Geolocation-based Sector Selection for Vehicle-to-Infrastructure 802.11ad Communication

Mateus Mattos[a], António Rodrigues[a], Rui Meireles[b,*], Ana Aguiar[a]

[a] *Instituto de Telecomunicações, FEUP DEEC, University of Porto, Portugal*
[b] *Computer Science Department, Vassar College, USA*

Abstract

To improve range, IEEE 802.11ad uses directional communication, the first step of which is to choose the best antenna configuration, known as sector. We studied the sector selection behavior of Commercial Off-The-Shelf (COTS) 802.11ad equipment in an experimental Vehicle-to-Infrastructure (V2I) communication scenario. First, we created a framework for mm-Wave data collection in mobile scenarios, highlighting the challenges, justifying our solutions, and making them available, thus flattening the path for future experimentation by others. We then proceeded to use said framework to collect data in a realistic V2I communication scenario. We performed two independent measurement campaigns in the same area. Analysis of the collected data revealed the following inefficiencies in the devices' sector selection algorithm: (i) a large number of sector selection attempts that do not result in a sector change; and (ii) a *"ping-pong"* effect in which a node oscillates between two sectors. With this in mind we propose an alternative antenna sector selection scheme that uses spatially-indexed historical performance data to pick the statistically-best sector for any given geolocation. A trace-based analysis showed that a position-based strategy can improve throughput by more than 10% for 30% of locations, and that gains can be as high as 60 %, in some instances.

Keywords: VANETs, V2I, mm-Wave, 802.11ad, directional communication

1. Introduction

The wide unlicensed spectrum available in the 60 GHz region enables very high data rates. But with high frequency also comes high attenuation. IEEE 802.11ad [1] tries to address this by concentrating transmission energy on a narrow beam directed at the receiver, i.e., it uses directional communication.

However, the forming and maintenance of these beams — known as beamforming and tracking, respectively — are challenging. This is especially true in the case of Vehicle-to-Infrastructure communication. Roadside Access Points (APs) are fixed, whilst vehicles move along roads. As depicted by Fig. 1, when a vehicle moves the angle between it and the AP will change, necessitating significant beam corrections. While mounting APs above the road can minimize angle changes, accurate beam tracking will always be a key factor in this application.

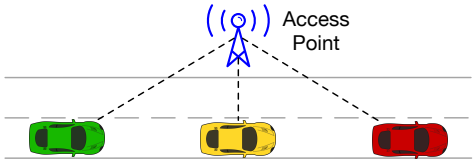

Figure 1: Change in antenna angles as vehicle passes by roadside access point.

802.11ad initiates beamforming by sweeping through a set of preset antenna configurations, or sectors, a process known as a Sector-Level Sweep (SLS). The best-performing sector, i.e., the one yielding the highest Signal-to-Noise Ratio (SNR), is chosen as a starting point for communication, possibly being subject to refinement later on.

We study the behavior of the sector selection algorithm employed by Commercial Off-The-Shelf (COTS) 802.11ad devices, under an experimental Vehicle-to-Infrastructure communication scenario. Our analysis uncovered inefficiencies such as a large number of sector sweeps that do not trigger a sector change, and oscillatory behavior where the selection repeatedly alternates between a pair of sectors, which we name *"ping-ponging"*.

In this paper, we study the efficacy of geolocation as a driver for sector selection, putting it forward as an alter-

*This work is a result of project FLOYD (POCI-01-0247-FEDER-045912), funded by the European Regional Development Fund (FEDER), through the Operational Competitiveness and Internationalization Programme (COMPETE 2020) and by Portuguese National Funds (OE), through Fundação para a Ciência e Tecnologia, I.P.; and UIDB/50008/2020, funded by the applicable financial framework (FCT/MCTES) (PIDDAC).

*Corresponding author.

Email addresses: mateus.o.mattos@gmail.com (Mateus Mattos), antonio.rodrig@fe.up.pt (António Rodrigues), rui.meireles@vassar.edu (Rui Meireles), anaa@fe.up.pt (Ana Aguiar)

native to the algorithm built into COTS 802.11ad devices. Geolocation of the two communicating nodes directly influences the angle between communicating antennas, as well as the fading environment surrounding the nodes, dictated by obstacles such as terrain and buildings. In turn, these factors directly influence channel quality, and hence the data rates that can be achieved. We propose using spatially-indexed historical network performance data to select antenna sectors for communication. This reduces the need for sweeping, thus increasing the amount of time available for data transmission. An analysis fed by experimental data traces indicates that this strategy can result in significant throughput gains for some locations.

In summary, we make the following contributions:

- Identify the challenges of collecting data using COTS mm-Wave devices in vehicular scenarios and present a framework to address them (§3.1-3.2).
- Capture time and spatially-indexed 802.11ad frames in a V2I environment, using COTS equipment and the proposed framework (§3.3-3.4).
- Analyze the efficiency of the sector selection algorithm employed by COTS devices (§4).
- Propose and evaluate a new geolocation-based sector selection algorithm for 802.11ad using data from two independent measurement campaigns (§5).

This paper is an extended version of a previous conference article [2]. Relative to it, we add two new contributions. First, we detail the challenges involved in collecting experimental data for vehicular mm-Wave networks using COTS devices, and present the framework we created to address them. Second, we have collected an additional independent dataset under similar conditions, which we use to test the reproducibility of our initial results.

2. Background and related work

802.11ad devices use phased antenna arrays, whose individual antenna gains can be controlled to achieve varying directionality. For two nodes to communicate, the direction of the antenna beam must be adapted to their relative positions — a process known as beamforming. Beamforming is split into two phases [1, 3]. First, both nodes perform Sector-Level Sweeps (SLSes) to choose an initial configuration (i.e., a sector) out of a set of predefined ones. A second, and optional, step — Beam Refinement Protocol (BRP) — can then be used to further fine tune the antenna gains, e.g., to further narrow the transmission and/or reception beams.

Our work focuses on the first phase, depicted in Fig. 2a. The initiator begins by sending a Sector SWeep (SSW) frame over each of the possible sectors, while the other node (known as the responder) listens using an omnidirectional pattern. The roles are then reversed, with the responder performing its sweep. The frames sent by the responder include feedback for the initiator, namely the transmit sector that yielded the strongest received signal. That will be the sector chosen by the initiator for future transmissions. Once the responder finishes, the initiator sends a feedback frame informing the responder of the best sector for it to transmit on. At this point, both initiator and responder have chosen transmit sectors.

The 802.11ad Medium Access Control (MAC) layer divides time into Beacon Intervals (BIs). As depicted in Fig. 2b, each BI is itself divided into a Beacon Header Interval (BHI), used for client association, and a Data Transmission Interval (DTI), used for data transfer [3].

The BHI is further subdivided into three parts. In the first, named Beacon Transmission Interval (BTI), the AP initiates an SLS, which all clients listen to. In the second, named Association Beamforming Training (A-BFT), each client performs its own responder sector sweep. Once this process is complete, the AP knows the best sector to use to reach each client, and each client knows the best sector to use to reach the AP. The third Beacon Header Interval phase, named Announcement Transmission Interval (ATI), is unrelated to sector selection. It is used to schedule transmission slots for the the Data Transmission Interval that follows it, and is optional.

The DTI is comprised of a variable number of scheduled Service Periods (SPs) and traditional 802.11 Contention-Based Access Periods (CBAPs). SLSes can be triggered during the DTI, by any node. The exact conditions that trigger a sector sweep may vary between different 802.11ad devices, i.e., it is implementation dependent. Our analysis

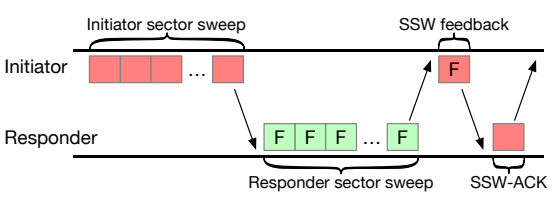

(a) Sector-Level Sweep phases.

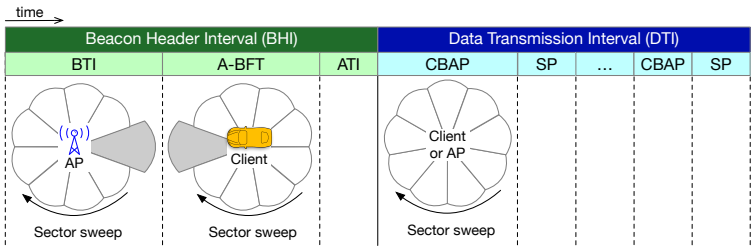

(b) Antenna sector selection timing within Beacon Interval (BI).

Figure 2: 802.11ad's Sector-Level Sweeping.

focuses on SLSes that occur during the DTI phase in a popular COTS device, the Talon AD7200.

The time required to perform a sector-level sweep increases linearly with the number of sectors probed. Early research has focused on reducing time complexity, and hence overhead. Two logarithmic-time strategies have been proposed. First, there is hierarchical beam searching [4]. In it, wider-beam sectors are tried first. Then, the best of them is subdivided into narrower ones, to be tried next. The process is repeated until good directionality is achieved. However, this requires multiple feedback rounds, introducing extra delay that can offset the original gains. Second, there is compressive tracking [5]. It enables beam alignment by having the AP send beacons using pseudo-random phases for each antenna element. Still, it requires nodes capable of measuring Channel State Information (CSI), which is seldom the case for COTS 802.11ad hardware. Rasekh *et al.* [6] propose a compressive tracking alternative that drops that requirement, at the cost of not distinguishing multipath components.

Instead of reducing SLS complexity, we focus on decreasing its frequency with an approach based on statistical historical performance that reduces the need for real-time measurements.

Applying 802.11ad to vehicular communication is challenging. Experimental studies [7, 8, 9] have reported short communication ranges ($\leq 20\,\mathrm{m}$) and frequent disconnections, highlighting the need for improved beam steering.

Focusing on V2I, Loch *et al.* [10] eschewed dynamic beamforming entirely, opting instead for fixed beam geometry determined by the relative orientation between the AP and the road it is on. Although effective, this strategy limits connection duration. Since most 802.11ad devices also support legacy 2.4/5 GHz Wi-Fi, Nitsche *et al.* [11] proposed using these lower-frequency, longer-range bands to perform out-of-band beam steering. Choi *et al.* [12] explored an 802.11p-based out-of-band solution for relative position and trajectory estimation in a vehicular setting. However, their evaluation was limited to simulations. In contrast, Muns *et al.* [13] implemented radar functionality in the 802.11ad band. This lets nodes estimate their relative positions and reduce the number of candidate sectors that need to be swept. Evaluation was simulation-based, and the strategy was shown to loose some of its effectiveness in strong multipath environments.

3. Experimental data collection

mm-Wave experimentation in vehicular environments is challenging. Data is produced at very high rates, logs need to be synchronised across devices, and high accuracy positioning is advisable. In this section, we describe the data collection framework that we created for this purpose.

3.1. Data collection framework requirements

Our goals were to analyze: i) the performance resulting from the use of different antenna sectors, for each location that a moving vehicle passes by, and ii) the efficiency and effectiveness of 802.11ad COTS devices' antenna sector selection algorithm. With this in mind we defined the following requirements for a data collection framework. First, it must be able to generate traffic at a fast enough rate to saturate the network. Second, it must be able to collect:

- **Network performance data:** For each data frame, we must collect its capture timestamp, length, transmission data rate, and, if received by its intended destination, received signal strength.
- **Antenna sector sweep and selection data:** For each SLS-related control frame, we must collect a capture timestamp, frame type (probe, feedback or feedback acknowledgment), transmission sector, and, for feedback frames, the chosen sector and its SNR.

Third, each piece of data must be annotated with:

- A geographical location, i.e., GPS coordinates.
- A timestamp, to allow the matching of data collected by different nodes. In order to do this, time must be synchronized across nodes.

Finally, the framework should provide a way to start and stop data collection in all nodes, as well as monitor its progress.

3.2. Data collection framework design

Fig. 3 depicts the architecture of the framework we created to fulfill the above requirements. There are two main nodes, a sender and a receiver, each comprised of multiple devices. The sender acts as the AP, and the receiver as the client. The sender contains a data producer application that continuously generates pseudo-random data. This data is then encapsulated in UDP segments and sent to the receiver over the 802.11ad link. UDP transport was used to avoid any potential influence from TCP's flow and congestion control schemes.

The router is connected to the data-producing computer through Gigabit Ethernet. The receiver side contains a complementary computer and application that consume the data. Given that 802.11ad supports multi-gigabit data rates, the Gigabit Ethernet connection constitutes a performance bottleneck. This compromise was made for two reasons. First, it reduces the load on the router's CPU, freeing it up for lower-level communication related tasks. Second, there are no publicly-available tools to build applications to run on the device's factory firmware, and the open source alternative, OpenWRT/LEDE, proved unstable when running in AP mode.

We now focus on how frames are captured. Normally, Wi-Fi control frames, such as those involved in SLSes, are entirely handled by the network interface and hence not visible to user-level processes running on the device. This behavior can be changed by setting the wireless interface

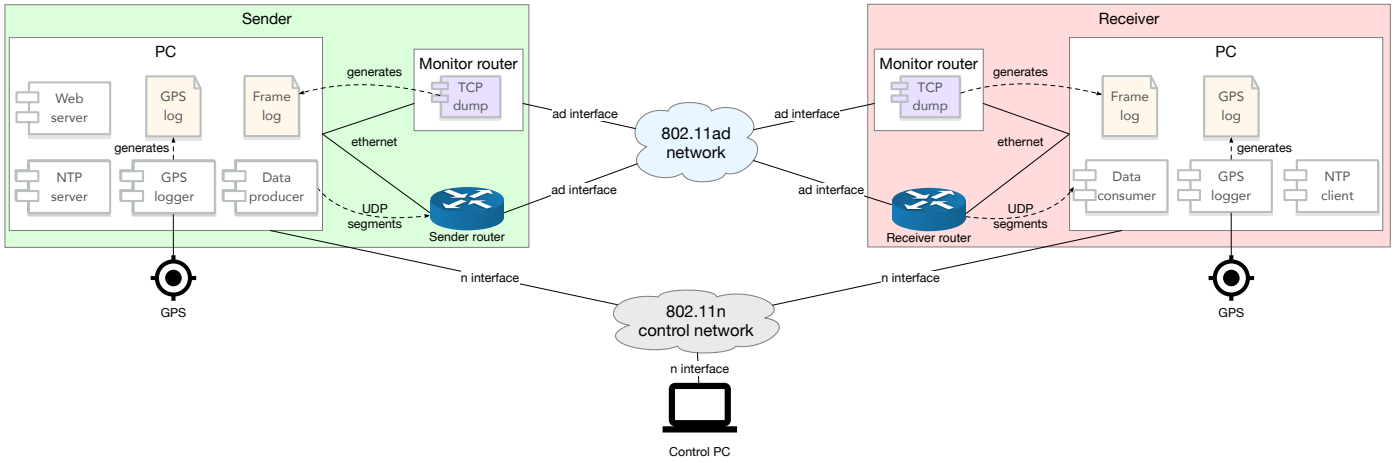

Figure 3: Data collection framework - physical architecture.

to promiscuous mode, which lets all frames through. However, when so configured, an interface can not be used for regular communication. As a consequence, the interface used to capture frames must be different from the one used for communication.

For this reason, our framework prescribes the use of monitor interfaces for frame capturing. To minimize the channel variation between interfaces, the monitor is to be placed as close to the communicating interface as possible. As shown in Fig. 3, ideally there should be one monitor for the sender and another for the receiver. Thus, each node consists of 3 devices: two 802.11ad routers, one for communication and another for monitoring, and one Personal Computer (PC). We used TP-Link Talon AD7200 802.11ad routers. The monitors run a Talon-specific firmware from the Talon Tools project [14]. This let us build and run `tcpdump` to capture frames, which is not possible on the factory firmware. Due to the routers' storage limitations, the capture logs are stored on the PC's drive, which is mounted using the Network File System (NFS) protocol.

In order to be able to annotate the captured frames with the geographical location they were sent or received from, each node is equipped with a GPS receiver. A logger application uses it to create a timestamp-indexed log of mobility-related data (latitude, longitude, altitude, speed, heading, etc). After the data is collected, a post-processing script uses each frame's timestamp to associate it with the last known GPS coordinates at the time the frame is recorded, i.e., no interpolation is used.

Since the analyses to be performed use timestamps to associate information captured by different nodes, it is important for clocks to be synchronized. The framework does this by running the Network Time Protocol (NTP) over a separate 802.11n network. This network is also used to control and monitor the collection process. The sender node gathers status information and makes it available through a web server that a control computer connects to. The computer also directs the start and end of the data collection process, which is automated through the execution of shell scripts. 802.11n's long range, compared with 802.11ad, makes it a good choice for this purpose, as nodes are able to travel hundreds of meters while remaining connected to the control network.

Note that this framework can scale up to support parallel data collection for multiple Wi-Fi standards (as done in [8]), by adding more data producer and consumer instances, along with the respective communication and monitor interfaces.

The code and scripts used to collect and process data using this framework, along with usage instructions, are available on GitHub [15].

3.3. Accuracy of GPS measurements

Our study involves the association of 802.11ad network performance data with the position of a moving vehicle over time. Since 802.11ad uses millimeter wavelengths its performance is sensitive to very minute position changes. As such, we require an highly accurate source of geographical location data: the greater the geolocation accuracy, the more accurate the downstream analysis can be.

One option is to use a high-accuracy GPS receiver such as the Trimble Pro Series 6H [18], which nominally achieves 50 cm precision for speeds up to 50 km/h. However, since such high-precision GPS devices are expensive and bulky, it may be more practical to use lower-accuracy receivers, such as those embedded in COTS smartphones. In order to better understand the real-world differences between these two classes of GPS devices, we assessed the accuracy of two COTS smartphones:

- LG G6 (launched in February 2017);
- Google Nexus 4 (launched in May 2013).

Our assessment was performed using the mobility patterns and suburban environment later described in §3.4. Fig. 4 shows the results, namely the positioning error in

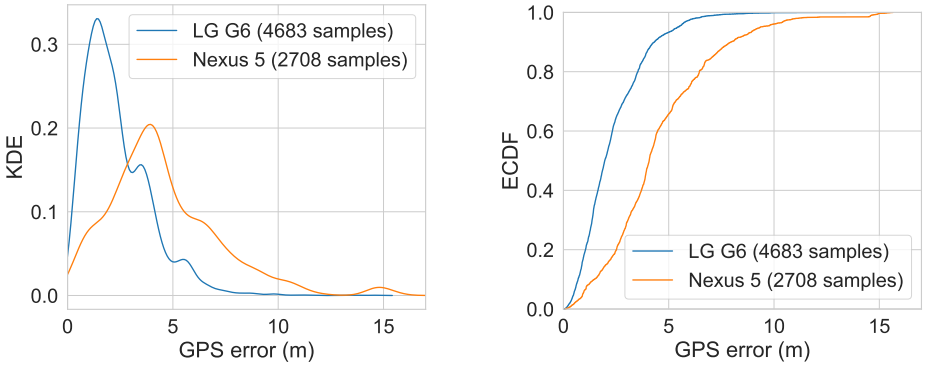

(a) Kernel Density Estimate (KDE) of GPS errors, in meters.

(b) Empirical Cumulative Distribution (ECDF) of GPS errors, in meters.

Figure 4: Smartphone GPS error distribution using using a Trimble 6H Pro Series as reference.

meters, taking the output of a Trimble 6H Pro Series high-precision GPS device as ground truth. The LG G6 exhibited a lower error (on average 2.33 m) when compared to the Nexus 4 (on average 4.54 m). We can thus conclude that smartphones can be used for geographical analyses with a granularity coarser than a couple of meters.

3.4. Data collection procedure

We now describe the scenario in which the data collection was performed. The setup, depicted in Fig. 5, was similar to the one used in [7]. The AP was placed on top of a vehicle parked at a corner of an intersection. A similarly equipped client vehicle then drove around the intersection while trying to download data from the AP.

We ran two sets of experiments using the same basic overall setup. One in July 2020 and another in October 2021, 15 months later. Topographical features stayed similar during this period, i.e., no new buildings were constructed in the immediate vicinity. Henceforth, the two sets of data shall be referred to as the *2020 dataset* and the *2021 dataset*.

The two datasets vary slightly in terms of mobility. In 2020, the client drove a circuit around the intersection, approaching it from all possible directions, at normal speeds for the road in question (i.e., up to 14 m/s, or 50 km/h).

In 2021 we used the same mobility pattern as in 2020, but added a second one for variety. In it, the client moved slowly (i.e., below 2 m/s, or 7 km/h) in a specific direction, either away or towards the AP. Fig. 5a depicts the eight different trajectories used. The low speeds enabled the collection of more spatially-dense data samples.

In all experiments, the AP sent pseudo-random data towards the client using UDP, at an application-level constant rate of ∼420 Mbps. 802.11ad frames were captured by a monitor node co-located with the mobile client, as per Fig. 5b. Due to lack of hardware, an AP-side monitor was not present. Geopositioning was provided by a high-accuracy Trimble Pro Series 6H [18] GPS receiver. Additional procedural details can be found in [19].

3.5. Collected data

We collected the following data:

- A 1 s-resolution client geolocation log.

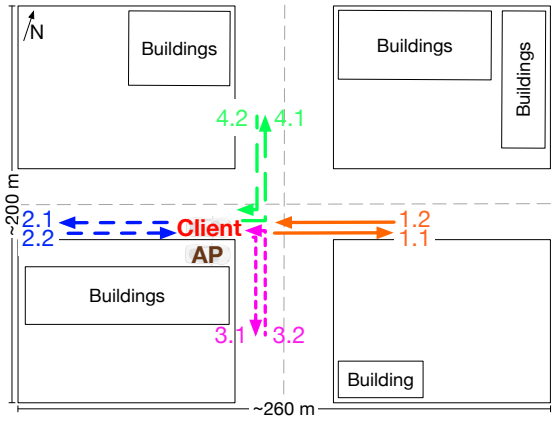

(a) Experiment environment and client mobility patterns.

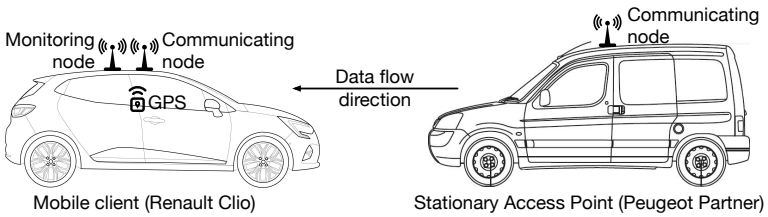

(b) Vehicles (to scale) and devices mounted on them. Application data was streamed from the parked AP to the mobile client. Vehicle blueprints courtesy of Renault UK[16] and BlueprintBox [17].

Figure 5: Experimental setup. The experiments took place at a residential-area intersection (coordinates 41.111929, -8.631083).

- A 1 ns-resolution timestamped capture of all frames received by the monitor node, which is co-located with the client node. This includes both data and control frames, such as the ones used to perform SLSes.
- A 1 s-resolution timestamped log of SLSes received by the client node, i.e., those performed to select the AP's transmit sector. This log includes the SNR observed for each sector at the client.
- Application-level throughput information from both the AP and client nodes.

The collection setup caused the following limitations:

- Since we had a single monitor node, co-located with the client, we were unable to capture all the frames sent by the AP, as some will inevitably have been lost on the way to the monitor. This makes it impossible to determine exactly how many frames were lost and how much time the AP spent trying to send data. We investigated the use of layer-2 frame sequence numbers to quantify losses, but their small 12-bit size means they roll over quickly — in the collected data, 75 % of sequence number *"wraparounds"* occurred within 500 ms or less. Therefore, accurately determining losses through sequence number analysis proved unfeasible.
- We were unable to match the SLSes captured by the monitor node with the SLS SNR records captured by the client. As such, we were not able to analyze the relationship between observed SNR and chosen sector — we expect the highest-SNR sector to be the one chosen, but could not confirm it. The reasons no match was possible are twofold: (i) the sets of SLSes captured by client and monitor are not the exact same, rendering a reception order-based matching unfeasible; and (ii) the monitor records per-frame timestamps at the kernel level, while the client records a single timestamp per SLS, at the user level, which introduces variable delay. Given the high frequency of SLSes, this proved to be a significant issue.
- With the tools available, i.e., `tcpdump`, we were unable to extract per-frame SNRs. Hence, we could not compare the distribution of data frame SNRs across sectors.

These limitations impacted the evaluation of our proposed antenna sector selection strategy, detailed in §5.

4. Default sector selection strategy analysis

This section takes the 2020 dataset described in §3 and analyzes the sector selections made by the COTS devices.

Initiator	#SLSes	#Sector selections			
		By Access Point		By client	
		Total	Inconseq.	Total	Inconseq.
AP	115,542	41,644	83%	34,048	66%
Client	15,414	5,280	81%	2,737	52%
All	130,956	46,924	82%	36,785	65%

Table 1: Sector-Level Sweep (SLS) statistics (2020 dataset).

The sector-level sweeps captured in this dataset are summarized in Tab. 1.

A successful SLS results in transmit sectors being selected for both initiator and responder (§2). The chosen sectors are included in the frames carrying feedback, i.e., the frames marked "F" in Fig. 2a. A significant portion of SLSes in our dataset were incomplete, meaning no feedback from the responder was captured — only the initiator's sweep.

cSome incomplete SLSes can be attributed to insufficient link budget. However, it is more difficult to explain why, frequently, the monitor was able to hear an SLS initiated by the AP, but failed to hear any feedback from the client, which was right next to it. We analyzed the correlation of this phenomenon with multiple variables, such as client location and speed, in an attempt to understand it, but results were inconclusive.

Consider the successful SLSes, i.e., the ones ending with a sector selection. Within these, we noticed two issues: (i) a large number of SLSes that do not trigger a sector change, and (ii) switching back and forth, or *"ping-ponging"*, between a pair of sectors.

Inconsequential SLSes: When an SLS does not trigger a sector change, i.e., the node elects to remain on the current sector, we call it *"inconsequential"*. As Tab. 1 shows, most SLSes were inconsequential, specially for the AP (82 % versus 65 % for the client).

The table excludes periodic SLSes done by the AP during the Beacon Header Interval to facilitate the discovery of new clients. As such, the higher percentage of inconsequential SLSes for the AP can potentially be explained by its physical installation. Since the AP was parked on the roadside, next to a building, the range of angles it could use for communicating with the client was limited. The client moved in multiple different directions (Fig. 5a), and was thus able to utilize a larger range of angles. The sector frequency histogram of Fig. 6a agrees with this. The Talon AD7200 devices feature a total of 64 antenna sectors, numbered 0 through 63. Sectors 20 and 24 alone accounted for ∼60 % of the AP's time. In contrast, the client's sector usage distribution was more uniform.

Interestingly, neither node used sectors 32 through 58, which indicates that the devices did not find their performance good for the range of transmitter-receiver angles involved in the experiments. Other sources have reported similar observations. Namely, Steinmetzer *et al.* [20] ran an extensive set of stationary experiments with the same Talon routers, and reported that sectors 32 through 60 were not used for transmission.

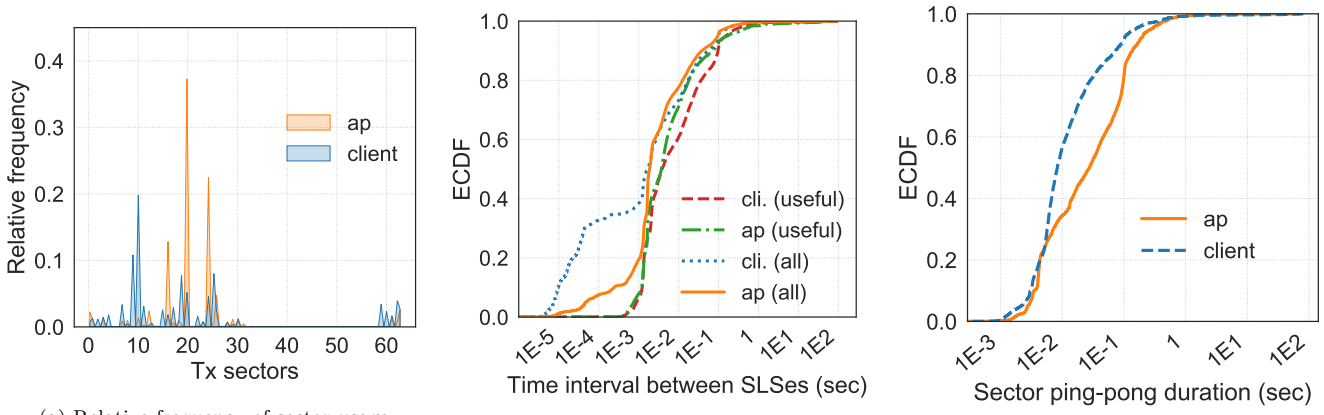

(a) Relative frequency of sector usage.

(b) Time between sector-level sweeps.

(c) Time between sector repetitions.

Figure 6: COTS devices' sector selection analysis results.

Fig. 6b shows the Empirical Cumulative Distribution (ECDF) for the time elapsed between consecutive sector sweeps, for all sweeps, and for consequential sweeps alone. If we focus on the client's sector choices, we can see that ∼40 % of all sweeps were separated by 1 ms or less. But only ∼6 % of consequential sweeps were that close. This tells us that many inconsequential sweeps, in fact around 40 % of them, were performed very soon after another sweep. To a lesser extent, the same phenomenon can be observed in the AP's SLSes.

Digging further, we discovered that 62 % of all inconsequential SLSes occurred when the client was less than 2 m away from the AP, and moving at a speed below 1 m/s. We could not find a clear reason for them, and since they were inconsequential, they were effectively unnecessary.

The median interval between consequential sweeps was around 5 ms for both client and AP. Only a residual amount of consecutive sweeps were separated by more than 500 ms.

We hypothesized that increased vehicle speed would lead to more frequent sector sweeping. However, when we plotted vehicle speed — both ground speed and angular, relative to the AP — versus SLS frequency we found no correlation between them.

The large number of inconsequential SLSes lead us to conclude that the sweeping strategy employed by the Talon routers was inefficient in the studied vehicular environment. The number of SLSes could be reduced significantly without loss in performance.

Sector *"ping-ponging"*: Within the consequential sweeps, we saw a pattern where a node would alternate between a pair of sectors, e.g., between sectors 20 and 24. We call this a *"ping-pong"*. To quantify it, we took each triplet of consecutive consequential sector selections — $\{s_1, s_2, s_3\}$ — and verified how often $s_1 = s_3$. *"Ping-ponging"* occurred in 5,431 out of 21,614, or 25.1 %, of triplets containing sector switches.

Fig. 6c shows the ECDF of the time elapsed between leaving and returning to a sector. The AP reverted 35 % of its sector changes back to the original sector within 10 ms. The client was even more extreme, reverting 60 % of its sector changes within the same period of time. Therefore, the effective utility of a good portion of *"ping-pongs"* seems limited. Just as with inconsequential sweeps, this is inefficient — the time could have been better used for data transmission.

5. Geolocation-based 802.11ad sector selection

5.1. Proposed strategy

Our ultimate goal is to maximize the amount of transferable data between network nodes. This can be accomplished by increasing two factors: (i) the data rate at which communication can occur, and (ii) the amount of time available for communication. Our proposal operates on both of these fronts.

In V2I communication geolocation determines antenna angles and majorly affects the characteristics of the fading environment, as it codifies the presence or absence of persistent obstacles such as terrain and buildings. With this in mind we investigate the possibility of using geolocation to select antenna sectors for 802.11ad communication. The idea is to use spatially-indexed historical network performance measurements to select the sector that has been shown to, statistically, yield the highest data rate. This will make more efficient use of the time available. Further, it reduces the need for sector-level sweeps, thus decreasing overhead and increasing the amount of time that can be used for communication.

The key assumption underlying this strategy is the stationarity of the distribution of achievable data rates as a function of $\langle sector, geolocation \rangle$ combination. This means that the likelihood of a sector being able to achieve a given data rate from a given location is assumed not to change with time.

The proposed strategy boils down to the following steps:

1. Discretize space into square cells, e.g., 1 × 1m, to allow for geolocation-based performance data aggregation.

2. Create a record of what sectors are used when the client is in each spatial cell, for how long, and the data rates and SNR levels they achieved.

3. For each spatial cell c, use the historical record to choose the statistically-best antenna sector to use when the client is in that particular cell.

Following, we describe how we evaluated the effect of geolocation-based antenna sector selection, what sector selection metrics we considered, and the results of their application.

5.2. Evaluation methodology

The focus of our evaluation was throughput. We compared the amount of data that was communicated in the datasets from §3 with the amount of data that could be transferred if a single sector, considered to be the best according to some metric (detailed later), was used for each cell.

We considered two variants, one where the frequency of sector level sweeps is unchanged, and one where they are eliminated, freeing extra time for data transmission. We begin with the variant that maintains SLSes unchanged, meaning the available communication time is unchanged as well. For each cell c, we:

1. Found the amount of data successfully communicated during the time the client spent in cell c: $D_{c,og}$.

2. Computed how long it took to transmit $D_{c,og}$ Bytes: Tx_c.

3. Estimated how much data could have been sent in Tx_c seconds if the sector deemed best, S, had been used during the entire transmission period: $D_{c,S}$.

4. Compared $D_{c,S}$ with $D_{c,og}$.

Let $\{f_1, f_2, ..., f_n\}$ be the frames exchanged with the client in cell c. The amount of data sent from/to c, $D_{c,og}$, is the sum of the frames' sizes. The total transmission time, Tx_c, is the sum of each frame's size divided by the rate it was sent at:

$$D_{c,og} = \sum_{i=1}^{n} size(f_i) \quad \text{and} \quad Tx_c = \sum_{i=1}^{n} \frac{size(f_i)}{rate(f_i)}. \quad (1)$$

Let $R_{c,S} = \{r_1, r_2, ..., r_m\}$ be the data rates that were used, historically, to send frames from/to cell c using sector S, the one chosen for the cell. The amount of data that can be sent using sector S in Tx_c seconds is the product of Tx_c by the data rate expected from S. The latter can be estimated as the average of rates $R_{c,S}$, weighted by the portion of time spent using each rate. We use time as a weight factor, rather than the amount of data sent, because we want to compute the expected data rate over a period of time. We can write:

$$D_{c,S} = Tx_c \times \bar{R}_{c,S} \quad \text{, with} \quad \bar{R}_{c,S} = \sum_{i=1}^{m} \frac{Tx_{c,S,r_i}}{Tx_{c,S}} \times r_i. \quad (2)$$

$Tx_{c,S}$ represents the time communicating using sector S in cell c, and Tx_{c,S,r_i}, the time spent using rate r_i under the same conditions. Both are computed similarly to Eq. 1, but including only the frames of interest.

We note this $D_{c,S}$ definition assumes that the time spent sending data at one rate can be used to send data at a different rate. Since each frame incurs an overhead in the form of a preamble and inter-frame spacing, this means we are assuming that the total number of frames remains constant regardless of rate. For that to be true, the frame size must be increased for higher data rates, as the higher the rate, the more data can be sent in a given amount of time. If this assumption does not hold, the communication time will be slightly overestimated.

Having defined the amount of data that can be sent for each (cell c, sector S) combination, $D_{c,S}$, the throughput gain expected from always using sector S for cell c can be computed as:

$$TG_{c,S} = \frac{D_{c,S} - D_{c,og}}{D_{c,og}}. \quad (3)$$

If sector selection is geolocation-based, we can also eliminate sector sweeps and use the extra time, $Tsls_c$, for data transmission. However, insufficient link budget may preclude data transfer during that period. In fact, one motivation for sector sweeps is trying to increase the link budget to acceptable levels. Hence, to be conservative, we scale $Tsls_c$ by the ratio of time spent successfully communicating data to the total amount of time spent in the cell, T_c. The total amount of transferable data when SLSes are disabled, $Dno_{c,S}$, is then:

$$Dno_{c,S} = D_{c,S} + Tsls_c \times \frac{Tx_c}{T_c} \times \bar{R}_{c,sm}. \quad (4)$$

Finally, $Dno_{c,S}$ can be compared with $D_{c,og}$ to find the throughput gain, similarly to Eq. 3.

5.3. Sector selection metrics

Geolocation-based sector selection hinges on the chosen sector being able to be used consistently for a given cell. If, in the dataset, a sector was only used in a cell for a very brief period, it is a leap to assume it will be usable in general. We consider these sectors to be outliers. In order to prevent them from being selected, we employed a Median Absolute Deviation (MAD) outlier detection scheme [21]. In short, we took the sector usage times for each cell, and dropped from contention any sector whose usage time was less than the median sector usage time by more than two times the average deviation from the median.

Once the MAD filter was applied, we considered three distinct strategies to pick the best sector for a given cell c:

Random: Pick a sector uniformly at random from the set of sectors that were able to communicate when the client was in cell c. Because this is a non-deterministic process, we repeated it 10,000 times per cell and averaged out the throughput results. The performance

of this strategy represents a lower bound by which others can be measured.

Median SNR: Pick the sector with the highest median SNR. If $SNR_{c,s}$ is the set of recorded SNR values for a sector s in cell c, the selected sector S_c will be:

$$S_c = \underset{s}{\text{argmax}}\ median(SNR_{c,s}). \quad (5)$$

Our dataset only contains SNR samples for sector sweep frames, so those are the ones used.

The use of SNR is motivated by its common use as a channel quality estimator and strong correlation with achievable data rates. We used the median and not the mean because the former is less sensitive to outliers.

Optimal: Since the goal is to maximize the amount of data transferred over a fixed period of time, the optimal strategy is to choose the sector that maximizes the mean data rate weighted by time. For a cell c, sector s combination, this is $\bar{R}_{c,s}$ from Eq. 2. Therefore, we choose sector:

$$S_c = \underset{s}{\text{argmax}}\ \bar{R}_{c,s}. \quad (6)$$

The performance of this metric represents an upper bound by which others can be judged.

5.4. Evaluation results

We analyze the change in throughput that can potentially be derived from geolocation-based sector selection on a per-cell basis. Since in our experiments data was downloaded from the AP to the client, we focus on the selection of the AP's transmit antenna sector. We did this for both the 2020 and 2021 datasets.

Let us first focus on the 2020 data. Fig. 7a shows the ECDFs of the per-cell throughput gain for the different selection metrics, with 1×1 m cells, and SLSes not eliminated, i.e., their time was not reclaimed for data transmission. Under such conditions, additional throughput can only be realized through an increase in mean data rate.

The optimal strategy performed markedly better than the others. Still, while there were no instances of decreased throughput, almost half the cells experienced no increase either. This indicates that, often, the combination of sectors used by the off-the-shelf beamforming algorithm could not be beaten by any single sector, and that therefore the algorithm performed well.

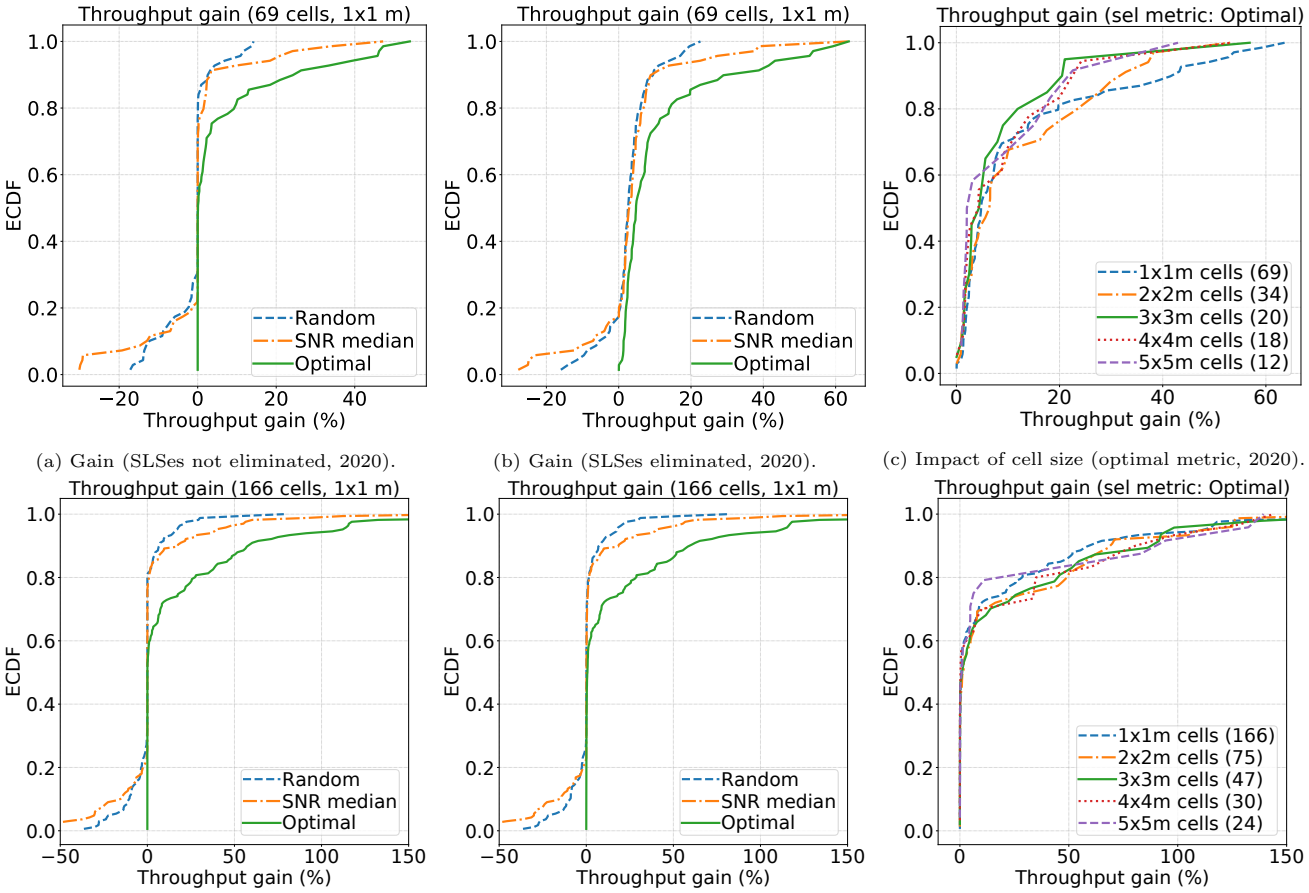

(a) Gain (SLSes not eliminated, 2020).
(b) Gain (SLSes eliminated, 2020).
(c) Impact of cell size (optimal metric, 2020).
(d) Gain (SLSes not eliminated, 2021).
(e) Gain (SLSes eliminated, 2021).
(f) Impact of cell size (optimal metric, 2021).

Figure 7: Geolocation-based sector selection per-cell throughput gain ECDFs, for both 2020 and 2021 datasets.

However, the distribution has a long tail, with 20 % of cells having increases of over 10 %, and a handful over 40 %. Thus, there is, under some circumstances, significant room for improvement over the off-the-shelf algorithm.

In the long run, the data rate achieved by the random sector selection metric is the average of all sector data rates, since they are all equally likely. This translated into a decrease in throughput for 40 % of cells, no change for 44 %, and a small increase for the remaining 16 %.

Interestingly, while the SNR median metric resulted in larger gains at the upper end of the distribution relative to the random metric, it also resulted in larger losses at the lower end. This leads us to believe that SNR samples taken from sector-level sweeps do not correlate as well as expected with the data rates used to send data frames.

Fig. 8 shows the relative frequency with which the different antenna sectors were chosen, for both the optimal and the SNR median metrics. Although they are quite similar, the distribution for the optimal metric skews more towards lower-numbered sectors (20 and below), relative to the SNR median one. It is also interesting to note that while sector 20 was the most popular AP transmit sector for the off-the-shelf algorithm (Fig. 6a), the geolocation-based schemes picked sector 16 most frequently. This tells us that, statistically, sector 16 yielded larger mean data rates more often than any other sector.

Consider now the same throughout gain, but for the 2021 dataset, which is shown in Fig. 7d. The overall shape of the distributions is very similar. The optimal scheme significantly outperformed the others. The percentage of cells for which significant gains can be realized is around 30 %, which is consistent with the 2020 results. The only real difference is that the 2021 distribution features a longer tail, with a small number of cells experiencing gains of over 100 %. Looking at the datasets in detail, the mobile client moved more slowly in the 2021 dataset. We hypothesize this may have resulted in better alignment and consequently larger differences between good and a bad sector choices.

If we leverage the fact that sector selection is based on geolocation to eliminate SLSes, additional time is freed up for communication. Fig. 9 quantifies this gain as a percentage of the original transmission time for the 2020 dataset. The median time gain was small: 3.6 %. The maximum was 15 %. 2021 results are similar.

Using this additional time for data transfer slightly increases throughput across the board, as depicted by Figs. 7b (2020) and 7e (2021).

Fig. 7c shows the optimal strategy's throughput gain for different cell sizes in the 2020 dataset, always assuming SLSes are eliminated. Although for most cells the difference is not large, cell size appears to be negatively correlated with throughput gain, i.e., larger cells lead to worse performance. The gap was largest at the distribution's tail. For example, the maximum observed gain was 63 % for 1×1 m cells, but only 43 % for 5×5 m cells.

Fig. 7f shows cell-size sensitivity for the 2021 dataset. As shown, this data set is less sensitive to cell size than the 2020 one. We have not found an explanation for this.

Finally, Fig. 10 shows the throughput gain as a function of the client's position relative to the AP, for both the 2020 and 2021 datasets. To allow for an easier comparison across datasets, cells are colored relative to the maximum throughput gain obtained within that particular dataset. For instance, the dark red cells represent the highest gains, regardless of what the actual absolute gain value was.

As expected, cells further away from the AP had smaller gains. If we focus on the optimal-metric results however, another pattern emerges. Namely, the cells aligned with the side and rear of the AP experienced the highest gains. This was consistent across datasets, despite the 2021 one having a much larger total number of cells.

The results presented so far were obtained using each dataset for both training and test. We also tried training on the 2020 dataset and testing on the 2021 one, but throughput performance was poor. Although the surrounding topography and communication equipment did not change significantly between the two years, the AP's location and orientation were not exactly the same. We hypothesize that this caused the 2020 data to not be capable

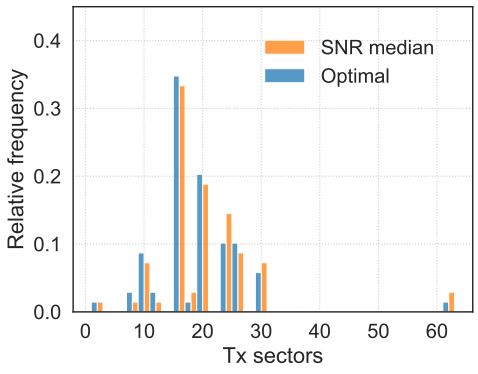

Figure 8: Geolocation-based sector selection frequency (2020 dataset).

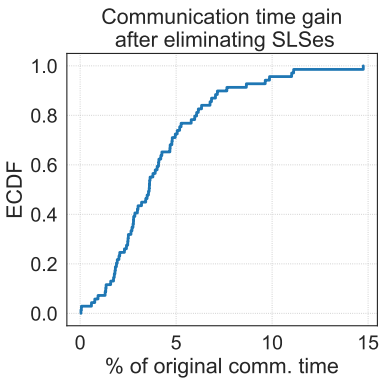

Figure 9: Time gained by eliminating SLSes, relative to the original transmission time (2020 dataset).

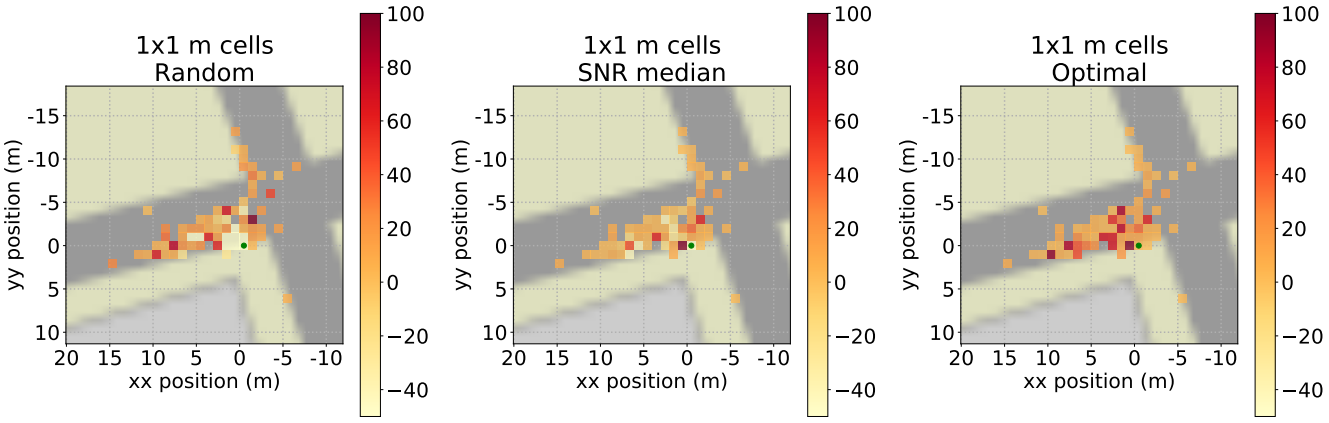

(a) 2020 dataset.

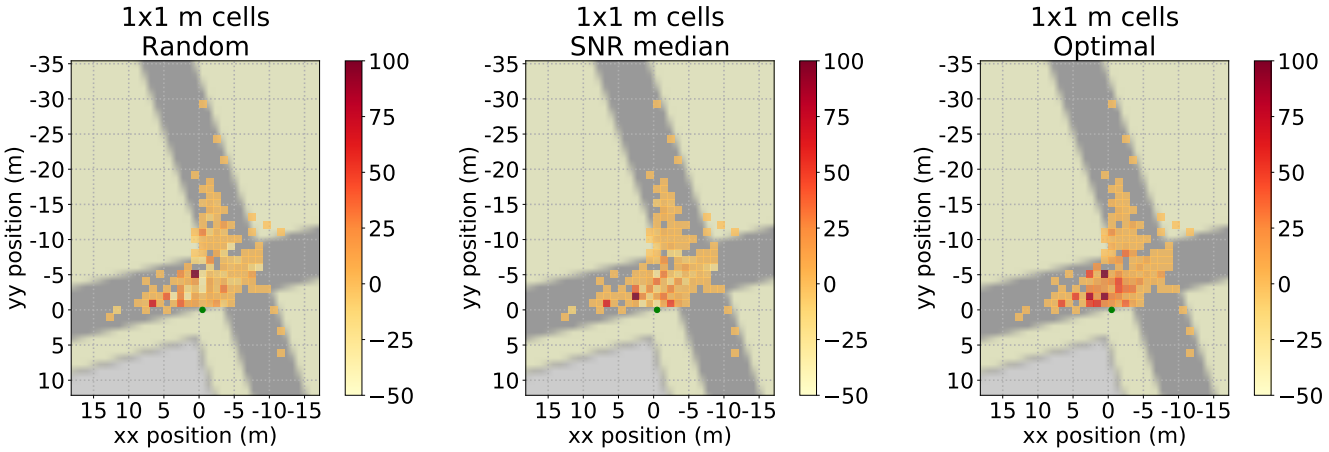

(b) 2021 dataset.

Figure 10: Geolocation-based sector selection throughput gain by spatial cell. Cell locations are shown relative to the AP's position, which is the origin of the coordinate system. The *xx* and *yy* axes represent the west-east and south-north directions, respectively. The plots are oriented with north at the top, south at the bottom, west to the left, and east to the right.

of predicting performance in the 2021 dataset. Due to the short wavelength being used, even slight repositioning of the AP requires new training.

We were however successful in splitting the 2020 dataset into two halves, and using one for training and another for testing. We did this by, for each cell, splitting the data by their timestamp. The earlier half was used for training, and the later half for testing. Fig. 11 shows the resulting throughput gain ECDF for different cell sizes, assuming the optimal metric is used and SLSes are eliminated. Results are largely similar to those in Fig. 7c, which reflects the use of the entire 2020 dataset for both training and testing. Throughput gains of at least 10 % can be observed for roughly 40 % of cells. Interestingly, the negative correlation between cell size and throughput observed earlier is not seen here. Also, the gain distribution has a longer tail, with a small number of cells seeing their throughput more than double.

Overall, these results validate our assumption that, if conditions are similar, the distribution of achievable data rates as a function of ⟨*sector*, *geolocation*⟩ combination persists through time, allowing prior measurements to in-

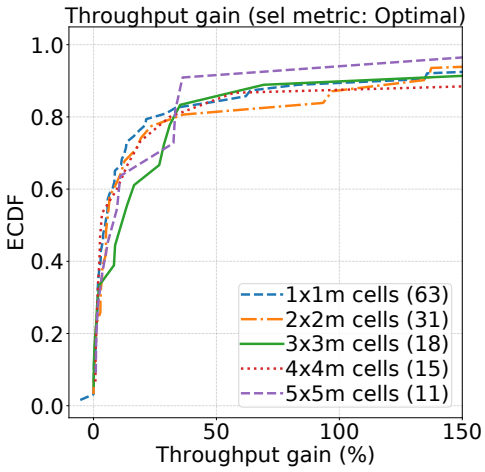

Figure 11: Geolocation-based sector selection per-cell throughput gain (optimal metric, 2020 split dataset).

form future antenna sector decisions.

6. Conclusions

In this work we focused on antenna sector selection for 802.11ad communication in a vehicular context. We presented a data collection framework and then used it to collect data regarding the beamforming behavior of 802.11ad COTS devices in a V2I communication scenario. Analysis of these data uncovered efficiency issues. Namely, that only a minority of SLSes resulted in an actual sector change, and that a significant portion of those changes caused *"ping-ponging"* between sectors.

We investigated how a geolocation-based sector selection scheme that picks the antenna sector that statistically performs the best for a given location can help increase performance, in such a scenario. The results of our trace-based evaluation show that significant gains of 10 % or more are possible for around 30 % of spatial cells. These results were observed in two separate datasets, collected more than one year apart. The gains derive from the use of higher data rates and the extra communication time gained by avoiding unnecessary sector sweeps. However, for most cells gains were minimal, and the observed communication range was short, which hints at the difficulty of using COTS 802.11ad devices in a V2I context.

In the future we would like to explore the practical feasibility of the proposed sector selection strategy through prototyping. Given that fine-grained spatially-indexed data is used for decision making, scalability is one of the concerns. We believe it can be addressed by having APs store the historical sector performance data, rather than clients. Since APs are fixed, this would limit the area of interest and consequently the amount of storage needed.

References

[1] 802.11 Working Group, IEEE 802.11ad-2012, Standard, IEEE (2012).
URL https://standards.ieee.org/standard/802_11ad-2012.html

[2] M. Mattos, A. Rodrigues, R. Meireles, A. Aguiar, Geolocation-based sector selection for vehicle-to-infrastructure 802.11ad communication, in: 2021 19th Mediterranean Communication and Computer Networking Conference (MedComNet), 2021, pp. 1–7. doi:10.1109/MedComNet52149.2021.9501244.

[3] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, J. C. Widmer, IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi, IEEE Communications Magazine 52 (12) (2014) 132–141. doi:10.1109/MCOM.2014.6979964.

[4] S. Hur, T. Kim, D. J. Love, J. V. Krogmeier, T. A. Thomas, A. Ghosh, Millimeter wave beamforming for wireless backhaul and access in small cell networks, IEEE Transactions on Communications 61 (10) (2013) 4391–4403. doi:10.1109/TCOMM.2013.090513.120848.

[5] Z. Marzi, D. Ramasamy, U. Madhow, Compressive channel estimation and tracking for large arrays in mm-wave picocells, IEEE Journal of Selected Topics in Signal Processing 10 (3) (2016) 514–527. doi:10.1109/JSTSP.2016.2520899.

[6] M. E. Rasekh, Z. Marzi, Y. Zhu, U. Madhow, H. Zheng, Noncoherent MmWave Path Tracking, in: 18th International Workshop on Mobile Computing Systems and Applications, HotMobile '17, ACM, 2017, p. 13–18. doi:10.1145/3032970.3032974.

[7] A. Aguiar, A. Rodrigues, M. Mattos, B. Maurício, R. Meireles, P. Steenkiste, Exploratory Study of 802.11ad Vehicular Links, in: 2020 Mediterranean Communication and Computer Networking Conference (MedComNet), 2020, pp. 1–4. doi:10.1109/MedComNet49392.2020.9191643.

[8] R. Meireles, A. Rodrigues, A. Stanciu, A. Aguiar, P. Steenkiste, Exploring Wi-Fi Network Diversity for Vehicle-To-Infrastructure Communication, in: 2020 IEEE Vehicular Networking Conference (VNC), 2020, pp. 1–8. doi:10.1109/VNC51378.2020.9318407.

[9] W. Kim, Experimental Demonstration of MmWave Vehicle-to-Vehicle Communications Using IEEE 802.11ad, Sensors 19 (9) (2019) 2057. doi:10.3390/s19092057.

[10] A. Loch, A. Asadi, G. H. Sim, J. Widmer, M. Hollick, mm-Wave on wheels: Practical 60 GHz vehicular communication without beam training, in: 2017 9th International Conference on Communication Systems and Networks (COMSNETS), 2017, pp. 1–8. doi:10.1109/COMSNETS.2017.7945351.

[11] T. Nitsche, A. B. Flores, E. W. Knightly, J. Widmer, Steering with eyes closed: Mm-wave beam steering without in-band measurement, in: 2015 IEEE Conference on Computer Communications (INFOCOM), 2015, pp. 2416–2424. doi:10.1109/INFOCOM.2015.7218630.

[12] J. Choi, V. Va, N. Gonzalez-Prelcic, R. Daniels, C. R. Bhat, R. W. Heath, Millimeter-Wave Vehicular Communication to Support Massive Automotive Sensing, IEEE Communications Magazine 54 (12) (2016) 160–167. doi:10.1109/MCOM.2016.1600071CM.

[13] G. R. Muns, K. V. Mishra, C. B. Guerra, Y. C. Eldar, K. R. Chowdhury, Beam Alignment and Tracking for Autonomous Vehicular Communication using IEEE 802.11ad-based Radar, in: IEEE INFOCOM 2019 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), 2019, pp. 535–540. doi:10.1109/INFCOMW.2019.8845121.

[14] D. Steinmetzer, D. Wegemer, M. Hollick, Talon tools: The framework for practical ieee 802.11ad research, (accessed on 2022-05-30) (2017).
URL https://seemoo.de/talon-tools

[15] A. Rodrigues, Vehicular Wi-Fi Data collection framework, (accessed on 2022-05-30) (2021).
URL https://github.com/adamiaonr/wifi-vehicles/tree/802.11ad

[16] Renault Clio Dimensions and Specifications, (accessed on 2022-05-30).
URL https://www.renault.co.uk/cars/clio/specifications.html

[17] 2005 Peugeot Partner blueprints, (accessed on 2022-05-30).
URL http://blueprintbox.com/details.php?image_id=20174

[18] Trimble, Trimble pro series datasheet, (accessed on 2022-05-30) (2012).
URL http://www.windenvironmental.com/Data-Sheets/Trimble-Pro%20Series-DS.pdf

[19] M. Mattos, Seleção de setores para 802.11ad baseado em geolocalização em redes veiculares, Master's thesis, Faculty of Engineering of the University of Porto (2020).

[20] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, M. Hollick, Compressive millimeter-wave sector selection in off-the-shelf ieee 802.11 ad devices, in: Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies, 2017, pp. 414–425. doi:10.1145/3143361.3143384.

[21] C. Leys, C. Ley, O. Klein, P. Bernard, L. Licata, Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median, Journal of Experimental Social Psychology 49 (4) (2013) 764–766. doi:10.1016/j.jesp.2013.03.013.