Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices

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

Daniel Wegemer
Secure Mobile Networking Lab
TU Darmstadt, Germany
dwegemer@seemoo.de

Matthias Schulz
Secure Mobile Networking Lab
TU Darmstadt, Germany
mschulz@seemoo.de

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

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

ABSTRACT

Achieving data-rates of multiple Gbps in 60 GHz millimeter-wave (mm-wave) communication systems requires efficient beam-steering algorithms. To find the optimal steering direction on IEEE 802.11ad compatible devices, state-of-the-art approaches sweep through all predefined antenna sectors. Recently, much more efficient alternatives, such as compressive path tracking, have been proposed, which scale well even with arrays with thousands of antenna elements. However, such have not yet been integrated into consumer devices. In this work, we adapt compressive path tracking for sector selection in off-the-shelf IEEE 802.11ad devices. In contrast to existing solutions, our compressive sector selection tolerates the imperfections of low-cost hardware, tracks beam directions in 3D and does not rely on pseudo-random beams. We implement our protocol on a commodity router, the TP-Link Talon AD7200, by modifying the sector sweep algorithm in the IEEE 802.11ad chip's firmware. In particular, we modify the firmware to obtain the signal strength of received frames and to select custom sectors. Using this extension, we precisely measure the device's sector patterns. We then select the best sector based on the measured patterns and sweep only through a subset of probing sectors. Our results demonstrate, that our protocol outperforms the existing sector sweep, increases stability, and speeds up the sector selection by factor 2.3.

ACM Reference format:
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 *Proceedings of CoNEXT '17: The 13th International Conference on emerging Networking EXperiments and Technologies, Incheon, Republic of Korea, December 12–15, 2017 (CoNEXT '17),* 12 pages.
DOI: 10.1145/3143361.3143384

CoNEXT '17, Incheon, Republic of Korea

978-1-4503-5422-6/17/12…$15.00
DOI: 10.1145/3143361.3143384

1 INTRODUCTION

Due to high path loss in the 60 GHz millimeter-wave (mm-wave) band, IEEE 802.11ad systems use directional antennas to focus their energy in the direction of the receiving node. To change the radiation patterns of these antennas and to enable beamsteering during runtime, manufacturers employ antenna arrays that offer a programmable interface to change gains and phases in discrete steps per antenna element. Searching all beam combinations results in a huge search space. Hence, practical IEEE 802.11ad systems use a set of predefined beam patterns, so-called sectors. A sector sweep (SSW) algorithm then selects the optimal sector to connect to another node. It performs an extensive search which is very time consuming as all sectors require probing. To improve performance, protocols may establish links with less overhead [35] or proactively switch to alternative beam alignments [31]. For instance, a hierarchical beam-search [15] probes sectors with wide beams first and continues with subsectors featuring narrow beams until it finds the best one. Compressive sensing techniques [18, 23–25], instead, do not evaluate all available antenna settings. Due to spatial similarities in different antenna sectors, these algorithms only probe a subset and determine the best selection by correlating the probes.

Recent publications on mm-wave beam training either limit themselves to theoretical approaches [18, 23, 24], or use simplified measurements to verify their solutions. Due to the lack of widespread mm-wave consumer hardware, many researchers base their work on prototyping platforms with directional horn antennas [13, 31, 34] or custom antenna arrays [25] that both exhibit a different behavior than commercial off-the-shelf devices. The low-cost components integrated in the latter cause imperfections and do not achieve the precision of laboratory equipment.

In this paper, we develop a new compressive sector selection protocol that extends existing compressive path tracking solutions with the 3D sector selection required for planar phased arrays and the imperfections of low-cost hardware elements. Instead of using random beams and theoretical beam patterns based on geometrical antenna layouts, we use the already well performing beam patterns defined as sectors in the IEEE 802.11ad chip's firmware. To benefit from their characteristics in our algorithm, we first measured them in three dimensions which allows us to optimally tune the pattern selection to our hardware and directly integrate it with the sector sweep in IEEE 802.11ad off-the-shelf devices.

Unfortunately, off-the-shelf mm-wave devices lack suitable interfaces to control the beamtraining. In low-cost hardware, the firmware of the wireless interface chips is a black box [19] and hinders researchers from directly manipulating the PHY and MAC layer parameters. Low-layer access is possible only within the limitations of the firmware interface [17]. However, to implement our compressive sector selection protocol in such devices we need to modify the sector sweep. To this end, we turn the Talon AD7200 router into a flexible research platform and provide access to both the router's firmware and its IEEE 802.11ad chip's firmware. We integrate custom firmware extensions to access the signal strength measurements of received sector sweep frames and to overwrite the sector selection in sector sweep feedback frames. With this framework, we measure three-dimensional patterns for all sectors predefined in the device and implement and evaluate our compressive sector selection protocol to minimize the number of probed sectors in off-the-shelf IEEE 802.11ad devices.

In the following, we summarize our contributions:

- We design, implement and evaluate a compressive sector selection protocol that extends existing solutions, integrates in the sector sweep and operates on off-the-shelf IEEE 802.11ad devices.
- We build a research platform based on the Talon AD7200 with open access to the system and the firmware of the IEEE 802.11ad chip that allows us to obtain signal strength measurement from received sector sweep frames and to select custom sectors in the sector sweep feedback.
- We obtain 3D radiation patterns for each sector predefined in the Talon AD7200 from an extensive measurement campaign in an anechoic chamber.

The remainder of this paper is structured as follows: In Section 2, we present our compressive path tracking algorithm. Section 3 discusses the necessary changes to the router's firmware and operating system. In Section 4, we analyze predefined sectors and provide their radiation patterns. Section 5 describes the effects of measurement deviations. In Section 6, we evaluate the performance of our protocol and discuss results in Section 7. Finally, we state related work in Section 8 and conclude this paper in Section 9.

2 COMPRESSIVE SECTOR SELECTION

In this section, we adopt a compressive path tracking algorithm for optimized sector selection on commodity devices. It integrates with the the sector sweep algorithm of IEEE 802.11ad devices.

2.1 Existing Solutions

To determine the best sectors for the communication between two devices, the IEEE 802.11ad standard [20, 36] employs the sector sweep algorithm. It works with predefined beam patterns, the antenna sectors. In this algorithm, two stations mutually exchange probing frames to determine the received signal strength per sector and then choose the one with maximum signal strength p_n for communication. The selected sector $\hat{n}$ out of N probed sectors is fed back to the transmitter:

$$\hat{n} = \underset{n}{\operatorname{argmax}}\, p_n. \tag{1}$$

As this approach scales linearly with the number of sectors in a training set, it is only appropriate for low numbers of sectors.

To reduce training complexity, compressive path tracking [18, 23–25] reduces the search space to a subset of all sectors in a training set and thereby provides higher efficiency with lower time overhead and a search complexity of $\log N$. Instead of probing all sectors, compressive path tracking as proposed by Rakesh et al. in [25] pseudo-randomly varies phase shifts of the antenna elements to produce random beam patterns. The receiver chooses the most dominant path to the transmitter by correlating the received signal strength with the expected beam patterns. Additional phase information even enables multi-path estimation as proposed in [18]. Such compressive algorithms achieve high efficiency and accuracy especially for antenna arrays with hundreds of elements.

Nevertheless, current solutions cannot directly be integrated in commodity devices, which contain low-cost array antennas with limited numbers of antenna elements. Due to the placement in a device chassis, surrounding materials may impair radiation. Hence, antenna manufacturers already provide a preselected set of antenna patterns with good radiation characteristics. Our preliminary experiments to apply random phase shifts substantially reduced the link quality between our devices under test, thus severely limiting the communication range. Given the minimum SNR requirements that are necessary for proper frame reception, choosing random phase shifts results in less accurate measurements on our low-cost hardware. Additionally, predicting paths between devices only in a two dimensional environment is insufficient, as the placement of communicating nodes is seldom limited to planes in practical setups. Hence, for this work, we extended existing compressive path tracking algorithms to cope with practical three-dimensional setups.

2.2 Our Sector Selection Algorithm

To cope with the limitations of commodity IEEE 802.11ad devices, we adapted the non-coherent path-tracking approach provided in [25]. Instead of relying on random beam patterns, we use the beam patterns defined by sectors that perform well with our device's antenna. Additionally, we extend the search space to three-dimensional patterns to accommodate realistic environments. Our approach works in two steps. First, we take a random subset of M out of N sectors supplied by the antenna to find the path a signal took to arrive at a receiver. We use this to determine the best of the original N sectors that optimizes the signal strength in the estimated direction. In the first step, we probe all M sectors resulting in M received signal strength values p_m, where $m \in [1, M]$. To find the sector with the highest expected signal strength according to expected three-dimensional beam patterns $\vec{x}_m(\phi, \theta)$, for azimuth angles ϕ and the elevation angles θ, we correlate a normalized vector of all received signal strength values $\vec{p}$ with a normalized vector of the corresponding expected beam pattern $\vec{x}(\phi, \theta)$:

$$W(\phi, \theta) = \left\langle \frac{\vec{p}}{||\vec{p}||}, \frac{\vec{x}(\phi, \theta)}{||\vec{x}(\phi, \theta)||} \right\rangle^2, \tag{2}$$

where $\langle \vec{u}, \vec{v} \rangle$ indicates the inner product of two vectors $\vec{v}$ and $\vec{u}$ and $||\vec{u}||$ is the norm. To estimate the path taken by a signal to arrive at the receiver, we determine the angle of arrival by maximizing the

correlation $W(\phi, \theta)$:

$$\hat{\phi}, \hat{\theta} = \underset{\phi,\theta}{\operatorname{argmax}}\, W(\phi, \theta). \tag{3}$$

Given a discrete grid of ϕ and θ, we find the angles $\hat{\phi}$ and $\hat{\theta}$ with maximum correlation numerically. For the second step to find the best out of all N sectors, we use the estimated angles $\hat{\phi}$ and $\hat{\theta}$ to select a sector with ID $\hat{n}$ that provides the strongest gain according to the measured beam patterns $x_n(\phi, \theta)$ with $n \in [1, N]$:

$$\hat{n} = \underset{n}{\operatorname{argmax}}\, x_n(\hat{\phi}, \hat{\theta}). \tag{4}$$

The number of available sectors N can thus be significantly larger than the number of probing sectors M.

2.3 Integration into Devices

To integrate our compressive sector selection into the sector sweep of commodity IEEE 802.11ad devices, we first need to gain access to the signal strength measurements of received sector sweep frames, determine the optimal sector and then insert the selected sector ID in the sector sweep feedback field. This requires our modifications to the Wi-Fi chip's firmware (described in Section 3) and precise knowledge on the antenna patterns of all available sectors (as determined in Section 4).

3 FRAMEWORK SETUP

As no suitable off-the-shelf IEEE 802.11ad testbeds exist, we developed a new open-source platform to perform mm-wave experiments on Talon AD7200 routers which contain Qualcomm's QCA9500 IEEE 802.11ad Wi-Fi chip. As the sector selection algorithm extended in this work runs in the Wi-Fi chip's firmware, we made it modifiable using the Nexmon framework [27]. In the following, we describe our new testbed.

3.1 Talon AD7200 as IEEE 802.11ad Platform

Currently, only a few off-the-shelf devices support IEEE 802.11ad. One of these devices is the Talon AD7200 router that contains Qualcomm's QCA9500 FullMAC IEEE 802.11ad Wi-Fi chip. It is connected to an antenna array with 32 antenna elements used for beam steering by changing phase shifts and amplitudes on each array element. The full control over the antenna as well as the handling of IEEE 802.11ad frames is encapsulated in the Wi-Fi chips proprietary firmware running on two ARC600 processor cores for real-time (ucode processor) and other MAC layer operations (firmware processor). To make the full potential of the Wi-Fi chip accessible to researchers, we first had to open up the router's operating system to get direct access to the wireless interface. To this end, we ported the Linux Embedded Development Environment (LEDE) from the Archer C2600 to the Talon AD7200 and extended it with a current version of the *wil6210* driver for Qualcomm's QCA9500 chip. Equipped with this platform, we can operate the Wi-Fi chip as access point or station as well as in monitor mode, but we are still limited to the features provided by the firmware.

3.2 Jailbreaking the IEEE 802.11ad Firmware

We use the Nexmon firmware patching framework [27] to modify the firmwares running on the two ARC600 processors. This framework allows writing patches in C instead of Assembly which eases patch development. By providing new attributes and pragmas as a GCC plugin, it allows to define where functions and variables should be placed in the patched firmware. However, using ARC600 processors complicates the patch creation as those processors contain separate read-only code and writable data regions (see Figure 1), while Nexmon assumes that the whole patched memory is writable. Fortunately, we discovered that code memory is also accessible at high memory addresses, where it is writable so that it can contain patches, where code and data sections are merged.

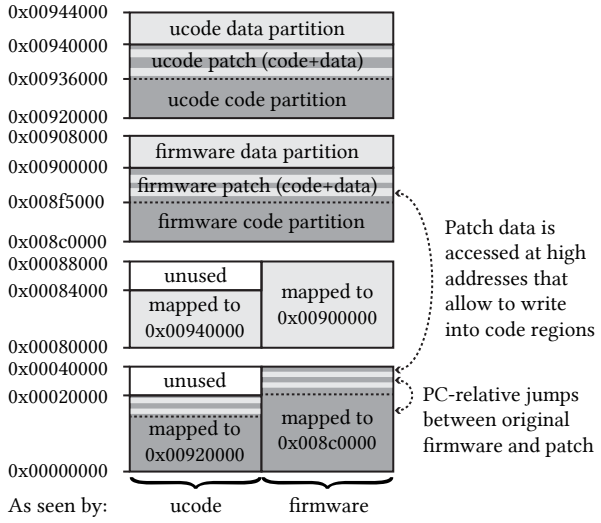

Figure 1: Memory layout of the QCA9500 IEEE 802.11ad Wi-Fi chip with two ARC600 processors (ucode and firmware) that have separate write-protected code and writable data memories at low addresses. All four memory regions are remapped into high addresses, where they are writable and accessible from the host.

Before writing patches, we first had to understand how the original firmware works to identify the parts handling the information we wanted to extract and modify. Without any source code and no function name strings in this firmware, this was a very tedious process. As the firmware delivered with the Talon AD7200 was outdated, we focused our analysis on firmware version "3.3.3.7759" extracted from Acer TravelMate notebooks that also runs well on the Talon AD7200.

To support the research community, we offer the LEDE firmware as well as the modified Nexmon framework for download [30]. This will allow other researchers to build upon our work and further improve the router's MAC and PHY mechanisms. In the next two sections, we describe the firmware patches we created to influence the handling of sector sweep frames.

3.3 Extracting Signal Strength Information

The default firmware neither provides access to sector sweep information, nor does it allow the modification of the sector selection. Our patches aim at enabling both features from the user space. Sector sweeps are handled in the ucode firmware which is hard

to analyze due to the lack of information such as strings that indicate what functions do. Instead, we had to match patterns of IEEE 802.11ad frames in the firmware memory to figure out where sector sweeps are handled. In Figure 2, we illustrate the original sector sweep handling in gray and our extensions to it in white. The original firmware uses received signal strength indicator (RSSI) and signal-to-noise ratio (SNR) measurements from received SSW frames to select the sector with the highest signal strength. We modified the firmware to extract both measurements for each sector sweep into a ring buffer that we can read from user space using our modified driver.

3.4 Selecting Custom Transmit Sectors

In the current state, we cannot directly select a sector for transmission at a node. Instead, we manipulate the selected sector ID in the feedback frames that are exchanged at the end of a sector sweep, as illustrated in Figure 2. To this end, we add a switch to the original implementation that allows to either use the sector ID selected by the original algorithm, or a custom value set from user space by sending a custom Wireless Module Interface (WMI) command. In either case, the sector ID is copied into the sector sweep feedback field of both SSW frames and SSW acknowledgment frames sent by the responder and SSW feedback frames sent by the initiator. As we can modify the feedback in all frames, we can control the sector selection at both the initiator and responder.

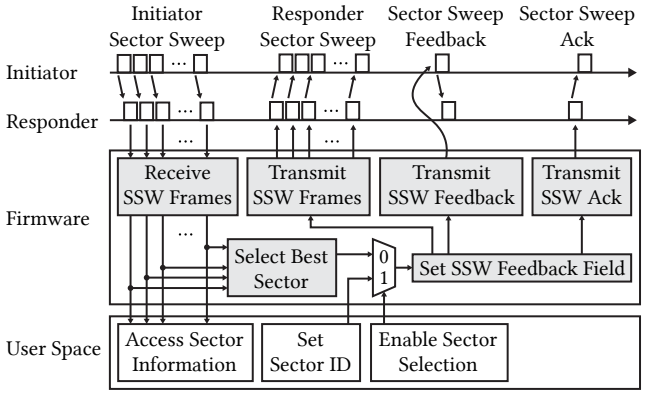

Figure 2: Sector sweep as responder with firmware extensions to access the received signal strength and select custom sectors from user space. White boxes indicate the extended functions, gray boxes the original behavior.

4 SECTOR PATTERN ANALYSIS

As presented in Section 2, our adopted compressive-path-tracking algorithm requires knowledge of the default beam patterns offered by the mm-wave antenna of the Talon AD7200. As we require the device-dependent antenna radiation characteristics for the Talon AD7200 platform, we cannot directly build on existing measurement campaigns which characterizing mm-wave propagation. Instead, we perform extensive measurements to determine the antenna gains in all three dimensions of all of the available antenna patterns that are predefined in the device. In this section, we present how we measured those beam patterns for sector sweep frames with our setup presented in Section 3.

4.1 Beacons and Sector Sweep

To initiate a communication between two IEEE 802.11ad nodes, both have to find an antenna sector pair that provides the best gain for their link. As Access points (APs) do not know the best sectors to advertise their existence to potential clients, they periodically transmit beacon frames successively over multiple sectors. After establishing a connection, IEEE 802.11ad nodes need to periodically update their sector selections to react to environmental changes or device movements that may impair the signal quality on a previously chosen sector through so-called sector sweeps. To figure out which sectors the Talon AD7200 uses for beaconing and sweeping, we deployed three of these devices in close proximity. We set up one of them in AP mode and another in managed mode and established a connection between both of them. We use the third device to capture all received beacon and sector sweep frames by operating it in monitor mode. To capture the frames, we employ tcpdump and analyze their content with Wireshark.

As defined in [36], both beacon and sector sweep frames contain a *sector sweep* field. It contains the sector ID of the sector used to transmit the current frame and a decreasing counter CDOWN that indicates the remaining frames in a burst. To analyze which sectors are used during beaconing and sweeping, we captured the sector IDs and the values of CDOWN and list them in Table 1. Both beaconing and sweeping use sector IDs from 1 to 31 and the ID 63. Sweeping, additionally, probes sectors 61 and 62. These are all the sector IDs potentially used for transmissions; sectors with IDs from 32 to 60 are undefined. In Table 1, we left out some slots which seem to be unused. For those, we never received any frames, independent of the monitor's position. We also observed that the sector sweeping settings stay constant over time. As no training is done for receive sectors, *the same (quasi omni-directional) sector is always used for reception.* Regarding timings, the AP triggers beacons every 102.4 ms [36] and sector sweeps at least once per second. Training transmit sectors in both directions takes on average 1.27 ms, consisting of 18.0 µs per sweep frame and additional 49.1 µs for initialization and feedback frames.

Table 1: IEEE 802.11ad devices transmit Beacon and Sector Sweep frames as bursts into multiple sections. CDOWN counts the remaining sectors in a burst, the rows indicate which sector IDs are used at a given CDOWN value.

CDOWN	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Beacon	-	63	-	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	-
Sweep	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	-	61	62	63

4.2 Beam Pattern Measurement Setup

Knowing the radiation patterns of the aforementioned sectors is crucial to make an optimal selection using a path-tracking algorithm. The shapes of antenna patterns highly depend on the antenna's geometries and objects in the surrounding. Also the packaging and placement of the antenna inside a device influences the radiation characteristics. Hence, to obtain device specific radiation patterns, we performed measurements with two Talon AD7200 devices (see Figure 3) in an anechoic chamber to omit disturbing reflections and multi-path effects. As shown in Figure 4, we mounted one device on a custom rotation head equipped with a step-motor with microstepping support to obtain a high rotation precision in the azimuth plane. To rotate the mm-wave antenna, we placed it directly on the rotation axis. Then, we placed a second Talon AD7200 in a distance of three meters, facing the rotating one. To measure the radiation patterns, we established a connection between both IEEE 802.11ad nodes and transmitted pings for 20 seconds to make them perform sector sweeps and keep the connection alive. Using our firmware patches introduced in Section 3.3 we extracted the received signal strength as SNR values. To control the testbed, we set up a 2.4 GHz network to remotely start experiments on the Talon routers from an ordinary computer by executing commands through an ssh connection (see Figure 3). The step-motor in the rotating head is controlled over USB. To automate experimentation and evaluate the collected measurement data, we used a combination of Python and MATLAB scripts. We present the results in the next section.

4.3 Measuring Antenna Sector Patterns

In our first experiment, we measured the sector patterns by rotating the first Talon router from −180° to 180° in steps of 0.9°, leaving the elevation angle at 0°. The SNR measurements extracted by our firmware patch are quantized in quarters of dB in a range from -7 to 12 dB. In Figure 5, we illustrate our measured radiation patterns for all 35 default sectors of the Talon AD7200's antenna array. To generate these patterns, we omitted obvious outliers, averaged over multiple measurements, and interpolated over gaps where we could not capture any frames due to misses in directions with low gains and decoding errors. To obtain the transmit sector patterns, we observed the sector sweep frames transmitted by the rotating device on the second fixed device. To measure the pattern of the receiving sector (Sector RX) we switch the roles of the two routers, so that the fixed node transmits sector sweep frames that are received and evaluated by the rotating node. To minimize the noise-floor in this pattern, we only considered frames transmitted on sector 63, as it has a strong unidirectional gain.

4.4 Discussing Antenna Sector Patterns

From these results, we learn that some transmit sectors such as 2, 8, 12, 20 24, and 63 provide a strong gain in one particular direction. Other sectors such as those with sector ID 13, 22, and 27 come with multiple, equal powered lobes or cover a wide range such as sector 26. Due to these strong variations, we do not provide beamwidths or sector steering angles. In the direction behind the antenna—for angles higher than ±120°—we observe distorted patterns. This is not surprising as the antenna array is partially blocked by a chip and shielded in this direction. Additionally, some sectors such as

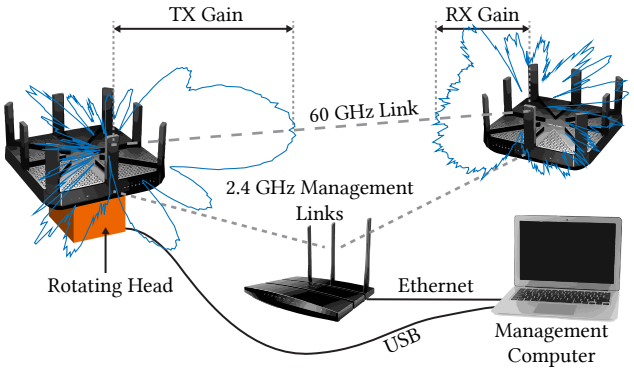

Figure 3: Schematic setup of experiment with two Talon routers and rotation head controlled by a computer over a management network.

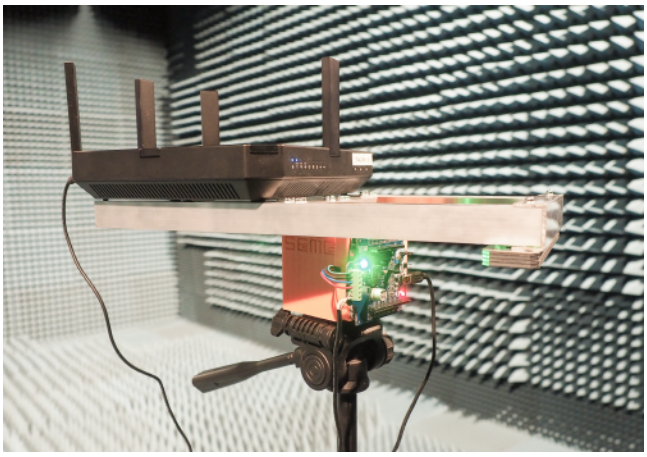

Figure 4: Experimental setup with a Talon router mounted on our custom rotation head in an anechoic chamber.

those with ID 5, 25, and 62 only exhibit low gains in all directions. It is likely that those have their maximum outside the evaluated azimuth plane. Therefore, we extended our measurements to also cover elevation angles.

4.5 Extending to 3D Patterns

To map the antenna sector patterns in a spherical space, we repeat our measurements and manually tilted the rotation head from 0° to 32.4° in steps of 3.6°. We used the same setup as before, but limited the azimuth angle to ±90° and performed SNR measurements every 1.8°. We illustrate our resulting patterns in Figure 6. For each sector, we plotted the measured value as heatmap over the azimuth and elevation angle. Comparing these results with those of the planar patterns in Section 4.3, we find that some sectors, such as Sector 5, with low gains in the plane exhibit stronger lobes at higher elevation angles. Sectors 25 and 62, however, still have low gain in the measured space. Sector 26, which exhibits a strong gain in the azimuth plane has significantly lower gains for higher elevation

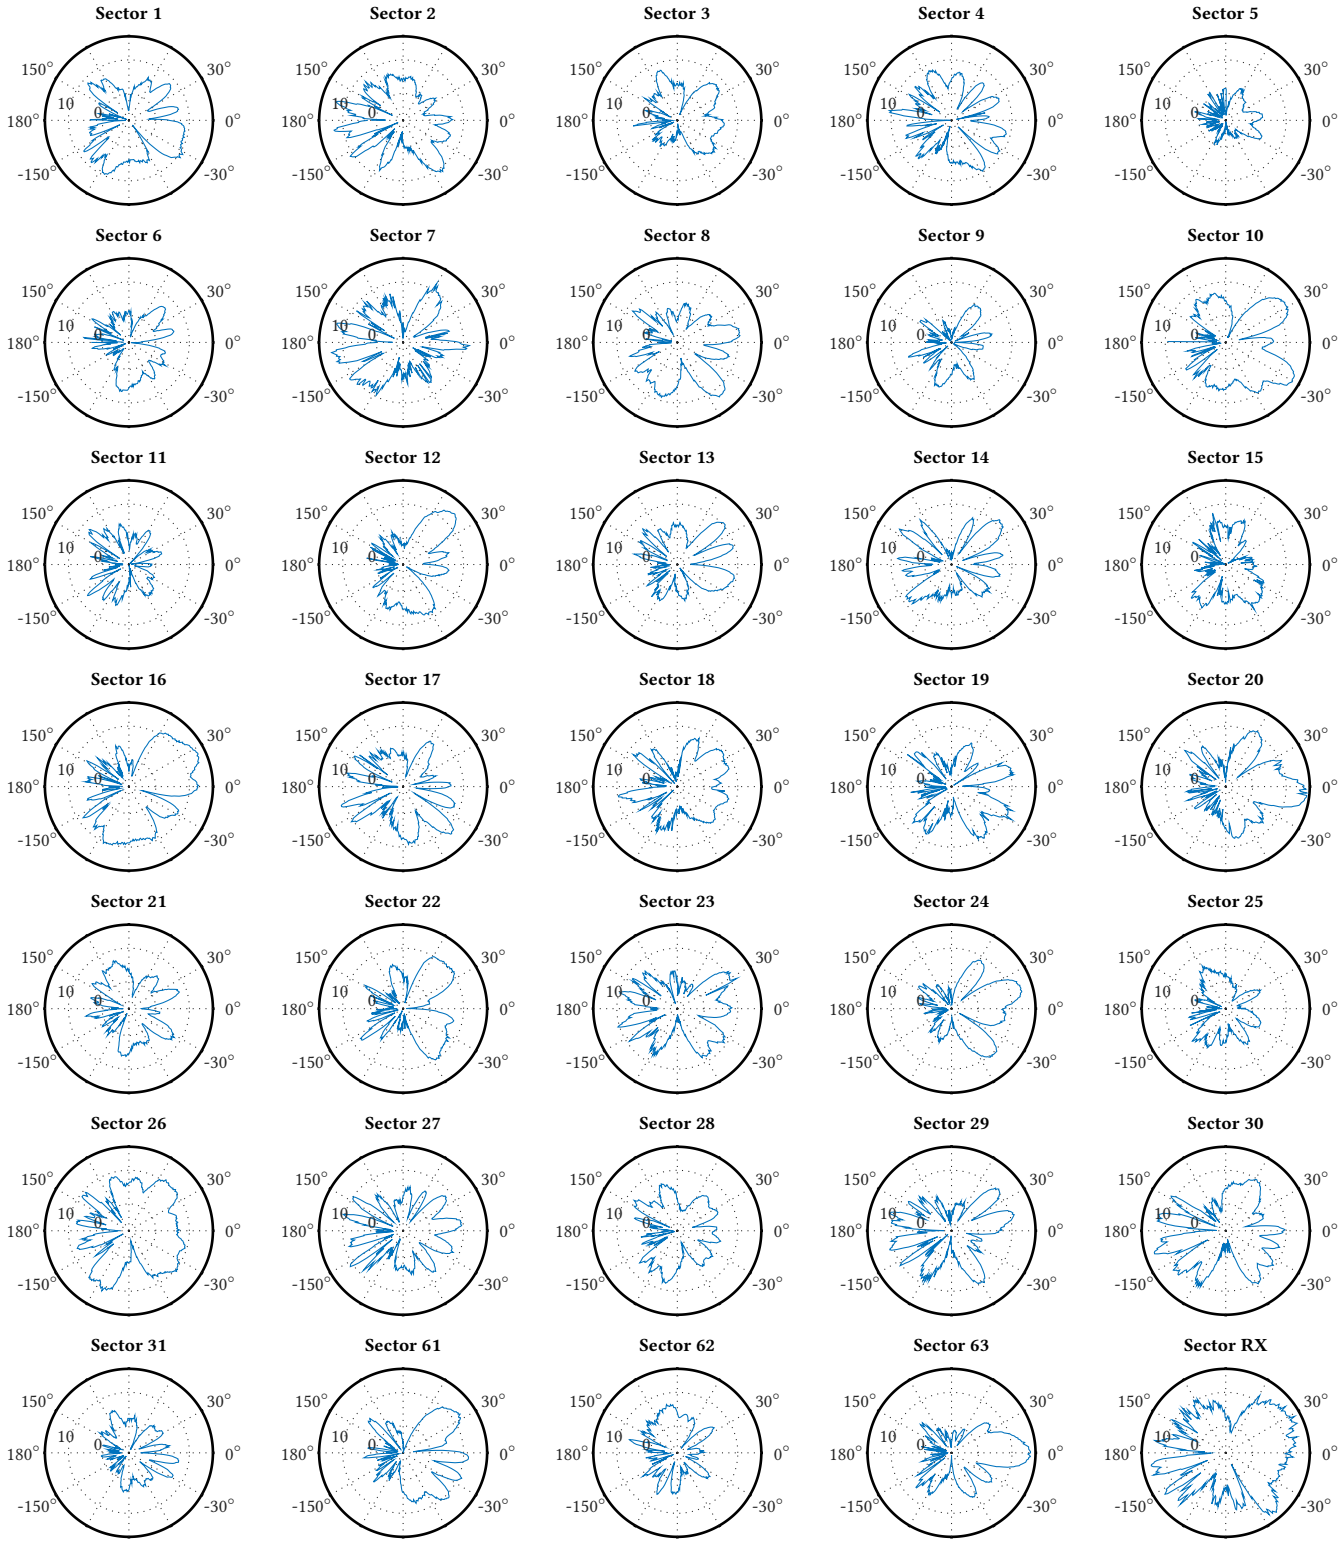

Figure 5: Measured SNR antenna patterns in azimuth plane of all sectors.

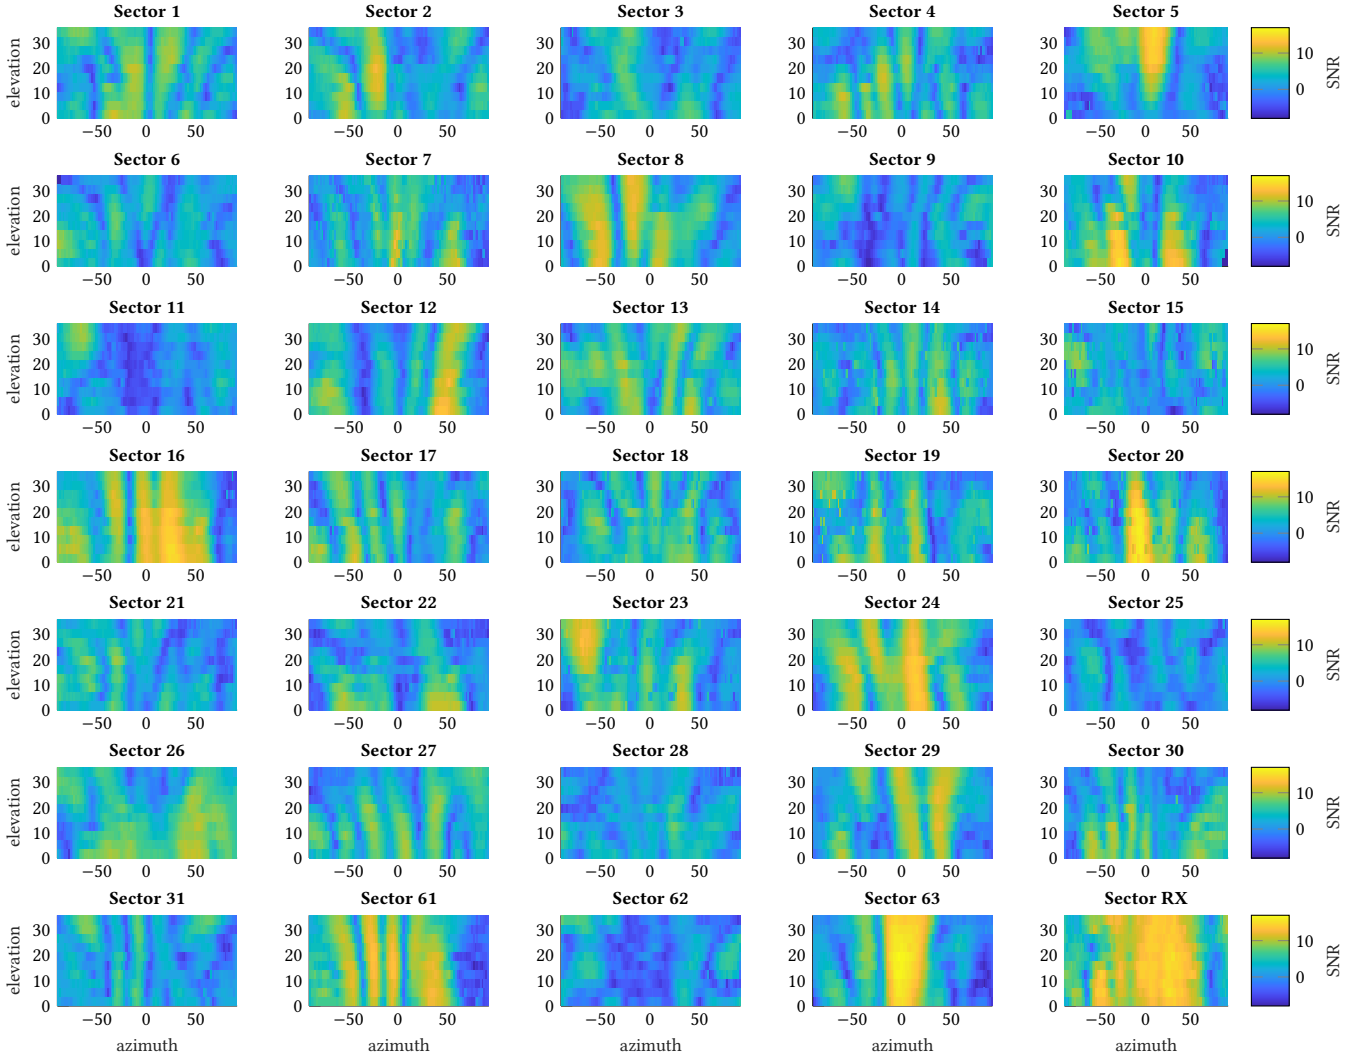

Figure 6: Measured spherical SNR antenna patterns over azimuth and elevation angles.

angles. It covers approximately a torus. Even though our measured patterns only provide an excerpt of spherical coverage, we obtained sufficient information for the main communication direction.

Our measurements, as presented above, capture the radiation characteristics for one particular device. Although we have confirmed that different devices exhibit similar patterns with slight variations, other Talon AD7200 devices might behave differently. All our measurement results can be found online [30]. The integration of the measured patterns in our compressive sector selection algorithm is described in the following section.

5 ESTIMATION WITH ERRORS

While the radiation patterns only need to be measured once, we now focus on evaluating the signal strength values collected in every sector sweep during the regular operation of an IEEE 802.11ad device. To enhance the quality of the patterns, we were able to average over multiple measurements, but when running our path tracking algorithm to select the optimal sector, we have to rely on the results we get from each sector sweep. Those results can be imprecise and fluctuating. Reading the signal strength values from the firmware as described in 3.3, we observed severe outliers. Especially channels with low gains resulted in high signal strength deviations and sometimes the firmware does not report any measurements at all. This behavior is also observable with the original firmware without any modifications. It impairs the original sector sweep algorithm and causes fluctuations in the sector selections.

Averaging over multiple measurements is, however, not feasible as sector selection algorithms need to react quickly to changes in the environment to adjust the selected sector accordingly. To cope with the effects of measurement fluctuations in our path tracking protocol, we decided to not only use SNR measurements of the sector sweep frames, but also RSSI measurements. Both are provided by the firmware and extractable with our patch, but internally they seem to be acquired differently as fluctuations are not observable in

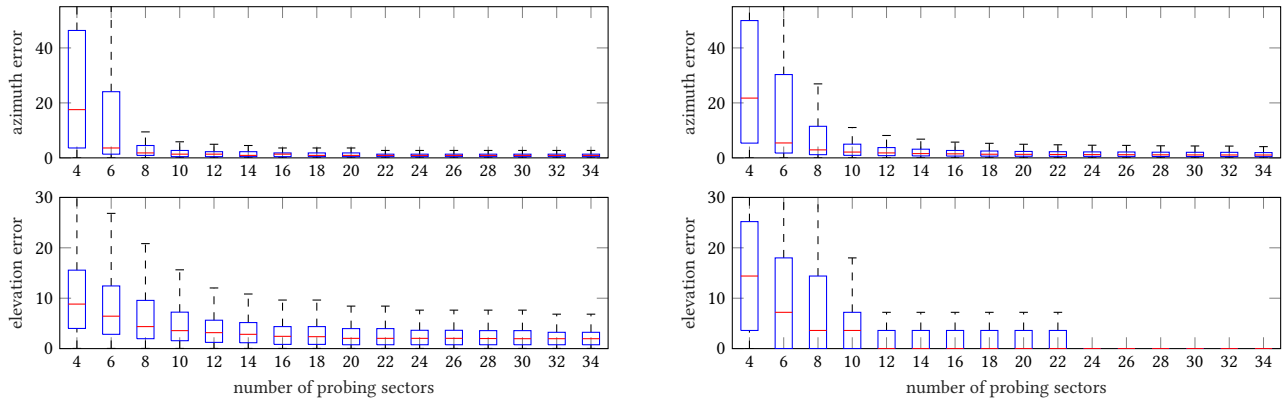

(a) Lab Environment **(b) Conference Room**

Figure 7: Angular estimation errors with compressive sector selection in lab environment and conference room allow for accurate estimation of spherical path direction with only a few probing sectors.

both values at the same time. Nevertheless, on average SNR and RSSI measurements are correlated. Thus, we extended the calculation of the estimation of the angle of arrival presented in Equation 3 by considering correlations with the measured SNR and RSSI values:

$$\hat{\phi}, \hat{\theta} = \underset{\phi,\theta}{\operatorname{argmax}}\, W_{\text{SNR}}(\phi,\theta) \cdot W_{\text{RSSI}}(\phi,\theta). \tag{5}$$

This effectively tolerates more outliers and increases the robustness against measurement deviations in either value. By probing only a subset of antenna sectors, we have a huge advantage over the original sector sweep algorithm: we naturally compensate missing measurements effectively. Assuming that the firmware misses to report the signal strength of the best sectors, the sector sweep algorithm would choose a less optimal one. With compressive path tracking, we still have good chances to find the optimal sector despite a decreased set of probing values.

6 PERFORMANCE EVALUATION

In the following, we present our evaluation results to demonstrate that our compressive sector selection algorithm is feasible for off-the-shelf IEEE 802.11ad devices and decreases the training time in comparison to the prevalent sector sweep. After describing our setup, we analyze the angular estimation accuracy, the impacts of sector selections and evaluate the overhead reduction and throughput.

6.1 Evaluation Setup

For evaluation purposes, we use the same setup as for obtaining the antenna patterns in Section 4.2 but take measurements in a lab environment and a conference room. In the lab environment, we place the two devices three meters apart, in the conference room six meters apart. The conference room contains a couple of potential reflectors such as white-boards so that we can expect reflections and multi-path effects. For both scenarios, we set the range of our rotation head to ±60°. In the lab environment, we tilt the rotation head in steps of 2° from 0° to 30° and use an azimuth resolution of 2.25°. In the conference room, we do not change the elevation angle, but increase the resolution of azimuth angles to 1.3°. These resolutions differ from those in the anechoic chamber (see Section 4.2) to demonstrate the general applicability and independence of a specific direction. Running experiments in both the lab environment and the conference room, we cover typical, indoor line-of-sight scenarios for IEEE 802.11ad communication.

In all of our experiments, we make the two devices perform sector sweeps with our firmware extension enabled and record the signal strength as SNR and RSSI value for each sweep and sector. We then perform offline analyses in MATLAB to assess the performance of our compressive sector selection (CSS) algorithm and the original sector sweep (SWP) algorithm. To assess the desired behavior of compressive estimation with a limited number of probing sectors, we only consider a variable number of random measurements in each sweep. Based on the computed outcome, we set the corresponding sector in the sector sweep feedback field to force devices to use our selected sector for data transmission.

6.2 Path Estimation Error

As our compressive sector selection approach relies on the path a mm-wave signal takes to arrive at its destination, we first evaluate the accuracy of the angle-of-arrival estimation by calculating the estimation error in azimuth and elevation direction. The error is the difference of the physical orientation and the estimated one. From the setup of our experiments, we know the physical orientation of the device mounted on top of our rotation head. To determine the estimated path, we perform our protocol as presented in Section 2.2 by computing the correlation according to Equation 5 and estimating the direction with Equation 3. Both errors in azimuth and elevation direction are handled independently, since we measured them with different resolution and accuracy.

Our results for both evaluation scenarios are shown in Figure 7, where boxes indicate the 50%, whiskers the 99% confidence bounds and the dash the median over all our measurements. A suitable approximation of the azimuth direction is achieved with only a few probing sectors. For example, with 10 probing sectors in the lab environment, the error stays with 99% confidence below 5.8° and

maximum median of 1.3°. With 20 probing sectors, upper bound decreases to 3.6° and the median becomes as precise as 0.9°. In the conference room setup the errors we observe are slightly higher due to worsened environmental conditions. A higher distance between the devices and stronger multi-paths make the correlation of measurements and estimated patterns less accurate. However, with 10 probing sectors, we still achieve an azimuth estimation error with a median of 2.1° that decreases to 1.3° by using 20 probing sectors. In elevation direction, we observe errors in both environments below 15° with 10 probing sectors and below 8.4° with 20 probing sectors. Given that we obtained the sector patterns in elevation direction with half the resolution than in azimuth direction, higher errors are expected. Also, we observe that elevation errors in the conference room are lower than in the lab environment. We expect some of the errors to be caused by manually tilting the rotation head. Despite of using a digital mechanic's level, we did not achieve a sub-degree precision in this direction. In the conference room environment we did not evaluate different elevation direction, hence, observe lower elevation error rates for a high number of probing sectors. In summary, we observe that with at least 12 probing sectors a suitable approximation of the signal path becomes possible.

6.3 Sector Selection Accuracy

Based on the path directions that are evaluated in the previous section, we select the sector with the highest gain in the path direction from our measurements. To investigate the accuracy of this selection we consider the SNR-loss and the selection stability. We state these results for both algorithms, compressive sector selection and the sector sweep, in the following.

The selection stability represents the time a selection algorithm spend in one particular sector. We determine this rate by evaluating the selection of our protocol and the sector sweep for each captured sweep in the conference room. For each physical path direction, we identify the sector that is selected most and count the occurrences. This number divided by the total number of evaluated sweeps provides the selection stability. In the assumption that each sweep has the same duration, the selection stability directly corresponds to the time spent in one particular sector. Spending more time in one sector indicates a higher stability. The selection stabilities, averaged over all evaluated path directions, are shown in Figure 8. For SSW, the results are constant as this algorithm always use the maximum number of probes. It achieves a stability of 73.9 % meaning that it's outcome is in average with 73.9 % one constant sector. With 26.1 % the SSW results in inconsistent selections. The sector sweep tends to switch between sectors due to measurement deviations. Multiple sectors with similar gains might be alternating for the highest measurement. With at least 13 probing sectors, compressive sector selection achieves a higher selection stability than the sector sweep. Applying all probing sectors that are available we stay in average of 94.7 % of the time in one particular sector. This demonstrates that compressive sector selection not only achieves the stability of the sector sweep with fewer probing sectors, but also reduces the amount of alternating sector selections.

However, the selection stability does not represent the quality of the selected sectors. To catch up on this, we additionally investigate the loss in SNR achieved by compressive sector selection and the

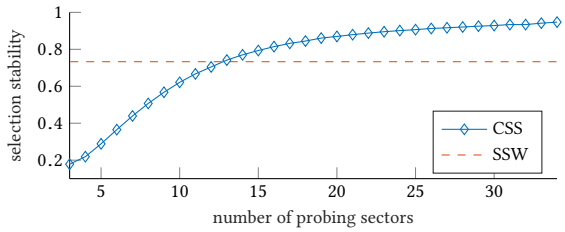

Figure 8: The selection stability illustrates the time spent in the most prominent sector. With at least 13 probing sectors our compressive sector selection (CSS) has less selection fluctuations than the sector sweep (SSW). Plotted results are averaged over all evaluated directions in the conference room.

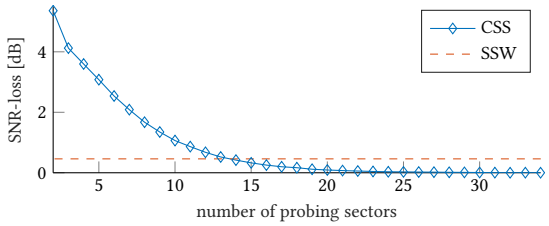

Figure 9: Average loss of SNR in compressive sector selection and sector sweep in dependency of the number of probing sectors.

sector sweep in comparison to the optimal achievable SNR. Doing so, we not only determine the sector selection but also the SNR value in each sweep. As optimal SNR, we consider the sector with the highest SNR as reported in the current and previous measurements. The SNR-loss is determined as difference towards this value. We average the SNR-loss over all evaluated directions and plot our results in dependency of the number of probing sectors in Figure 9. In this comparison, the sector sweep achieves a reasonable SNR in average just 0.5 dB below the optimum. Our protocol already achieves considerable results with only a few probing sectors: 6 probes are sufficient to find a sector with an average 2.5 dB below optimum. This might be insufficient to achieve high data-rates but already allows for simple data exchange. Greater values than those of the SSW are achieved with at least 14 probing sectors. With about 20 probing sectors the achieved SNR approaches the optimal one quite well. These values exhibit high variations over different directions that are not encountered in the averaged results. Nevertheless, 14 probing sectors are sufficient to make a sector selection that provides the same SNR as the one obtained by the sector sweep.

6.4 Overhead Reduction and Throughput

The previous results revealed that with at least 14 probing sectors, compressive sector selection achieves a low estimation error as well as an SNR and selection stability in the order of that of the sector sweep. Based on these results, we evaluate the overhead reduction and the throughput that is achieved with only 14 probing sectors.

The training time directly depends on the number of probing sectors. As found in Section 4.1, we assume that probing a single sector takes 18.0 µs and the feedback and acknowledgment add another overhead of 49.1 µs. The resulting training time in dependency of the number of probing sectors is illustrated in Figure 10. Specifically, with only 14 probing sectors, the time to complete a mutual training of transmit sectors could be performed in 0.55 ms. In comparison to the original sector sweep that take 1.27 ms, this effectively speeds up the beam-training by factor 2.3.

To assess the throughput, we determine the selected sectors as done before for each sweep. With this sector set in the sector sweep feedback field, we generate random TCP payload and measure the application-layer data-rate by running iPerf3[1]. Even in static scenarios, the devices trigger the beam training approximately once every second. We average the throughput that we measured over 10 seconds over all selected sectors to take into account the impacts of suboptimal selections. We take measurements in the conference room with the rotation head steered to −45°, 0° and 45°. As shown in Figure 11, the expected throughput of compressive sector selection is with 1.48 Gbps, 1.51 Gbps, and 1.50 Gbps only slightly better than that of the SSW. In practice, these differences might barely be recognizable but yet show the additional performance gain we achieve from higher stability.

Theoretically, a decreased training time, leaves more time for data transmission and thereby increases the throughput. This aspect is not covered in our evaluation yet. The provided throughput comparisons are performed with equal sweep duration. We leave adjusting the duration of the sector sweep for future work.

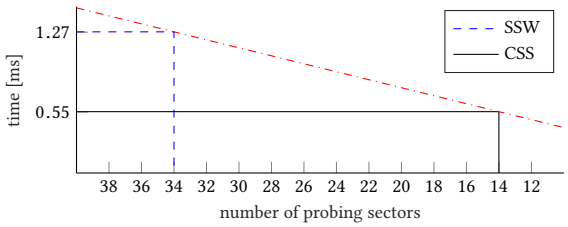

Figure 10: Required time to perform a mutual training in dependency of the number of probing sectors to illustrate that compressive sector selection (CSS) completes the training in 0.55 ms by using 14 probing sectors while the sector sweep (SSW) takes 1.27 ms.

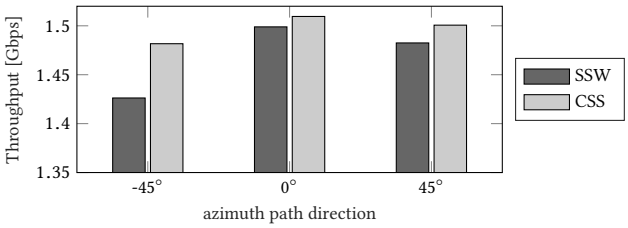

Figure 11: Additionally expected throughput gain from increases stability in compressive sector selection with 14 probing sectors in comparison with the sector sweep.

6.5 Summary

In this evaluation, we demonstrated the feasibility of compressive sector selection on IEEE 802.11ad off-the-shelf devices. In our protocol 14 out of 34 probing sectors are sufficient to outperform the sector sweep. Compressive sector selection estimates the path direction with high accuracy and error of only a few degree. The achieved signal strength and throughput are in the same order as those of the sector sweep. With our approach, we increase the stability of sector selections and therewith spend more time transmitting in the optimal sector. By using only 14 probing sectors, compressive sector selection reduces the mutual training time from 1.27 ms to 0.55 ms. However, we support a variable number of sectors, so that our system adjusts to arbitrary settings.

7 DISCUSSION

As presented in the evaluation section, our compressive sector selection algorithm is the first working implementation of compressed path tracking that runs on off-the-shelf 60 GHz mm-wave hardware. Our results show that probing only 14 out of 34 possible sectors is sufficient to speed up the sweeping process by a factor of 2.3. With only two nodes per room, the gain in throughput is low. However, if we consider dense mm-wave node deployments, we need to keep in mind that each sector sweep performed by a pair of nodes pollutes the whole mm-wave channel in all directions. This reduces the benefit of using mm-wave hardware to communicate with many stations in parallel over directional links. Thinking about mobile setups, users need to be tracked over time. The shorter the sweeping time, the more often a sweep can be performed without degrading the throughput too much. Hence, our approach is best suited to increase the performance and frequency of sweeping in emerging mm-wave applications.

While current mm-wave devices, such as the Talon AD7200 router, use only 34 different sectors, future generations are likely to demand higher directivities and more fine-grained beam control. Such requirements could be addressed by increasing the number of implemented and predefined sectors. The Talon AD7200' phased antenna array with 32 individual antenna elements already suffices to refine the beams. However, increasing the number of sectors adds additional overhead to the training process. While with few sectors, the sweep completes in suitable time, it becomes inefficient with a large set of probing sectors. Solutions such as our compressive sector selection overcome this limiting factor and scale well with high number of sectors. For example, with our approach we could significantly increase the number of available sectors while keeping the number of probes as low as in the current sweep. As a result, more precise beam patterns could be efficiently selected without adding additional training time overhead.

Further improvements are achievable from adaptively controlling the number of sectors that are probed in the sweep. For example, in static scenarios, few probes are sufficient to validate the current antenna settings. Whenever a node starts moving, the number of probes may increase to keep track of the movement. Instead of

[1] iPerf3 network bandwidth measurement tool accessible at https://iperf.fr

applying a random selection, predefined probing sectors might provide further benefits. Given an approximate signal direction, sectors with low gain or similar gains in that direction are unlikely to perform efficient. However, choosing the most-efficient probing sectors turned out to be highly context specific. Still, compressive sector selection is suited for such scenarios by keeping the number of probes as well as the selection of sectors a variable parameter.

Research with practical mm-wave systems is currently getting started and we hope that our Talon AD7200 research framework encourages other researchers to explore the boundaries of this low-cost off-the-shelf hardware and conduct own experiments. To this end, we make our tools, in particular the adapted Nexmon framework, the LEDE sources, and the measured sector patterns, publicly available [30]. However, current firmwares impede researchers from rapidly prototyping new algorithms on cheap and widely available hardware. We hope that chip manufacturers may open their interfaces to allow a better integration and low-level access to their hardware that enables the evaluation of new protocols.

8 RELATED WORK

Optimal algorithms for sector selection and beam maintenance in mm-wave communication systems are heavily discussed in the research community. Phased array antennas achieve beamforming by weighting individual antenna elements which enables new beamforming and training concepts [16]. To increase the efficiency of beamtraining algorithms, Wang et al. found out that the number of probing beams must be only twice the number of antenna elements to estimate the path direction [35]. Hierarchical codebook structures of beams with different beamwidths allow to iteratively refine the beamtraining [2, 15]. While their complexity to find the best matching beam is $\log N$, hierarchical protocols increase the communication overhead to feedback their results after each round. Compressive sensing based path tracking protocols achieve the same complexity but only require one training round. As described in [18, 23, 24], such methods select pseudo-random beams for probing and derive the path based on magnitude and phase of the received signal. As phase information are seldom accessible on practical systems, Rasekh et al. extend compressive path tracking by only relying on signal strength information in [25]. Efficient training algorithms alone are, however, insufficient to handle highly directional mm-wave connections threatened by mobility and blockage. Adaptive beam maintenance ensures connectivity in scenarios with changing environmental conditions. To reduce the risk of disconnections, Haider et al. adaptively controls data-rates and beamwidths during communications[13]. Their approach allows to widen the beam and lower the data-rate on demand. The authors in [31] proactively predict the quality of multiple paths to switch beams, whenever the signal quality degrades. Nitsche et al. [21] use out-of-band channel measurements from 2.4 or 5 GHz networks to predict the signal path before establishing a mm-wave link. Ali et al. [1] follows a similar approach combined with compressed sensing techniques to keep track of the optimal beam-steering direction.

In Multiple Input Multiple Output (MIMO) systems with multiple antenna arrays driven by individual RF chains, beamforming techniques allow to increase the gain of spatial re-usability [5, 6, 22, 28, 29] to serve multiple users simultaneously. Adaptive algorithms with hierarchical codebook structures allow for efficient channel estimation [2, 3] which also operates in frequency selective wide-band channels [4]. Structured compressive sensing approaches such as [11, 12] enable to estimate the wireless channel of massive MIMO networks with few probes efficiently. The authors in [10] apply compressive sensing to select the correct beams in multi-user MIMO with low interference. Theoretical considerations of the advantages of compressed sensing in multi-path sparsity have been formalized in [7]. The coding based beamforming scheme proposed in [33] allows in-packet training and continuously switches beams during transmissions.

Due to the lack of wide-spread mm-wave interfaces in consumer hardware, researchers often base their work on expensive prototyping platforms with directional horn antennas [13, 31, 34] or custom antenna arrays [25, 37]. Both platform types exhibit a behavior that is different from commercial off-the-shelf devices and often lack standard compliant IEEE 802.11ad implementations. While common firmwares for WiFi chips in the lower frequency bands could be modified to support custom MAC-layer protocols [8, 9] or extract channel state information (CSI) [14], this feature is not yet available for any IEEE 802.11ad chip. Recent work has shown that some low-layer statistics can be obtained from wireless docking stations [17, 19] but are strongly limited by the firmware interface. In [26, 32], the authors perform practical evaluations with commodity routers, similar to the Talon AD7200, and extended the device driver to capture extended statistics. Nevertheless, all work on commodity devices lack encompassing access to low-layer parameters, which is required to implement custom beam training protocols.

9 CONCLUSION

In this work, we implemented a compressive sector selection protocol that integrates with the sector sweep on off-the-shelf IEEE 802.11ad devices. To run our protocol on practical devices, we turned the Talon AD7200 router into an extensible mm-wave research platform. In particular, we ported LEDE to the Talon AD7200 and adopted the Nexmon firmware patching framework for the IEEE 802.11ad Wi-Fi chip. We extended the firmware of this chip to get access to sector sweep information and made overwriting the transmit sector possible. Through extensive measurements with this platform in an anechoic chamber, we generated three-dimensional sector patterns for each of the 35 predefined sectors on the Talon AD7200 router. These patterns, used for compressive sector selection, allow us to estimate the signal direction with low error. Thereby, we significantly decrease the number of sectors that need to be probed in order to perform an optimal sector selection. Our results show that the time to search for the optimal sector in a mutual beamtraining scenario can be reduced from 1.27 ms needed for a complete sector sweep to 0.55 ms. Hence, we speed up the beamtraining on off-the-shelf devices by a factor of 2.3. Our system design directly integrates with the sector sweep, and therefore operates orthogonal to other MAC and network layer optimizations.

To allow other researchers to benefit from our results and the IEEE 802.11ad research platform, we make our measured sector patterns, the source code of the LEDE system, as well as the adopted Nexmon framework publicly accessible.

ACKNOWLEDGEMENTS

This work has been funded by the German Research Foundation (DFG) in the Collaborative Research Centers (CRCs) 1119 "CROSSING — Cryptography-Based Security Solutions" and 1053 "MAKI — Multi-Mechanism-Adaptation for the Future Internet". This work also has been the funded by the German Federal Ministry of Education and Research (BMBF) and the State of Hesse within CRISP-DA, and the Hessian LOEWE excellence initiative within NICER. Further, this work was supported by the European Research Council grant ERC CoG 617721, the Ramon y Cajal grant from the Spanish Ministry of Economy and Competitiveness RYC-2012-10788, the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919), and the Alexander von Humboldt Foundation.

REFERENCES

[1] Anum Ali and Robert W Heath Jr. 2017. Compressed Beam-Selection In Millimeter Wave Systems With Out-of-band Partial Support Information. In *Proc. of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE.

[2] 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* 8, 5 (2014), 831–846.

[3] Ahmed Alkhateeb, Omar El Ayach, Geert Leus, and Robert W Heath. 2014. Single-sided adaptive estimation of multi-path millimeter wave channels. In *Proc. of the 15th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*. IEEE, 125–129.

[4] Ahmed Alkhateeb and Robert W Heath. 2016. Frequency Selective Hybrid Precoding for Limited Feedback Millimeter Wave Systems. *IEEE Transactions on Communications* 64, 5 (2016), 1801–1818.

[5] Omar El Ayach, Robert W Heath, Shadi Abu-Surra, Sridhar Rajagopal, and Zhouyue Pi. 2012. Low complexity precoding for large millimeter wave MIMO systems. In *Proc. of the International Conference on Communications (ICC)*. IEEE, 3724–3729.

[6] Omar El Ayach, Robert W Heath, Shadi Abu-Surra, Sridhar Rajagopal, and Zhouyue Pi. 2012. The capacity optimality of beam steering in large millimeter wave MIMO systems. In *Proc of the 13th Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*. IEEE, 100–104.

[7] Waheed U Bajwa, Jarvis Haupt, Akbar M Sayeed, and Robert Nowak. 2010. Compressed Channel Sensing: A New Approach to Estimating Sparse Multipath Channels. *Proc. IEEE* 98, 6 (2010), 1058–1076.

[8] Daniel S. Berger, Francesco Gringoli, Nicolò Facchi, and Ivan Martinovic. 2014. Gaining insight on friendly jamming in a real-world IEEE 802.11 network. In *ACM Conference on Security and Privacy in Wireless & Mobile Networks (WiSec) 2014*. Oxford, United Kingdom.

[9] Daniel S. Berger, Francesco Gringoli, Nicolò Facchi, Ivan Martinovic, and Jens B. Schmitt. 2016. Friendly Jamming on Access Points: Analysis and Real-World Measurements. *IEEE Transactions on Wireless Communications* 15, 9 (2016), 6189–6202.

[10] Jinho Choi. 2015. Beam Selection in mm-Wave Multiuser MIMO Systems Using Compressive Sensing. *IEEE Transactions on Communications* 63, 8 (2015), 2936–2947.

[11] Zhen Gao, Linglong Dai, and Zhaocheng Wang. 2016. Channel estimation for mmWave massive MIMO based access and backhaul in ultra-dense network. In *Proc. of the International Conference on Communications (ICC)*. IEEE, 1–6.

[12] 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.

[13] Muhammad Kumail Haider and Edward W Knightly. 2016. Mobility resilience and overhead constrained adaptation in directional 60 GHz WLANs: protocol design and system implementation. In *Proc. of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*. ACM Press, New York, USA, 61–70.

[14] Daniel Halperin, Wenjun Hu, Anmol Sheth, and David Wetherall. 2011. Tool Release: Gathering 802.11n Traces with Channel State Information. *ACM SIGCOMM CCR* 41, 1 (Jan. 2011), 53.

[15] 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), 4391 – 4403.

[16] Shajahan Kutty and Debarati Sen. 2016. Beamforming for Millimeter Wave Communications: An Inclusive Survey. *IEEE Communications Surveys & Tutorials* 18, 2 (2016), 949–973.

[17] Adrian Loch, Guillermo Bielsa, and Joerg Widmer. 2016. Practical Lower Layer 60 GHz Measurements Using Commercial Off-The-Shelf Hardware. In *Proc. of the 10th International Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WiNTECH)*. ACM, 9–16.

[18] 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), 514 – 527.

[19] 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 *Proc. of the 11th Conference on Emerging Networking Experiments and Technologie (CoNext)*. ACM, 1–13.

[20] 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.

[21] Thomas Nitsche, Adriana B Flores, and Edward W Knightly. 2015. Steering with eyes closed: mm-wave beam steering without in-band measurement. In *Proc. of the Conference on Computer Communications (INFOCOM)*. IEEE, 2416–2424.

[22] Minyoung Park and Praveen Gopalakrishnan. 2009. Analysis on spatial reuse and interference in 60-GHz wireless networks. *IEEE Journal on Selected Areas in Communications* 27, 8 (2009), 1443–1452.

[23] Dinesh Ramasamy, Sriram Venkateswaran, and Upamanyu Madhow. 2012. Compressive adaptation of large steerable arrays. In *Proc. of the Information Theory and Applications Workshop (ITA)*. IEEE, 234–239.

[24] Dinesh Ramasamy, Sriram Venkateswaran, and Upamanyu Madhow. 2012. Compressive tracking with 1000-element arrays: A framework for multi-Gbps mm wave cellular downlinks. In *Proc. of the 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 690–697.

[25] Maryam Eslami Rasekh, Zhinus Marzi, Yanzi Zhu, Upamanyu Madhow, and Haitao Zheng. 2017. Noncoherent mmWave Path Tracking. In *Proc. of the 18th International Workshop on Mobile Computing Systems and Applications*. ACM, New York, USA, 13–18.

[26] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2017. WiFi-Assisted 60 GHz Wireless Networks. In *Proc. of the 23rd Annual International Conference and Mobile Computing and Networking (MobiCom)*.

[27] Matthias Schulz, Daniel Wegemer, and Matthias Hollick. 2017. Nexmon: The C-based Firmware Patching Framework. https://nexmon.org. (2017).

[28] Jaspreet Singh and Sudhir Ramakrishna. 2014. On the feasibility of beamforming in millimeter wave communication systems with multiple antenna arrays. In *Proc. of the Global Communications Conference (GLOBECOM)*. IEEE, 3802–3808.

[29] 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.

[30] Daniel Steinmetzer, Daniel Wegemer, and Matthias Hollick. 2017. Talon Tools: The Framework for Practical IEEE 802.11ad Research. (2017). https://seemoo.de/talon-tools/

[31] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: enabling robust 60 GHz links under blockage. In *Proc. of the 13th Usenix Conference on Networked Systems Design and Implementation (NSDI)*. Microsoft Research, USENIX Association, 193–206.

[32] Teng Wei and Xinyu Zhang. 2017. Pose Information Assisted 60 GHz Networks: Towards Seamless Coverage and Mobility Support. In *Proc. of the 23rd Annual International Conference and Mobile Computing and Networking (MobiCom)*.

[33] Y Ming Tsang, Ada S Y Poon, and Sateesh Addepalli. 2011. Coding the Beams: Improving Beamforming Training in mmWave Communication System. In *Proc. of the Global Communications Conference (GLOBECOM)*. IEEE, 1–6.

[34] Vignesh Venkateswaran, Xinyu Zhang, Sanjib Sur, Vignesh Venkateswaran, and Parmesh Ramanathan. 2015. *60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling*. A Link-Level Profiling, Vol. 43. ACM, New York, USA.

[35] Junyi Wang, Zhou Lan, Chang-Woo Pyo, Tuncer Baykas, Chin-Sean Sum, M Azizur Rahman, Jing Gao, Ryuhei Funada, Fumihide Kojima, Hiroshi Harada, and Shuzo Kato. 2009. Beam codebook based beamforming protocol for multi-Gbps millimeter-wave WPAN systems. *IEEE Journal on Selected Areas in Communications* 27, 8 (2009), 1–6.

[36] IEEE 802.11 working group. 2012. Standard for Information technology–Telecommunications and information exchange between systems–Local and metropolitan area networks–Specific requirements-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band. (Dec. 2012).

[37] Jialiang Zhang, Xinyu Zhang, Pushkar Kulkarni, and Parameswaran Ramanathan. 2016. OpenMili: A 60 GHz Software Radio Platform with a Reconfigurable Phased-array Antenna. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking (MobiCom)*. ACM, New York, NY, USA, 162–175. DOI: https://doi.org/10.1145/2973750.2973760