

Millimeter Wave Communications: From Point-to-Point Links to Agile Network Connections

Omid Abari^{†1} Haitham Hassanieh^{†2} Michael Rodriguez¹ Dina Katabi¹

¹ MIT CSAIL ²UIUC

{abari, mrod, dina}@csail.mit.edu haitham@illinois.edu

[†]Co-primary Authors

ABSTRACT

Millimeter wave (mmWave) technologies promise to revolutionize wireless networks by enabling multi-gigabit data rates. However, they suffer from high attenuation, and hence have to use highly directional antennas to focus their power on the receiver. Existing radios have to scan the space to find the best alignment between the transmitter's and receiver's beams, a process that takes up to a few seconds. This delay is problematic in a network setting, where the base station needs to quickly switch between users and accommodate mobile clients.

We present Agile-Link, the first mmWave beam steering system that is demonstrated to find the correct beam alignment without scanning the space. Instead of scanning, Agile-Link hashes the beam directions using a few carefully chosen hash functions. It then identifies the correct alignment by tracking how the energy changes across different hash functions. Our results show that Agile-Link reduces beam steering delay by orders of magnitude.

1. INTRODUCTION

The ever-increasing demand for mobile and wireless data has placed a huge strain on today's WiFi and cellular networks [9, 12, 33]. Millimeter wave (mmWave) frequency bands address this problem by offering multi-GHz of unlicensed bandwidth – 200× more than the bandwidth allocated to today's WiFi and cellular networks [23, 25]. Further, mmWave radio hardware has recently become commercially viable [26, 11, 19]. This led to multiple demonstrations of point-to-point mmWave communication links [15, 38, 31]. These advances have generated much excitement about the role that mmWave technology can play in future

wireless networks, and have led to mmWave communication being declared as a central component in next-generation (5G) cellular networks [25, 16, 19]. It has also led to multiple mmWave standards including IEEE 802.11ad for wireless LANs [18] and IEEE 802.15.3c for wireless PANs [17].

However, a key challenge has to be addressed before mmWave links can be integrated into cellular or 802.11 networks. mmWave signals attenuate quickly with distance, so they need to use highly directional antennas to focus their power. Due to the narrow beam of the antennas, communication is possible only when the transmitter's and receiver's beams are well-aligned. First generation mmWave radios used horn antennas, which require mechanical steering to identify the best beam alignment. More advanced mmWave radios use phased-array antennas, which can be steered electronically. Still, current phased array mmWave radios require multiple seconds to scan the space with their beams to find the best alignment [39]. Taking a long delay before aligning the beams may be acceptable in today's fixed point-to-point links. However, such a long delay hampers the deployment of mmWave links in cellular (or 802.11) networks, where a base station has to quickly switch between users and accommodate mobile clients.

So, how is beam steering done in mmWave phased-arrays? Since the wavelength is very small (a few millimeters), a small phased array, the size of a credit card, can have tens or hundreds of antennas, leading to a very narrow beam, as shown in Fig.2. Beam steering is done in the analog domain using phase shifters, which add a controllable phase to each antenna. Identifying the best beam alignment is equivalent to identifying the correct phase setting for all phase shifters on both the transmitter and receiver. This is done by sequentially trying different phase shifts (i.e., different beams) and measuring the received signal power. The best beam alignment maximizes the power.

Trying all possible beam directions incurs excessive delay though. Indeed, existing products can take seconds to converge [39]. Thus, multiple proposals have been introduced to optimize the steering time. In particular, the 802.11ad standard proposes to set the transmitter's beam pattern to a quasi-omnidirectional shape, while the receiver scans the space for the best signal direction. The process is then reversed to have the transmitter scan the space while keeping

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 the author(s) 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.

HotNets-XV, November 09 - 10, 2016, Atlanta, GA, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4661-0/16/11...\$15.00

DOI: <http://dx.doi.org/10.1145/3005745.3005766>




Figure 1—Illustrative example of Agile-Link’s algorithm: The figure shows how the algorithm recovers the direction of the signal of arrival (i.e. 60° in this example) in three steps. The algorithm first hashes the spatial directions into bins (step b and c). Then, it gives votes to all directions that hash into the bin which has energy. Finally, it picks the direction with the highest number of votes (step d).

the receiver quasi-omnidirectional [18, 37]. While this design reduces the steering delay, it comes at the cost of worse SNR and it still requires each node to sequentially scan the whole space of beam directions, continuing to incur significant delay in practice [30]. Further, none of the other proposals for improving the steering delay have been evaluated with mmWave phased arrays, and the vast majority of them are purely theoretical.

This paper presents Agile-Link, the first mmWave beam steering system that is evaluated on phased arrays and demonstrated to find the correct beam alignment without sequentially scanning the space.

Agile-Link’s design relies on a combination of smart hashes and voting. Specifically, there are only a few paths that the mmWave signal can take between the transmitter and receiver [25, 4]. Thus, instead of trying all beam directions, Agile-Link works by hashing spatial directions into bins, where each bin collects energy from a large number of directions. Agile-Link can then ignore all bins that have no energy and focus on those with high energy because they contain the correct signal alignment. Agile-Link uses a few carefully-crafted hash functions, which allows it to quickly identify the best beam alignment by observing how the energy changes across bins, from one hash function to another.

Designing appropriate hash functions for this problem is challenging. First, while in theory there are many good hash functions that one could apply to the signal, in practice we are allowed to change *only* the phases on the phase shifters (see Fig. 2). We are neither allowed to manipulate the magnitude of the signal on the individual antennas, nor to turn off some antennas. This renders many of the standard hashing techniques useless and significantly constrains the space of hash functions. Second, the design of the hash functions has to deal with the possibility of signals along different directions being hashed to the same bins, combining destructively and canceling out. In §3, we develop a steering algorithm that addresses these challenges and quickly identifies the optimal steering direction.

We evaluate Agile-Link using mmWave radios, each equipped with a phased array that has 8 antennas. We also

use simulations to explore its scaling behavior to large arrays with hundreds of antennas, which are expected in the future [8]. We compare Agile-Link with two baselines: an exhaustive scan of the space to find the best beams, and the quasi-omnidirectional search proposed in the 802.11ad standard. Our evaluation reveals the following findings. In comparison with the exhaustive search, Agile-Link reduces the search time by one to three orders of magnitude, for array sizes that range from 8 antennas to 256 antennas. In comparison to the quasi-omnidirectional search, Agile-Link reduces the delay by $1.5\times$ to $10\times$, for the same range of array sizes. Thus, we believe that Agile-Link provides an important step towards practical mmWave networks.

2. ILLUSTRATIVE EXAMPLE

Agile-Link’s algorithm is best understood through a high-level example. Consider an antenna array with $N = 16$ antenna elements and the case where the received signal has a dominant path along the spatial direction of 60° as shown in Fig. 1(a). Agile-Link’s goal is to discover this 60° direction from which the signal arrives in order to steer its beam that direction and maximize the SNR.

Agile-Link starts by hashing the spatial directions into bins, where each bin collects energy from multiple directions in space. For example, if we hash the space into 4 bins, then each bin should collect energy from $N/4 = 4$ different directions. This is achieved by steering the antenna array to beam along 4 different directions as shown in Fig. 1(b). A bin will contain energy only if the signal’s path matches a direction that is hashed to the bin. Hence, in Fig. 1(b) only the first bin, highlighted in red, will contain energy since it captures the signal along the 60° direction. This significantly reduces the search space to the directions within this bin.

However, we cannot simply iterate on these candidate directions. Multiple paths can collide in the same bin and potentially cancel each other –i.e., the RF waves along these paths can sum up destructively. Thus, there is a probability that a bin may have negligible energy though it does contain the directions of real signal paths. To avoid missing directions of the signal with high SNR, we need to randomize the




Figure 2—Millimeter Wave Phased Arrays. mmWave phased arrays use phase shifters to control the phase on each antenna and steer the beam.

hashing. Agile-Link repeats the hashing while randomizing the directions that fall into the same bin as shown in Fig. 1(c). This ensures that if two paths collide in the first hashing, they will not continue to collide and cancel each other.

After repeating the random hashing, Agile-Link uses a voting based scheme to discover the direction of the signal that has energy. Specifically, in the first hashing in Fig. 1(b), the first bin contained energy and hence, the directions that fall into that bin get a vote. In the second hashing in Fig. 1(c), the third bin contained energy and the directions in that bin get votes as shown in Fig. 1(d). However, only the direction along 60° will get a vote from both hashing functions, so Agile-Link can quickly recover it as shown in Fig. 1(e), without having to scan all possible directions.¹

In order to be able to implement the above algorithm, however, Agile-Link needs to steer the antenna array to beam along any random set of directions. This is very challenging in practice since we are only allowed to change the phases on the phase shifters (see Fig. 2). We are neither allowed to manipulate the magnitude of the signal on the individual antennas, nor to turn some antennas off. This significantly constrains the space of possible beam patterns which we can create to hash the directions to bins. In the following section, we describe in detail how to address this problem and create beam patterns that can capture different directions simply by setting the phase of the phase shifters.

3. Agile-Link

This section describes Agile-Link in detail. For clarity, we will describe the problem and the algorithm assuming only the receiver has an antenna array, while the transmitter has an omni-directional antenna that transmits in all directions. The extension to the case where both transmitter and receiver have antenna arrays is achieved simply by applying the algorithm on both sides.

3.1 Formalizing the Problem

At mmWave frequencies, the wireless channel follows a geometric model [3]. Specifically, for a receiver with N an-

¹Note that in more general settings, Agile-Link can use more than just two hashing functions. In that case, directions that do not have energy are unlikely to get votes whereas directions that have energy will get a lot of votes allowing Agile-Link to quickly recover them.

tenna elements that are equally spaced by a distance $d = \lambda/2$, where λ is the wavelength, the wireless channel at the n -th antenna element in the array can be written as:

$$\mathbf{h}_n = \sum_k^K \alpha_k e^{j(2\pi \cos(\theta_k)nd/\lambda + \psi_k)}, \quad (1)$$

where K is the total number of paths and α_k , ψ_k , and θ_k are respectively the attenuation, phase, and direction of the signal along the k -th path. By replacing $\cos(\theta_k) = f$, the above becomes a standard Fourier transform equation where antenna elements are analogous to time samples and signal directions are analogous to frequencies. Thus, we can rewrite the above equation as:

$$\mathbf{h} = \mathbf{F}'\mathbf{x}, \quad (2)$$

where \mathbf{h} is an $N \times 1$ vector representing the channels on the antenna elements, \mathbf{F}' is the inverse Fourier transform matrix and \mathbf{x} is a vector representing the channel across different directions in space, *i.e.*, $\mathbf{x}_{\cos(\theta_k)} = \alpha_k e^{j\psi_k}$. \mathbf{x} is sparse since only a few directions will have energy. Hence, by taking a Fourier transform across the channels of the antenna elements in the array and calculating the energy in the elements of \mathbf{x} , we can recover the directions of the paths that the wireless signal takes.

Unfortunately, we cannot measure the channel on each antenna element in Equations 1 or 2. This is because, in a phased array, we can only measure the combined signal after the phase shifters, as shown in Fig. 2. Say ϕ_n is the phase set to the phase shifter on the n -th antenna element. Let \mathbf{a} be a $1 \times N$ vector such that $\mathbf{a}_n = e^{j\phi_n}$, then each measurement we collect of the channel can be written as: $y = \mathbf{a}\mathbf{F}'\mathbf{x}$. We can repeat this measurement M times for different settings of phase shifters to obtain a vector of measurements:

$$\mathbf{y} = \mathbf{A}\mathbf{F}'\mathbf{x} \quad (3)$$

The above is a standard sparse recovery equation; we know from sparse recovery theory that we can recover the vector \mathbf{x} using only $M = K \log N/K$ measurements. However, trying to recover \mathbf{x} by applying standard sparse recovery to these measurements does not work in practice. These measurements are taken over time and due to CFO (carrier frequency offset) between the transmitter and receiver, the measurements will accumulate a random unknown phase which will corrupt the phase of the measurements. Hence, when reasoning across multiple measurements, we can only use the magnitude of these measurements. Agile-Link designs an algorithm that recovers the directions of the paths using only the magnitude of the measurements.

3.2 Agile-Link's Algorithm

As described earlier, Agile-Link works in two stages. First, by randomly hashing the space into bins such that each bin collects power from a range of directions. Then, by using a voting mechanism to recover the directions that have the energy. Below, we describe these two stages in detail.

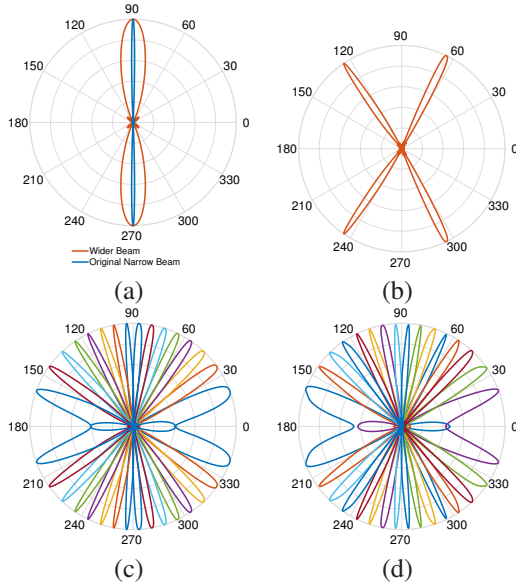


Figure 3—Hashing Beam Patterns (a) Using a wider beam to hash the direction cannot be done since it requires changing the size of the array. (b) Using Agile-Link’s hash function to hash well-spread directions into a bin (c) All the bins of Agile-Link’s hash function, where the same color corresponds to directions in the same bin. (d) A random permutation of the hash function in (c).

A. Hashing Spatial Directions into Bins: As described earlier, creating wider beams that hash the space into bins is difficult since we can only control the phase on each antenna *i.e.*, in the above model, we can only set the phases of the vector \mathbf{a} . Each setting of this vector \mathbf{a} will create a different beam pattern and the magnitude squared of each measurement $|y|^2$ will correspond to the energy in the regions covered by the beam pattern. Thus, Agile-Link’s goal is to find settings of the vector \mathbf{a} that (1) create good beam patterns that can hash the space into bins and cover all regions in space, and (2) create beam patterns that can be randomized to change the different directions that hash to the same bins.

Ideally, we would like to set the vector \mathbf{a} to create a wider beam as shown in Fig. 3(a) to hash several directions into a bin. This, however, will require changing the size of the array, which is not feasible as described earlier. Instead, Agile-Link divides the antenna array into sub-arrays and makes each sub-array beam toward a different direction. Specifically, the vector \mathbf{a} is divided into R segments, each of length N/R *i.e.*, $\mathbf{a}_{[1:N/R]}$, $\mathbf{a}_{[N/R+1:2N/R]}$, \dots , $\mathbf{a}_{[(R-1)N/R:N]}$. Each segment then sets its beam towards a different direction. The sub-beam created by each segment will be larger than the beam created by the full array by a factor of R , and will cover R different directions. Further, since there are R such sub-beams and each is R times larger than the basic beam created by the full array, the beam created by this setting of the vector \mathbf{a} will cover R^2 directions. Now, if we wish to hash the space of directions into B bins, then the total number of directions covered by each bin will be N/B and thus, $R = \sqrt{N/B}$.

But how do we set the direction in which each sub-array should direct its beam? The naïve solution would be to try

to recreate the wide beam in Fig. 3(a) by directing the beams of the sub-arrays to nearby or consecutive directions. However, this creates bad beams because the side-lobes of each sub-array will now sum up with the main-lobe of the next sub-array and significantly reduce its power (the sum is in the complex domain). In fact, the best beams are obtained when the sub-arrays direct their beams in well spaced directions so that the leakage from their side-lobes is minimized. Fig 3(b) shows an example when we have two sub-arrays of half the size of the full array direct their beams 60° apart. In this case, the beam pattern will hash directions that are 60° apart into the same bin. By shifting the direction, we can then create all different bins as shown in Fig 3(c) where each color corresponds to a bin in the hashing function.

The question that remains is how do we randomize these beam patterns (*i.e.*, how do we randomize the beams that hash directions to bins)? To do that, we leverage a nice property of the Fourier transform that says we can randomly permute the output of the Fourier transform by randomly permuting its input samples and modulating their phase. Specifically, consider a vector \mathbf{z} of length N and its Fourier Transform $\hat{\mathbf{z}}$, then given a random σ , invertible modulo N , and a random $\beta < N$, we can permute the input samples according to:

$$\mathbf{z}'(t) = \mathbf{z}(\sigma t \bmod N) \times e^{-2\pi j\beta t/N} \quad (4)$$

This will result in a random permutation of the frequencies according to:

$$\hat{\mathbf{z}}'(f) = \hat{\mathbf{z}}(\sigma f + \beta \bmod N). \quad (5)$$

Thus, by randomly permuting the input samples, we can randomly permute the frequencies. Recall from Eq. 2 that in our setting, input samples correspond to the channel \mathbf{h} on the antenna elements of the array and frequencies correspond to directions of the signal \mathbf{x} . Unfortunately, as described earlier, we do not have access to the antenna elements of the array to perform this permutation. Luckily, however, we can shift this permutation from the antenna elements to the phase shifters. Recall that each measurement is obtained by $y = \mathbf{a}_{1 \times N} \mathbf{h}_{N \times 1}$ and hence we can perform the permutation on \mathbf{a} instead of \mathbf{h} . Thus, we will set the phase of the n -th phase shifter to a new phase $\phi'(n) = \phi(\sigma n \bmod N) - 2\pi\beta n/N$. This creates a random permutation of the directions and will randomize which directions hash to which bins. Fig. 3(d) shows an example of this where the directions that hash to the same bin are shown in the same color and are different from the hashing in Fig. 3(c) —*i.e.*, the hashing in Fig. 3(d) is a permutation of the hashing in Fig. 3(c), over the space of directions.

B. Recovering the Directions of the Paths After hashing the spatial directions into bins, Agile-Link discovers the actual directions of the signal using a voting based scheme, where each bin gives votes to all directions that hash into that bin. After few random hashes, the directions that have energy will collect the largest number of votes which allows Agile-Link to recover them. Unfortunately, directly applying this voting approach does not work well in practice. Mainly because the

side-lobes of the beams create leakage between the bins and thus a strong path in one bin can leak energy into other bins which corrupts the voting process. To overcome this problem, Agile-Link uses a form of soft voting and takes into account the leakage between the bins.

Specifically, Agile-Link models the beam patterns shown in Fig. 3 as a probability distribution $P(\theta)$ that indicates the probability that the received signal arrived from direction θ . If we hash into B bins, we will have B such patterns and collect B measurements $\mathbf{y}_{1 \times B}$ corresponding to the bins. After taking the magnitude squared of each measurement and normalizing their power, we can compute the probability that the energy of the signal is coming from direction θ as:

$$\Pr\{\theta\} = \frac{1}{\|\mathbf{y}\|_2^2} \sum_{b=1}^B |\mathbf{y}_b|^2 \times P_b(\theta), \quad (6)$$

where $P_b(\theta)$ is the beam pattern corresponding to the b -th bin, $\|\mathbf{y}\|_2^2$ is the L2 norm squared of \mathbf{y} , and the summation is taken over the probabilities since the direction of the signal can fall into any of the bins. After performing a few random hashes, we can compute the probability the energy of the signal is coming from direction θ as:

$$\Pr\{\theta\} = \prod_{l=1}^L \frac{1}{\|\mathbf{y}_l\|_2^2} \sum_{b=1}^{B_l} |\mathbf{y}_{b,l}|^2 \times P_{b,l}(\theta), \quad (7)$$

where L is the total number of random hashes performed by Agile-Link and the product is taken over the probabilities since the direction of arrival has to be large in every single hashing. Finally, the K directions of θ that have the highest probabilities will correspond to the directions of the K paths that the signal traverses. The best beam alignment is then chosen to be the direction of the path that delivers the maximum energy.

3.3 Measurement Complexity

Agile-Link’s algorithm takes B measurements for each hashing of the space, where B is the number of bins. In the case where both transmitter and receiver have antenna arrays, Agile-Link takes $B \times B$ measurements. The hashing is repeated L times, for a total number of measurements $L \times B^2$. Since the channel is sparse with only K paths, we only need to set $B = O(K)$ to ensure a small number of path collisions in the same bin. We also set $L = \log N$ to ensure that the probability of missing a path is polynomially small. Thus, Agile-Link’s algorithm requires $O(K^2 \log N)$ measurements, which scales sub-linearly with the size of the array. Hence, for large N , it delivers significant gains over the 802.11ad standard, as we will show in §4.

While we described the algorithm in the context of 1D antenna arrays, the algorithm holds for 2D arrays as well. We simply need to apply the hash function along both dimensions of the array. For an $N \times N$ antenna array, the complexity will simply be $O(K^2 \log N^2)$, scaling sub-linearly with the total number of antennas in the array.




Figure 4— Recovered Directions (a) the direction of arrival that Agile-Link recovers at the receiver and (b) the direction of departure that is recovered at the transmitter when the transmitter is placed 120° relative to the receiver’s array and the receiver is placed 50° relative to the receiver’s array.




Figure 5— Platform (a) The phased array and (b) the mmWave radio we used to perform our experiments.

4. RESULTS

In order to evaluate Agile-Link’s ability to identify the best beam alignment quickly, we ran experiments in an office/lab area with standard furniture and multipath effects. We used millimeter radios operating in the 24 GHz ISM band and equipped with a phased array that has 8 antennas as shown in Fig. 5 [2].

Fig. 4 shows the signal direction recovered by Agile-Link on the transmitter and receiver side when we set the transmitter to be at an angle of 120° with respect to the receiver’s array and the receiver at an angle of 50° with respect to the transmitter’s array. As can be seen, Agile-Link can accurately identify the direction of arrival of the signal at the receiver as 120° and the direction of departure of the signal from the transmitter as 50° .

Next we would like to evaluate the gain in latency that Agile-Link delivers. However, since our radio has a fixed array size, we cannot empirically measure how this gain scales for larger arrays. Hence, we perform extensive simulations to compute this gain for larger arrays and we use our empirical results from our 8-antenna array to find the delay for this array size. We compare against two baselines.

- **Exhaustive Search:** In this case, the transmitter and the receiver each uses $2N$ different beams to scan the different directions where N is the number of antennas. This takes $4N^2$ measurements. Then, the combination of transmitter and receiver beams that delivered the maximum power is picked as the direction of the signal.
- **802.11ad Standard:** In this case, the transmitter sets its antenna array to a quasi-omnidirectional mode while the receiver scans $2N$ directions of the beam. This is followed by the receiver setting its antenna array to a quasi-omnidirectional mode while the transmitter scans $2N$ beam directions. Then, the γ transmit and receive




Figure 6— Beam Searching Latency The reduction in search time for Agile-Link compared to 802.11ad and exhaustive search.

beams that delivered the highest power are tested against each other, *i.e.*, γ^2 combinations are tried. The combination that delivers the maximum power is then picked for beamforming.

Fig. 6 plots the reduction in latency that Agile-Link achieves over exhaustive search and the standard. The figure shows that, for an 8-antenna phased array, Agile-Link can reduce the search time by $5.3\times$ and $1.2\times$ compared to exhaustive search and the 802.11ad standard, respectively. The gain however increases quickly as the number of antennas increase. This is due to the fact that the search time is directly proportional to the number of measurements collected by each scheme. Recall that, exhaustive search requires a quadratic number of measurements as a function of the array size and the standard is linear in the antenna array size since it uses $4N + \gamma^2$ measurements. However, Agile-Link is sub-linear in the array size and uses only $K^2 \log N$ measurements, as described earlier.² Thus, the gain of Agile-Link over both exhaustive search and 802.11ad increases very fast: for arrays of size 256, it is $10\times$ better than the 802.11ad standard and multiple orders of magnitude better than exhaustive search.

5. RELATED WORK

The related work can be classified into the following areas:

Practical mmWave Phased Array Systems: Working implementations of phased array mmWave systems have been limited to industry, with very few known examples, such as Qualcomm’s 28 GHz demo [8], Samsung’s 28 GHz prototype [28], and 60 GHz products from two startups: Wilocity and SiBeam [1, 29].³ However, none of these systems present a steering algorithm or measurements of steering delays. In fact, current products are designed for static links [29, 1]; they take a long time to steer the beam and are not suitable for mobile or multi-user networks [39].

²We set K to 4 since most empirical measurement studies [25, 4, 30, 31] show that at mmWave frequencies the channel has only 2 to 3 paths.

³Note that there are other mmWave products on the market [19, 22]. However, they do not support phased arrays and require the use of horn antennas. Further, while the circuits community produced several VLSI chips for mmWave phased arrays [27], these chips have not been demonstrated to work as part of a full-fledged mmWave communication system.

Point-to-Point mmWave Communication: Recent interest in mmWave communication has led to a lot of demonstrations of point-to-point links for Data Centers applications [15, 38, 10], as well as cellular picocells and WiFi applications [39, 31, 30]. These implementations mainly focus on using horn antennas to direct the beam, which require mechanical steering and are not suitable for non-static links or multi-user networks.

Simulation-Based Beam Searching Methods: There is a large body of theoretical work that proposes more efficient beam searching algorithms. Most of this work proposes enhancements on the 802.11ad standard, employing hierarchical beams to speed up the search [21, 5, 36, 20, 34, 37, 32]. However, in practice, hierarchical search requires feedback from the receiver to guide the transmitter at every stage of the hierarchy, which incurs significant protocol delay. Furthermore, hierarchy-based algorithms do not work in the worst case. Different paths can combine destructively and cancel each other at any level of the hierarchy, resulting in lost paths. Agile-Link’s algorithm, however, randomizes that hashing of the path directions in order to avoid such worst case scenarios, as we described in §3.

Sparse Recovery: Past theoretical work proposes using compressive sensing to reduce the number of measurements needed to discover the right alignment of the beam [24, 14, 13]. This approach, however, does not work with practical hardware because it ignores CFO (Carrier Frequency Offset), which corrupts the phase of the measurements. In contrast, Agile-Link’s algorithm relies only on the magnitude of the measurements to recover the correct beam alignment, and hence does not suffer due to CFO, as we explained in §3.

Our work is also related to massive MIMO systems at GHz frequencies [35, 7, 6]. These systems, however, connect each antenna to its own TX/RX chain, and hence can immediately measure the channel at each antenna. At mmWave frequencies, we can only measure the combined signal from all antennas in the array, since they are connected to one TX/RX chain, as shown in Fig 2. Further, using many TX/RX chains makes the system huge and power hungry, which is not suitable for mobile devices and access points.

6. CONCLUSION

In this paper, we have presented Agile-Link, the first phased array mmWave system that is capable of fast beam steering. Agile-Link delivers a new algorithm that finds the correct alignment of the beams between a transmitter and a receiver orders of magnitude faster than existing radios that have to scan the entire space to find the best alignment. This process currently takes up to a few seconds, which is impractical for dynamic and multi-user networks where the direction of alignment is constantly changing. Finally, the high data rates that mmWave communication can deliver makes it an indispensable part of future cellular networks and wireless LANs. We believe Agile-Link brings us closer towards practical mmWave networks.

Acknowledgments: We thank the NETMIT group and the reviewers for their insightful comments. This research is supported by NSF and HKUST. We thank the members of the MIT Center for their interest and support.

7. REFERENCES

- [1] Wilocity 802.11ad Multi-Gigabit Wireless Chipset. <http://wilocity.com>, 2013.
- [2] O. Abari, H. Hassanieh, M. Rodreguiz, and D. Katabi. Poster: A Millimeter Wave Software Defined Radio Platform with Phased Arrays. In *MOBICOM*, 2016.
- [3] A. Alkhateeb, O. El Ayach, G. Leus, and R. W. Heath. Channel estimation and hybrid precoding for millimeter wave cellular systems. *Selected Topics in Signal Processing, IEEE Journal of*, 2014.
- [4] C. R. Anderson and T. S. Rappaport. In-Building Wideband Partition Loss Measurements at 2.5 and 60 GHz. *IEEE Transactions on Wireless Communications*, 3(3), May 2004.
- [5] D. C. Araújo, A. L. de Almeida, J. Axnas, and J. Mota. Channel estimation for millimeter-wave very-large mimo systems. In *Signal Processing Conference (EUSIPCO), 2014 Proceedings of the 22nd European*, pages 81–85. IEEE, 2014.
- [6] W. U. Bajwa, J. Haupt, A. M. Sayeed, and R. Nowak. Compressed Channel Sensing: A New Approach to Estimating Sparse Multipath Channels. *Proceedings of the IEEE*, 98(6), June 2010.
- [7] C. R. Berger, Z. Wang, J. Huang, and S. Zhou. Application of Compressive Sensing to Sparse Channel Estimation. *IEEE Communications Magazine*, November 2010.
- [8] M. Branda. Qualcomm Research demonstrates robust mmWave design for 5G. Qualcomm Technologies Inc., November 2015.
- [9] Cisco. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013.
- [10] Y. Cui, S. Xiao, X. Wang, Z. Yang, C. Zhu, X. Li, L. Yang, and N. Ge. Diamond: Nesting the Data Center Network with Wireless Rings in 3D Space. In *NSDI*, 2016.
- [11] C. Doan, S. Emami, D. Sobel, A. Niknejad, , and R. Brodersen. Design considerations for 60 GHz CMOS radios. *IEEE Communication Magazine*, 42(12):132.
- [12] Ericsson. Traffic and market data report, 2011.
- [13] B. Gao, Z. Xiao, C. Zhang, D. Jin, and L. Zeng. Joint SNR and Channel Estimation for 60 GHz Systems using Compressed Sensing. In *WCNC*, 2013.
- [14] B. Gao, Z. Xiao, C. Zhang, D. Jin, and L. Zeng. Sparse/dense channel estimation with non-zero tap detection for 60-GHz beam training. *IET Communications*, 8(11):2044–2053, 2014.
- [15] D. Halperin, S. Kandula, J. Padhye, P. Bahl, and D. Wetherall. Augmenting Data Center Networks with Multi-Gigabit Wireless Links. In *ACM SIGCOMM*, 2011.
- [16] S. Han, C. I. Z. Xu, and C. Rowell. Large-Scale Antenna Systems with Hybrid Analog and Digital Beamforming for Millimeter Wave 5G. *IEEE Communications Magazine*, January 2015.
- [17] IEEE Standards Association. IEEE Standards 802.15.3c-2009: Millimeter-wave-based Alternate Physical Layer Extension, 2009.
- [18] IEEE Standards Association. IEEE Standards 802.11ad-2012: Enhancements for Very High Throughput in the 60 GHz Band, 2012.
- [19] J. Kilpatrick, R. Shergill, and M. Sinha. 60 GHz Line of Sight Backhaul Links Ready to Boost Cellular Capacity. Analog Devices Inc.
- [20] J. Kim and A. F. Molisch. Fast Millimeter-Wave Beam Training with Receive Beamforming. *Journal of Communications and Networks*, 16(5), October 2014.
- [21] B. Li, Z. Zhou, W. Zou, X. Sun, and G. Du. On the Efficient Beam-Forming Training for 60GHz Wireless Personal Area Networks. *IEEE Transactions on Wireless Communications*, 12(2), February 2013.
- [22] Pasternack Enterprises Inc. 60 GHz Transmitter/Receiver Development System. www.pasternack.com.
- [23] Z. Pi and F. Khan. An introduction to millimeter-wave mobile broadband systems. *Communications Magazine, IEEE*, 2011.
- [24] D. Ramasamy, S. Venkateswaran, and U. Madhow. Compressive tracking with 1000-element arrays: A framework for multi-gbps mm wave cellular downlinks. In *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*. IEEE, 2012.
- [25] S. Rangan, T. S. Rappaport, and E. Erkip. Millimeter-wave cellular wireless networks: Potentials and challenges. *Proceedings of the IEEE*, 2014.
- [26] T. S. Rappaport, J. N. Murdock, and F. Gutierrez. State of the art in 60GHz integrated circuits and systems for wireless communications. *Proceedings of the IEEE*, 2011.
- [27] G. Rebeiz. Millimeter-wave SiGe RFICs for large-scale phased-arrays. In *IEEE Bipolar/BiCMOS Circuits and Technology Meeting (BCTM)*, 2014.
- [28] W. Roh, J.-Y. Seol, J. Park, B. Lee, J. Lee, Y. Kim, J. Cho, K. Cheun, and F. Aryanfar. Millimeter-Wave Beamforming as an Enabling Technology for 5G Cellular Communications: Theoretical Feasibility and Prototype Results. *IEEE Communications Magazine*, February 2014.
- [29] SiBeam, Lattice Semiconductor. www.sibeam.com.
- [30] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *SIGMETRICS*, 2015.
- [31] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *NSDI*, 2016.
- [32] Y. M. Tsang, A. S. Y. Poon, and S. Addepalli. Coding the Beams: Improving Beamforming Training in mmWave Communication System. In *IEEE GLOBECOM*, 2011.
- [33] UMTS Forum. Mobile traffic forecasts: 2010-2020 report, 2011.
- [34] J. Wang, Z. Lan, C.-W. Pyo, T. Baykas, C.-S. Sum, M. A. Rahman, J. Gao, R. Funada, F. Kojima, H. Harada, and S. Kato. Beam Codebook Based Beamforming Protocol for Multi-Gbps Millimeter-Wave WPAN Systems. *IEEE Journal of Selected Areas in Communications*, 27(8), October 2009.
- [35] X. Xie, E. Chai, X. Zhang, K. Sundaresan, A. Khojastepour, and S. Rangarajan. Hekaton: Efficient and Practical Large-Scale MIMO. In *MOBICOM*, 2015.
- [36] W. Yuan, S. M. D. Armour, and A. Doufexi. An Efficient and Low-complexity Beam Training Technique for mmWave Communication. In *PIMRC*, 2015.
- [37] L. Zhou and Y. Ohashi. Efficient Codebook-Based MIMO Beamforming for Millimeter-Wave WLANs. In *PIMRC*, 2012.
- [38] X. Zhou, Z. Zhang, Y. Zhu, Y. Li, S. Kumar, A. Vahdat, B. Y. Zhao, and H. Zheng. Mirror Mirror on the Ceiling: Flexible Wireless Links for Data Centers. In *ACM SIGCOMM*, 2012.
- [39] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng. Demystifying 60GHz Outdoor Picocells. In *MOBICOM*, 2014.