Noncoherent mmWave Path Tracking

Maryam Eslami Rasekh, Zhinus Marzi, Yanzi Zhu, Upamanyu Madhow, Haitao Zheng
University of California, Santa Barbara, CA, 93106
{merasekh, zh_marzi, madhow} @ece.ucsb.edu, {yanzi, htzheng}@cs.ucsb.edu

ABSTRACT

Millimeter (mm) wave picocellular networks have the potential for providing the 1000X capacity increase required to keep up with the explosive growth of mobile data. However, maintaining beams towards mobile users and adapting to frequent blockage, requires efficient, dynamic path tracking algorithms. In this paper, we develop and experimentally demonstrate a novel noncoherent compressive strategy for this problem, and compare it with conventional hierarchical and exhaustive beam scanning. To the best of our knowledge, this is the first experimental demonstration of practical, scalable path estimation for mmWave/60GHz picocells. Our results indicate the feasibility of sub-second path tracking with low overhead on today's mmWave hardware, and open up a rich space for design of 5G mmWave networks.

1. INTRODUCTION

Both industry and academia have been exploring new 5G communication techniques in the millimeter (mm) wave band [21, 20] as a means of providing a quantum leap in cellular capacity. Recent work has experimentally demonstrated the feasiblity of outdoor mmWave picocellular systems in 60GHz [28], where low-cost base stations deployed on lampposts can connect outdoor mobile users by forming concentrated beams towards them[1]. Thanks to the tiny wavelength of mmWave bands, a very large number of antenna elements can be packed into both compact picocellular base stations and mobile devices.[2] They enable highly directional transmissions, which not only achieve very high data rate (*e.g.* multi-Gbps) in the downlink direction, but also minimize interference to other transmissions. A follow-up work [12] analyzed network-wide interference patterns, showing that, when carefully designed, mmWave outdoor picocells can easily deliver the required orders of magnitude increase in mobile capacity.

While prior work [28, 12] demonstrates the vast potential of mmWave picocells, it is based on an idealized assumption: for each user, nearby basestations can always identify the path(s) towards the user, and hence can form narrow (or pencil-like) beam towards the user. Path discovery and tracking, therefore, are critical requirements for mmWave picocells to function.

In practice, however, path tracking for mmWave picocells is challenging due to three key requirements. *First*, since mmWave transmissions are highly directional and easily blocked, outdoor path tracking needs to occur at sub-second time scales, *e.g.* every 100ms, even for pedestrian mobility [28]. *Second*, the tracking design must support the use of very large antenna arrays and scale in the number of mobile users. *Finally*, in order to be incorporated into commodity mobile devices like smartphones, the design should be low-cost and not rely on complex hardware modules.

Considering these requirements, existing approaches on user path tracking cannot be applied to mmWave picocells. They either require excessive amount of user feedback thus cannot scale with the number of users, *e.g.* the *hierarchical beam scanning* defined by the 802.11ad indoor 60GHz networking [2], or demand a large number of beacon measurements which translate into high latency and overhead, *e.g.* the *exhaustive scanning* scheme [9]. The recently proposed *compressive user tracking* scheme [13, 17, 18] overcomes the above two challenges by applying compressive estimation, but requires phase coherence across beacons that is not supported by commodity hardware[3].

Noncoherent Compressive Path Tracking. We propose and demonstrate a novel *noncoherent* compressive tracking scheme. Our design follows the same beacon broadcast scheme as the prior coherent design [13]. But the key difference is that, unlike the coherent design which requires knowledge of both amplitude and phase on the beacons, our design only requires the signal strength value (RSS). This relaxation eliminates the need for maintaining coherence across beacons, allowing immediate deployment on commodity mmWave hardware, and direct comparison with scan-based techniques, which also only require RSS.

The price we pay for the lack of phase information is that it becomes harder to estimate multiple paths, whereas coherent techniques can do so using greedy interference cancellation techniques [13, 11]. Instead, our current noncoherent design estimates the dominant path between each base station and user pair, assuming that

[1]The current 60GHz picocellular design focuses on downlink transmissions where a base station using an electronically steerable antenna array to beam towards mobile receivers in its picocell. The uplink (user to basestation) can leverage existing WiFi or LTE radios on the mobile device and the basestation.

[2]A 64-element (e.g., 8×8) array at 60 GHz is the size of a postage stamp, and a 1024-element (32×32) array is palm-sized.

[3]Standard communication receivers employ coherent processing *within* each packet, but maintaining coherence *across* packets requires tight coordination between transmitter and receiver, and more sophisticated processing, than is possible with commodity 60GHz hardware.

HotMobile '17, February 21-22, 2017, Sonoma, CA, USA

DOI: http://dx.doi.org/10.1145/3032970.3032974

other paths, if present, are substantially weaker. We argue that identifying the dominant path suffices for robust beamforming, since our overhead calculations show that it is possible to track users at sub-second time scales with very small beacon overhead ($<1ms$), so that we can quickly identify and switch beams if needed, to handle abrupt changes in the dominant path (*e.g.*, when a strong LoS path gets blocked). As ongoing work, we are improving our noncoherent design to recover multiple paths using joint estimation techniques.

Using a 60GHz phased array testbed, we performed initial validation on our design. Each testbed radio operates under the 802.11ad standard, and has a 8×16 phased array (128 elements), with effectively 16 horizontally steerable elements (each consisting of 8 vertical elements with fixed phase relationships). Experimental results show that our noncoherent design achieves accuracy similar to that of exhaustive beam scanning but at 50-75% less overhead. We also use simulations to show that this advantage increases sharply with the number of steerable elements (99% less overhead for 1000 elements). To the best of our knowledge, our work is the first to demonstrate that a low cost mmWave pico base station can discover and maintain, at sub-second time scales, knowledge of dominant paths to all mobile users in its vicinity.

Summary and Moving Forward. We explore the feasibility of real-time noncoherent path tracking for mmWave picocell networks. Such tracking only relies on relatively slow feedback from mobile receivers (possibly over LTE or WiFi uplink rather than mmWave uplink). This is critical for ensuring robust connectivity in the face of mobility and blockage, and for base station coordination for handoff and interference management. By establishing the basic feasibility of real-time user tracking, our work motivates research into a rich set of design issues for mmWave picocells, which we discuss in §6. While our noncoherent technique provides an immediate solution for commodity hardware that outperforms conventional scan-based techniques, we also plan to explore efficient coherent implementations, given the performance advantages offered by coherent compressive estimation. Significant strides have been made in even more challenging problems in coherence and fine-grained estimation for lower carrier frequencies over the past few years [8, 15, 7, 16, 5], hence we are hopeful that similar advances, enabled by accompanying protocol changes and sophisticated signal processing across packets, could greatly enhance the capabilities of mmWave systems, despite the higher phase noise at such frequencies.

2. MMWAVE PICOCELL PATH TRACKING

To provide context, we describe in this section the key requirements of user path tracking for mmWave picocell networks, and discuss the drawbacks of existing solutions.

2.1 Key Requirements

Low Latency. Path tracking must operate in real-time to identify the available path(s) to each user. Since mmWave transmissions are both highly directional and easily blocked by obstacles, the spatial channel between a base station and a mobile user can be highly dynamic. As shown in Figure 1, when a pedestrian user turns and faces in a different direction, her own body can block the path being currently used; either a different path or paths to a different base station must be used. Similarly, blockage by other (moving) pedestrians also occurs over similar time constants. Therefore, the base stations must be prepared to adapt their narrow beams, and to coordinate among themselves in serving users, at sub-second time scales (*e.g.*, once per 100 msec as shown by [28]).

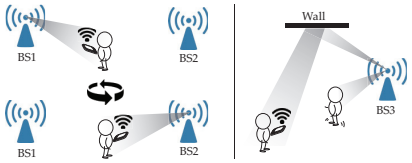

Figure 1: Available path changes as a user rotates her body or gets blocked by another pedestrian.

Supporting Large Arrays and Multiple Users. The overhead of path tracking must scale with both the number of users and the number of antenna elements. To create highly directional transmissions (or pencil-like beams), both mmWave picocell base stations and mobile devices use radios with a large number of antenna elements. Today, our testbed hardware already uses a 128-element phased array, and future hardware can easily go up to 1024-element (32×32) arrays. Furthermore, while indoor mmWave communications are mostly point to point, each outdoor picocell will serve multiple users. Thus supporting large arrays and multiple users is a critical requirement for path tracking.

Minimizing Hardware Cost. Since hardware design at the mmWave band is costly and complex, it is important that networking functionalities do not rely on complex custom hardware modules. This leads to two implications for path tracking. First, it should use coarse phase control for each antenna element to simplify hardware design, enabling scaling to a large number of elements. Second, in order to accommodate standard IEEE 802.11ad technology in commodity hardware (including our testbed from Facebook's Terragraph project [1]), it should not require tracking carrier phase across beacon packets. As a simple example, if the transmitter and receiver's local oscillators are offset by 10 parts per million at 60 GHz (i.e., by 0.6 MHz), then we need to have control of beacon packet launches to a granularity of about 100 nanoseconds in order to hope to track carrier phase across beacons, which is not available with standard IEEE 802.11ad implementations.

2.2 Why Existing Solutions Fail

Next we discuss existing solutions for user path tracking, and show that they fail to meet the above requirements. Figure 2 provides a high-level view of the three key schemes, and summarizes their key properties when applied to outdoor mmWave picocells.

Hierarchical Beam Scanning. This is the approach adopted in the IEEE 802.11ad indoor 60GHz standard [2, 23]. To discover the dominant paths to a specific user, the base station starts with a coarse partition of the search space, and refines it iteratively based on feedback from the user in the previous iteration. Specifically, it first sends beacons with broad beams, and then narrows the beams as it keeps receiving feedback from the target user. For each single user, hierarchical scanning may appear as efficient because it partitions the search space for beam directions efficiently using a tree, thus the number of beams scales as $\log N$, where N is the number of steerable antenna elements at the base station. Recent work also adapts the design to use both WiFi and 60GHz (horn antenna) scanning [14].

However, when applied to outdoor picocells which serve multiple users per cell, it faces several significant disadvantages, listed in order of importance. *First*, the scheme requires multiple feedback packets, and even more crucially, the decision on which beam to transmit the next beacon depends on, and hence must wait for, feedback regarding previous transmitted beacons. Therefore, it incurs high overhead and delays. *Second*, it does not scale well as the number of users grows, since different users at different locations may require different sequences of beacons to be sent. *Third*,

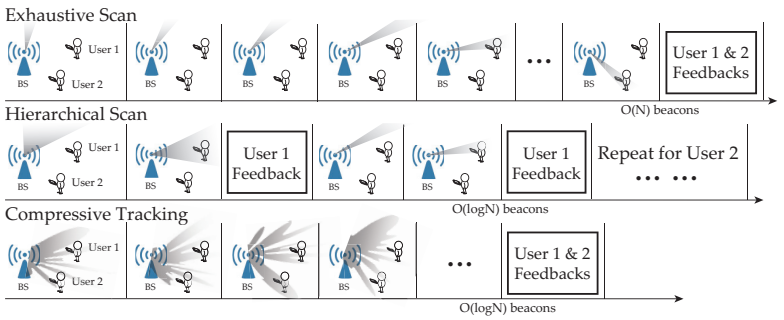

	Hier. scan	Exh. scan	Comp. tracking	**Our goal**
scale w/ # of users	Yes	No	Yes	**Yes**
scale w/ N	O(N)	O(logN)	0(logN)	**O(logN)**
feedback per user	O(logN)	1	1	**1**
require coherence	No	No	Yes	**No**
support coarse phase	Fragile	Fragile	Yes (2-bit)	**Yes (2-bit)**

Figure 2: Comparing 3 methods for beacon-based path tracking. N: # of steerable antenna elements.

sending a coarser beam requires that fewer antenna elements be used, thus compromising the range and reliability. Multiple paths within a coarse beam can interfere with each other, leading to performance deterioration. Cancellation techniques could alleviate the latter problem [6], but require coherence across channel measurements, and have not been experimentally validated.

Exhaustive Beam Scanning. A simpler alternative is exhaustive beam scanning, where the base station employs a fine-grained partition of the search space up front, sending out beacons using narrow beams [9]. It can be shown that the number of beacons required in this strategy scales linearly with N, the number of antenna elements (at the base station). This means significant tracking overhead (and latency) when using large arrays. Finally, we note that both the hierarchical and exhaustive scanning schemes require fine-grained phase control to generate the required beam patterns, which becomes costly and difficult to implement for large arrays.

Compressive Tracking. The third scheme is a *compressive* approach [17, 18, 13]: the base station sends beacons using pseudorandom phases assigned to each antenna element. Each compressive beacon covers the entire search space, but the beam pattern exhibits (pseudo)random peaks, and is therefore very different from either broad beam (as used at the beginning of a hierarchical scheme) or a highly directional pattern. By putting together the feedback corresponding to multiple compressive beacons, one can infer the directions of the dominant paths. In a nutshell, the compressive approach combines the best features of hierarchical and exhaustive scanning: the number of beacons scales as $\log N$, a common set of beacons can be used for all users in the cell, and each user consolidates its feedback into a single packet at the end of all the beacons. It also supports coarse phase control, because the pseudorandom beams can be implemented by randomly sampling among four phases $\{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\}$ for 2-bit control.

Unfortunately, current compressive designs [17, 18, 13] all require coherence across beacons, which is unavailable in commodity mmWave platforms. In particular, we have confirmed experimentally that these designs fail on our 60 GHz testbed, which does not support such coherence.

3. NONCOHERENT PATH TRACKING

We propose a new, noncoherent compressive tracking design, which maintains the key properties of compressive tracking but eliminates the need for coherence across beacons. Specifically, our design follows the same pseudorandom beacon broadcast of the coherent design [13], but only requires user feeding back the signal strength value (RSS) of the beacons and not the phase. This makes the compressive architecture directly competitive with existing hierarchical and exhaustive scans, which also only employ RSS. We show that despite the absence of phase information, the pseudorandom variations across beacon RSS values already carry sufficient information on the most dominant path. Thus, we can estimate the path by correlating the beacon RSS vector with the expected RSS patterns for different angles of departure (which can be computed based on knowledge of the array geometry and the pseudorandom beacon phases), and picking the one with the largest correlation magnitude. We now describe our design in more detail.

3.1 Models of Compressive Path Estimation

Consider an N-element regular linear array with inter-element spacing d at the mmWave picocell base station. A direction of departure θ from the broadside maps to a spatial frequency $\omega = \frac{2\pi d \sin\theta}{\lambda}$, with the corresponding array response represented by $\mathbf{a}(\omega) = [1, e^{j\omega}, e^{j2\omega}, ..., e^{j(N-1)\omega}]^T$. In each tracking interval, the compressive tracking design will broadcast M pseudorandom beacons, $i = 1, ...M$, produced by assigning beamforming weights of $\mathbf{b}_i = [e^{j\phi_{i1}}, e^{j\phi_{i2}}, ..., e^{j\phi_{iN}}]$ to the N elements of the array, where ϕ_{in}'s are independent and identically distributed (i.i.d.) random variables, chosen from a uniform distribution over $(0, 2\pi)$[4]. The response to the ith beacon at spatial frequency ω is

$$f_i(\omega) = \mathbf{b}_i^T \mathbf{a}(\omega) = \sum_{n=0}^{N-1} e^{j\phi_{in}} e^{jn\omega}. \quad (1)$$

Consider the simple scenario where the channel between the base station and the user consists of a single path with complex amplitude h_0, angle of departure θ_0, and corresponding spatial frequency ω_0. Ideally, the *coherent* measurement for beacon i by a user is given by $z_i = h_0 f_i(\omega_0) + v_i$, where $\{v_i\}$ are i.i.d. complex-valued Gaussian noise samples, $v_i \sim \mathcal{CN}(0, 2\sigma^2)$. In practice, lack of precise synchronization between the transmitter and receiver local oscillators and the resulting independent time drifts translate into a random phase rotation of the observed channel response. Thus, the receiver actually sees $z_i e^{j\phi_i}$, where $\{\phi_i\}$ is the random, unpredictable phase rotation for each i. Such lack of coherence across beacons destroys phase information, hence the receiver must employ noncoherent measurements involving only the amplitudes:

$$y_i = |z_i| = |h_0 f_i(\omega_0) + v_i| \quad (2)$$

3.2 Noncoherent Estimation Design

It can be proved, under a high SNR approximation to the observation statistics, that the optimal (maximum likelihood) estimate is given by intuitively pleasing rule of choosing the spatial frequency whose nominal RSS pattern best matches the observed RSS pattern (proof omitted due to space limitations). Specifically, denote the nominal RSS pattern for spatial frequency ω by $\mathbf{x}(\omega) = [|f_1(\omega)|, |f_2(\omega)|, ..., |f_M(\omega)|]^T$. The cost function to be maximized is the squared normalized inner product of these nominal

[4]The available phase values depend on the granularity of phase control, *e.g.* $\{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\}$ for 2-bit control.

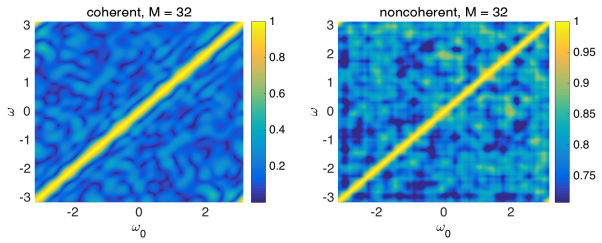

Figure 3: Correlation of beacon responses of different spatial frequencies using coherent and noncoherent measurements, with 32 beacons.

RSS patterns with the measured RSS pattern: $J(\omega) = \langle \frac{\mathbf{y}}{||\mathbf{y}||}, \frac{\mathbf{x}(\omega)}{||\mathbf{x}(\omega)||} \rangle^2$, where, for vectors $\mathbf{u}$, $\mathbf{v}$, $\langle \mathbf{u}, \mathbf{v} \rangle = \mathbf{u}^T\mathbf{v}$ denotes inner product, and $||\mathbf{u}|| = \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle}$ denotes the norm. Our estimate of the spatial frequency of the most dominant path is given by

$$\hat{\omega_0} = \underset{\omega}{\text{argmax}} \ J(\omega) \tag{3}$$

We solve this optimization problem by first evaluating the cost function on a discrete grid, and then refining around the maximizing grid point via Newton pursuit.

Path Estimation Accuracy. We now provide insight into the relative efficacy of the noncoherent and coherent frameworks, showing that, despite some performance degradation due to lack of coherence, the noncoherent framework can be expected to provide an accurate estimate for the dominant path. Ignoring noise, the RSS observations $\mathbf{y}$ are proportional to $\mathbf{x}(\omega_0)$, so that the variations across ω are captured by the normalized correlations: $K(\omega, \omega_0) = \left| \langle \frac{\mathbf{x}(\omega_0)}{||\mathbf{x}(\omega_0)||}, \frac{\mathbf{x}(\omega)}{||\mathbf{x}(\omega)||} \rangle \right|$. Clearly, we will have a peak at $\omega = \omega_0$, but large local maxima at other values of ω would imply higher vulnerability to noise. We can also define analogous normalized correlations for (idealized) coherent measurements.

Figure 3 shows these normalized correlations $K(\omega, \omega_0)$ for both coherent and noncoherent measurements with 32 beacons. Clearly, coherent measurements provide better suppression of such undesired local maxima, but noncoherent measurements also provide enough discrimination between the desired $\omega = \omega_0$ line and undesirable local maxima.

3.3 Estimating Multiple Paths

mmWave channels are often well characterized by a single dominant path - usually the LOS path or a strong reflection when LOS is blocked - and several other weaker paths. In this case, it is easy to show (detailed mathematical models omitted due to lack of space) that the contribution of the other paths to noncoherent measurements appears as additional noise, so that optimal estimation based on the above single-path model can be expected to perform well. We have found this to be the case in all our experiments on a typical university campus (see §4).

For the case of multipath channels with more than one strong path, the additional paths cannot be treated as noise and the multiple strong paths can be jointly estimated. For the idealized coherent measurement model, effective greedy techniques based on coherent successive cancellation of estimated paths have been developed ([13] and references therein) but not experimentally validated due to the difficulty of attaining coherence across beacons. For noncoherent measurements, such cancellation is difficult, and estimating K distinct paths requires joint estimation of $3K$ real-valued parameters: the spatial frequency and complex amplitude (real and imaginary component) of each of the K paths. Developing efficient solutions for this problem is our ongoing work.

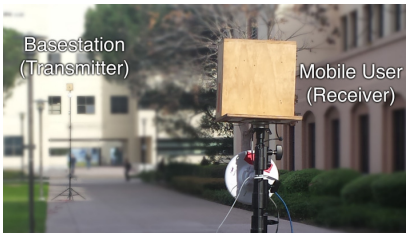

Figure 4: Our 60GHz phased array testbed emulating picocells.

4. INITIAL VALIDATION

Using a 60GHz phased array testbed, provided by the Facebook Terragraph project [1], we perform initial validation of our noncoherent scheme in terms of accuracy and overhead. Overall, our results show that our noncoherent design achieves accuracy similar to that of exhaustive beam scanning but at 50% smaller overhead. We also use simulations to show that this advantage increases sharply with the number of steerable antenna elements. Thus using commodity 802.11ad hardware, our design can lead to accurate user tracking at sub-second time scales even for very large arrays.

The 60GHz Phased Array Testbed. Our testbed consists of two 60GHz radios, each operating according to the 802.11ad single carrier mode [2]. To emulate a picocellular network setting, we place the transmitter (base station) at a height of 4 m, and mount the receiver on a mobile rack at 1.5 m emulating a user holding a smartphone (see Figure 4). Each antenna consists of an 8-by-16 rectangular array, with effectively 16 horizontally steerable elements (each consisting of 8 vertical elements with fixed phase relationships). The phase control uses a 4-bit quantization. Each radio is also paired with a WiFi radio for uplink feedback and system control. The base station broadcasts each beacon as a 802.11ad packet with zero payload. The receiver (user) operates in the omnidirectional mode to receive beacons.

We performed experiments at various outdoor locations on our campus (open space, pathways near trees and buildings, etc). In most scenarios we could only find one dominant path. When placing antennas close to relatively smooth building surfaces we obtained two-path scenarios. We ran path estimation in three different ranges of 50, 100, and 200 meters, and observed similar results.

4.1 Path Estimation Accuracy

We first evaluate whether our noncoherent compressive design can estimate the dominant path accurately. To obtain the ground truth path direction(s), we perform the exhaustive beam scanning of the azimuth angle with 64 beacons in the angular span of $(-45°, 45°)$ (the valid angular span for our hardware configuration), and obtain the corresponding RSS value per beacon. Under the same propagation environment, we run the compressive estimation using 5-35 beacons. From the user returned beacon RSS values, we generate the likelihood function $J(\omega)$ and identify the peak as the most dominant path.

Single Path Scenarios. Figure 5 shows an example result of exhaustive scanning (RSS vs. spatial frequency ω) with 64 beacons and the likelihood function $J(\omega)$ from noncoherent compressive estimation with 16 and 32 beacons. We see that the peaks of $J(\omega)$ coincide closely with the peak of the RSS from exhaustive scan, demonstrating the accuracy of compressive estimation.

Next, we plot in Figure 6(a) the median and 95th percentile of the path estimation error (absolute spatial frequency error $|\omega_{est} - \omega_{true}|$), and the resulting beamforming loss (relative signal strength loss when the base station beamforms to the user in the estimated direction), across all measurements. Here we also examine the impact of phase control granularity, 4-bit (maximum available on the

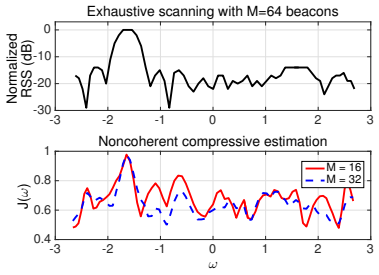

Figure 5: Exhaustive scan measurements vs. likelihood curve obtained by compressive measurements.

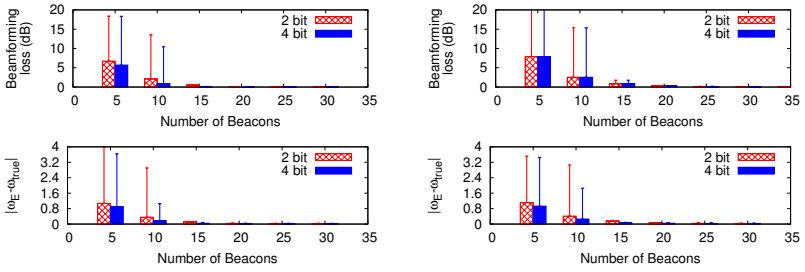

(a) 1-path scenarios (b) 2-path scenarios

Figure 6: Mean & 95th percentile of beamforming loss and absolute spatial frequency error of noncoherent estimation vs. the number of beacons.

testbed) and 2-bit. We make two key observations. *First*, with just 15 beacons, our design achieves accuracy similar to that of exhaustive beam scanning with 32 or 64 beacons. *Second*, the estimation error and beamforming loss are similar for 2-bit and 4-bit phase control, confirming that our design can be implemented on low-cost hardware with coarser phase control.

2-Path Scenarios. In this case, the existence of a large flat surface (building wall or windows) provides modest reflection to the user. But the reflection path is always more than 8dB weaker than the LoS path (due to reflection loss). Results in Figure 6(b) show that our noncoherent design can always detect the most dominant path at high accuracy.

4.2 Tracking Overhead and Latency

The overhead of path tracking includes two parts: the time taken by base stations to send beacons, and the time taken by users to feedback beacon RSS reports. The *beacon overhead* depends on the number of beacons and the beacon spacing, since each beacon is short ($2\mu s$). For our current 802.11ad hardware, the beacon spacing is bounded by the beam dwelling time ($\leq 25\ \mu s$). Thus, the overhead for sending 16 beacons is 0.4 ms, when we adhere to the 802.11ad standard. The *feedback overhead* depends on the number of users and the wireless technique and protocol used for uplink. When using WiFi (802.11a/b), it takes on average 3.6 ms for one user and $30ms$ for five users (due to uplink contention). This delay can be significantly reduced by improving the uplink design (which we leave to future work).

To evaluate the overhead when using large arrays, we use simulations to derive the minimum beacons required to ensure at most $3dB$ beamforming loss, as a function of the number of steerable antenna elements (N). Figure 7 shows that the beacon overhead for exhaustive scan scales linearly with N and that of our design scales nicely with $log(N)$. For example, supporting arrays with 1000 steerable elements will require 22 beacons for our design and yet 1800 beacons for exhaustive scan. Using 802.11ad hardware, our design leads to 0.55 ms beacon overhead, or 0.55% overhead for a 100ms frame, while exhaustive scan faces 45 ms overhead.

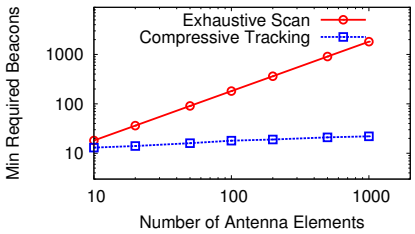

Figure 7: Beacon overhead vs. # of antenna elements

5. RELATED WORK

There has been tremendous progress over the past decade in developing commercially viable mmWave hardware, and products based on the IEEE 802.11ad indoor 60 GHz networking standard (also backed by the WiGig alliance) are available on the market. Directional MAC protocols supporting RF beamforming have been included in 802.11ad, and existing products are able to automatically steer around blockage [28]. However, these do not scale to the rapid changes in outdoor mobile environments.

More recently, there is significant interest in *outdoor* mmWave links for flexible wireless backhaul [25] and base-to-mobile communication [20]. Recent works report propagation measurements supporting the feasibility of mmWave cellular (focusing on ranges of 200m, in bands other than 60 GHz) [19, 20, 3, 21]. As discussed earlier, the basic feasibility and massive capacity gains of 60 GHz picocells with shorter links were established by [28, 12],

In addition to the prior work on coherent compressive tracking [17, 18, 13], other prior theoretical work on mmWave channel estimation includes [22, 6]. [22] explores one-bit compressive estimation (and does not attempt to super-resolve the channel as in [13]), while [6] employs hierarchical search with interference cancellation across paths. These approaches also rely on coherent measurements, like [17, 18, 13].

A recent work [4] addresses fast path tracking without phase coherence by designing beacons that each sense a subset of possible angles, and then voting on the respective angles to identify the direction of the strongest path. Due to interference from side lobes, they performed "soft voting" that is similar to the maximum likelihood method defined by Eq. 3. While the number of beacons required for successful estimation is not reported, our simulations found that comparing with our approach, [4] requires significantly more beacons to operate. For instance, using the same 16-element setup as our experiments, [4] needs 140 beacons for 90% accurate estimation (*i.e.* within $3dB$ of optimal beamforming), while ours requires 12. Their increased overhead is due to the beacon design that results in beam patterns with wide and relatively uniform peaks, thus limiting the information provided by each beacon measurement relative to randomly generated compressive beacons.

Recent experimental research in 60 GHz networking includes [10, 27, 28, 29, 26, 14, 30, 23, 24]. They all employ horn antennas, unlike the electronically steerable phased arrays used in our testbed that allow real-time beam steering and tracking. For quasi-stationary links, [24] allows fast beam adaptation during blockage, but requires full beam scanning (hierachical or exhaustive) to identify paths. Our work differs by targeting mobile links where the set of potential paths change dynamically and must be discovered at much finer time scales.

6. CONCLUSION AND OPEN PROBLEMS

Our proposed noncoherent framework enables immediate deployment of compressive tracking with commodity IEEE 802.11ad hardware, since it requires the same information (i.e., RSS) as conventional scan-based techniques. Using this framework, we have provided the first experimental demonstration of compressive mmWave path estimation, showing that it provides accurate estimates at significantly lower overhead than standard scanning. This motivates a sustained effort in pursuing compressive architectures, with two major thrusts as follows. The first is to extend the noncoherent framework to estimating multiple paths, and to understand it more deeply at a theoretical level, relating it to the literature on phase retrieval. The second is to pursue the custom hardware and protocol changes required for obtaining the performance advantage provided by coherent compressive measurements, and evaluating experimentally whether the required phase coherence across beacons, and the gains promised by theoretical studies such as [13], can be attained despite the higher phase noise at mmWave frequencies.

Our experiments also provide confidence in the feasibility of mmWave path tracking to allow abstractions for higher layer network design, including interactions with TCP. For example, we may assume that each base station will have estimates of the dominant paths to all users in its vicinity, updated on a sub-second basis (e.g., every 100 ms). Under this assumption, how well can we serve pedestrian users in crowded urban environments, where paths may be blocked by the user's body or other obstacles? How well can we track vehicular mobility for antennas (a) on the vehicle roof, or (b) hand-held inside the vehicle? How can base stations use path estimates to coordinate for increasing robustness to blockage, or managing interference by (a) choosing which user to serve, and (b) applying interference suppression techniques (e.g., forming nulls towards one user while directing beams at another)? A comprehensive research agenda for answering such questions is the natural next step towards realizing mmWave picocellular networks delivering orders of magnitude capacity gains over existing systems.

Acknowledgments

We thank the anonymous reviewers for their helpful comments, and the Facebook Terragraph team for providing the 60 GHz testbed on which these experiments were carried out. This project was supported by NSF grants CNS-1518812 and CNS-1317153, and by a gift from Facebook.

7. REFERENCES

[1] Facebook Terragraph and Project ARIES. https://code.facebook.com/posts/1072680049445290/introducing-facebook-s-new-terrestrial-connectivity-systems-terragraph-and-project-aries/.

[2] IEEE 802.11 Task Group AD. http://www.ieee802.org/11/Reports/tgad_update.htm.

[3] OnQ blog. https://www.qualcomm.com/news/onq/2015/11/19/qualcomm-research-demonstrates-robust-mmwave-design-5g.

[4] ABARI, O., HASSANIEH, H., RODREGUEZ, M., AND KATABI, D. Millimeter wave communications: From point-to-point links to agile network connections. In *Proc. of HotNets* (2016).

[5] ABARI, O., RAHUL, H., KATABI, D., AND PANT, M. Airshare: Distributed coherent transmission made seamless. In *Proc. of INFOCOM* (2015).

[6] ALKHATEEB, A., AYACH, O. E., LEUS, G., AND HEATH, R. W. Channel estimation and hybrid precoding for millimeter wave cellular systems. *IEEE Journal of Selected Topics in Signal Processing 8*, 5 (2014).

[7] BALAN, H. V., ROGALIN, R., MICHALOLIAKOS, A., PSOUNIS, K., AND CAIRE, G. Airsync: Enabling distributed multiuser mimo with full spatial multiplexing. *IEEE/ACM Transactions on Networking (TON) 21*, 6 (2013).

[8] BIDIGARE, P., PRUESSING, S., RAEMAN, D., SCHERBER, D., MADHOW, U., AND MUDUMBAI, R. Initial over-the-air performance assessment of ranging and clock synchronization using radio frequency signal exchange. In *Proc. of SSP* (2012).

[9] CELIK, N., ISKANDER, M. F., EMRICK, R., FRANSON, S. J., AND HOLMES, J. Implementation and experimental verification of a smart antenna system operating at 60 ghz band. *IEEE Transactions on Antennas and Propagation 56*, 9 (2008).

[10] HALPERIN, D., ET AL. Augmenting data center networks with multi-gigabit wireless links. In *Proc. of SIGCOMM* (2011).

[11] MAMANDIPOOR, B., RAMASAMY, D., AND MADHOW, U. Frequency estimation for a mixture of sinusoids: A near-optimal sequential approach. In *Proc. of GlobalSIP* (2015).

[12] MARZI, Z., MADHOW, U., AND ZHENG, H. Interference analysis for mm-wave picocells. In *Proc. of GLOBECOM* (2015).

[13] MARZI, Z., RAMASAMY, D., AND MADHOW, U. Compressive channel estimation and tracking for large arrays in mm wave picocells. *IEEE Journal of Selected Topics in Signal Processing PP*, 99 (2016).

[14] NITSCHE, T., FLORES, A. B., KNIGHTLY, E. W., AND WIDMER, J. Steering with eyes closed: Mm-wave beam steering without in-band measurement. In *Proc. of INFOCOM* (2015).

[15] QUITIN, F., RAHMAN, M. M. U., MUDUMBAI, R., AND MADHOW, U. A scalable architecture for distributed transmit beamforming with commodity radios: Design and proof of concept. *Wireless Communications, IEEE Transactions on 12*, 3 (2013).

[16] RAHUL, H., KUMAR, S. S., AND KATABI, D. Megamimo: Scaling wireless capacity with user demand. In *Proc. of SIGCOMM* (2012).

[17] RAMASAMY, D., VENKATESWARAN, S., AND MADHOW, U. Compressive adaptation of large steerable arrays. In *Proc. of ITA* (2012).

[18] RAMASAMY, D., VENKATESWARAN, S., AND MADHOW, U. Compressive tracking with 1000-element arrays: a framework for multi-gbps mm wave cellular downlinks. In *Proc. of Allerton* (2012).

[19] RANGAN, S., RAPPAPORT, T., AND ERKIP, E. Millimeter-wave cellular wireless networks: Potentials and challenges. *Proc. of the IEEE 102*, 3 (2014).

[20] RAPPAPORT, T., SUN, S., MAYZUS, R., ZHAO, H., AZAR, Y., WANG, K., WONG, G., SCHULZ, J., SAMIMI, M., AND GUTIERREZ, F. Millimeter wave mobile communications for 5g cellular: It will work! *IEEE Access 1* (2013).

[21] ROH, W., SEOL, J.-Y., PARK, J., LEE, B., LEE, J., KIM, Y., CHO, J., CHEUN, K., AND ARYANFAR, F. Millimeter-wave beamforming as an enabling technology for 5g cellular communications: theoretical feasibility and prototype results. *IEEE Communications Magazine 52*, 2 (2014).

[22] RUSU, C., MENDEZ-RIAL, R., GONZALEZ-PRELCIC, N., AND HEATH, R. W. Adaptive one-bit compressive sensing with application to low-precision receivers at mmwave. In *Proc. of GLOBECOM* (2015).

[23] SUR, S., VENKATESWARAN, V., ZHANG, X., AND RAMANATHAN, P. 60 ghz indoor networking through flexible beams: A link-level profiling. In *ACM SIGMETRICS Performance Evaluation Review* (2015), vol. 43, ACM, pp. 71–84.

[24] SUR, S., ZHANG, X., RAMANATHAN, P., AND CHANDRA, R. Beamspy: enabling robust 60 ghz links under blockage. In *Proc. of NSDI* (2016).

[25] TAORI, R., AND SRIDHARAN, A. In-band, point to multi-point, mm-wave backhaul for 5g networks. In *Proc. of ICC* (2014).

[26] WEI, T., AND ZHANG, X. mtrack: High-precision passive tracking using millimeter wave radios. In *Proc. of MobiCom* (2015).

[27] ZHOU, X., ZHANG, Z., ZHU, Y., LI, Y., KUMAR, S., VAHDAT, A., ZHAO, B. Y., AND ZHENG, H. Mirror mirror on the ceiling: Flexible wireless links for data centers. In *Proc. of SIGCOMM* (2012).

[28] ZHU, Y., ZHANG, Z., MARZI, Z., NELSON, C., MADHOW, U., ZHAO, B. Y., AND ZHENG, H. Demystifying 60ghz outdoor picocells. In *Proc. of MobiCom* (2014).

[29] ZHU, Y., ZHOU, X., ZHANG, Z., ZHOU, L., VAHDAT, A., ZHAO, B. Y., AND ZHENG, H. Cutting the cord: a robust wireless facilities network for data centers. In *Proc. of MobiCom* (2014).

[30] ZHU, Y., ZHU, Y., ZHAO, B. Y., AND ZHENG, H. Reusing 60GHz radios for mobile radar imaging. In *Proc. of MobiCom* (2015).