

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

by

Michael Rodriguez

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Masters of Engineering in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2017

© Massachusetts Institute of Technology 2017. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
June, 2017

Certified by.....
Dina Katabi
Professor
Thesis Supervisor

Accepted by
Christopher Terman
Chairman, Department Committee on Graduate Theses

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

by

Michael Rodriguez

Submitted to the Department of Electrical Engineering and Computer Science
on June, 2017, in partial fulfillment of the
requirements for the degree of
Masters of Engineering in Electrical Engineering

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.

This research encompasses the implementation and testing of Agile-link, the first mmWave beam steering system that is implemented and evaluated on phased arrays, and 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. Two major limitations are addressed in this research. First is the issue of delays in scanning and the second is the accuracy of the beams. Here we propose, implement and examine solutions to these two major issues. Our results show that not only does Agile-link create accurate phase shifted beams, but, it also reduces beam steering delay by orders of magnitude.

Thesis Supervisor: Dina Katabi

Title: Professor

Acknowledgments

The work for this thesis was done in collaboration with Omid Abari and Haitham Hasanieh under the supervision of Dina Katabi. I thank them for their guidance and patience. I am grateful for the opportunity to have learned from and worked with them.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 8 |
| 2 | Millimeter Wave Primer | 13 |
| 3 | Related Work | 17 |
| 4 | Agile-Link Overview | 21 |
| 5 | Agile-Link Software Radio Platform | 24 |
| 6 | Phased Array Calibration | 29 |
| 7 | Experimental Results | 33 |
| 7.1 | Agile-Link's Beam Searching Performance | 33 |
| 7.2 | Agile-Link's Platform Performance | 39 |
| 7.3 | Importance of Calibration | 41 |
| 8 | Conclusion | 44 |

List of Figures

| | | |
|-----|---|----|
| 1-1 | Millimeter Wave Phased Arrays The figure shows mmWave phased arrays which use phase shifters to control the phase on each antenna and steer the beam. | 10 |
| 2-1 | Beam Scanning The figure shows that a phased array with N antennas has to try 2N beams in order to cover different directions in space. In this figure N = 8. | 14 |
| 5-1 | Agile-Link's Platform The figure shows (a) the mmWave radio platform we built to operate as a daughterboard for the USRP software radio and (b) the phased array with 8 antenna elements and 8 phase shifters. | 25 |
| 5-2 | Agile-Link's Architecture. The figure shows block diagrams for both Agile-Link's transmitter and receiver. Agile-Link uses a heterodyne architecture where the mmWave signal is first taken into an intermediate frequency, before the I and Q components are separated. | 26 |
| 6-1 | Phased array radiation pattern. The figure plots the phased array radiation pattern for (a) before and (b) after calibration when the beam was steered at 60 degrees. The figure shows that without calibration, the phased array has an unwanted large sidelobe at 80 degrees. The calibration improves the radiation pattern of the phased array by minimizing the sidelobe. | 30 |
| 6-2 | Phase Shifters' Performance. The figure shows the amount of phase shift introduced by each phase shifter (mounted on Agile-Link's phased array) versus the value of its 8-bits control line. Phase shifters performs differently once they are mounted on the board. | 31 |

| | | | |
|-----|---|--|----|
| 7-1 | Beam Accuracy with a Single Path | The figure shows the SNR loss due to beam misalignment for Agile-Link, standard and exhaustive search. | 37 |
| 7-2 | Beam Accuracy with Multipath | The figure shows the SNR loss of standard and Agile-Link with respect to the exhaustive search. | 37 |
| 7-3 | Beam Searching Latency | The figure plots the reduction in search time for Agile-Link compared to standard and exhaustive search. | 38 |
| 7-4 | Agile-Link Coverage | Effective SNR at the receiver versus distance between receiver and transmitter. | 40 |
| 7-5 | 256 QAM and 16 QAM constellation. | Figure shows that Agile-Link provides a full OFDM PHY capable of delivering up to 256 QAM and 16QAM for short and long distances, respectively. | 40 |
| 7-6 | Agile-Link MU-MIMO Network Throughput Gain | The figure shows that Agile-Link design is flexible and can be used to provide mmWave MU-MIMO. | 41 |
| 7-7 | Performance of Phased Array Calibration. | CDF of sidelobe level relative to the main beam with and without calibration. The figure shows that Agile-Link's calibration significantly reduces the radiation outside of the main beam. | 42 |

Chapter 1

Introduction

The ever increasing demand for mobile and wireless data has placed a huge strain on today's WiFi and cellular networks [9, 15, 44]. 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 [28, 31]. Further, mmWave radio hardware has recently become commercially viable [33, 12, 24]. This led to multiple demonstrations of point-to-point mmWave communication links [18, 51, 39]. These advances have generated much excitement about the role that mmWave technology can play in future wireless networks, and led to mmWave communication being declared as a central component in next-generation (5G) cellular networks [31, 19, 24]. It has also led to multiple mmWave standards including IEEE 802.11ad for wireless LANs [23] and IEEE 802.11c for wireless PANs [22].s

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; hence 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 [52]. 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.1-1. 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. Indeed, existing products can take seconds to converge [52]. 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 the receiver quasi-omnidirectional. While this design reduces the steering delay, it still requires each node to sequentially scan the whole space of beam directions and hence continues to incur significant delay in practice [38, 52]. 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 thesis introduces Agile-link, the first mmWave beam steering system that is implemented and evaluated on phased arrays, and demonstrated to find the correct beam alignment without sequentially scanning the space. We focus on the hardware platform and then address key issues related to using phased array antennas for hashing.

Agile-link’s design relies on a combination of smart hashes and voting. Specifically,

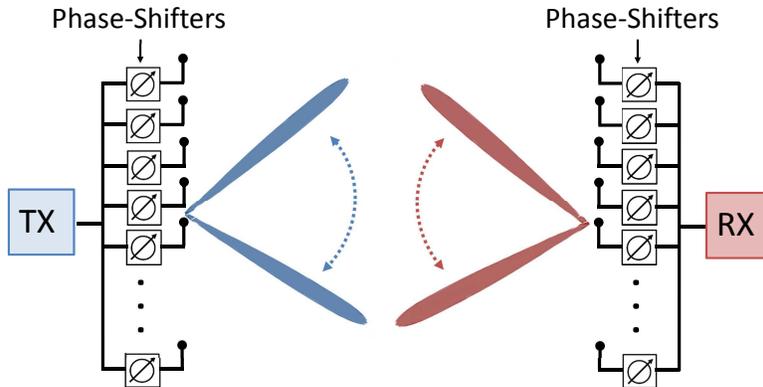


Figure 1-1: **Millimeter Wave Phased Arrays** The figure shows mmWave phased arrays which use phase shifters to control the phase on each antenna and steer the beam.

there are only a few paths that the mmWave signal can take between the transmitter and receiver [31, 3]. 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. 1-1). We are neither allowed to manipulate the magnitude of the signal on the individual antennas nor allowed 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. The specific hashing techniques are outside the scope of this thesis as they were developed by a colleague.

We have built a full-fledged mmWave radio capable of fast beam steering. This mmWave radio operates as a daughterboard for the USRP software radio. This enables easy manipulation of mmWave signals using standard GNU-radio software, and

helps bring mmWave to the GNU-radio community. The platform operates in the ISM band at 24GHz. Each phased array has 8 antennas, but the same device can have multiple such phased arrays. The design enables all USRP-GNU functions to be performed in mmWave frequencies. For example, one can control the bandwidth of the signal, change the center frequency, and coordinate multiple USRPs with a shared clock to act as a MIMO device.

We have also developed calibration techniques for the phased-array antennas which eliminate side-lobes when steering the signal to different angles. This greatly reduces the unwanted side effects of steering. These results were used to evaluate Agile-link's hashing algorithm and were also used as a benchmark for calibration quality.

We evaluate Agile-Link using our SDR mmWave radio. 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 1.5x to 10x, for the same range of array sizes.
- The improved steering delay of the quasi-omnidirectional search with respect to the exhaustive search comes at the cost of worse SNRs. In particular, the 90th percentile SNR loss in multipath scenarios is 12 dB. In contrast, Agile-link's SNR loss in multipath scenarios is only 2 dB.
- We use our Agile-link platform to study the performance of mmWave phased array radios operating in the 24GHz ISM band. Our results show that these radios can support delivering up to 256 QAM modulation and can easily operate at distances that exceed 100 meters.

We believe the tests and implementation strategies described here show that Agile-link provides an important leap towards practical mmWave networks.

Chapter 2

Millimeter Wave Primer

Millimeter wave systems have to use large antenna arrays to compensate for the high pathloss and attenuation that wireless signals experience at mmWave frequencies. Unlike traditional wireless systems having a digital TX/RX chain per antenna in mmWave is prohibitively expensive because of the large number of antennas needed to beamform the signal. As a result, mmWave phased arrays use analog phase shifters that control the phase on each antenna and allow us to steer the beam as shown in Fig. 1-1. This creates a challenge: since we do not have a digital chain per antenna, we cannot estimate the channel on each antenna separately like we currently do in traditional wireless systems. Hence, we cannot immediately discover the direction where the strongest power is arriving from and steer the beam in that direction. Instead, mmWave systems have to search for the correct beam alignment of the transmitter and the receiver.

Discovering the correct beam alignment is expensive. An exhaustive scan of the different directions of the beams can take a long time. Specifically, for a phased array with N antennas, each of the transmitter and receiver has to try $2N$ beams in order to cover the different directions in space [38], as shown in Fig. 2-1.¹ This requires $4N^2$ measurements to scan the entire space and align the beam of the transmitter with that of the receiver. Since each measurement requires a packet transmission, the

¹The number is $2N$ to ensure there is overlap between the beams; otherwise the best alignment could be at the edge between two beam positions.

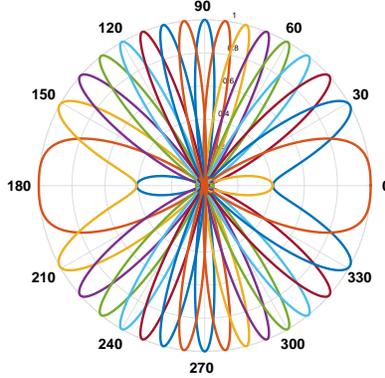


Figure 2-1: **Beam Scanning** The figure shows that a phased array with N antennas has to try $2N$ beams in order to cover different directions in space. In this figure $N = 8$.

process of establishing a link can take an extremely long time [38, 52].

Further, beam alignment has to be done often because of the low coherence time at mmWave frequencies. Specifically, the wavelength is very small at mmWave frequencies, and hence even the slightest motion or change in the environment can significantly change the channel leading to a channel coherence time of sub-millisecond [31]. This is orders of magnitude smaller than that of today’s cellular and WiFi networks. Hence, this expensive process of beam alignment has to be done quite often which is very detrimental to the throughput in cellular systems and wireless LANs.

Several mmWave measurement studies show that mmWave channels are very sparse with at most 2 to 3 paths between a transmitter and receiver [31, 3, 38, 39]. Given that there are only $K \ll N$ paths between a transmitter and receiver, there are at most K directions from which the signal can leave the transmitter and K directions from which it can arrive at the receiver. The 802.11ad standard leverages this structure to speed up the search process. While there are several variations of the standard presented in different documents, at a high level, the standard uses three stages in order to establish a link between an access point (AP) and a client [23, 50]:²

1. **Sector Level Sweep (SLS):** In this stage, the AP sets its receiver beam pattern to a quasi-omnidirectional beam ie, a very wide beam that covers an almost omni directional range. The client then scans and transmits using all its $2N$ beam patterns. Next, this process is repeated with the client setting its

²For simplicity, we assume both AP and client have N antennas.

receiver antenna to quasi-omnidirectional and the AP sweeping through its $2N$ transmit beams. At the end of this stage, the AP and client each pick the γ transmit directions that deliver the largest power.

2. **Multiple sector ID Detection (MID):** In this stage, the scanning roles are reversed. The AP sets its transmit antenna to quasi-omnidirectional while the client sweeps through its receive beams. Then, the client sets its transmit antenna to quasi-omnidirectional while the AP sweeps through its receiver beams. At the end of this stage, the AP and client each pick the γ receive directions that deliver the largest power. Since the direction of the paths that the signal traverses is the same for transmission and reception, there should be no difference between the output of the SLS stage and this stage. In fact, this stage is optional in the standard (see section 9.35.3.1 in [23]) and is included in case of imperfections in the quasi-omnidirectional receive patterns.
3. **Beam Combining (BC):** In this stage, each of the γ best transmit/receive directions at the AP are tried with each of the γ transmit/receive at the client. Hence, γ^2 combinations are tested and the combination of transmit and receive beam directions that deliver the maximum power is then selected and used for beamforming during the data transmission. This stage is also optional in the standard. However, unlike the MID stage, it is highly needed in practice because of multipath i.e., we cannot set $\gamma = 1$ and skip this stage.

This process reduces the number of measurements to $8N + 2\gamma^2$ (or $4N + \gamma^2$ in case MID is skipped) which is $O(N + K^2)$ since $\gamma = O(K)$ where K is the number of paths between transmitter and receiver. This still requires significant time as demonstrated in [52, 38]. It also suffers in multipath scenarios as we will show in Chapter 7.

Chapter 3

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 [7], Samsung’s 28 GHz prototype [35], and 60 GHz products from two startups: Wilocity and SiBeam [1, 36].¹ However, none of these systems present a steering algorithm or measurements of steering delays. In fact, current products are designed for static links [36, 1] and hence take a long time to steer the beam and are not suitable for mobile or multi-user networks [52].

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 [18, 51, 11] as well as cellular picocells and WiFi applications [52, 39, 38]. These implementations mainly focus on using horn antennas to direct the beam. Horn antennas, however, require mechanical steering and are not suitable for non-static links or multi-user networks. In addition to using horn antennas, the work in [52] uses the Wilocity chipsets found on Dell Latitude 6430U laptops and shows that they take multiple seconds to align the beam. Similarly, the work in [38] uses horn antennas to evaluate the 802.11ad standard and shows that it can also take multiple

¹Note that there are other mmWave products on the market [24, 27]. 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 [34], these chips have not been demonstrated to work as part of a full-fledged mmWave communication system.

seconds to align the beam in the right direction which confirms the need for fast beam searching algorithms. The closest to our work is [39] which tries to avoid expensive beam searching by predicting the next best beam to switch to in the case where the current beam gets blocked due to human motion. The work, however, focuses on static links and assumes that signal propagation paths are known and measured a priori which is not the case in dynamic and mobile networks.

mmWave Measurement Studies and Channel Profiling: Our work is also related to several measurement studies that use horn antennas to profile the wireless channels in mmWave networks both indoors and outdoors [38, 31, 3, 40, 48, 41, 37, 45, 32, 10]. These studies confirm the sparsity of the wireless channel at mmWave frequencies showing only 2 to 3 paths between transmitter and receiver. They also emphasize the large overhead of beam searching and hence motivate the need for faster algorithms.

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 standard which use hierarchical beams to speed up the search [26, 4, 49, 25, 46, 50, 42]. However, hierarchical beams are hard to generate using fixed antenna arrays [29, 2]. The work in [2, 14] leverages compressive sensing to optimize the signal power and create good hierarchical beams. This, however requires changing the hardware architecture and connecting all the antenna elements to around 10 to 15 TX/RX chains to achieve good beams (See Fig.5 in [2]). Such architecture does not apply to any existing mmWave hardware and significantly complicates the hardware design and increases the cost [19]. 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 and hence these paths will be lost. Agile-link's algorithm randomizes that hashing of the directions of the paths in order to avoid such worst case scenarios.

Sparse Recovery: Our work is also related to sparse recovery such as compressive sensing [13, 8] and the sparse Fourier transform [21, 20] algorithms. But, these al-

gorithms do not lend well in practice to mmWave systems because of the hardware constraints and architecture of mmWave systems. Specifically, past theoretical work proposes using compressive sensing to reduce the number of measurements needed to discover the right alignment of the beam [30, 17, 16]. This approach does not work with practical hardware because it ignores CFO (Carrier Frequency Offset) between the transmitter and receiver. Specifically, since the measurements occur over time, they accumulate CFO which corrupts the phase of the measurements and hence diverge from the underlying model. 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 because of CFO.

Our work is also related to massive MIMO systems at GHz frequencies (i.e., frequencies similar to today’s WiFi and cellular systems), where sparse recovery is again used to reduce the overhead for discovering the angle of arrival of the signal [47, 6, 5]. These systems are intrinsically different from mmWave phased array because each antenna is connected to its own TX/RX chain, and the channel can be immediately estimated at each antenna. Hence, with N antennas, these systems can immediately retrieve N measurements from a single packet transmission. 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 1-1. Hence, each measurement requires a packet transmission and time on the channel which significantly reduces the throughput.

Chapter 4

Agile-Link Overview

Agile-link’s goal is to discover the spatial directions from which the signal arrives at the receiver or departs from the transmitter in order to steer its beam and maximize the SNR. Since the paths between transmitter and receiver at mmWave frequencies are very sparse [31, 39], there are only a few directions in which the signal has energy. Hence, there is no point scanning and collecting measurements from all possible directions. Agile-link instead leverages the sparsity to minimize the number of measurements needed to discover the directions with highest SNR.

At a high-level, Agile-link works by hashing spatial directions into bins, where each bin collects energy from a wider range of directions in space. (For example, one bin may collect energy from directions 10° , 55° , 100° , and 145° .) A bin will contain energy only if one or more of the signal’s paths matches a direction that is hashed to the bin. Since mmWave signals have two or three paths [31, 3], only a few of the bins will have energy, which will significantly reduce the search space to the directions within those bins.

However, 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. Hence, we need to randomize the hashing; otherwise, we might miss the directions with high SNR. Agile-link repeats the hashing while randomizing the directions that fall into the same bin in order to ensure that if two

paths collide in the first hashing, they will not continue to collide and cancel each other. After repeating the random hashing a few times, Agile-link uses a voting based scheme to discover the directions of the signal that have energy. Specifically, for each hashing, a direction will get a vote if it falls in a bin that has energy. Since the signal is sparse and the hashing is randomized, directions that have energy are likely to get a lot of votes whereas directions that do not have energy are unlikely to get a lot of votes. Thus, the directions that have energy will have the highest votes which allows Agile-link to discover them quickly.

But how should we hash the spatial directions into bins and how do we randomize the hashing? In theory, there are many good hash functions that one could apply to the signal to generate and randomize the hash. In practice, we are allowed to change *only* the phases on the phase shifters (see Fig. 1-1). We are neither allowed to manipulate the magnitude of the signal on the individual antennas nor allowed to turn some antennas off. This renders many of the standard hashing techniques useless and significantly constrains the space of possible beam patterns which we can create to hash the directions to bins.

Chapter 5

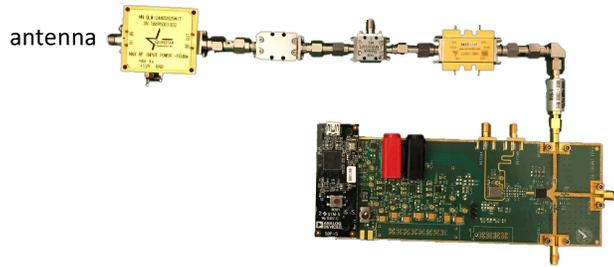
Agile-Link Software Radio Platform

We have designed and built a full-fledged mmWave radio capable of fast beam steering, as shown in Fig.5-1. The radio operates in the new 24GHz ISM band. Its physical layer supports a full OFDM stack up to 256 QAM. Our radio platform addresses critical system and design issues that are described below.

(a) Heterodyne Architecture: mmWave hardware is significantly more expensive than GHz hardware. Thus, we advocate a heterodyne architecture where the mmWave signal is first taken into an intermediate frequency of a few GHz, before the I and Q (real & imaginary) components are separated. Such a design reduces the number of components that need to operate at very high frequencies (e.g., mixers, filters, etc) and replaces them with components that operate at a few GHz, which are much cheaper.

The architecture of Agile-Link's receiver is shown in 5-2(b). The first block is a mmWave phased array which allows us to steer the beam electronically. The phased array consists of 8 antenna elements where each element is connected to a phase shifter component. The outputs of the phase shifters are combined and fed to a single mmWave front-end. Then, the mmWave front-end downconverts the mmWave signal to an intermediate frequency (IF) and feeds it to the daughterboard on the USRP which samples it and passes the digitized samples to the UHD driver.

The mmWave front-end consists of a low-noise amplifier (LNA) to amplify the received signal, a band-pass filter to remove the out-of-band noise and interference



(a) mmWave Radio



(b) Phased Array

Figure 5-1: **Agile-Link’s Platform** The figure shows (a) the mmWave radio platform we built to operate as a daughterboard for the USRP software radio and (b) the phased array with 8 antenna elements and 8 phase shifters.

signals, and a mixer to downconvert the mmWave signal to the IF signal. The main step of our front-end design is the generation of a high-frequency local oscillator (LO), which is used to generate the mmWave carrier. This signal can be obtained from a phase locked loop (PLL) working at mmWave frequencies. Unfortunately, to the best of our knowledge, such PLLs are not available commercially. To overcome this difficulty, we use a component called a frequency doubler. Specifically, instead of using a mmWave PLL, we use a PLL working at much lower frequency and feed its output to a frequency doubler to generate an LO signal at a mmWave frequency.

The mmWave transmitter architecture mirrors that of the receiver to send data at 24Ghz. Data from the USRP is upconverted with a mixer to create the 24Ghz signal. This signal is transmitted through the phased array antenna as shown in 5-2(a).

(b) Integration with GNU-Radio: To support flexible development, we designed our platform as a daughterboard for the USRP software radio. This enables easy manipulation of mmWave signals using standard GNU-radio software, and helps bring mmWave to the GNU-radio community. For example, one can use typical GNU-radio functions to transmit and receive, control the bandwidth of the signal, and change the

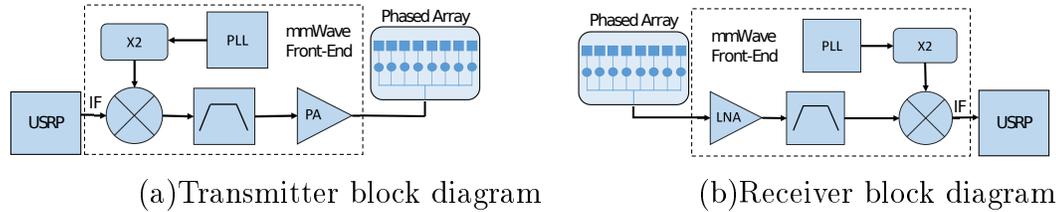


Figure 5-2: **Agile-Link's Architecture.** The figure shows block diagrams for both Agile-Link's transmitter and receiver. Agile-Link uses a heterodyne architecture where the mmWave signal is first taken into an intermediate frequency, before the I and Q components are separated.

center frequency. One can also use our radio to support mmWave MIMO functions, in a manner similar to how one builds a MIMO USRP node –i.e., by coordinating multiple USRPs with a shared clock to act as a MIMO device.

(c) Phased Array Calibration: mmWave phased arrays require a one-time calibration. To see the importance of such calibration, Fig.6-1 plots patterns of Agile-Link's phased array before and after calibration when the beam was steered at 60 degrees. The figure shows that without calibration, the phased array has an unwanted large sidelobe at 80 degrees. Such a sidelobe reduces the power beamed to the receiver and can further cause interference to other mmWave connections.

The need for calibration stems from the non-linearity of phase shifters. Specifically, phase shifters are analog components used to change the phase of an RF signal. The phase shift introduced by a phase shifter is a function of its control voltage. This function is typically provided by the manufacturer as a plot which shows the phase shift for each control voltage. However, once these phase shifters are mounted on the phased array board, they perform differently due to finite size of the antenna array, variation in antenna's feeding network, etc. Hence, it is required to calibrate individual phase shifters after mounting them on the board. To do so, we fix the input of all phase shifters except one, which we vary to scan the whole range of input. We empirically observe the phase shift resulting from each input and create a table that maps a phase shifter's input to the resulting phase.

It is important to realize that each individual phase shifter has to be calibrated, and results in a different calibration table. Fig. 6-2 plots the calibration functions for eight phase shifter in our array. The figure shows that for the same control value,

phase shifters may have up to a 100 degree difference in the amount of phase shift they introduce. Hence, instead of the specification provided by manufacturer, we use the empirical calibration tables to adjust the phase shifters' control voltage.

(d) Parts Used: We implemented the design in Fig. 5-2 using off-the-shelf components. For the mmWave low-noise amplifier (LNA) and power amplifier (PA), we use Hittite HMC-C020 and Quinstar QLW-2440, respectively. For the mmWave mixer, we use Marki M1R-0726MS. To generate local oscillator (LO) signals, we use Analog Devices ADF5355 PLL and Hittite HMC-C035 frequency doubler. We use USRP X310 as an IF and baseband signal processing unit. The phased array is designed using HFSS software and fabricated on printed circuit board (PCB) using Rogers substrate. The phased array includes 8 antenna elements separated by $\frac{\lambda}{2}$, where each element is connected to a Hittite HMC-933 analog phase shifter. We use Analog Device AD7228 digital-to-analog converters (DAC) and Arduino Due micro-controller board to digitally control the phase shifters.

(e) Implementation Details: The phased-array antenna is implemented using an Arduino Due and an 8 output DAC. Each output of the DAC is connected to one of the phase-shifters on the phased array antenna. This way the Arduino can control each phase shifter individually. The Arduino writes a value between 0 to 255 to the DAC which, in turn, sends a voltage between 0 and 5V, respectively, to each phase shifter. The amount of phase shift as a function of input voltage is shown in Fig. 6-2. Using the calibration values from the figure we are able to steer the beam to the desired locations.

Chapter 6

Phased Array Calibration

Similar to phased arrays working at a few GHz, mmWave phased arrays also require a one-time calibration. To see the importance of such calibration, Fig.6-1 plots patterns of Agile-Link's phased array before and after calibration when the beam was steered at 60 degrees. The figure shows that without calibration, the phased array has a unwanted large sidelobe at 80 degrees. This means that a significant amount of power is radiated at some angles which is not desirable. This non-ideality is due to the fact that there is a big difference between the specifications of a phase shifter and how it performs in practice once mounted on the radio board.

Phase shifters are RF components used to change the phase of RF signal. The amount of the phase shift introduced by a phase shifter is a function of their control voltage. This function is typically provided by manufacturer as a plot which shows the phase shift for each control voltage. However, once these phase shifters are mounted on the phased array board, they perform differently due to finite size of the antenna array, variation in antenna's feeding network, etc. Hence , it is required to calibrate individual phase shifter after mounting them on the board.

To calibrate the phase shifters, one can look at the output of individual phase shifter and see how it performs (i.e. finding the introduced phase shift for every control voltage). In our case, since the output of multiple such phase shifters are combined and connected to a single Tx/Rx chain (see Fig. 5-2), we cannot directly isolate the output signal of one phase shifter from another. Instead, we use a method

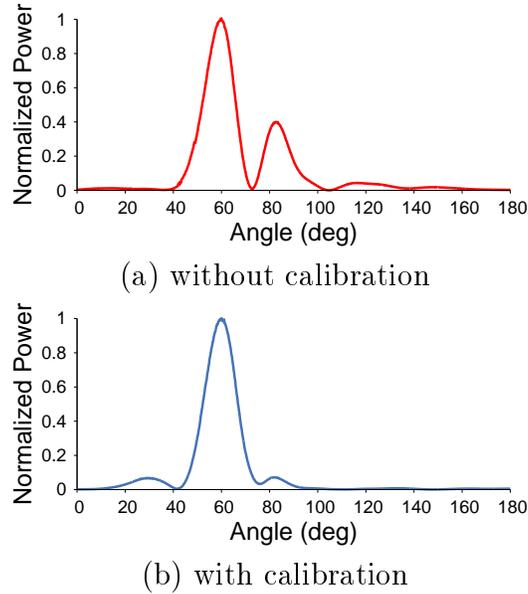


Figure 6-1: **Phased array radiation pattern.** The figure plots the phased array radiation pattern for (a) before and (b) after calibration when the beam was steered at 60 degrees. The figure shows that without calibration, the phased array has an unwanted large sidelobe at 80 degrees. The calibration improves the radiation pattern of the phased array by minimizing the sidelobe.

called Rotating Element Electric Field Vector (REV) to isolate the impact of one phase shifter on the signal of phased array. Specifically, we measure the amplitude variation of the phased array signal while the control value of only one phase shifter is incrementally shifted from 0 to 255 (i.e. max of 8'b) with the other phase shifters held constant. The signal of a phased array is given by the sum of the signal of phase shifters. Hence, when we change the phase of a phase shifter, the summed signal varies as the signal of the phase shifter rotates. We measure the amplitude variation of the sum signal and determine the phase of the phase shifter for each control voltage.

We use the process described above to calibrate the Agile-Link's phase shifters. Fig. 6-2 demonstrates the performance of individual phase shifters mounted on Agile-Link's phased array. The figure plots the phase shift introduced by each phase shifter versus the value of its 8-bits control line. As can be seen, for the same control value, phase shifters may have up to 100 degrees difference in the amount of phase shift they introduce. Hence, instead of the specification provided by manufacturer, we use this plot to adjust the phase shifters' control voltages to generate correct phase shifts.

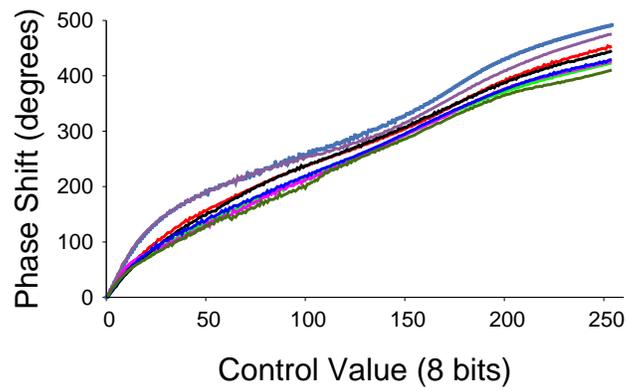


Figure 6-2: **Phase Shifters' Performance.** The figure shows the amount of phase shift introduced by each phase shifter (mounted on Agile-Link's phased array) versus the value of its 8-bits control line. Phase shifters performs differently once they are mounted on the board.

Chapter 7

Experimental Results

We evaluate Agile-Link’s performance using both indoor and outdoor experiments. For indoor scenarios, we ran experiments in an office/lab area with standard furniture and multipath effects. We also ran experiments in an anechoic chamber, where we can accurately measure the ground truth. The anechoic chamber walls are covered with RF absorbers to eliminate multipath and isolate the space from exterior interference. This isolation is necessary to measure the ground truth path traveled by the signal without having RF reflections. Outdoor scenarios involved setting up the system in an open-air park area. We ran our experiments in the 24 GHz ISM band. We use a total of 4 Agile-Link radios, each of them equipped with an 8-antenna phased array. In some experiments, two USRP radios are connected via an external clock to create a MIMO base station. The radios are moved around in the lab and the outdoor space to cover various client locations.

7.1 Agile-Link’s Beam Searching Performance

We start by evaluating Agile-Link’s ability to identify the best beam alignment quickly. Below we describe the baselines we compare against, the metrics we use, the experiments we ran and our results.

A. Baselines:

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 as was described in Chapter 2. This takes $4N^2$ measurements. Then, the combination of transmitter and receiver beams that delivered the maximum power is picked and used for beamforming during data transmission.
- **802.11ad Standard:** Recall from Chapter 2 that 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 $2N$ beams. Then, the γ transmit and receive 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.

B. Metrics:

We evaluate the performance of Agile-Link’s beam searching algorithm along two axes. The first is the accuracy in detecting and aligning the beams of the receiver and transmitter. In this case, our metric is the SNR loss versus the optimal alignment, i.e., how much SNR could we have gained had we known the ground truth. We calculate this metric by measuring the SNR achieved by our beam alignment and subtract it from the SNR achieved by the optimal alignment.

$$SNR_{loss} = SNR_{optimal} - SNR_{Agile-Link} \quad (7.1)$$

The lower the SNR loss, the higher our accuracy in detecting the direction of the signal. In order to measure the optimal SNR, we ran experiments in an anechoic chamber where there is no multipath and we can accurately measure the ground truth direction of the signal and align the beams along those directions. We also ran experiments in multi-path rich environments. In this case, since we do not know the ground truth, we compute the SNR loss metric relative to the exhaustive search

baseline described above.

$$SNR_{loss} = SNR_{Exhaustive} - SNR_{Agile-Link} \quad (7.2)$$

The second axis is the latency in identifying the correct beam alignment. Here, we compute two metrics. The gain verses the exhaustive search baseline and the gain verses the 802.11ad standard.

C. Beam Alignment Accuracy vs. the Ground Truth:

As described above, we first run the experiments in an anechoic chamber, where there is only one strong path that we can accurately measure as our ground truth. For each experiment, we place Agile-Link’s transmitter and receiver at two different locations. We then change the orientation of the transmitter’s and receiver’s antenna arrays with respect to each other. Since there is only a single line-of-sight path in the anechoic chamber, this path will appear at a different direction at the transmitter and at the receiver depending on the orientation of the antenna arrays. Hence, this allows us to test any combination of directions from which the strongest path can leave the transmitter and arrive at the receiver. For each setting, the transmitter then transmits packets consisting of OFDM symbols. The receiver receives these packets, and computes the directions of the best beam alignment. We then steer the beams based on the output of the alignment and measure the SNR achieved by this alignment. We repeat each run with Agile-Link’s beam searching algorithm, the exhaustive search and the 802.11ad standard. In order to calculate the optimal SNR, we align the beams along the direct line of sight since it is the only path and then try very fine grained adjustment to lock in on the best SNR.

Fig. 7-1 plots a CDF of the SNR loss for Agile-Link’s beam searching scheme, the exhaustive search and the 802.11ad standard. The figure reveals two interesting points:

- The figure shows that Agile-Link performs better than the two baselines in that it has minimal SNR loss. While all schemes have a median SNR loss below 1dB, the 90th percentile SNR loss for both exhaustive search and the standard is

3.95dB which is higher than the 1.89dB SNR loss of Agile-Link. This is because the standard and exhaustive search choose to steer using the best beam from the set of $2N$ beams which they tested. However, this does not cover all possible directions. In this case, they will end up picking the closest beam which is not necessarily optimal. SNR loss is further exasperated by the fact that this can happen on both sides i.e., the transmitter and the receiver. In contrast, Agile-Link uses the beams as probability distributions and picks the angle that maximizes the probability of beaming toward the strongest path. Thus, Agile-Link can discover the direction of the path beyond the $2N$ directions used by exhaustive search and the standard.

- The figure also shows that standard and exhaustive search have similar performance. This might seem surprising since one may expect exhaustive search to find a better beam alignment since it spends more time searching the space. However, it is important to recall that the standard differs from the exhaustive search only in the first stage where it uses a quasi-omnidirectional beams to limit the search space to a few top candidates. In the final stage, the standard tries all possible combinations of these candidate beams. Since there is only one path in this experiment, as long as the best beam is picked as one of the candidate beams in the first stage, the standard will converge to the same beam alignment as the exhaustive search. In the next section we will show that this does not hold in multipath settings.

D. Beam Alignment Accuracy in Multipath Environments:

We repeat the above experiments in an office/lab area where, due to multipath, the signal can arrive from different directions. In this case, we do not have the ground truth for the direction of strongest path and hence we measure the SNR loss relative to the exhaustive search baseline. Note that since exhaustive search exhaustively tries all possible combinations of directions, it is not sensitive to multipath and maintains its performance.

Fig. 7-2 plots a CDF of SNR loss for Agile-Link and the standard with respect to

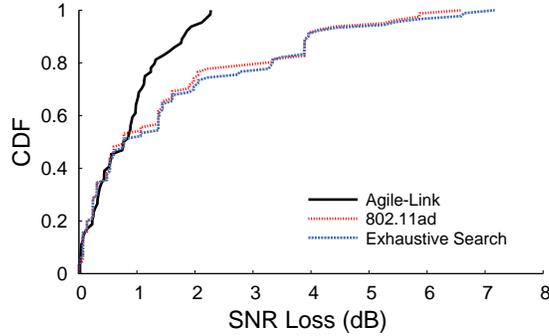


Figure 7-1: **Beam Accuracy with a Single Path** The figure shows the SNR loss due to beam misalignment for Agile-Link, standard and exhaustive search.

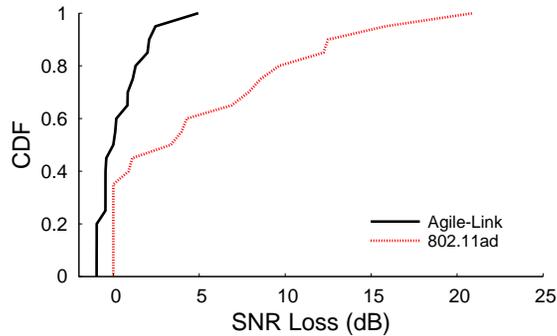


Figure 7-2: **Beam Accuracy with Multipath** The figure shows the SNR loss of standard and Agile-Link with respect to the exhaustive search.

the exhaustive search. The figure shows that the standard performs much worse in multipath scenarios. Specifically, instead of having a similar SNR to the exhaustive search as before, the median and 90th percentile SNR loss (with respect to exhaustive search) are 4dB and 12.5dB, respectively. This is due to the fact that the standard is using its phased array as a quasi-omnidirectional antenna and hence the multiple paths can combine destructively and get lost. Further, due to imperfections in the quasi-omnidirectional patterns, some paths can get attenuated and hence the standard can easily choose the wrong direction to align its beam. In contrast, Agile-Link performs well even in the presence of multipath. Specifically, the median and 90th percentile SNR loss with respect to exhaustive search are 0.1dB and 2.4dB, respectively. Finally, the figure also shows that, in some cases, the Agile-Link SNR loss with respect to exhaustive search is negative. This is because in some case, Agile-Link performs better than exhaustive search for the same reasons described above.

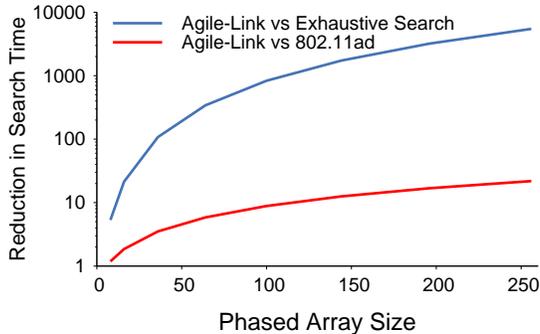


Figure 7-3: **Beam Searching Latency** The figure plots the reduction in search time for Agile-Link compared to standard and exhaustive search.

E. Beam Alignment Latency:

Next we would like to evaluate the gain in latency that Agile-Link delivers over the two baselines. 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.

Fig. 7-3 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 standard, respectively. The gain increases quickly as the number of antennas increase. This is because 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 since it uses $4N^2$ measurements. The standard is linear in the antenna array size since it uses $4N + \gamma^2$ measurements as explained earlier. Agile-Link is sublinear in the array size and uses only $K^2 \log N$ measurements as described earlier.¹ Thus, the gain of Agile-Link over exhaustive search and the standard increases very fast. For arrays of size 256 is $10\times$ better than the standard and multiple orders of magnitude better than exhaustive search.

¹We set K to 4 since most empirical measurement studies [31, 3, 38, 39] show that at mmWave frequencies the channel has only 2 to 3 paths.

7.2 Agile-Link’s Platform Performance

Agile-Link is not only a beam-steering system. It is a full-fledged steerable mmWave phased-array platform. Much of the previous work on mmWave measurements is performed using horn antennas that emulate mmWave phased-array. Thus, in this section, we use Agile-Link to study the performance of mmWave communication with actual phased arrays. Our results give insights to the performance of wireless communications in the 24GHz ISM band. They also demonstrate the high performance of Agile-Link’s platform and its flexibility that allows us to extend it to enable Multi-User MIMO.

A. Agile-Link’s Performance:

We first evaluate Agile-Link’s ability in enabling high data rates and long range communication using phased arrays. We measure the effective SNR of the received signal for different distances between Agile-Link’s receiver and transmitter. Both transmitter and receiver are using phased array antennas and the transmit power complies with FCC part15. At each distance, we run 30 different measurements where we transmit OFDM packets. We then decode the packets and calculate the effective SNR at the receiver side. Fig. 7-4 shows the effective SNR at the receiver side versus the distance between transmitter and receiver ranging from 2.5 m to 100 meters. The figure shows that Agile-Link provides SNR of more than 30 dB for distances smaller than 10 m. As expected, the SNR degrades as the distance increases. However, even at 100 meters, Agile-Link enables an SNR of 17 dB which is sufficient for relatively dense modulations such as 16 QAM [43].²

Fig. 7-5 zooms in on the OFDM modulations at two different distances. The figure shows the constellation for 256 QAM and 16 QAM signals received at 2.5 and 100 meters, respectively. This provides visual evidence that the receiver can accurately decode the received signal, even for very dense constellations like 256 QAM and hence can deliver very high data rates.

²Note that, while one would expect higher SNR at closer distances, the increase in SNR is limited by the dynamic range of the USRP.

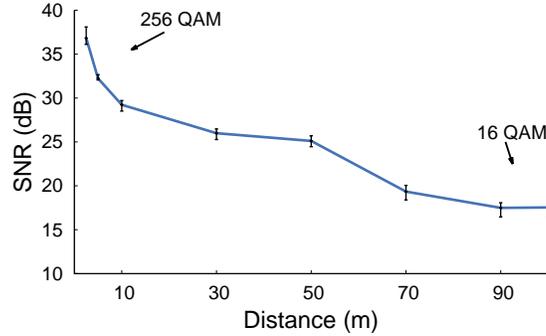


Figure 7-4: **Agile-Link Coverage** Effective SNR at the receiver versus distance between receiver and transmitter.

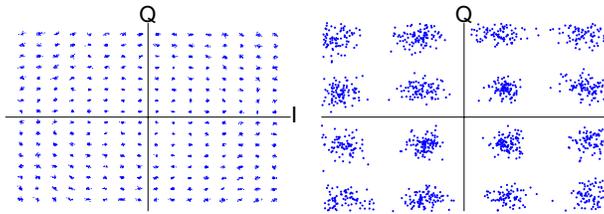


Figure 7-5: **256 QAM and 16 QAM constellation.** Figure shows that Agile-Link provides a full OFDM PHY capable of delivering up to 256 QAM and 16QAM for short and long distances, respectively.

B. Agile-Link's Extension MU-MIMO:

Agile-Link provides a flexible platform that operates as a daughterboard for the USRP software radios. One can use multiple Agile-Link radios to support mmWave MU-MIMO. This is done in a similar manner to how one builds a MIMO USRP node at lower frequencies. We simply connect two USRPs each equipped with an Agile-Link daughterboard to an external shared clock and we connect the Agile-Link radios to the single PLL described in Chapter 5. This gives us a mmWave MIMO node with two chains. We use this node as an access point (AP) transmitter and we implement MU-MIMO enabling it to send two streams simultaneously to two independent clients. We vary the positions of the clients in our office space and measure the gain in throughput of using Agile-Link's MU-MIMO over a mmWave system with no MIMO capability where the access point transmits only one stream at any point in time. Fig. 7-6 shows the gain of Agile-Link's MU-MIMO system. The figure shows that MU-MIMO has increased the network throughput by an average of $1.6\times$.

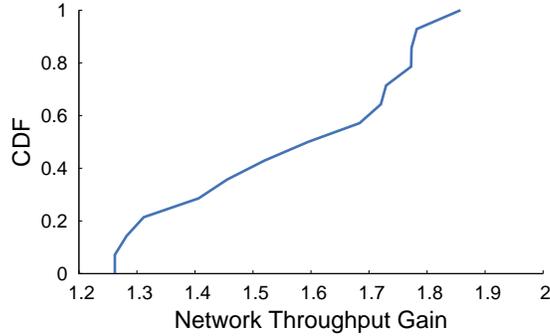


Figure 7-6: **Agile-Link MU-MIMO Network Throughput Gain** The figure shows that Agile-Link design is flexible and can be used to provide mmWave MU-MIMO.

7.3 Importance of Calibration

Finally, we would like to emphasize the importance of calibration. In Chapter 6, we described how Agile-Link calibrates the phase shifters. Here, we evaluate the impact of such calibration on improving the phased array radiation pattern. To do so, we measure and compare the pattern of the phased array before and after the calibration process. We run 120 experiments for a variety of beam directions on 4 different phased arrays (two transmitter and two receiver arrays). We run these experiments in an anechoic chamber designed for antenna measurement. To calibrate the phased array on the transmitter, we use a horn antenna with a narrow beam at the receiver side facing the transmitter. The phased array is mounted on a pole equipped with a precise step motor. We setup the phase shifters of the array to steer at a specific angle, then we rotate the phased array antenna from 0-180 degrees while the horn antenna is receiving the transmitted signal. We measure the received power at each angle which gives us the radiation pattern. We perform the same experiments to calibrate the receiver’s phased arrays while the horn antenna is used for transmitting.

For each measured pattern, we calculate the sidelobe level (SLL), which is the sidelobe power relative to the peak power of the main beam. Higher SLL numbers imply that the antenna array leaks more power outside the main beam, which impedes its beamforming and creates interference. Fig. 7-7 plots a CDF of SLL for our phased arrays before and after calibration. The figure shows that without calibration, the median SLL is -4.48dB, while calibration reduces it to -8.76dB. These results suggest

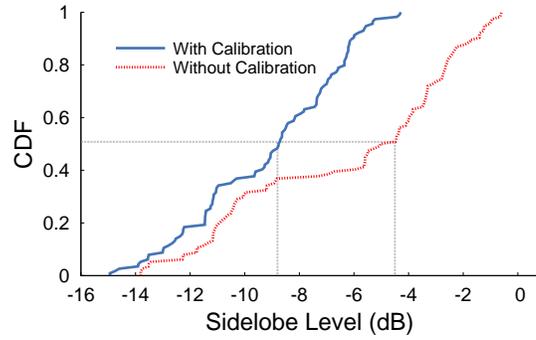


Figure 7-7: **Performance of Phased Array Calibration.** CDF of sidelobe level relative to the main beam with and without calibration. The figure shows that Agile-Link’s calibration significantly reduces the radiation outside of the main beam.

that calibration improves the directionality of the phased array’s beam which improves communication range and data rate, and reduces interference.

Chapter 8

Conclusion

In this thesis, we have presented Agile-Link, the first implemented 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. Agile-Link also delivers a full-fledged mmWave radio platform which operates as a daughterboard for the USRP software radio and helps bring mmWave to the GNU-radio community.

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, and its software radio platform opens up mmWave research to the networking community.

Bibliography

- [1] Wilocity 802.11ad Multi-Gigabit Wireless Chipset. <http://wilocity.com>, 2013.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] C. R. Berger, Z. Wang, J. Huang, and S. Zhou. Application of Compressive Sensing to Sparse Channel Estimation. *IEEE Communications Magazine*, November 2010.
- [7] M. Branda. Qualcomm Research demonstrates robust mmWave design for 5G. Qualcomm Technologies Inc., November 2015.
- [8] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52:489 – 509, 2006.
- [9] Cisco. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013.

- [10] S. Collonge, G. Zaharia, , and G. Zein. Influence of the human activity on wide-band characteristics of the 60 ghz indoor radio channel. *IEEE Transactions on Wireless Communications*, 3(6), 2004.
- [11] 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.
- [12] C. Doan, S. Emami, D. Sobel, A. Niknejad, , and R. Brodersen. Design considerations for 60 GHz CMOS radios. *IEEE Communication Magazine*, 42(12):132.
- [13] D. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289 – 1306, 2006.
- [14] M. E. Eltayeb, A. Alkhateeb, R. W. Heath, and T. Y. Al-Naffouri. Opportunistic Beam Training with Hybrid Analog/Digital Codebooks for mmWave Systems. In *GLOBESIP*, 2015.
- [15] Ericsson. Traffic and market data report, 2011.
- [16] 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.
- [17] 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.
- [18] D. Halperin, S. Kandula, J. Padhye, P. Bahl, and D. Wetherall. Augmenting Data Center Networks with Multi-Gigabit Wireless Links. In *ACM SIGCOMM*, 2011.
- [19] 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.
- [20] H. Hassanieh, P. Indyk, D. Katabi, and E. Price. Nearly optimal sparse fourier transform. In *STOC*, 2012.
- [21] H. Hassanieh, P. Indyk, D. Katabi, and E. Price. Simple and practical algorithm for sparse FFT. In *SODA*, 2012.
- [22] IEEE Standards Association. IEEE Standards 802.15.3c-2009: Millimeter-wave-

- based Alternate Physical Layer Extension, 2009.
- [23] IEEE Standards Association. IEEE Standards 802.11ad-2012: Enhancements for Very High Throughput in the 60 GHz Band, 2012.
 - [24] J. Kilpatrick, R. Shergill, and M. Sinha. 60 GHz Line of Sight Backhaul Links Ready to Boost Cellular Capacity. Analog Devices Inc.
 - [25] J. Kim and A. F. Molisch. Fast Millimeter-Wave Beam Training with Receive Beamforming. *Journal of Communications and Networks*, 16(5), October 2014.
 - [26] 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.
 - [27] Pasternack Enterprises Inc. 60 GHz Transmitter/Receiver Development System. www.pasternack.com.
 - [28] Z. Pi and F. Khan. An introduction to millimeter-wave mobile broadband systems. *Communications Magazine, IEEE*, 2011.
 - [29] K. Ramachandran, N. Prasad, K. Hosoya, K. Maruhashi, and S. Rangarajan. Adaptive Beamforming for 60 GHz Radios: Challenges and Preliminary Solutions. In *mmCom'10, Workshop on Millimeter Wave Communications*, 2010.
 - [30] 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.
 - [31] S. Rangan, T. S. Rappaport, and E. Erkip. Millimeter-wave cellular wireless networks: Potentials and challenges. *Proceedings of the IEEE*, 2014.
 - [32] T. Rappaport, F. Gutierrez, E. Ben-Dor, J. Murdock, Y. Qiao, and J. Tamir. Broadband Millimeter-Wave Propagation Measurements and Models Using Adaptive-Beam Antennas for Outdoor Urban Cellular Communications.
 - [33] 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.
 - [34] G. Rebeiz. Millimeter-wave SiGe RFICs for large-scale phased-arrays. In *IEEE*

Bipolar/BiCMOS Circuits and Technology Meeting (BCTM), 2014.

- [35] 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.
- [36] SiBeam, Lattice Semiconductor. www.sibeam.com.
- [37] P. F. M. Smulders. Statistical Characterization of 60-GHz Indoor Radio Channels. *IEEE Transactions on Antennas and Propagation*, 57(10), 2009.
- [38] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *SIGMETRICS*, 2015.
- [39] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *NSDI*, 2016.
- [40] Sylvain Collonge and Gheorghe Zaharia and Ghais El-Zein. Influence of the Human Activity on Wide-Band Characteristics of the 60 GHz Indoor Radio Channel. *IEEE Transactions on Wireless Communications*, 3(6), November 2004.
- [41] X. Tie, K. Ramachandran, and R. Mahindra. On 60 GHz Wireless Link Performance in Indoor Environments. In *PAM*, 2011.
- [42] Y. M. Tsang, A. S. Y. Poon, and S. Addepalli. Coding the Beams: Improving Beamforming Training in mmWave Communication System. In *IEEE GLOBECOM*, 2011.
- [43] D. Tse and P. Vishwanath. *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- [44] UMTS Forum. Mobile traffic forecasts: 2010-2020 report, 2011.
- [45] E. J. Violette, R. H. Espeland, and G. R. Hand. Millimeter-wave Urban and Suburban Propagation Measurements Using Narrow and Wide Bandwidth Channel Probes. NASA Tech Report.
- [46] 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 Jour-*

nal of Selected Areas in Communications, 27(8), October 2009.

- [47] X. Xie, E. Chai, X. Zhang, K. Sundaresan, A. Khojastepour, and S. Rangarajan. Hekaton: Efficient and Practical Large-Scale MIMO. In *MOBICOM*, 2015.
- [48] H. Xu, V. Kukshya, and T. S. Rappaport. Spatial and Temporal Characteristics of 60-GHz Indoor Channels. *IEEE Journal on Selected Areas in Communications*, 20(3), April 2002.
- [49] W. Yuan, S. M. D. Armour, and A. Doufexi. An Efficient and Low-complexity Beam Training Technique for mmWave Communication. In *PIMRC*, 2015.
- [50] L. Zhou and Y. Ohashi. Efficient Codebook-Based MIMO Beamforming for Millimeter-Wave WLANs. In *PIMRC*, 2012.
- [51] 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.
- [52] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng. Demystifying 60GHz Outdoor Picocells. In *MOBICOM*, 2014.