figure 6(a) shows that with just $\delta = 4$ probing refinement the median SNR loss can be contained within 0.39 dB and worse case by 2.87 dB. With 8 additional probings, the 90th percentile SNR loss is contained within 1.5 dB. The effect from this loss can be overcome by the rate (MCS) adaptation algorithm since the average SNR difference between two MCS options is typically bounded by 1.5 dB [16]. Unless otherwise mentioned, we will use $\delta = 8$ for the rest of our experiments.

As expected, the SNR loss under *multi path* cases is higher. Specifically, Figure 6(b) shows that the median SNR loss is 3.3 dB, while the worse case SNR loss can go upto 17.4 dB with $4K$ probings. The key reason is that it becomes harder to estimate the multipath profile under the limitations of off-the-shelf platforms as we described in Section 3.2.2. The error quickly drops, however, with additional refinements. With the additional 8 probings, the average SNR loss is only 0.79 dB and 90th percentile loss is within 2.2 dB. In all, *UbiG*'s fast beam alignment protocol yields very close to optimal result in both *single* and *multi path* environments.

UbiG's accurate beam alignment allows for near-optimal rate selection compared to the ground-truth, as shown in Figure 6(c) and (d) for both *single* and *multi path* settings. Rate under-selection can still occur under low SNR conditions, because the channel measurements become noisier, yielding less accurate prediction from the optimization model. The average rate loss under the fast beam alignment, however, is only 6% compared to the ground-truth.

Effectiveness of SSIM probing: Recall that *UbiG* leverages the SSIM metric between beam patterns to decide which $4K$ beams to probe to extract the 4-tuple properties of the K paths for beam alignment. Figure 8(a) shows the effectiveness of SSIM probing compared to a random probing scheme. While the median SNR loss under $4K$ probe remains the same under both SSIM and random probing, SSIM improves the worse case (last 10-percentile SNR loss). Under single path scenarios, the worse case SNR loss with SSIM probing can be upto 7.7 dB, while random probing can exacerbate this SNR loss upto 17.42 dB. This loss is amplified under multi path cases where

⁵Multiple spikes with less than 1% energy of the first-arrival-path are still considered as a *single path* case.




Figure 6: Effectiveness of *UbiG*'s fast beam alignment: CDF of (a) SNR loss under *single path* scenarios. (b) Under *multi paths*. (c) Selected rates comparison with ground-truth under *single path*. (d) Under *multi paths*.




Figure 9: Performance scalability of *UbiG*'s fast beam alignment under three SNR conditions: (1) Poor (< 5 dB). (2) Medium (5 ~ 15 dB). (3) Good (> 15 dB). Error bar shows 90th percentile SNR loss.




Figure 8: Effectiveness of SSIM probing: (a) Worse case (last 10 percentile) improvement under 4K SSIM vs. random probing. (b) Additional probing beyond 4K to contain SNR loss under 3 dB.

the worse SNR loss under random probing can be upto 40.24 dB! Because, under random probing the correlation of the measured channel of $4K$ can not be reduced and thus the optimization solver yields poor result in aligning the best beam direction. Note that, this SNR loss can be reduced under both scenarios by employing additional δ refinements as well. Under random probing, however, it needs more refinement to reduce the SNR loss. Figure 8(b) shows that under both single and multi path cases, SSIM probing can confine the SNR loss within 3 dB with 8 additional probings. In contrast, to achieve the same performance, the random approach needs upto 14 additional probings.

Performance scalability: Our experimental platform supports a small phased-array (4×8), and can generate only 64

beams. Thus, we resort to a large-scale simulation to understand the scalability of the fast beam alignment protocol. We build a ray tracing simulator to represent a large $50 \times 50 m^2$ space with multiple specular and point reflectors that generate a multipath rich environment. The simulator uses the classical Saleh-Valenzuela multi-path channel model [22]. We mounted the AP at one corner of the indoor area and placed the user device in 300 different locations. To compare the results under different sized phased-array antennas, we vary the number of antenna elements from 32 to 512 in a 2D antenna array. For each of the antenna array, we follow a standard codebook [4] to generate 64 to 1024 beam patterns. For each device location, the AP steers all its beam directions to measure the channel and SNR condition at the user device. Then, we apply *UbiG*'s fast beam alignment algorithm⁶ to predict the best beam. In order to understand the behavior of the alignment protocol under various SNR conditions, we bucketize the measurements into three conditions: (1) "Poor" SNR: best beam can achieve maximum SNR up to 5 dB. (2) "Medium" SNR: best beam can achieve SNR between 5-15 dB. (3) "Good" SNR: best beam can achieve SNR beyond 15 dB. Figure 9 compares the SNR loss of *UbiG*'s fast beam alignment over the *exhaustive search*.

The simulation shows multiple interesting results. Under "medium" SNR conditions (Figure 9(b)), *UbiG*'s fast beam alignment can already approach the optimal result even with 1024 beams. Whereas a small $4K$ probing can still have 10.9 dB SNR loss

⁶ K is set to 3, in consistent with prior empirical studies [3, 39, 43].




Figure 11: Minimum number of APs to probe to find the best AP. Mean and std are calculated from 10 random deployment in $10 \times 7 m^2$ area for each density setting.

under 1024 beams, this quickly converge to only 1.6 dB with additional 8 probings. Under “good” SNR conditions, the SNR loss with 1024 beams can be even smaller, down to 0.99 dB. Under “poor” SNR condition (best beam SNR < 5 dB), however, *UbiG* can not align the best beam with high accuracy for large phased-arrays. While the mean SNR loss can be under 2.3 dB for a 256-beam phased-array with the fast beam alignment, this loss can be very high under 1024 beams with 90th percentile upto 26.2 dB! Because, under “poor” SNR condition, the channel measurement becomes highly noisy and gives noisy input to *UbiG*’s channel pruning (Algorithm 1) and optimization solver (Section 3.2). The error under “poor” SNR, however, has minimal impact on *UbiG*’s overall goal of ubiquitous Gbps connectivity. Under poor SNR conditions, *UbiG* will rapidly switch the user device to a new AP that provides better SNR.

Effectiveness of AP ranking: *UbiG* leverages cooperation among APs to enable ubiquitous Gbps connectivity. Specifically, *UbiG* first ranks the APs to identify the candidate APs that could potentially support Gbps connectivity when the primary AP is blocked. We first evaluate the efficacy of the AP ranking scheme through an example scenario. We deploy 10 APs in the $10 \times 7 m^2$ indoor area (Figure 7) to emulate an ultra-dense deployment and connect a user device to one of them (AP_2). We then leverage the channel measurements from the fast beam alignment of the primary AP_2 and employ the AP ranking algorithm (Section 4.1) to predict the beam strength distribution profile of all 10 APs. Figure 10(a) and (b) show the ground-truth and predicted beam distribution profile from all 10 APs to the user device. Each curve represents the relative strength of all the beams from the AP to the user device. The ground-truth measurement in Figure 10(a) shows AP_5 to be the best one, since it can establish the strongest beam to the user device. While *UbiG*’s predicted beam strength distribution profile in Figure 10(b) does not match with measured profile exactly in absolute sense, the prediction correctly ranks AP_5 as the best one (also c.f., Figure 10(c)). Another key point to note here is that, the best AP prediction is achieved using only channel measurements from the fast beam alignment of

the AP_2 where the user is currently connected to.

We further evaluate *UbiG*’s AP ranking under different AP densities (3-10 APs in our floorplan area). For each AP deployment, we place the user device in 10 random locations and identify the ground-truth best AP through extensive measurement search. We then compare the effectiveness of *UbiG*’s AP ranking by utilizing the channel measurement from the fast beam alignment of one of the APs. The results in Figure 11 show that *UbiG*’s AP ranking scheme can find the best AP for each setting. The best AP always lies within the first 3 APs of *UbiG*’s ranked order, even under an ultra-dense deployments of 10 APs.

6.2 TCP over *UbiG*

Next, we evaluate the direct benefit of *UbiG*’s fast beam alignment algorithm to upper-layer protocols, in particular TCP which is known to be sensitive to link dynamics. Since we are unable to run the optimization solver of Equation (5) in real-time due to hardware constraints, we first calculate the latency of the beam alignment algorithm and then emulate the impact of such latency on TCP. Specifically, we modify the kernel of our platform to drop packets at the 60 GHz radio during the beam alignment phase, which mimics the effect of beam alignment time on TCP⁷. Figure 13(a) shows the end-to-end latency of a TCP flow from an AP to a user (5m away). As expected, the TCP latency is strictly affected by the underlying beam alignment latency, growing upto 764.43 ms for 512 beams when running IEEE 802.11ad. This is clearly unacceptable for real-time applications. The TCP end-to-end latency further grows with the number of users in the network, and can go upto 2 seconds. *UbiG*’s fast beam alignment, in contrast, can limit the end-to-end latency to less than 20 ms.

Figure 12 provides a microscopic exposition of how the beam alignment latency is amplified at the TCP level. We observe that *UbiG* has no effect on the TCP congestion window (cwnd) adaptation⁸ for either 64 or 1024 beams (Figure 12(a, c)). For IEEE 802.11ad, however, even tens of ms of alignment latency for 64 beams can disrupt the cwnd growth (Figure 12(b)). The effect is even worse for 1024 beams. Figure 12(d) shows that the cwnd completely shuts down after 3.2 s of beam alignment due to TCP time-out, and takes additional 6.6 s to recoup and start growing back. Finally, TCP’s cwnd dynamics directly affect the TCP throughput. Our *iPerf* measurement (Figure 14(a, b)) shows that *UbiG* maintains a relatively stable TCP throughput of around 2 Gbps. In contrast, several seconds of TCP disruptions can occur while running the IEEE 802.11ad beam alignment.

⁷During beam alignment no data packets are transmitted.

⁸We use default TCP CUBIC[20] in our OpenWrt AP platform.

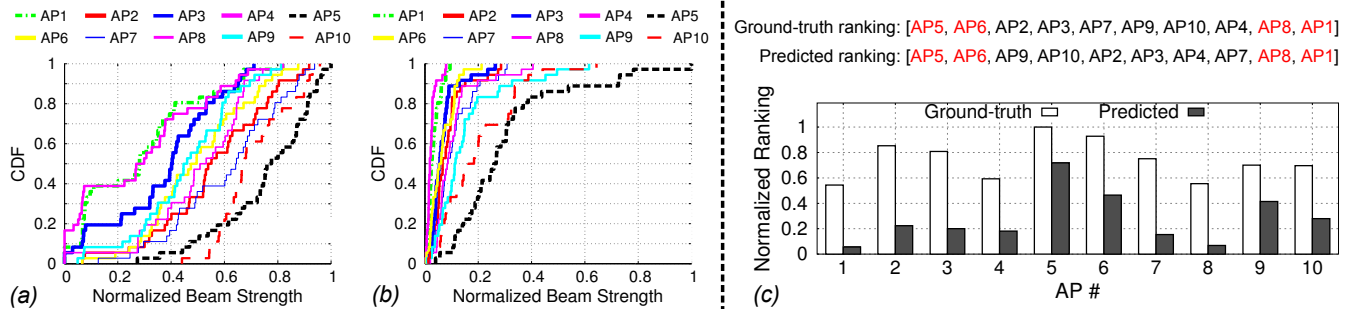


Figure 10: Effectiveness of AP ranking: (a) Ground-truth relative beam strength distribution profile for 10 AP deployment in the vicinity of a user device. (b) Predicted distribution profile, based on channel measurements from the fast beam alignment of the AP_2 that the user is connected to. (c) Ranking order of the 10 APs shows that *UbiG* can pick the best AP, i.e., AP_5 , even if the absolute profiles between (a) and (b) do not exactly match.




Figure 12: Effect of beam alignment latency on TCP adaptation: (a) 64 beams with *UbiG*. (b) 64 beams with IEEE 802.11ad; Congestion window is disrupted even under 10s of ms alignment latency. (c) 1024 beams with *UbiG*. (d) 1024 beams with IEEE 802.11ad; Congestion window entirely shuts-down 3.2 seconds after beam-searching is triggered. Then it takes additional 6.6 seconds to recoup and start growing.




Figure 13: TCP end-to-end latency: (a) Single user. (b) Latency per user when AP is serving 5 users.




Figure 14: *iPerf* TCP throughput under *UbiG* and IEEE 802.11ad for (a) 64 beams. (b) 1024 beams.

7 RELATED WORK

Beam alignment algorithms: The standard *multi-level* beam searching used by the COTS IEEE 802.11ad devices can take hundreds of ms to few seconds to find the best beam alignment [41, 54]. State-of-the-art compressed sensing solutions [23, 30, 31] have improved scalability by requiring only $\mathcal{O}(K \log(N))$ probes, for N beam options and K channel paths. These proposals, unfortunately, require phase-coherence across measurements or customized phased-array antennas [2, 15], which are not supported by commodity platforms.

Recent non-coherent designs can be implemented on commodity platforms [7, 36]. They use coarse SNR (and not channel response) measurements, however, and hence they have limitations in aligning the beams in multipath mmWave channels. BeamSpy [43] allows to infer the best beam alignment with fast search, but works only for quasi-stationary links. The system proposed in [21] allows for zero overhead beam alignment by implementing a two-lobe beam pattern at the receiver. The system, however, requires specific antenna hardware support.

It is unclear how it performs with the imperfect beams supported by commodity devices (*c.f.*, Figure 2). Finally, a few out-of-band solutions *e.g.*, [26, 41, 48] require additional hardware, such as MIMO Wi-Fi radios, location sensors and depth cameras, to identify the best beam. Different from all existing works, *UbiG* is able to identify the best beam in dynamic mmWave settings with a very few probe packets. It works with today's cheap COTS mmWave hardware and scales well for future platforms supporting very large phased-array antennas.

AP switching: Many AP switching solutions have been proposed in lower-frequency Wi-Fi and cellular networks (*e.g.*, [8, 25, 29, 40, 51]), but the challenges remain largely unexplored in mmWave networks. The concept of “picocell” switching was recently discussed in [36, 54], but the path tracking from several “picocell” base-stations will likely overwhelm the link connection during continuous movement. Pia, a pose information assisted AP switching mechanism, has been recently proposed to make mmWave networks robust [48]. It requires precise location and orientation sensing capability on mobile devices which is available only on specialized mobile hardware (*e.g.*, certain VR headsets), however. Further, the FoV based AP switching solution does not work when the FoV of the mobile devices itself changes during partial blockage or mobility. *UbiG* is the first system that provides Gbps mmWave connectivity in a dense realistic mmWave deployment with a scalable AP-switching solution and requires no out-of-band sensors.

8 CONCLUSION

In this paper, we design and implement the key building blocks of a scalable and ubiquitous Gbps mmWave wireless network, which can provide latency guarantees, even using today's cheap COTS hardware. *First*, we introduce an efficient beam alignment algorithm that can identify the best beam for a mmWave link with guaranteed latency, even when the devices are equipped with large phased-array antennas. *Second*, we devise an infrastructure-side predictive AP switching solution to guarantee consistent Gbps connectivity under device mobility and blockage. Our experimental results show that *UbiG* performs close to an “Oracle” solution which instantaneously knows the best AP-user pair and the strongest communication beam. We believe the algorithms in *UbiG* can be the key building blocks for the next-generation 5G networks consisting of dense mmWave deployments.

ACKNOWLEDGEMENT

We sincerely thank the anonymous reviewers and shepherd for their valuable comments and feedback. Sanjib Sur and Xinyu Zhang were partially supported by the NSF under Grant CNS-1506657, CNS-1518728, and CNS-1350039 during this work.

REFERENCES

- [1] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. 2016. Cutting the Cord in Virtual Reality. In *Proc. of ACM HotNets-XV*.
- [2] Omid Abari, Haitham Hassanieh, Michael Rodriguez, and Dina Katabi. 2016. Millimeter Wave Communications: From Point-to-Point Links to Agile Network Connections. In *ACM HotNets-XV*.
- [3] C.R. Anderson and T.S. Rappaport. 2004. In-Building Wideband Partition Loss Measurements at 2.5 and 60 GHz. *IEEE Transactions on Wireless Communications* 3, 3 (2004).
- [4] IEEE Standards Association. 2009. IEEE Standards 802.15.3c-2009: Millimeter-wave-based Alternate Physical Layer Extension.
- [5] T. Tony Cai and Lie Wang. 2011. Orthogonal Matching Pursuit for Sparse Signal Recovery With Noise. *IEEE Transactions On Information Theory* 57, 7 (2011), 4680–4688.
- [6] Cisco Systems, Inc. 2016. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016-2021 White Paper. <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html>.
- [7] Daniel Steinmetzer and Daniel Wegemer and Matthias Schulz and Joerg Widmer and Matthias Hollick. 2017. Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices. In *Proc. of ACM CoNEXT*.
- [8] Kristian Evensen, Andreas Petlund, Haakon Riiser, Paul Vigmostad, Dominik Kaspar, Carsten Griwodz, and Pal Halvorsen. 2011. Mobile Video Streaming Using Location-based Network Prediction and Transparent Handover. In *International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*.
- [9] Facebook, Inc. 2016. Introducing Facebook's new terrestrial connectivity systems — Terragraph and Project ARIES. <https://code.facebook.com/posts/1072680049445290/introducing-facebook-s-new-terrestrial-connectivity-systems-terragraph-and-project-aries/>.
- [10] Federal Communications Commission. 2015. FCC Promotes Higher Frequency Spectrum for Future Wireless Technology. <https://apps.fcc.gov>.
- [11] Fierce Wireless. 2015. Verizon applauds FCC chairman's move to 5G spectrum. <http://www.fiercewireless.com/tech/story/verizon-applauds-fcc-chairmans-move-5g-spectrum/2015-10-05>.
- [12] Vladimir B. Gantovnik, Zafer Gurdal, Layne T. Watson, and Christine M. Anderson-Cook. 2005. A Genetic Algorithm for Mixed Integer Nonlinear Programming Problems Using Separate Constraint Approximations. *The American Institute of Aeronautics and Astronautics (AIAA)* 43, 8 (2005), 1844–1849.
- [13] Daniel Halperin, Srikanth Kandula, Jitendra Padhye, Paramvir Bahl, and David Wetherall. 2011. Augmenting Data Center Networks with Multi-gigabit Wireless Links. In *ACM SIGCOMM*.
- [14] R. C. Hansen. 2009. *Phased Array Antennas*. John Wiley & Sons, Inc.
- [15] Haitham Hassanieh, Omid Abari, Michael Rodriguez, Mohammed Abdelghany, Dina Katabi, and Piotr Indyk. 2018. Fast Millimeter Wave Beam Alignment. In *Proc. of ACM SIGCOMM*.
- [16] IEEE Standards Association. 2012. IEEE Standards 802.11ad-2012, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band. <http://standards.ieee.org/findstds/standard/802.11ad-2012.html>.
- [17] IEEE Standards Association. 2017. IEEE Standards 802.11ay, Enhancements Throughput for Operation in License-Exempt Bands above 45 GHz. http://www.ieee802.org/11/Reports/tgay_update.htm.
- [18] Qualcomm Incorporated. 2016. Exploring the Potential of mmWave for 5G Mobile Access.
- [19] Qualcomm Incorporated. 2018. Wilocity Kernel Driver. <https://github.com/torvalds/linux/tree/master/drivers/net/wireless/ath/wil6210>.

- [20] Injong Rhee and Lisong Xu and Sangtae Ha and Alexander Zimmermann and Lars Eggert and Richard Scheffenegeger. 2015. CUBIC for Fast Long-Distance Networks. In *Internet-Draft*.
- [21] Loch, Adrian and Assasa, Hany and Palacios, Joan and Widmer, Joerg and Suys, Hans and Debaillie, Björn. 2017. Zero Overhead Device Tracking in 60 GHz Wireless Networks Using Multi-Lobe Beam Patterns. In *Proc. of ACM CoNEXT*.
- [22] A. Maltsev, R. Maslennikov, A. Sevastyanov, A. Lomayev, and A. Khoryaev. 2010. Statistical Channel Model for 60 GHz WLAN Systems in Conference Room Environment. In *In Proceedings of the Fourth European Conference on Antennas and Propagation (EuCAP)*.
- [23] Zhinus Marzi, Dinesh Ramasamy, and Upamanyu Madhow. 2016. Compressive Channel Estimation and Tracking for Large Arrays in mm-Wave Picocells. *IEEE J. Sel. Topics Signal Processing* 10, 3 (2016), 514–527.
- [24] Meru Networks. 2005. Virtual Cells: The Only Scalable Multi-Channel Deployment.
- [25] Anthony J Nicholson, Yatin Chawathe, Mike Y Chen, Brian D Noble, and David Wetherall. 2006. Improved Access Point Selection. In *In Proceedings of the 4th ACM International Conference on Mobile Systems, Applications and Services*.
- [26] Thomas Nitsche, Adriana B. Flores, Edward W. Knightly, and Joerg Widmer. 2015. Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement. In *IEEE INFOCOM*.
- [27] OpenWrt. 2017. OpenWrt Wireless Freedom. <https://openwrt.org/>.
- [28] Park, Minyoung and Gopalakrishnan, Praveen. 2009. Analysis on Spatial Reuse and Interference in 60-GHz Wireless Networks. *IEEE J.Sel. A. Commun.* 27, 8 (2009), 1443–1452.
- [29] V. Pichapati, H. Kowshik, A. P. Subramanian, R. Kokku, and M. Chetlur. 2014. Location Assisted Handoffs in Dense Cellular Networks. In *IEEE International Conference on Sensing, Communication, and Networking (SECON)*.
- [30] Dinesh Ramasamy, Sriram Venkateswaran, and Upamanyu Madhow. 2012. Compressive Adaptation of Large Steerable Arrays. In *Information Theory and Applications Workshop (ITA)*.
- [31] Dinesh Ramasamy, Sriram Venkateswaran, and Upamanyu Madhow. 2012. Compressive Tracking with 1000-element Arrays: A Framework for Multi-Gbps MM-wave Cellular Downlinks. In *Allerton Conference on Communication, Control and Computing*.
- [32] T.S. Rappaport, F. Gutierrez, E. Ben-Dor, J.N. Murdock, Yijun Qiao, and J.I. Tamir. 2013. Broadband Millimeter-Wave Propagation Measurements and Models Using Adaptive-Beam Antennas for Outdoor Urban Cellular Communications. *IEEE Transactions on Antennas and Propagation* 61, 4 (2013).
- [33] T.S. Rappaport, Shu Sun, R. Mayzus, Hang Zhao, Y. Azar, K. Wang, G.N. Wong, J.K. Schulz, M. Samimi, and F. Gutierrez. 2013. Millimeter Wave Mobile Communications for 5G Cellular: It Will Work! *IEEE Access* 1 (2013).
- [34] Theodore S. Rappaport, Eshar Ben-Dor, James N. Murdock, and Yijun Qiao. 2012. 38 GHz and 60 GHz Angle-Dependent Propagation for Cellular and Peer-to-Peer Wireless Communications. In *IEEE ICC*.
- [35] Theodore S. Rappaport, Robert W. Heath Jr., Robert C. Daniels, and James N. Murdock. 2014. *Millimeter Wave Wireless Communications*. Prentice Hall.
- [36] Maryam Eslami Rasekh, Zhinus Marzi, Yanzi Zhu, Upamanyu Madhow, and Haitao Zheng. 2017. Noncoherent mmWave Path Tracking. In *ACM HotMobile*.
- [37] Mathew K. Samimi and Theodore S. Rappaport. 2015. 3-D Statistical Channel Model for Millimeter-Wave Outdoor Mobile Broadband Communications. In *Prof. of IEEE ICC*.
- [38] Ioulia Skliarova and António B.Ferrari. 2002. FPGA-based Implementation of Genetic Algorithm for the Traveling Salesman Problem and Its Industrial Application. In *IEA/AIE*.
- [39] Peter F. M. Smulders. 2009. Statistical Characterization of 60-GHz Indoor Radio Channels. *IEEE Transactions on Antennas and Propagation* 57, 10 (2009).
- [40] Zhenyu Song, Longfei Shangguan, and Kyle Jamieson. 2017. Wi-Fi Goes to Town: Rapid Picocell Switching for Wireless Transit Networks. In *ACM SIGCOMM*.
- [41] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2017. WiFi-Assisted 60 GHz Wireless Networks. In *Proc. of ACM MobiCom*.
- [42] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 2015. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *Proc. of ACM SIGMETRICS*.
- [43] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *Proc. of USENIX NSDI*.
- [44] The White House. 2016. Unlocking the Promise of Broadband for All Americans. <https://www.whitehouse.gov/blog/2016/07/15/unlocking-promise-broadband-gains-all-americans>.
- [45] Michalis Vavouras, Kyprianos Papadimitriou, and Ioannis Papaefstathiou. 2009. High-Speed FPGA-based Implementations of a Genetic Algorithm. In *International Symposium on Systems, Architectures, Modeling, and Simulation*.
- [46] Junyi Wang, Zhou Lan, Chin-Sean Sum, Chang-Woo Pyo, Jing Gao, Tuncer Baykas, Azizur Rahman, Ryuhei Funada, Fumihide Kojima, Ismail Lakkis, Hiroshi Harada, and Shuzo Kato. 2009. Beamforming Codebook Design and Performance Evaluation for 60GHz Wideband WPANs. In *IEEE VTC*.
- [47] Zhou Wang, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P. Simoncelli. 2004. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing* 13, 4 (2004).
- [48] Teng Wei and Xinyu Zhang. 2017. Pose Information Assisted 60 GHz Networks: Towards Seamless Coverage and Mobility Support. In *Proc. of ACM MobiCom*.
- [49] Hao Xu, Vikas Kukshya, and Theodore S. Rappaport. 2002. Spatial and Temporal Characteristics of 60-GHz Indoor Channels. *IEEE Journal on Selected Areas in Communications* 20, 3 (2002).
- [50] ZDNet. 2017. Ericsson, Qualcomm, and Vodafone Trial 5G New Radio for Unified 5G. In <https://www.qualcomm.com/news/releases/2017/02/26/ericsson-qualcomm-and-vodafone-trial-5g-new-radio-unified-5g>.
- [51] Yunze Zeng, Ioannis Pefkianakis, Kyu-Han Kim, and Prasant Mohapatra. 2017. MU-MIMO-Aware AP Selection for 802.11ac Networks. In *ACM Mobihoc*.
- [52] Jialiang Zhang, Xinyu Zhang, Pushkar Kulkarni, and Parameswaran Ramanathan. 2016. OpenMili: A 60 GHz Software Radio Platform with a Reconfigurable Phased-Array Antenna. In *ACM MobiCom*.
- [53] Xia Zhou, Zengbin Zhang, Yibo Zhu, Yubo Li, Saipriya Kumar, Amin Vahdat, Ben Y. Zhao, and Haitao Zheng. 2012. Mirror Mirror on the Ceiling: Flexible Wireless Links for Data Centers. In *ACM SIGCOMM*.
- [54] Yibo Zhu, Zengbin Zhang, Zhinus Marzi, Chris Nelson, Upamanyu Madhow, Ben Y. Zhao, and Haitao Zheng. 2014. Demystifying 60GHz Outdoor Picocells. In *Proc. of ACM MobiCom*.