

Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices

Joan Palacios*
IMDEA Networks Institute
Madrid, Spain
joan.palacios@imdea.org

Daniel Steinmetzer*
Secure Mobile Networking Lab
TU Darmstadt, Germany
dsteinmetzer@seemoo.de

Adrian Loch
IMDEA Networks Institute
Madrid, Spain
adrian.loch@imdea.org

Matthias Hollick
Secure Mobile Networking Lab
TU Darmstadt, Germany
mhollick@seemoo.de

Joerg Widmer
IMDEA Networks Institute
Madrid, Spain
joerg.widmer@imdea.org

ABSTRACT

Beamforming is vital to overcome the high attenuation in wireless millimeter-wave networks. It enables nodes to steer their antennas in the direction of communication. To cope with complexity and overhead, the IEEE 802.11ad standard uses a sector codebook with distinct steering directions. In current off-the-shelf devices, we find codebooks with generic pre-defined beam patterns. While this approach is simple and robust, the antenna modules that are typically deployed in such devices are capable of generating much more precise antenna beams. In this paper, we adaptively adjust the sector codebook of IEEE 802.11ad devices to optimize the transmit beam patterns for the current channel. To achieve this, we propose a mechanism to extract full channel state information (CSI) regarding phase and magnitude from coarse signal strength readings on off-the-shelf IEEE 802.11ad devices. Since such devices do not expose the CSI directly, we generate a codebook with phase-shifted probing beams that enables us to obtain the CSI by combining strategically selected magnitude measurements. Using this CSI, transmitters dynamically compute a transmit beam pattern that maximizes the signal strength at the receiver. Thereby, we automatically exploit reflectors in the environment and improve the received signal quality. Our implementation of this mechanism on off-the-shelf devices demonstrates that adaptive codebook optimization achieves a significantly higher throughput of about a factor of two in typical real-world scenarios.

*These authors contributed equally to this work.

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

© 2018 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *The 24th Annual International Conference on Mobile Computing and Networking (MobiCom '18), October 29–November 2, 2018, New Delhi, India*, <https://doi.org/10.1145/3241539.3241576>.

ACM Reference Format:

Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. 2018. Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices. In *The 24th Annual International Conference on Mobile Computing and Networking (MobiCom '18), October 29–November 2, 2018, New Delhi, India*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3241539.3241576>

1 INTRODUCTION

Designing wireless communication systems that operate at very high frequencies such as the 60 GHz band is technically challenging. For instance, phased antenna arrays are hard to manufacture [13, 24] and phase noise plays a significant role [5, 37]. As a result, system designers often resort to simple yet sub-optimal solutions. This is the case for certain mechanisms such as analog beamforming in the IEEE 802.11ad standard for communication at 60 GHz. Beamforming is crucial for such systems since devices must use directional communication to overcome the very high path loss in the millimeter-wave band and reach the receiver. However, the standard is limited to a basic beamforming mechanism based on a codebook of generic beam patterns. Such patterns are envisioned to have the shape of uniform sectors and cover the entire azimuth range of the device. That is, instead of beamforming towards a specific direction, devices choose the beam pattern out of their codebook which provides the highest gain in that direction.

Beam training, also called Sector Level Sweep (SLS) in infrastructure-based 802.11ad networks, works as follows. The Access Point (AP) transmits beacon messages using each of its available beam patterns sequentially, while the station (STA) listens with a quasi-omnidirectional beam pattern. After that, the STA repeats the same process but includes in each of its messages the identifier of the beam pattern that it received best from the AP. Finally, the AP replies with the identifier of the best beam pattern of the STA in a dedicated

control message. While this mechanism is straightforward, it clearly does not exploit the full potential of the antenna array of IEEE 802.11ad devices. First, none of the available beam patterns in the codebook may steer exactly towards the receiver. Second, strong reflectors in the environment may remain unused even though they could contribute to the received signal strength. Third, existing reflections may result in destructive interference at the receiver, causing significant harm to the communication. If devices were to adapt their beam patterns to the specific environment in which they operate, they could easily mitigate the above issues. However, this requires full Channel State Information (CSI) at the transmitter, which is particularly challenging to obtain in millimeter-wave systems. The reasons are twofold. First, the feedback overhead of full CSI is significant. Second, devices must probe enough orthogonal beam patterns to cover all of the dimensions of the channel. The channel information extracted from all of the probes must be coherent in phase. In particular, achieving such phase coherence over different probes in consumer-grade devices is non-trivial.

In this paper, we design a mechanism that enables full CSI extraction on consumer-grade Commercial Off-The-Shelf (COTS) IEEE 802.11ad devices using only non-coherent signal-to-noise ratio (SNR) measurements. It allows us to fully exploit the capabilities of phased antenna arrays in such devices. Our Adaptive Codebook Optimization (ACO) mechanism probes the channel using carefully engineered beam patterns that allow devices to extract both amplitude and phase information from simple SNR readings. Obtaining the SNR is required by the IEEE 802.11ad standard to select the best beam during the sector sweep. Hence, measurements of this metric are supported on all IEEE 802.11ad devices which makes ACO widely applicable. Specifically, ACO measures the *relative phase* among the elements of the antenna array of a node. Figure 1 depicts a toy example of ACO for an antenna array with two-bit phase shifters per antenna element. ACO defines one of the elements of the array as a reference R . To measure the relative phase to an element E , it switches on R and E , keeping all other elements off. It then transmits one probe frame for each of the four possible values of the phase shifter at E . As shown in Figure 1, the complex gains of R (blue) and E (red) sum up (sum shown in black) for each of the four probes. However, the receiver only observes the resulting amplitude in terms of received signal strength. Since the four probes cover a 2π phase range, the four powers must lie on a sinusoidal curve. Our key insight is that the initial phase of that curve is directly related to the relative phase among R and E . ACO exploits this property using Fourier analysis to reconstruct the sinusoidal curve and compute the phase shift δ . Figure 1 shows an example for $\delta = \pi/4$. Repeating the process for each element of the antenna array, ACO can compute the full CSI.

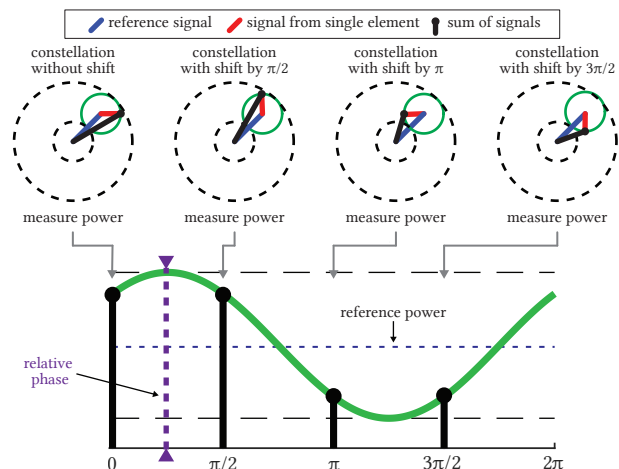


Figure 1: ACO toy example. The upper part shows the signal constellations with four different phase shifts at a single antenna element, whereas the lower part illustrates how ACO derives the relative phase using a sinusoidal curve.

The above technique is the core of ACO. Using the computed CSI, ACO derives beam patterns that maximize the SNR, exploit reflections, and prevent destructive interference. However, the toy example in Figure 1 may raise a number of concerns which we discuss below. In this paper, we develop ACO to address these concerns effectively.

- (1) **Overhead.** ACO requires at least four probes per antenna element, which increases the probing overhead compared to a codebook of generic beam patterns. However, in practice, only a subset of the antennas of the array has a significant impact. We show that obtaining the CSI for just a subspace of the channel is sufficient in most cases. As a result, (a) the probing overhead is similar to that for pre-defined generic beam patterns (i.e., IEEE 802.11ad SLS), and (b) the feedback overhead is limited.
- (2) **Performance.** ACO transmits using only two antenna elements simultaneously during the probing phase. For long links, the resulting transmit power may not be sufficient to reach the receiver, and thus the CSI measurement may fail. To address this issue, we extend the basic approach in Figure 1 to use not just single antenna elements but entire beam patterns. This enables ACO to operate also on low SNR links.
- (3) **Accuracy.** In our toy example, ACO probes each possible value of the phase shifters. The resulting overhead is limited for two-bit phase shifters but would increase significantly for higher-resolution hardware. However, ACO only needs four samples to reconstruct the sinusoidal curve in Figure 1. To keep the overhead low,

ACO only probes four phase shifts even if more combinations are possible.

We implement ACO on commodity COTS hardware. To this end, we gain full access to the beamforming control of the TP-Link Talon AD7200 60 GHz router. This device features a Qualcomm QCA9500 chipset that implements the IEEE 802.11ad standard. We disassemble the phased antenna array of the device to understand its structure and experimentally reconstruct the antenna weighting network. This enables us to set arbitrary beam patterns on the antenna. However, *ACO itself does not require any information about the antenna structure to operate*. Obtaining the weighting network of the QCA9500 60 GHz solution is a valuable contribution to the community in itself since Qualcomm's chipset has become a major research platform in the field [34, 39]. We modify the sector sweep behavior of the router to include the custom beam patterns that enable ACO. We embed ACO in the regular operation of the router, that is, *we do not obtain our results in post-processing* but as part of the normal data transmissions in our testbed. This enables us to quantify the performance of ACO in terms of SNR, data rate, and throughput using the Transmission Control Protocol (TCP). Our contributions are as follows:

- We design ACO, a method to obtain full CSI in terms of phase and amplitude using only SNR measurements. The key to our method is obtaining SNR values that resemble a discrete Fourier decomposition.
- We implement our method on commodity 60 GHz hardware. To this end, we disassemble a COTS 60 GHz router and modify its sector sweep behavior.
- We evaluate ACO in a real-world office environment and show that it achieves on average $2.5\times$ higher SNR as well as up to $2\times$ higher throughput.

2 ADAPTIVE CODEBOOK OPTIMIZATION

In the following, we define the architecture and the operation of ACO. For illustration purposes, we consider an infrastructure-based IEEE 802.11ad network. However, ACO is not limited to such networks but is applicable to any millimeter-wave link. In our case, both the AP and the STA can benefit from ACO for beamforming.

Figure 2 depicts a schematic overview of the operation of ACO, which consists of two phases: initialization and continuous adaptation. In the initialization phase, devices discover each other and establish a connection using the default IEEE 802.11ad procedure based on generic beam patterns. After that, ACO exhaustively probes all of the antenna elements of the array. This one-time overhead allows ACO to obtain the full CSI and to determine which of the elements contribute most to it. Based on this information, ACO restricts

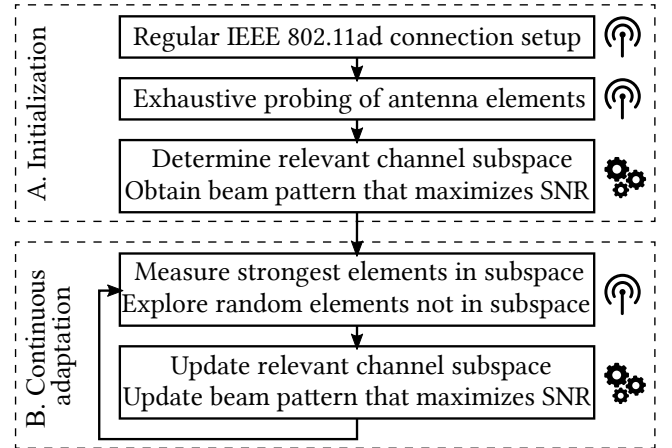


Figure 2: ACO operation with processing (gears symbol) and communication (antenna symbol) steps.

the following channel measurement to those elements. This reduces the probing overhead significantly and allows ACO to focus on the subspace of the channel which is relevant for a certain physical environment. To determine which and how many antenna elements N to use, ACO maximizes the expression in Equation 1, where a_k is the amplitude of antenna k and antenna elements are ordered by strength.

$$N = \arg \max_n \frac{(\sum_{k=1}^n a_k)^2}{n} \quad (1)$$

To complete the initialization phase, ACO obtains the beam pattern that maximizes the SNR within the computed subspace and uses it for regular communication. After that, ACO proceeds to the continuous adaptation phase, which allows ACO to readjust the beam pattern as the CSI and the environment change during on-going communication. Similarly to IEEE 802.11ad, ACO probes the channel periodically but limits the channel measurement to the strongest antenna elements and thus to the current subspace. However, the subset of relevant antenna elements may change if the environment changes due to, e.g., user mobility. To adapt to such changes, ACO divides its periodic channel probing sweep into a measurement part and an exploration part. The channel measurement is performed over the subspace of strongest antennas, whereas the exploration part probes a subset of randomly chosen antenna elements to identify potential changes in the relevant channel subspace. As shown in Figure 3, ACO can choose the overall duration of the sweep such that its duration is equivalent to a regular IEEE 802.11ad sector sweep. Still, if full CSI is desired, ACO can easily be reconfigured to probe all antenna elements at the expense of higher overhead. After each periodic sweep, ACO recomputes the beam pattern that maximizes the SNR and data communication

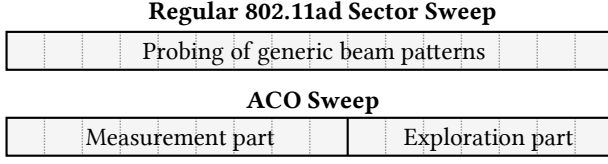


Figure 3: ACO sweep compared to the IEEE 802.11ad sector sweep.

continues. The complexity of the processing steps shown in Figure 2 is negligible due to the very high computational efficiency of our method (see Section 4). This enables ACO to operate both on powerful APs as well as resource-constrained STAs.

3 SYSTEM MODEL

To formally define the operation of ACO, we introduce our system model in this section. Specifically, we define the following antenna and channel models.

3.1 Antenna Model

For both the transmitter and the receiver, we consider a phased antenna array that implements analog beamforming. That is, all antenna elements are connected via a network of amplifiers and phase shifters to a single radio frequency (RF) chain. This type of beamforming is widely used in COTS 60 GHz hardware since more advanced architectures such as digital or hybrid beamforming are not cost-efficient. We allow for both regular and irregular antenna element layouts.

3.2 Channel Model

Given a transmitter TX that sends a signal x towards a receiver RX, the received signal can be expressed as

$$\mathbf{y} = \mathbf{c}^H \mathbf{H} \mathbf{p} \mathbf{x} + \mathbf{c}^H \mathbf{N}, \quad (2)$$

where \mathbf{c} and \mathbf{p} are the complex gains of the antennas at RX and TX respectively, and \mathbf{N} is uncorrelated additive white noise received at each antenna with variance σ^2 . The channel \mathbf{H} reflects the propagation effects. We consider the geometric channel model that is widely used in the millimeter-wave literature. Given a set of L different paths from the transmitter to the receiver, we can represent the channel as:

$$\mathbf{H} = \sqrt{P_t} \sum_{l=1}^L \alpha_l \mathbf{a}_{\text{RX}}(\mathbf{v}_l) \mathbf{a}_{\text{TX}}(\mathbf{w}_l)^H, \quad (3)$$

where P_t is the transmit power, L is the number of paths, α_l is the complex gain, \mathbf{v}_l , \mathbf{w}_l are unitary vectors of arrival and departure for the l^{th} path, and $\mathbf{a}_{\text{TX}} : \mathbf{S} \rightarrow \mathbb{C}^{N_{\text{TX}} \times 1}$ as well as $\mathbf{a}_{\text{RX}} : \mathbf{S} \rightarrow \mathbb{C}^{N_{\text{RX}} \times 1}$ are the steering vectors at TX and RX respectively. We assume no prior knowledge regarding the steering vectors \mathbf{a}_{TX} and \mathbf{a}_{RX} because they are typically

unknown for COTS devices. Further, we denote the receiving channel for a fixed \mathbf{c} as $\mathbf{h}_{\text{RX}} = \mathbf{c}^H \mathbf{H}$. Thus, the received signal can be written as $\mathbf{y} = \mathbf{h}_{\text{RX}} \mathbf{p} \mathbf{x} + \mathbf{c}^H \mathbf{N}$. Conversely, the transmitting channel for a fixed \mathbf{p} is $\mathbf{h}_{\text{TX}} = \mathbf{H} \mathbf{p}$ and hence $\mathbf{y} = \mathbf{c}^H \mathbf{h}_{\text{TX}} \mathbf{x} + \mathbf{c}^H \mathbf{N}$. Based on these definitions, we can compute the received signal power as $|\mathbf{h}_{\text{RX}} \mathbf{p}|^2$ or $|\mathbf{c}^H \mathbf{h}_{\text{TX}}|^2$, and the received noise power as $\|\mathbf{c}\|^2 \sigma^2$. In Table 1, we list the mathematical symbols that are used throughout this paper.

While COTS devices may provide SNR or RSSI values, they typically do not expose detailed signal power measurements. Since $\|\mathbf{c}\|$ is a constant value, we define $\|\mathbf{c}\|^2 \sigma^2 = 1$ as noise power reference. Using this definition along with the well-known formula of the SNR in logarithmic units, we obtain that the received signal power is $|\mathbf{h}_{\text{RX}} \mathbf{p}|^2 = 10^{\frac{\text{SNR}}{10}}$.

4 CHANNEL MEASUREMENT ON OFF-THE-SHELF DEVICES

Based on the system model introduced in Section 3, in the following, we detail the design of the core component of ACO, that is, retrieving complex gains from simple signal strength readings. This is particularly challenging because commodity COTS devices use cost-efficient components that do not allow for exact amplitude and phase measurements. To address this limitation, we transmit multiple probes with different phase shifts and use this information to reconstruct the full CSI, as sketched in Section 1. Figure 1 summarizes this procedure.

To enable the above measurement, ACO needs to retrieve the complex gain of each antenna \mathbf{h}_{RX} . However, obtaining the actual value of \mathbf{h}_{RX} is impossible because any set of measurements explained by \mathbf{h}_{RX} can also be explained by $\alpha \mathbf{h}_{\text{RX}}$ where α is any complex number with unitary modulus. Thus, ACO computes \mathbf{h}_{RX} relative to a given \mathbf{p} in terms of phase shift. That is, we assume that $\arg(\mathbf{h}_{\text{RX}} \mathbf{p}) = 0$. This

Table 1: List of used symbols.

Symbol	Description
k	Antenna index
\mathbf{H}	Channel matrix
\mathbf{h}_{TX}	Transmission channel matrix
\mathbf{h}_{RX}	Reception channel matrix
\mathbf{c}	Reception beam pattern
\mathbf{p}	Transmission beam pattern
\mathbf{p}'	Unnormalized transmission beam pattern
\mathbf{x}	Signal
\mathbf{N}	Noise matrix
m	Measurement index
\mathbf{a}	Measurements matrix, where each $[\mathbf{a}]_m$ corresponds to \mathbf{p}_m
\mathbf{f}	Measurements matrix discrete Fourier transform
Γ	$ \mathbf{f} _1$
Δ	$ \mathbf{f} _2$

is not a problem because we only need the relative phases among antenna elements to measure path interactions and compute beamforming weights that maximize the SNR. Further, determining the actual value of \mathbf{h}_{RX} is not useful in practice because the phase of \mathbf{h}_{RX} can change by a full π rad if the length of the link changes by half a wavelength, which for 60 GHz is less than two millimeters. That is, any slight movement would invalidate the result. Fortunately, ACO only requires *relative* phase shifts.

4.1 Complex Gain Retrieval

The intuitive explanation in Figure 1 depicts how ACO measures relative phase shifts among individual antenna elements. In Section 4.1.1 we formalize this procedure. Using such individual antenna elements translates into estimating the channel with quasi-omnidirectional beam patterns. This approach is highly effective on high SNR links and is immune to sudden changes in the environment. For instance, if the alignment among devices changes abruptly, using quasi-omnidirectional beam patterns allows ACO to react instantly to the change. Since millimeter-wave communication is typically restricted in range (e.g., within a room), the high SNR case is common.

Nevertheless, we extend ACO in Section 4.1.2 such that it can also operate in low SNR cases. To this end, we generalize the approach in Figure 1 such that the reference signal is generated not only by a single antenna element but via a directional beam pattern. This beam pattern provides a higher SNR than the omnidirectional reference, but only allows to retrieve the complex gain in specific directions. The underlying concept is identical but the analytical definition changes. We formalize this approach in Section 4.1.2. Further details on the mathematical derivation of both approaches can be found in the accompanying technical report [26].

4.1.1 Quasi-Omnidirectional Estimation. We identify each of the elements on the antenna array using index k . Further, we define the index of the reference antenna R (see Section 1) as \bar{k} . We divide the procedure to retrieve the complex gain into the measurement of (a) the amplitude and (b) the phase.

(a) Amplitude Measurement. Measuring $|\mathbf{h}_{\text{RX}}|_k \forall k$ is straightforward since ACO only needs to transmit probes switching on one element at a time, as defined below.

$$[\mathbf{p}']_{k'} = \begin{cases} 0 & k' \neq k \\ 1 & k' = k \end{cases}, \quad \mathbf{p} = \frac{\mathbf{p}'}{\|\mathbf{p}'\|}. \quad (4)$$

In the remainder of this subsection, we use the prime symbol to refer to the non-normalized version of a beam pattern. As per Section 3.2, ACO obtains the amplitude as $|\mathbf{h}_{\text{RX}}|_k|^2 = |\mathbf{h}_{\text{RX}}\mathbf{p}'|^2$. Based on the result, ACO defines as reference element \bar{k} the element with the largest amplitude.

(b) Phase Measurement. Next, we need to obtain the phase shift of each element, except for the reference \bar{k} , which by definition has zero phase. That is, we obtain $\arg[\mathbf{h}_{\text{RX}}]_k, \forall k \neq \bar{k}$. Since we measure the relative phase of element $k \neq \bar{k}$ with respect to \bar{k} , we need to enable both antennas k and \bar{k} . We then probe the phase values $[\mathbf{p}']_k \in \{1, i, -1, -i\}$ while $[\mathbf{p}']_{\bar{k}} = 1$ and all other components are 0. We then group the measurements $|\mathbf{h}_{\text{RX}}\mathbf{p}_m|^2$ in a vector of four components $\mathbf{a} \in \mathbb{R}^{4 \times 1}$ where $[\mathbf{p}'_m]_k$ has the values 1, i , -1 , $-i$ for $[\mathbf{a}]_1, [\mathbf{a}]_2, [\mathbf{a}]_3, [\mathbf{a}]_4$, respectively. That is, each element in \mathbf{a} contains the sum of two signals—the signal transmitted by the reference antenna element \bar{k} and the signal transmitted by measured antenna element k for one out of four possible phase shifts m . We can reformulate this as in Equation 5.

$$\begin{aligned} [\mathbf{a}]_m &= |\mathbf{h}_{\text{RX}}\mathbf{p}_m|^2 = \left| \frac{1}{\sqrt{2}}([\mathbf{h}_{\text{RX}}]_{\bar{k}} + [\mathbf{h}_{\text{RX}}]_k e^{(m-1)\frac{\pi}{2}i}) \right|^2 \\ &= \frac{1}{2}(|\mathbf{h}_{\text{RX}}]_{\bar{k}}|^2 + |\mathbf{h}_{\text{RX}}]_k|^2 e^{(\arg[\mathbf{h}_{\text{RX}}]_k + (m-1)\frac{\pi}{2})i})^2 \\ &= \Gamma + 2\Delta \cos(\arg[\mathbf{h}_{\text{RX}}]_k + (m-1)\frac{\pi}{2}) \\ &= \Gamma + \Delta e^{(\arg[\mathbf{h}_{\text{RX}}]_k + (m-1)\frac{\pi}{2})i} + \Delta e^{-(\arg[\mathbf{h}_{\text{RX}}]_k + (m-1)\frac{\pi}{2})i} \\ &= \Gamma + \Delta e^{\arg[\mathbf{h}_{\text{RX}}]_k i} e^{(m-1)\frac{\pi}{2}i} + \Delta e^{-\arg[\mathbf{h}_{\text{RX}}]_k i} e^{-(m-1)\frac{\pi}{2}i} \end{aligned}, \quad (5)$$

for $\Gamma = \frac{|\mathbf{h}_{\text{RX}}]_{\bar{k}}|^2 + |\mathbf{h}_{\text{RX}}]_k|^2}{2}$ and $\Delta = \frac{|\mathbf{h}_{\text{RX}}]_{\bar{k}}| |\mathbf{h}_{\text{RX}}]_k|}{2}$. This is a wave decomposition expression. Hence, the fast Fourier transform of \mathbf{a} defined as $\mathbf{f} = \text{FFT}(\mathbf{a})$ has the expression

$$\mathbf{f} = [\Gamma, \Delta e^{\arg[\mathbf{h}_{\text{RX}}]_k i}, 0, \Delta e^{-\arg[\mathbf{h}_{\text{RX}}]_k i}]^T. \quad (6)$$

This means that the phase shift of antenna k with respect to \bar{k} is $\arg[\mathbf{f}]_2$, i.e., the phase of the second element of the fast Fourier transform of the received signal strength vector \mathbf{a} .

The above approach is feasible whenever the transmitter can reach the receiver using the quasi-omnidirectional beam pattern given by switching on only the reference antenna element \bar{k} . The stronger \bar{k} , the more accurate is the estimation of the phase shift. As described in Section 2, we define a measurement codebook that contains the required phase shifts to compute the phase for all of the elements of the antenna. The number of probes in the ACO sweep (see Figure 3) is directly related to N , which is the overall number of antenna elements probed in the measurement and the exploration parts of the sweep.

- For antenna element \bar{k} , we only need one probe because we define it as a phase reference and thus we only measure its amplitude. Its relative phase is zero.
- For all other antenna elements, we need one probe to measure its amplitude and four additional probes to obtain its relative phase with respect to \bar{k} .

Thus, the ACO sweep contains a total of $N + 4(N - 1)$ probes to measure the complex gain of N antenna elements.

4.1.2 Directional Estimation. In the following, we discuss an extension of our approach in Section 4.1.1 for the low SNR case. Intuitively, instead of using a single antenna element

as a reference, we resort to a beam pattern \mathbf{p} . Due to the higher directionality compared to the quasi-omnidirectional case, this results in much higher received signal strength during the channel estimation. Beam pattern \mathbf{p} is typically one that the transmitter used earlier for communication. For instance, ACO could use one of the generic beam patterns of the device. Based on this, \mathbf{p} is suitable as a reference if \mathbf{p} has higher directionality than any single antenna beam pattern even when switching off that antenna. That is, if we set any k and define $\bar{\mathbf{p}}$ as $[\bar{\mathbf{p}}']_{k'} = [\mathbf{p}']_{k'}$ for any $k' \neq k$ and $[\bar{\mathbf{p}}']_k = 0$, then $|\mathbf{h}_{\text{RX}}\bar{\mathbf{p}}| > |\mathbf{h}_{\text{RX}}\mathbf{k}|$. This condition is likely to hold for a generic beam pattern selected for communication since it is very probable that such a beam pattern has a higher gain than a quasi-omnidirectional one. Since ACO uses $\mathbf{h}_{\text{RX}}\mathbf{p}$ as the phase reference, we define its phase to be zero, that is, $\arg(\mathbf{h}_{\text{RX}}\mathbf{p}) = 0$. Based on the above definitions, we compute the amplitude and the relative phase for antenna element k . We distinguish two cases: (a) a simple case when antenna element k is switched off in the reference \mathbf{p} , and (b) an analytically harder case when k is on in \mathbf{p} .

(a) *Antenna element k is off in \mathbf{p} .* Since element k is off, $[\mathbf{p}]_k = 0$. We define \mathbf{p}_m such that $[\mathbf{p}]_{k'} = [\mathbf{p}_m]_{k'} \forall k' \neq k$ with $[\mathbf{p}_m]_k = 1, i, -1, -i$ for $m = 1, 2, 3, 4$, respectively. That is, we probe all possible phase shifts for antenna k . Next, we follow the same approach as in Equation 5 but substitute $[\mathbf{h}_{\text{RX}}]_{\bar{k}}$ with $\mathbf{h}_{\text{RX}}\mathbf{p}$. The result for the phase is equivalent, that is, the relative phase shift of antenna element k is again $\arg[\mathbf{f}]_2$. However, the amplitude is harder to measure since we cannot measure each antenna element individually. To compute it, we resort to Γ and Δ , which for the directional case are as defined in Equation 7.

$$\Gamma = \frac{(\mathbf{h}_{\text{RX}}\mathbf{p})^2 + |\mathbf{h}_{\text{RX}}]_k|^2}{(\|\mathbf{p}'\|^2 + 1)} \quad \Delta = \frac{(\mathbf{h}_{\text{RX}}\mathbf{p})[\mathbf{h}_{\text{RX}}]_k}{(\|\mathbf{p}'\|^2 + 1)} \quad (7)$$

Per Equation 6, both Γ and Δ are known since we can obtain them directly from the fast Fourier transform of the amplitude measurements \mathbf{a} . Thus, Equation 7 is a quadratic system of equations from which we can compute $|\mathbf{h}_{\text{RX}}]_k|$. Further, since $\mathbf{h}_{\text{RX}}\mathbf{p} \in \mathbf{R} > |\mathbf{h}_{\text{RX}}]_k|$, the system has two real-valued solutions as shown in Equation 8.

$$\begin{aligned} \mathbf{h}_{\text{RX}}\mathbf{p} + |\mathbf{h}_{\text{RX}}]_k| &= \sqrt{(\Gamma + 2\Delta)(\|\mathbf{p}'\|^2 + 1)} \\ \mathbf{h}_{\text{RX}}\mathbf{p} - |\mathbf{h}_{\text{RX}}]_k| &= \sqrt{(\Gamma - 2\Delta)(\|\mathbf{p}'\|^2 + 1)} \end{aligned} \quad (8)$$

Rearranging Equation 8 and substituting Γ and Δ with the coefficients of the fast Fourier transform of \mathbf{a} , we obtain the amplitude of the complex gain as in Equation 9.

$$|\mathbf{h}_{\text{RX}}]_k| = \frac{\sqrt{\|\mathbf{p}'\|^2 + 1}}{2} (\sqrt{[\mathbf{f}]_1 + 2|[\mathbf{f}]_2|} - \sqrt{[\mathbf{f}]_1 - 2|[\mathbf{f}]_2|}). \quad (9)$$

(b) *Antenna element k is on in \mathbf{p} .* Since element k is on, $[\mathbf{p}]_k \neq 0$. We again probe all possible phase shifts for antenna

k , that is, we set $[\mathbf{p}'_m]_k = 1, i, -1, -i$ for $m = 1, 2, 3, 4$, group measurements into a variable \mathbf{a} , and compute $\mathbf{f} = \text{FFT}(\mathbf{a})$. We obtain the amplitude and phase of antenna element k following an approach along the lines of (a). Since the detailed analysis is similar but more protracted than in (a), we do not discuss the details but directly give the analytical expression for the amplitude in Equation 10 and the one for the relative phase in Equation 11.

$$|\mathbf{h}_{\text{RX}}]_k| = \frac{\|\mathbf{p}'\|}{2} (\sqrt{[\mathbf{f}]_1 + 2|[\mathbf{f}]_2|} - \sqrt{[\mathbf{f}]_1 - 2|[\mathbf{f}]_2|}) \quad (10)$$

$$\arg[\mathbf{h}_{\text{RX}}]_k = \arg[\mathbf{f}]_2 - \arg\left(\gamma + |\mathbf{h}_{\text{RX}}]_k|[\mathbf{p}'_k]_k e^{\arg[\mathbf{f}]_2}\right). \quad (11)$$

In Equation 11, parameter γ is defined as in Equation 12 and is used to subtract the phase contribution to the beam pattern \mathbf{p} of the antenna that we are measuring.

$$\gamma = \frac{\|\mathbf{p}'\|}{2} (\sqrt{[\mathbf{f}]_1 + 2|[\mathbf{f}]_2|} + \sqrt{[\mathbf{f}]_1 - 2|[\mathbf{f}]_2|}). \quad (12)$$

Similarly to the quasi-omnidirectional case in Section 4.1.1, we define a measurement codebook that contains the required phase shifts to compute the complex gains of the relevant antenna elements. However, in this case, the number of required probes depends on whether each antenna element is on or off in the reference beam pattern \mathbf{p} .

- If the element is on in \mathbf{p} , we need to measure \mathbf{p} and three additional probes with changing phase shifts.
- If the element is off in \mathbf{p} , we need to transmit four probes to compute its relative phase and amplitude.

Thus, the ACO sweep contains a total of $1 + 3N_{\text{on}} + 4N_{\text{off}}$ probes, where $N_{\text{on}}/N_{\text{off}}$ is the number of antenna elements that are on/off in reference \mathbf{p} , respectively.

4.2 Channel Matrix Computation

Using the methods introduced in Section 4.1, ACO can compute \mathbf{h}_{RX} and analogously \mathbf{h}_{TX} . That is, it can obtain the receiving and the transmitting channel. In the following, we formulate how both can be combined to obtain the full CSI \mathbf{H} of the measured subspace. Since both \mathbf{h}_{RX} and \mathbf{h}_{TX} are relative values, we multiply them with unitary complex values δ and λ to model the unknown relative phase shift, obtaining $\mathbf{h}_{\text{RX}} = \delta \mathbf{c}^H \mathbf{H}$ and $\mathbf{h}_{\text{TX}} = \lambda \mathbf{H} \mathbf{p}$. For each antenna element $k \in 1, 2, \dots, N$ in the subspace, we rearrange the measured complex gain of the transmit and the receive channel into the matrices \mathbf{G}_{TX} and \mathbf{G}_{RX} , respectively. Specifically, we stack them as $[\mathbf{G}_{\text{RX}}]_{k,:} = \delta_k [\mathbf{W}]_{:,k}^H \mathbf{H}$ and $[\mathbf{G}_{\text{TX}}]_{:,k} = \lambda_k \mathbf{H} [\mathbf{W}]_{:,k}$, where \mathbf{W} is a matrix that we define as $\mathbf{W} \in \mathbb{C}^{N \times N}$. Based on this definition, we can formulate \mathbf{G}_{RX} and \mathbf{G}_{TX} as follows:

$$\mathbf{G}_{\text{RX}} = \mathbf{D}_\delta \mathbf{W}^H \mathbf{H} \quad (13)$$

$$\mathbf{G}_{\text{TX}} = \mathbf{H} \mathbf{W} \mathbf{D}_\lambda, \quad (14)$$

where $[\mathbf{D}_\delta]_{k,k} = \delta_k$ and $[\mathbf{D}_\lambda]_{k,k} = \lambda_k$ are diagonal matrices. Next, we relate the above formulation to the singular-value decomposition (SVD) of the channel. To this end, we define $\mathbf{U}\mathbf{S}\mathbf{V}^H = \mathbf{H}$, $\mathbf{U}_{\text{RX}}\mathbf{S}_{\text{RX}}\mathbf{V}_{\text{RX}}^H = \mathbf{H}$, and $\mathbf{U}_{\text{TX}}\mathbf{S}_{\text{TX}}\mathbf{V}_{\text{TX}}^H = \mathbf{H}$ to be the SVD of \mathbf{H} , \mathbf{G}_{RX} and \mathbf{G}_{TX} respectively. Since \mathbf{D}_δ , \mathbf{D}_λ are unitary diagonal matrices and thus orthogonal, this means that $\mathbf{D}_\delta\mathbf{W}^H\mathbf{U}$ and $\mathbf{W}^H\mathbf{D}_\lambda^H\mathbf{V}$ are orthogonal if and only if \mathbf{W} is orthogonal. Unless \mathbf{H} contains two equal non-zero singular values, which is a zero-probability event, we have that its economic SVD decomposition is unique and translates into the following expressions:

$$\mathbf{S} = \mathbf{S}_{\text{RX}} = \mathbf{S}_{\text{TX}} \quad (15)$$

$$\mathbf{V}_{\text{RX}} = \mathbf{V} \quad \mathbf{U}_{\text{TX}} = \mathbf{U} \quad (16)$$

$$\mathbf{U}_{\text{RX}} = \mathbf{D}_\delta\mathbf{W}^H\mathbf{U} \quad \mathbf{V}_{\text{TX}} = \mathbf{W}^H\mathbf{D}_\lambda^H\mathbf{V} . \quad (17)$$

Thus, we can compute the CSI as $\mathbf{H} = \mathbf{U}_{\text{TX}}\sqrt{\mathbf{S}_{\text{RX}}\mathbf{S}_{\text{TX}}}\mathbf{V}_{\text{RX}}$, where the square root is computed element-wise. For the above to be valid, \mathbf{W} must be an orthogonal matrix such that each column $[\mathbf{W}]_{:,k}$ is a feasible analog configuration. To this end, we can use a Hadamard matrix with entries $\{-1, 1\}$ that has been normalized to satisfy power constraints.

4.3 SNR Maximizing Beam Patterns

Based on the measured CSI, ACO computes the beam pattern \mathbf{p} that maximizes the SNR. As discussed in Section 3.2, the received signal strength is $|\mathbf{h}_{\text{RX}}\mathbf{p}|^2$. Thus, the beam pattern \mathbf{p} that maximizes the expression is derived from the equality condition of the Schwartz inequality as $\mathbf{p} = \mathbf{h}_{\text{RX}}^H / \|\mathbf{h}_{\text{RX}}\|$. Due to hardware constraints, we round \mathbf{p} to match the device resolution. Since \mathbf{h}_{RX} is known, ACO can compute the optimal \mathbf{p} as explained before. Since ACO only requires an FFT and simple element-wise matrix operations, its additional computational complexity is minimal. Thus, ACO computes the optimal beam pattern in a highly efficient manner and can be executed even on highly constrained hardware components.

5 EXPERIMENT PLATFORM

We implement and evaluate our scheme on off-the-shelf IEEE 802.11ad devices that incorporate phased antenna arrays. We obtain control over the antenna array to adjust steering parameters, without requiring any hardware modifications. In the following, we describe the system architecture of our platform, the utilized antenna module, and how to control the antenna beam steering.

5.1 System Architecture

Only a few consumer devices can communicate in the 60 GHz band and support IEEE 802.11ad. One of them is the TP-Link Talon AD7200 tri-band router that uses a Qualcomm QCA9500 IEEE 802.11ad Wi-Fi chip. It features an antenna array with 32 elements that are individually controllable in

phase and amplitude. Control over the antenna is encapsulated in the firmware of the Wi-Fi chip, and only limited access is exposed to the host operating system and device driver. To obtain access to the antenna steering properties, as well as the IEEE 802.11ad beam training operation, we utilize the framework for practical IEEE 802.11ad research proposed in [33, 34]. This framework comes with a customized LEDE/OpenWrt image to provide open access to the operating system, network interfaces, and device drivers. Using the binary firmware patching framework (nexmon-arc) [33], additional features can be directly implemented in the firmware of the Wi-Fi chip.

5.2 Sector Level Sweep

The SLS is the standard beam training mechanism in IEEE 802.11ad and implemented in the firmware of the QCA9500 chip. It uses a set of predefined antenna configurations, the so-called sectors. These sectors are designed such that their beam patterns cover the entire azimuth to allow for an effective scanning of devices in proximity. We briefly discussed the operation of the SLS in Section 1. Through firmware patching, we can extract the SNR and RSSI of received sector sweep frames and report back custom feedback in the sector sweep feedback frames.

5.3 Antenna Module

To select antenna steering parameters, we need to obtain a thorough understanding of the antenna and its steering capabilities. Qualcomm's QCA9500 IEEE 802.11ad Wi-Fi chip consists of two modules, 1) a baseband IC that takes care of the signal and frame processing, and 2) an antenna module with an additional RF-IC that drives the antenna elements and controls the radiation characteristics. This modular design allows to flexibly place the antenna module at proper locations inside a device chassis to minimize radiation impairments. Both modules, the antenna and baseband IC, are connected with a coaxial cable for bi-directional transfer of modulated data, control, and clock signals, as well as power supply for the antenna module. In transmit mode, the antenna chip mixes up the modulated data signal from intermediate frequency (IF) to the desired RF channel in the 60 GHz band. In receive mode, the RF signal is mixed down to IF again. All antenna elements in the array are driven by an antenna weighting network, which is adjusted by the external control signal. Currently, only a single RF-chain is available. The antenna weighting network consists of an antenna switch, eight distribution amplifiers, 32 edge amplifiers and 32 phase shifters. All of these are controllable from within the firmware running on the baseband chip. To adjust the antenna steering, the following six 32-bit parameters are available:

- psh_hi: phase shift values for ant. chains [15-0]
- psh_lo: phase shift values for ant. chains [31-16]
- etype0: edge amp. bit 0 for all ant. chains [31-0]
- etype1: edge amp. bit 1 for all ant. chains [31-0]
- etype2: edge amp. bit 2 for all ant. chains [31-0]
- dtype_swch_off: dist. amp. values (3 bits each) + X16 switch bits

A discrete configuration of the antenna with these parameters refers to a so-called sector. The current firmware in version 4.1.0.55 supports 64 different transmit sectors out of which 35 are defined and used in beam training, as well as a single receive sector. The definitions for these sectors are stored in an antenna steering codebook.

The antenna steering codebook and individual sector configurations can be either changed directly in the memory of the firmware image (hardcoded) or dynamically adjusted during runtime. For the latter, the driver exposes specific netlink vendor commands. By changing the antenna parameters in the sector configurations, we can change the radiation patterns and gains—a crucial prerequisite to implement adaptive codebook optimization on off-the-shelf devices.

5.4 Reconstructing the Antenna Layout

Unfortunately, no public documentation for the antenna module or the baseband chip was available at the time of writing. Hence, we first had to analyze the internal structure of the antenna elements experimentally to create specific beam patterns. Each phase shifter in the weighting network is driven by two consecutive bits from either psh_hi or psh_lo. The edge amplifiers use a single bit from each of etype0, etype1, and etype2, while the distribution amplifiers consume three consecutive bits from dtype_swch_off. The most significant bits in dtype_swch_off are used to drive the antenna switch. We developed a parser to extract the parameters for each element which allows us to easily obtain and adjust the gain and phase of each antenna element.

By modifying individual parameters, we reconstructed the structure of the antenna weighting network experimentally. For example, by iterating over all pairwise combinations of distribution and edge amplifiers and setting all others to zero, we found that the first distribution amplifier drives the first four edge amplifiers, the second distribution amplifier drives the next four edge amplifiers, and so on. An antenna is only active if both the corresponding distribution and edge amplifier are set to non-zero. Similarly, we verified which phase shifter bits belong to which antenna chain by changing single values and monitoring the received signal strength at an unmodified device. The resulting antenna weighting structure as revealed in our experiments is shown in Figure 4.

By disassembling the antenna from the chip and shielding all except one antenna element, we reconstructed the

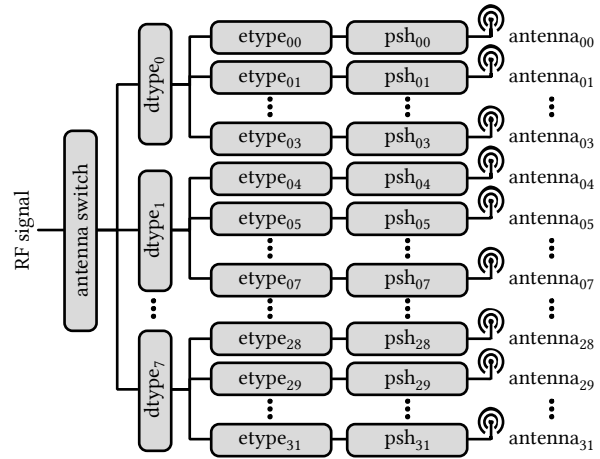


Figure 4: Experimentally reconstructed antenna weighting network for beam steering inside the QCA6310 antenna IC, which consists of an antenna switch, 8 distribution amplifiers (dtype), 32 edge amplifiers (etype) and 32 phase shifters (psh) to drive 32 antenna elements individually.




Figure 5: Phased array antenna chip disassembled from a Talon AD7200. Annotations show the antenna elements ID of the patch antennas. 14 additional dipole antennas are invisible from the surface.

3-dimensional layout of the antenna array as well as the physical element positions as shown in Figure 5. The antenna array consists of 12 patch antennas that are located on the front surface of the module in a 2 × 6 matrix shape. On the back side, there are only six patch antennas since the IC blocks the rest of the surface. The remaining 14 antennas are dipole antennas and directed towards the sides of the Printed Circuit Board (PCB). As a result, the three-dimensional layout, as well as its asymmetric assembly, lead to irregular beam patterns as already discussed in [21, 34].

This know-how of the antenna layout and the capability to control the weighting network is crucial to understand

how to optimize the beam patterns in the codebook. It is also a valuable resource for other researchers using this platform.

6 PRACTICAL IMPLEMENTATION

We implement ACO on the TP-Link Talon AD7200 router introduced in the previous section.¹ We first discuss in Section 6.1 the ACO implementation itself and then provide an overview on our experiment setup in Section 6.2.

6.1 Prototype Operation

Our prototype consists of an IEEE 802.11ad AP that serves one or multiple STAs. We use Talon devices to realize both AP and STAs, since the prototyping framework in [34] allows us to set them to either of those modes. For ease of experimentation, we control the prototype from a laptop which is external to the 60 GHz network. All of the devices that are part of the prototype are connected to a common 2.4 GHz Wi-Fi network that we use to control the experiments. Using this control network, the laptop instructs the individual STAs to connect to the AP, generates traffic in the IEEE 802.11ad network, and collects statistics. For simplicity, we also run certain components of ACO on the laptop, such as the reconfiguration of the codebooks via the interfaces described in Section 5.3. However, these portions of the code could as well run directly on the Talon devices since ACO has very low computational requirements.

The operation of our prototype is as follows. First, the control laptop loads the default codebooks of the Talon devices on both the AP and the STAs. Then, it triggers the STAs to establish a connection to the AP using the regular IEEE 802.11ad procedure. After that, ACO loads an initial measurement codebook on the devices and collects the resulting SNR readings to compute the transmit beam pattern that maximizes the SNR for each link. We carry out the process for AP and STAs, since we analyze both up-link and down-link traffic in Section 6.2 (and to ensure that acknowledgments are also sent with the optimized beam patterns). To specifically optimize the beam patterns under low SNR conditions, we employ the directional estimation scheme described in Section 4.1.2. Finally, the control laptop re-configures all devices with the computed communication beam patterns and generates traffic in the network. In particular, we use `iperf3` to generate ten seconds of down-link traffic and ten seconds of up-link traffic on each link. For each experiment, we repeat this measurement twenty times while we collect statistics such as TCP throughput, physical layer rate, and signal quality.

¹The ACO implementation is available at: <http://wireless.networks.imdea.org/downloads-software-and-traces>




Figure 6: Experiment setup in indoor office scenario.

6.2 Experiment Setup

We evaluate ACO in an open-plan office environment (see Figure 6). The area has a size of 18×7.5 meters and is surrounded by twelve individual offices. We deploy one AP below the ceiling next to the desks in the open-plan area. Further, we place STAs at 24 different positions within the room. Out of the 24 locations, 12 correspond to the desks. For these cases, we orient the STAs such that they mimic the usual placement of a user's laptop. Next, we select six further positions in an area with no desks but still in line of sight (LOS) of the AP. This enables us to study whether the desks and the separators among them play a significant role. Finally, we select six further STA locations behind an isolated wall that separates the entry area from the open-plan office space. Those locations are non line of sight (NLOS), which means that STAs must use reflections to reach the AP.

7 EVALUATION

We evaluate the performance of ACO in the real-world office scenario described in Section 6.2. Since we focus on the performance of individual links, we only activate one STA at a time. We obtain the achieved SNR, the TCP throughput, IQ constellations, and the expected bitrate.

7.1 SNR Maximization

As discussed in Section 4.3, we design ACO to maximize the SNR. In the first part of our experiments, we evaluate the SNR gains that ACO achieves in comparison to IEEE 802.11ad operation with the default antenna patterns. For both our optimized beams and the default ones, we measure the SNR at all locations of our testbed for the up- and down-link. For ease of comparison, we report the SNR on a linear scale throughout this section. As shown in Figure 7, we achieve a SNR of 8.16 for the up-link and 7.53 for the down-link when using the generic beam patterns. With ACO, we increase the average linear SNR to 22.06 for the up-link and 19.78 for the down-link, respectively. The cumulative distribution function (CDF) of the default beam pattern reveals that 90% of the measurements have an SNR that is lower than 18. With

ACOs optimized beams, we achieve an SNR higher than this in about 45% of the measurements. The floor-maps in Figure 8 illustrate the SNR gains that we achieve at the individual measurement points in our environment. We find that the distance between AP and STA does not influence ACOs performance. However, the variance among different locations is high. In some spots, ACO achieves gains of only 50% while in others the gains are over 500%. Some locations can be easily reached via reflectors while others cannot. Overall, we achieve an average gain of 168.61% on the up-link and 159.34% on the down-link.

To validate the high SNR gains, we additionally analyze the signals at the physical layer. Unfortunately, evaluating the accuracy of the CSI and how it degrades with SNR is impossible on the COTS devices themselves. This leaves us to compare the achievable signal constellations with those of the default codebook. To this end, we place a Sivers IMA FC2221V/01 V-band down-converter with a horn antenna at the location of the receiver and capture the raw samples using a Keysight DSOS254A oscilloscope. The bandwidth of the oscilloscope is sufficient to capture the full IEEE 802.11ad signal, which we decode using the Keysight Wideband Waveform Center. This enables us to obtain the IQ constellations of data frames transmitted both with the generic beam patterns and with ACO beam patterns. We show an example in Figure 9. Both frames in that example were transmitted with MCS8, which uses a $\pi/2$ -QPSK constellation. The ACO beam patterns show a visible accuracy improvement and lead to constellations that are significantly less noisy than those of the default beam pattern. This results in fewer symbol errors, which in turn improves the stability of the link.

7.2 Throughput Improvement

The large SNR improvement in Section 7.1 enables our Talon devices to switch to a higher MCS. As a result, we obtain significant TCP throughput gains when generating iPerf3 traffic in the network. Figure 10 depicts the CDF of the achieved TCP throughput at a LOS location in our testbed. We achieve gains for both up-link and down-link. For the former, the average gain of ACO is 58% while for the latter it is 102%. That is, we *double* the throughput in the down-link. In Figure 10, we also show the distribution of the MCS for that location. As expected, the Talon devices are able to choose high modulations much more often when using our ACO beam patterns compared to the generic ones. For instance, ACO enables MCS8 on the down-link more than 60% of the time whereas generic beam patterns only reach such a high MCS 18% of the time.

In Figure 11 we perform a similar analysis for one of the NLOS locations behind the isolated wall in our testbed (see




Figure 7: Average SNR and CDF of our ACO mechanism and the default IEEE 802.11ad operation for up-link and down-link measurements.

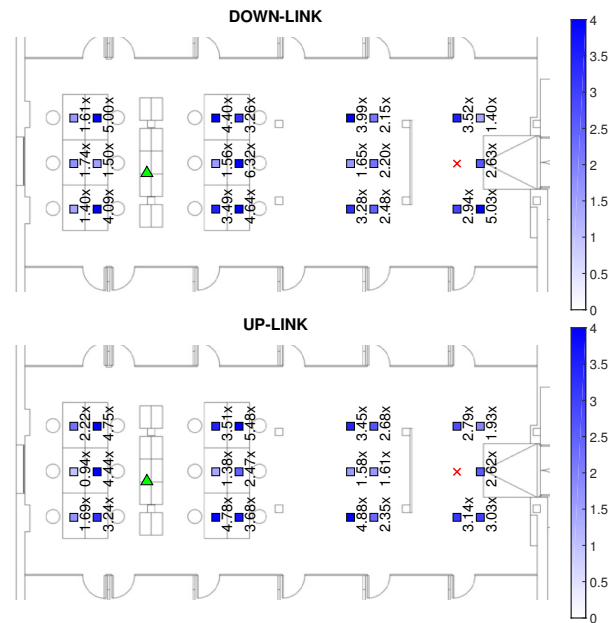


Figure 8: Map of the average SNR gains of optimized beams. Each square represents a STA location, crosses indicate blind spots, while the AP is indicated by the green triangle.




Figure 9: Constellation diagrams of MCS8 encoded frames transmitted with generic (left) and optimized beam patterns (right).

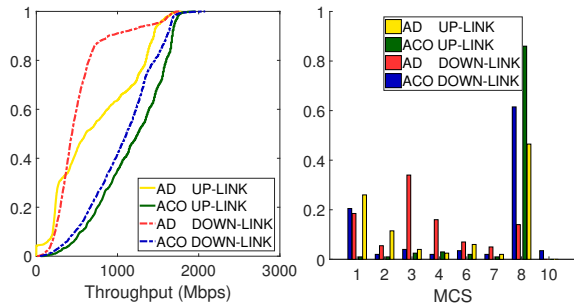


Figure 10: LOS throughput CDF and transmission MCS histogram.

Section 6). Due to the wall, the AP and the STA communicate via a reflection. Our results show that the generic beamforming mechanism in IEEE 802.11ad performs poorly in such a scenario, achieving on average only 42.74 Mbps in the up-link and 314.54 Mbps in the down-link. In contrast, ACO achieves on average about 855 Mbps on the same link in both directions. The MCS distribution in Figure 11 reflects this behavior, showing that IEEE 802.11ad operates mostly in the range of MCS1 to MCS3, whereas ACO enables rates beyond MCS4 in most cases. We conclude that ACO is highly beneficial for IEEE 802.11ad wireless networks and can achieve order-of-magnitude throughput gains.

7.3 Expected Bitrate

Although we consistently achieve very high SNR gains at all of the locations in our testbed (see Section 7.1), we only observe equivalent throughput gains at a subset of those locations. To understand this effect, we delved into the operation of the Talon AD7200 router. We found that in many cases the router does not switch to a higher modulation even though the SNR improves by a factor of three or more. In particular, we observed that the rate adaptation algorithm very often does not exceed MCS8 even though the router implements up to MCS12. As a result, TCP throughput is limited to about 1.6 Gbps. Figure 12 shows that the throughput saturates at this value for all of the locations in our testbed, irrespective of the length of the link and other factors that could influence signal quality. In exhaustive measurements, we found that the rate adaptation mechanism of the router is configured such that MCS configurations beyond MCS8 are disfavored. This is a coherent design decision for the Talon hardware that comes with certain limitations due to its inexpensive production. For example, the Gigabit Ethernet interface limits traffic to one Gigabit per second. The 60 GHz interface is not the bottleneck and already saturates the rest of the system even with an MCS below the maximum. For this reason, it seems that the manufacturer selected a conservative rate adaptation scheme that does not switch

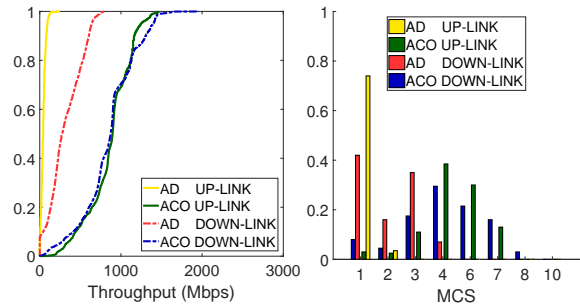


Figure 11: NLOS throughput CDF and transmission MCS histogram.




Figure 12: Measured throughput as a function of the measured SNR.

to high MCS values even when the SNR of the link is very high. We expect next-generation IEEE 802.11ad/ay devices to provide better support for such multi-Gigabit per second modulations in terms of both handling high volumes of traffic and generating cleaner radio-frequency signals that enable higher modulations such as 16-QAM. Hence, we conclude that the throughput saturation that we observe in Figure 12 is due to limitations of the hardware that prevent stable operation of the router at such high MCSs.

As a result of the above hardware limitations, the benefits of ACO only become visible in our testbed when a link using the generic beam patterns of IEEE 802.11ad operates below MCS8. Otherwise, the SNR improvement of ACO does not translate into an effective throughput gain. To illustrate the benefits that ACO would achieve on hardware that can process higher data throughput, in the following we compute the bitrate according to the expected MCS for a given SNR. To this end, we build on the signal quality thresholds recommended in the IEEE 802.11ad standard [15] as listed in Table 2. We transform the measured SNRs to sensitivity thresholds, lookup the recommended MCS, and, thus, obtain the expected bitrates and cross-validate them with our measurements.

In Figure 13, we show the average gain in terms of the expected bitrate. On both the up-link and the down-link ACO

Table 2: Recommended MCS selection, required receive sensitivity, and bitrate in IEEE 802.11ad [15].

MCS	Sensitivity	Bitrate
0	-78 dBm	27.5 Mbps
1	-68 dBm	385.0 Mbps
2	-66 dBm	770.0 Mbps
3	-65 dBm	962.5 Mbps
4	-64 dBm	1155.0 Mbps
5	-62 dBm	1251.3 Mbps
6	-63 dBm	1540.0 Mbps
7	-62 dBm	1925.0 Mbps
8	-61 dBm	2310.0 Mbps
9	-59 dBm	2502.5 Mbps
10	-55 dBm	3080.0 Mbps
11	-54 dBm	3850.0 Mbps
12	-53 dBm	4620.0 Mbps

obtains on average about 50% higher bitrate, but Figure 14 shows that gains reach up to 2.54x. We achieve significant gains across most of the testbed. In a few cases, the achieved bitrate gain in Figure 14 is limited compared to the large SNR gain for the same location in Figure 8. This occurs in particularly challenging locations, such as the desks which are very close to the AP. Since the AP hangs from the ceiling, the vertical angle of the link towards the STAs located at those desks is very steep. Due to the layout of the antenna array (see Section 5.3), the vertical steering capability of the devices is limited. While those links are inherently weak, ACO still provides a substantial SNR improvement even in such challenging scenarios and enables the system to reach a higher MCS in most cases.

7.4 Channel Probing Overhead

ACO requires additional SNR feedback. Instead of feeding back only the ID of the best antenna beam pattern as with conventional beam training, the SNRs of all the beam patterns that were probed are required. Furthermore, the overhead to acquire full channel state information would be higher than that of conventional IEEE 802.11ad training. However, measuring the full CSI for all antenna elements is usually not necessary. Our approach can flexibly adjust the number of active antennas or probing elements based on how rapidly the sub-channel changes, allowing us to reduce the overhead to achieve a certain accuracy. According to [34], probing a single sector in the sector sweep takes about 18.0 μ s. Probing all $N = 32$ antenna elements with ACO would need $N + 4(N - 1) = 156$ probes (see Section 4.1.1), which results in a total probing time of about 2.8 ms. With 13 active antenna elements, we perform the beam training as fast as the IEEE 802.11ad sector sweep with 64 sectors in about 1.2 ms. As not all antenna elements contribute to good beam patterns, learning a sub-space of the channel with less




Figure 13: Average expected bitrate and CDF of our ACO mechanism and the default IEEE 802.11ad operation for up- and down-link measurements.

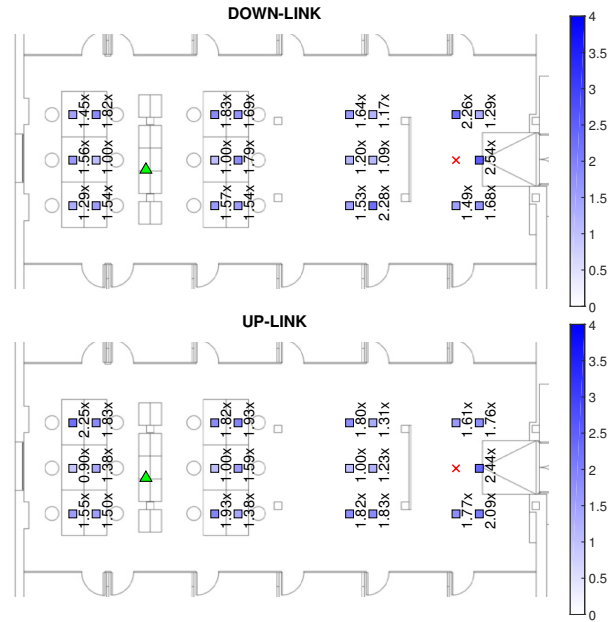


Figure 14: Map of the average expected bitrate gains of optimized beams. Each square represents a STA location, crosses indicate blind spots, while the AP is indicated by the green triangle.

active antenna elements is usually sufficient and provides a good trade-off between overhead and performance.

This flexible overhead also allows ACO to adapt to mobility which requires to train the beams continuously. By adapting the number of active antennas in the array, ACO can select a training overhead that matches the dynamics of the devices and environment. Specifically, ACO only needs to probe a small sub-space of the channel in case the movement is small.

8 RELATED WORK

The IEEE 802.11ad standard [15, 22] deploys a two-stage beamforming protocol that consists of the SLS and an optional beam refinement phase (BRP). The SLS operates with predefined beam patterns, so-called sectors, and trains the antenna steering by sweeping through all available sectors while sending probe frames. In the BRP, devices continuously adapt their antenna sector to compensate for mobility and small channel distortions. In IEEE 802.11ay—the upcoming millimeter-wave standard for Wi-Fi networks and the successor of IEEE 802.11ad—several refinements are introduced to make beam training more efficient with shorter frames and multiple channels [6, 9]. For a low number of sectors, the sector sweep completes in a reasonable amount of time. However, its complexity increases linearly with the number of probed sectors. Thus, either the number sector is low and thus beam patterns become wide and irregular, or a high training overhead must be taken into account. Common COTS devices typically use a fixed set of sectors and do not implement the BRP [34].

Several alternatives have been proposed that allow devices to steer their antennas efficiently. State-of-the-art protocols and techniques are surveyed in [17]. Hierarchical beam training with codebook structures of different width allows to refine the beam patterns iteratively [1, 2, 14, 25]. Organizing beam patterns in a tree hierarchy achieves logarithmic search complexity but introduces an additional communication overhead due to the required feedback for each of the multiple probing rounds. Compressive sensing based approaches derive the CSI by selecting pseudo-random probing beams [20, 27, 28] and determine an optimal beam after only one training round. However, complex channel measurements are not available on common COTS devices. Thus, [29] proposed a variant based on non-coherent measurements to determine the line of sight path without phase information. [34] integrated such a variant of compressive sensing with the sector sweep on COTS devices. Most beam-steering approaches use codebooks that consists of beams with specific geometrical contexts such as beam-width or beam-direction. To generate such beams, it is required to understand the antenna radiation characteristics, especially since the wavelength of mm-waves is of the order of millimeters. Small variations in the electrical components create significant phase shifts or irregular absorption. To tackle these impairments, antenna calibration under controlled conditions is required and must be incorporated in the codebook. Our approach, in contrast, does not rely on antenna calibration. We optimize the beam-steering based on current channel conditions without any geometrical properties.

In multi-antenna scenarios with multiple RF-chains, additional performance gains are expected. Compressive sensing

allows to estimate the channel with only a few probes efficiently [7, 8] and selects proper beams for multiple users [4]. Exploiting multiple antennas also improves spatial reuse [31, 32]. With multiple RF-chains, hybrid beamforming allows to take channel measurements in parallel and decreases the overhead or increases the channel estimation quality. Unfortunately, at the time of writing, no COTS device supports multiple RF-chains or hybrid beamforming.

Algorithms such as [36] improve resilience to blockage by instantaneously predicting the availability of alternative paths. The mobility of devices is considered in [42], in which correlations of channel profiles at nearby locations allow to reconstruct the channel profile when devices start to move. Doing so, the proposed protocol continuously realigns the links of mobile devices without explicit channel scanning. In [40], E-MI senses the environment and traces back the LOS and NLOS paths between devices by applying ray-tracing techniques. This allows to resolve the significant reflections in the paths, predict the quality of different links, and optimize the deployment of APs. Moreover, adding low pilot overhead, as done in [3], increases the channel estimation accuracy. Beam steering algorithms also benefit from multi-band connectivity. Wi-Fi connections in the lower frequency bands such as 2.4 and 5 GHz can assist millimeter-wave link adaptation [35]. Angular estimation of the communication direction in the lower frequency bands can already provide a coarse direction to steer the millimeter-wave beam [23]. Tracking of devices without probing overhead is possible by explicitly generating multi-lobe beam patterns that detect if devices move out of the beam's coverage area [18]. To lower the risk of outage during mobility, the authors of [10] propose to continuously adjust the width of the beam pattern as well as the selected modulation and coding schemes.

While millimeter-wave research is well advanced, widespread practical testbed and evaluation systems are rare. Available IEEE 802.11ad chipsets expose a single RF-chain, provide limited control, and only allow to extract a few low-layer parameters [19, 21]. As a result, researchers tend to create evaluation platforms that are equipped with directional horn antennas [10, 11, 36, 38] or custom phased antenna arrays [29, 41]. Due to the high amount of hardware customization, such platforms complicate the reproducibility of results. They typically lack IEEE 802.11ad compliant implementations, and their performance and behavior differ fundamentally from that of COTS devices. Commercial prototyping platforms such as that in [30] allow to evaluate different protocols but are not affordable for larger deployments and are impractical to analyze realistic millimeter-wave use-cases. In prior work, we extended the device driver to control the Wi-Fi interface in COTS devices [35, 39] and added new features to the firmware running on the IEEE 802.11ad Wi-Fi chip [34] to extend the access to low-layer

statistics and perform practical experiments. However, all of these works based on COTS devices lack access to the full CSI of the millimeter-wave channel.

Accessing CSI on Wi-Fi systems in the lower frequency bands is comparatively easier. These systems apply digital beamforming which allows for full channel estimation. Channel measurements are, for example, required to compensate the channel effects and apply beamforming in multi-antenna systems. Prototyping platforms such as WARP [16] allow estimating the complex channel gains with reasonable overhead in multi-antenna systems. The CSI can also be directly extracted from some specific COTS wireless interface cards [12]. Unfortunately, such approaches require community projects and a detailed understanding of the wireless chip. Notably, even for Wi-Fi that has been around for two decades, only a few chipsets to date provide CSI that is accessible to researchers. Manufacturers of commodity Wi-Fi equipment typically keep their systems closed. As of now, no such access to CSI is available for any IEEE 802.11ad device.

9 CONCLUSIONS

Obtaining full CSI on millimeter-wave devices enables beamforming to achieve higher directionality, allow for more accurate beam steering, and exploit multi-path effects caused by reflectors and obstacles in the environment. Unfortunately, this requires complex channel measurements which are not provided by commercial off-the-shelf devices. In this work, we propose Adaptive Codebook Optimization (ACO), a mechanism that extracts full CSI on such devices from non-coherent measurements. In particular, we engineer beam patterns with constant phase shifts that allow extracting CSI both in terms of phase and magnitude from simple SNR readings. As such readings are available on most devices, our approach is easily portable to any hardware that provides access to SNR information. We design ACO to incur the same or less probing overhead than IEEE 802.11ad. Moreover, we use the obtained CSI to compute beam patterns that adapt to the environment and maximize the SNR. In this way, we automatically avoid destructive interference and choose the best available path between transmitter and receiver. We implement ACO on TP-Link Talon AD 7200 tri-band routers by obtaining full access to the beamforming control of the integrated 32-element phased antenna array. To this end, we disassembled the phased antenna array and reconstructed the antenna weighting network experimentally. Our evaluations in a real-world office environment show that ACO increases the SNR by factor 2.5 and achieves a 2x higher TCP throughput. To support the community and allow other researchers to benefit from our results, we make our framework and firmware patches [33] as well as the source code of our actual implementation publicly available.

ACKNOWLEDGMENTS

This work has been supported in part by the ERC project SEARCHLIGHT grant no. 617721, the Ramon y Cajal grant RYC-2012-10788, the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919), the German Federal Ministry of Education and Research (BMBF), the State of Hesse within CRISP-DA, the Hessian LOEWE excellence initiative within NICER, and the German Research Foundation (DFG) within the Collaborative Research Center (CRC) 1053 “MAKI – Multi-Mechanism-Adaptation for the Future Internet”.

REFERENCES

- [1] Ahmed Alkhateeb, Omar El Ayach, and Geert Leus. 2014. Channel Estimation and Hybrid Precoding for Millimeter Wave Cellular Systems. *IEEE Journal of Selected Topics in Signal Processing* (2014).
- [2] Ahmed Alkhateeb, Omar El Ayach, Geert Leus, and Robert W. Heath. 2014. Single-Sided Adaptive Estimation of Multi-Path Millimeter Wave Channels. In *International Workshop on Signal Processing Advances in Wireless Communications (SPAWC) 2014*. IEEE, 125–129.
- [3] Daniel C. Araújo, André L. F. de Almeida, Johan Axnäs, and Joao C. M. Mota. 2014. Channel Estimation for Millimeter-Wave Very-Large MIMO Systems. *IEEE Journal of Selected Topics in Signal Processing* (2014), 81–85.
- [4] Jinho Choi. 2015. Beam Selection in mm-Wave Multiuser MIMO Systems Using Compressive Sensing. *IEEE Transactions on Communications* 63, 8 (2015), 2936–2947.
- [5] Aditya Dhananjay. 2015. *Iris: Mitigating Phase Noise in Millimeter Wave OFDM Systems*. Ph.D. Dissertation. New York University.
- [6] Alecsander Eitan and Carlos Cordeiro. 2016. *Short SSW Format for 11ay (IEEE 802.11-16/0416-01-00)*.
- [7] Zhen Gao, Linglong Dai, and Zhaocheng Wang. 2016. Channel Estimation for mmWave Massive MIMO Based Access and Backhaul in Ultra-Dense Network. In *International Conference on Communications (ICC) 2016*. IEEE, 1–6.
- [8] Zhen Gao, Chen Hu, Linglong Dai, and Zhaocheng Wang. 2016. Channel Estimation for Millimeter-Wave Massive MIMO With Hybrid Precoding Over Frequency-Selective Fading Channels. *IEEE Communications Letters* 20, 6 (2016), 1259–1262.
- [9] Yasaman Ghasempour, Claudio R. C. M. da Silva, Carlos Cordeiro, and Edward W. Knightly. 2017. IEEE 802.11ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi. *IEEE Vehicular Technology Conference (VTC) Fall 2017* 55, 12 (Dec. 2017), 186–192.
- [10] Muhammad Kumail Haider and Edward W. Knightly. 2016. Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc) 2016*. ACM Press, New York, New York, USA, 61–70.
- [11] Muhammad Kumail Haider and Edward W. Knightly. 2018. iTrack: Tracking Indicator LEDs on APs to Bootstrap mmWave Beam Acquisition and Steering. *International Workshop on Mobile Computing Systems Applications (HotMobile) 2018* (2018), 107–112.
- [12] Daniel Halperin, Wenjun Hu, Anmol Sheth, and David Wetherall. 2011. Tool Release: Gathering 802.11n Traces with Channel State Information. *ACM SIGCOMM Computer Communication Review* 41, 1 (Jan. 2011), 53–53.
- [13] Wonbin Hong, Kwang-Hyun Baek, Youngju Lee, Yoongeon Kim, and Seung-Tae Ko. 2014. Study and Prototyping of Practically Large-scale

- mmWave Antenna Systems for 5G Cellular Devices. *IEEE Communications Magazine* 52, 9 (2014).
- [14] Sooyoung Hur, Taejoon Kim, David J. Love, James V. Krogmeier, and Timothy A. Thomas. 2013. Millimeter Wave Beamforming for Wireless Backhaul and Access in Small Cell Networks. *IEEE Transactions on Communications* 61, 10 (Oct. 2013).
- [15] IEEE Standards Association. 2014. IEEE Std 802.11ad-2012: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band. ISO/IEC/IEEE 8802-11:2012/Amd.3:2014(E). (2014).
- [16] Mango Communications Inc. 2018. WARP Project. (2018). <http://warpproject.org>
- [17] Shajahan Kutty and Debarati Sen. 2016. Beamforming for Millimeter Wave Communications: An Inclusive Survey. *IEEE Communications Surveys & Tutorials* 18, 2 (2016), 949–973.
- [18] Adrian Loch, Hany Assasa, Joan Palacios, Joerg Widmer, Hans Suys, and Björn Debaillie. 2017. Zero Overhead Device Tracking in 60 GHz Wireless Networks using Multi-Lobe Beam Patterns. In *International Conference on emerging Networking EXperiments and Technologies (CoNEXT) 2017*. ACM, Incheon, Republic of Korea, 224–237.
- [19] Adrian Loch, Guillermo Bielsa, and Joerg Widmer. 2016. Practical Lower Layer 60 GHz Measurements Using Commercial Off-The-Shelf Hardware. *International Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WiNTECH)* (2016), 9–16.
- [20] Zhinus Marzi, Dinesh Ramasamy, and Upamanyu Madhow. 2016. Compressive Channel Estimation and Tracking for Large Arrays in Mm-Wave Picocells. *IEEE Journal of Selected Topics in Signal Processing* 10, 3 (April 2016).
- [21] Thomas Nitsche, Guillermo Bielsa, Irene Tejado, Adrian Loch, and Joerg Widmer. 2015. Boon and Bane of 60 GHz Networks: Practical Insights Into Beamforming, Interference, and Frame Level Operation. In *International Conference on emerging Networking EXperiments and Technologies (CoNEXT) 2015*.
- [22] Thomas Nitsche, Carlos Cordeiro, Adriana B. Flores, Edward W. Knightly, Eldad Perahia, and Joerg Widmer. 2014. IEEE 802.11ad: Directional 60 GHz Communication for Multi-Gigabit-Per-Second Wi-Fi. *IEEE Communications Magazine* 52, 12 (Dec. 2014), 132–141.
- [23] Thomas Nitsche, Adriana B. Flores, and Edward W. Knightly. 2015. Steering with Eyes Closed: Mm-Wave Beam Steering Without in-Band Measurement. *IEEE Conference on Computer Communications (INFOCOM) 2015* (2015), 2416–2424.
- [24] Yong Niu, Yong Li, Depeng Jin, Li Su, and Athanasios V. Vasilakos. 2015. A Survey of Millimeter Wave Communications (mmWave) for 5G: Opportunities and Challenges. *Wireless Networks* 21, 8 (2015).
- [25] Song Noh, Michael D. Zoltowski, and David J. Love. 2017. Multi-Resolution Codebook and Adaptive Beamforming Sequence Design for Millimeter Wave Beam Alignment. *IEEE Transactions on Wireless Communications* 16, 9 (2017), 5689–5701. <https://doi.org/10.1109/TWC.2017.2713357>
- [26] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. 2018. *Addendum to Adaptive Codebook Optimization for Beam Training on Off-The-Shelf IEEE 802.11ad Devices*. Technical Report TR-IMDEA-Networks-2018-1. IMDEA Networks.
- [27] Dinesh Ramasamy, Sriram Venkateswaran, and Upamanyu Madhow. 2012. Compressive Adaptation of Large Steerable Arrays. In *Information Theory and Applications Workshop (ITA) 2012*. IEEE, 234–239.
- [28] Dinesh Ramasamy, Sriram Venkateswaran, and Upamanyu Madhow. 2012. Compressive Tracking with 1000-Element Arrays: a Framework for Multi-Gbps Mm Wave Cellular Downlinks. In *Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 690–697.
- [29] Maryam Eslami Rasekh, Zhinus Marzi, Yanzi Zhu, Upamanyu Madhow, and Haitao Zheng. 2017. Noncoherent mmWave Path Tracking. In *International Workshop on Mobile Computing Systems and Applications (HotMobile) 2017*. ACM, New York, USA, 13–18.
- [30] Swetank Kumar Saha, Yasaman Ghasempour, Muhammad Kumail Haider, Tariq Siddiqui, Paulo De Melo, Neerad Somanchi, Luke Zakrajsek, Arjun Singh, Owen Torres, Daniel Uvaydov, Josep Miquel Jornet, Edward W. Knightly, Dimitrios Koutsonikolas, Dimitris Pados, and Zhi Sun. 2017. X60: A Programmable Testbed for Wideband 60 GHz WLANs with Phased Arrays. In *Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WinTech) 2017*.
- [31] Jaspreet Singh and Sudhir Ramakrishna. 2014. On the Feasibility of Beamforming in Millimeter Wave Communication Systems with Multiple Antenna Arrays. In *IEEE Global Communications Conference (GLOBECOM) 2014*. IEEE, 3802–3808.
- [32] Jaspreet Singh and Sudhir Ramakrishna. 2015. On the Feasibility of Codebook-Based Beamforming in Millimeter Wave Systems With Multiple Antenna Arrays. *IEEE Transactions on Wireless Communications* 14, 5 (2015), 2670–2683.
- [33] Daniel Steinmetzer, Daniel Wegemer, and Matthias Hollick. 2017. Talon Tools: The Framework for Practical IEEE 802.11ad Research. (2017). <https://seemoo.de/talon-tools/>
- [34] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. 2017. Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices. In *International Conference on emerging Networking EXperiments and Technologies (CoNEXT) 2017*. ACM Press, Incheon, Republic of Korea, 414–425.
- [35] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2017. WiFi-Assisted 60 GHz Wireless Networks. In *International Conference on Mobile Computing and Networking (MobiCom) 2017*.
- [36] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*.
- [37] Timothy A. Thomas, Mark Cudak, and Tom Kovarik. 2015. Blind phase noise mitigation for a 72 GHz millimeter wave system. In *2015 IEEE International Conference on Communications (ICC)*.
- [38] Vignesh Venkateswaran, Xinyu Zhang, Sanjib Sur, Vignesh Venkateswaran, and Parmesh Ramanathan. 2015. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *ACM SIGMETRICS Performance Evaluation Review*. ACM, New York, New York, USA, 71–84.
- [39] Teng Wei and Xinyu Zhang. 2017. Pose Information Assisted 60 GHz Networks: Towards Seamless Coverage and Mobility Support. In *ACM International Conference on Mobile Computing and Networking (MobiCom'17)*.
- [40] Teng Wei, Anfu Zhou, and Xinyu Zhang. 2017. Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI) 2017*.
- [41] Jialiang Zhang, Xinyu Zhang, Pushkar Kulkarni, and Parameswaran Ramanathan. 2016. OpenMili: A 60 GHz Software Radio with a Programmable Phased-array Antenna: Demo. In *International Conference on Mobile Computing and Networking (MobiCom)*.
- [42] Anfu Zhou, Xinyu Zhang, and Huadong Ma. 2017. Beam-Forecast: Facilitating Mobile 60 GHz Networks via Model-Driven Beam Steering. In *IEEE Conference on Computer Communications (INFOCOM) 2017*. IEEE, 1–9.