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

Mateus Mattos*, António Rodrigues*, Rui Meireles† and Ana Aguiar*

*Instituto de Telecomunicações, FEUP DEEC, University of Porto, Portugal
†Computer Science Department, Vassar College, USA

***Abstract*—To improve range, 802.11ad uses directional communication, the first step of which is to choose the best antenna configuration, or 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. Analysis of the collected data revealed the following inefficiencies: (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 studied 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 such a strategy can potentially improve throughput by up to 60 %, depending on location.**

***Index Terms*—VANETs, V2I, mm-Wave, 802.11ad, directional communication**

I. Introduction

The wide unlicensed spectrum available in the 60 GHz region enables very high data rates. But with high frequency also comes high attenuation. 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. Especially so in the case of vehicular communication, where there is a significant degree of mobility involved.

In 802.11ad, beamforming is initiated by sweeping through a set of preset antenna configurations, or sectors, a process known as a Sector-Level Sweep (SLS). The best sector 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 (V2I) 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.

Geolocation 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 study the efficacy of geolocation as a driver for sector selection, putting it forward as an alternative to the algorithm built into COTS 802.11ad devices. We propose to use spatially-indexed historical network performance data to select communication sectors. 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 particular locations.

In summary, we make the following contributions:

- Capture time and spatially-indexed 802.11ad frames in a V2I environment, using COTS equipment (§III).
- Analyze the efficiency of the sector selection algorithm employed by COTS devices (§IV).
- Propose and evaluate a new geolocation-based 802.11ad sector selection algorithm (§V).

II. Background and Related Work

802.11ad devices use phased antenna arrays, whose individual antenna gains can be controlled to achieve directionality. Beamforming is split into two parts [1], [4]. 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 optional phase – Beam Refinement Protocol (BRP) – can then be used to further fine tune the antenna gains.

Our work focuses on the first phase, depicted in Fig. 1. 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

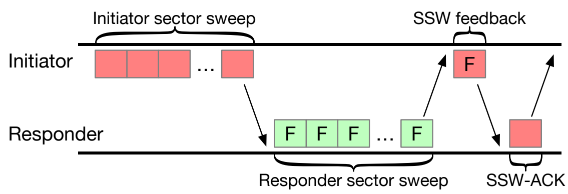

Fig. 1: 802.11ad Sector-Level Sweep process.

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).

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 divides medium access into two periodic intervals: the Beacon Header Interval (BHI), used for client association, and the Data Transmission Interval (DTI), used for data transfer [4]. The Access Point (AP) always initiates an SLS during the BHI to ease client association, and reserves channel time to allow responses from prospective clients. SLSes can also be triggered during the DTI, by any node. Our analysis focuses on SLSes that occur during this phase.

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 [5], in which 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 [6]. 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.mu Rasekh *et al.* [7] 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 poses significant challenges. Experimental studies [8]–[10] 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.* [11] 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 somewhat limits connection duration. Since most 802.11ad devices also support legacy $2.4/5\,\mathrm{GHz}$ Wi-Fi, Nitsche *et al.* [12] proposed using these lower-frequency, longer-range bands to perform out-of-band beam steering. Choi *et al.* [13] 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.* [14] 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 also simulation-based, and the strategy looses some of its effectiveness in strong-multipath environments.

III. Experimental data collection

A. Collection procedure

We conducted experiments with COTS 802.11ad devices – TP-Link Talon AD7200 routers – to collect beamforming data in a V2I scenario. The setup, depicted in Fig. 2, was similar to the one used in [8]. A router, acting as an 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. In the first set the client moved slowly (i.e., below $2\,\mathrm{m/s}$, or $7\,\mathrm{km/h}$) in a specific direction, either away or towards the AP. This enabled the collection of spatially-dense data samples. Fig. 2a depicts the eight different trajectories used in this first set of experiments.

In the second set, 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 $14\,\mathrm{m/s}$, or $50\,\mathrm{km/h}$). This let us collect data at higher speeds.

In all experiments the AP sent pseudo-random data towards the client, at an application-level constant rate of $\sim 420\,\mathrm{Mbps}$.

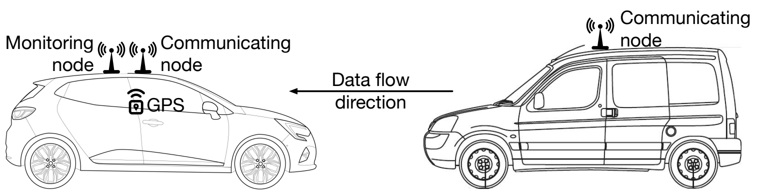

(a) Vehicles (to scale) and devices mounted on them. Application data was streamed from the parked AP to the mobile client. Vehicle blueprints courtesy of Renault UK [2] and BlueprintBox [3].

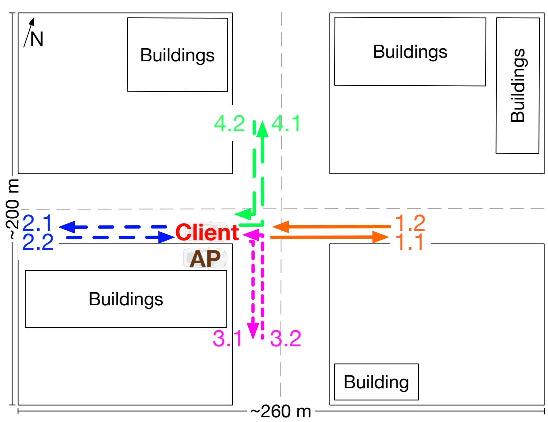

(b) Experiment environment and client mobility patterns.

Fig. 2: Experimental setup. The experiments took place at a residential-area intersection (coordinates 41.111929, -8.631083).

UDP transport was used to avoid any potential influence from TCP's flow and congestion control schemes.

802.11ad frames were captured with `tcpdump` by a third Talon AD7200, configured in monitor mode and co-located with the mobile client, as per Fig. 2a. All routers ran a version of the LEDE operating system, as well as modified `wil6210` drivers and firmware [15]. Geopositioning was provided by a high-accuracy Trimble Pro Series 6H [16] GPS receiver.

All devices were controlled and monitored through a separate 2.4 GHz 802.11n control network. This network was also used to synchronize all device clocks via NTP, allowing for timestamp-based fusion of the collected data.

Additional procedural details can be found in [17].

B. Collected data

We collected the following data:

- A 1 s-resolution client geolocation log.
- 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, making it impossible to accurately determine losses.
- 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 can 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 (`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 geolocation-based sector selection strategy, detailed in §V.

IV. Default Sector Selection Strategy Analysis

The sector-level sweeps captured in the dataset from §III are summarized below:

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%

A successful SLS results in transmit sectors being selected for both initiator and responder (§II). The chosen sectors are included in the frames carrying feedback, i.e., the frames marked F in Fig. 1. A significant portion of SLSes in our dataset were incomplete, meaning no feedback from the responder was captured – only the initiator's sweep.

Some 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 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 the table above shows, most SLSes were inconsequential, specially for the AP (82 % versus 65 % for the client).

The table excludes periodic SLSes used by the AP 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. 2b), and was thus able to utilize a larger range of angles. The sector frequency histogram of Fig. 3a agrees with this. Sectors 20 and 24 alone accounted for ~60 % of the AP's time. In contrast, the client used a much larger variety of antenna sectors.

Fig. 3b 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.

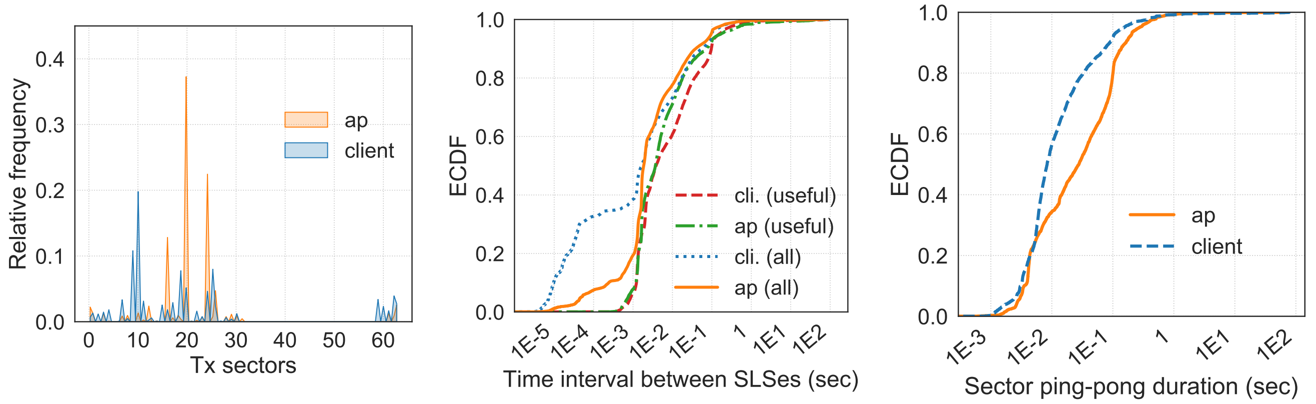

(a) Relative frequency of sector usage.

(b) Time between sector-level sweeps.

(c) Time between sector repetitions.

Fig. 3: COTS devices' sector selection analysis results.

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

The high number and frequency 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-pongs"*: 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. 3c 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 $10\,\mathrm{ms}$. 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.

V. Geolocation-based 802.11ad sector selection

A. 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 sector and geolocation combination. I.e., that the likelihood of a sector being able to achieve a given data rate from a given location does not change with time.

The proposed strategy boils down to the following steps:

1) Discretize space into square cells, e.g., $1 \times 1\mathrm{m}$, 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.

B. Evaluation methodology

The focus of our evaluation was throughput. We compared the amount of data that was communicated in the dataset from §III 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 start with the variant that maintains SLSes unchanged, meaning the available communication time is also unchanged.

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 to be the 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:

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

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 the 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:

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

$Tx_{c,S}$ 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.

Now the throughput gain expected from always using sector S for cell c can be computed as:

$$TG_{c,S} = \frac{D_{c,S} - D_{c,og}}{D_{c,og}}. \tag{3}$$

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:

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

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

C. 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 [18]. 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 different 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. I.e., 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:

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

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:

$$S_c = \underset{s}{\operatorname{argmax}} \ \bar{R}_{c,s} \tag{6}$$

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

D. 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.

Fig. 4a shows the ECDFs of the per-cell throughput gain for the different selection metrics, when 1×1 m cells were used, and SLSes not eliminated. I.e., their time was not reclaimed for data transmission. In this situation, 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 a substantial margin through the use of a single sector, and that therefore the algorithm performed well.

However, the distribution has a long tail, with 20 % of cells experiencing increases of over 10 %, and a handful over 40 %.

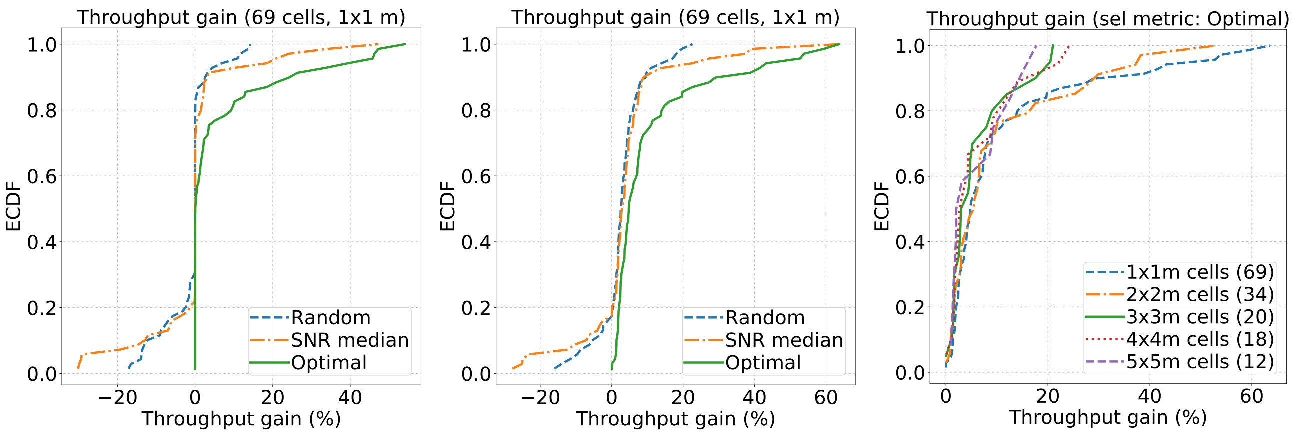

(a) Gain when SLSes are not eliminated. (b) Gain when SLSes are eliminated. (c) Impact of cell size (optimal metric).

Fig. 4: Geolocation-based sector selection per-cell throughput gain Empirical Cumulative Distributions (ECDF).

This shows that, under some circumstances, there is 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. 5 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. 3a), 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.

If we leverage the fact that sector selection is based on geolocation to eliminate SLSes, additional time is freed up for communication. Fig. 6 quantifies this gain as a percentage of the original transmission time. The median time gain was small: $3.6\,\%$. The maximum was $15\,\%$. Using this additional time for data transfer boosts throughput across the board, as depicted by Fig. 4b. Using the optimal selection strategy, all cells exhibited a positive throughput gain, with the median being $4.8\,\%$, and the maximum $63\,\%$.

Fig. 4c shows the optimal strategy's throughput gain for different cell sizes, 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. E.g., the maximum observed gain was $63\,\%$ for 1×1 m cells, but only $43\,\%$ for 5×5 m cells.

These results make intuitive sense since the larger the cell, the larger the diversity of antenna angles and radio obstacles that fall onto a single cell. From this we can conclude that the higher the accuracy of the geopositioning system used, the

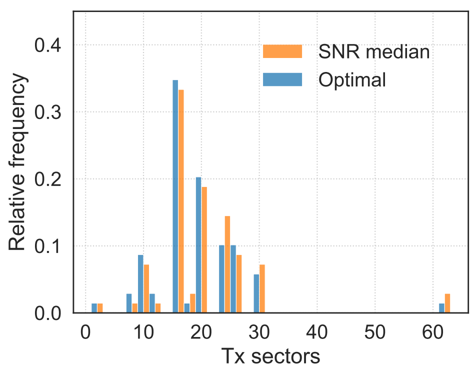

Fig. 5: Geolocation-based sector selection frequency.

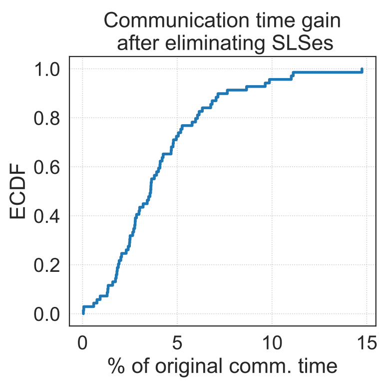

Fig. 6: Time gained by eliminating SLSes, relative to the original transmission time.

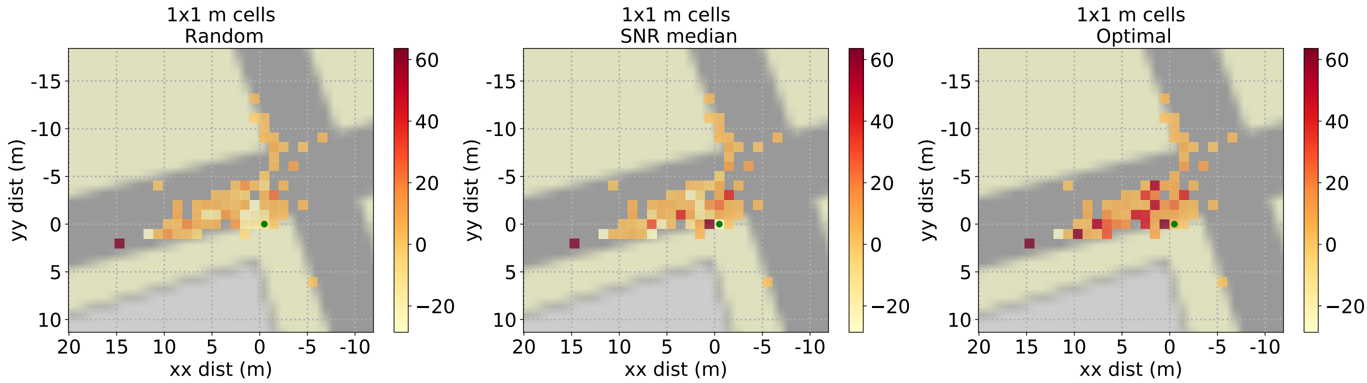

Fig. 7: Geolocation-based sector selection throughput gain by spatial cell.

better geolocation-based sector selection will perform.

Finally, Fig. 7 shows the throughput gain as a function of the client's position relative to the AP. The cells aligned with the side and rear of the AP experienced the highest gains, particularly in the optimal selection metric case.

VI. Conclusions

We collected 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 performed 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 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. Given that fine-grained spatially-indexed data is used in the decision making process, scalability is one of the concerns. We believe it can be addressed by having APs store the historical sector performance data, rather than clients. This would limit the area of interest and consequently the amount of storage needed.

References

[1] 802.11 Working Group, "IEEE 802.11ad-2012," IEEE, Standard, 2012.
[2] "Renault Clio Dimensions and Specifications," Renault United Kingdom. [Online]. Available: https://www.renault.co.uk/cars/clio/specifications.html (accessed on 2021-03-05).
[3] "2005 Peugeot Partner blueprints," BlueprintBox. [Online]. Available: http://blueprintbox.com/details.php?image_id=20174 (accessed on 2021-03-05).
[4] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and J. C. Widmer, "IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi," *IEEE Communications Magazine*, vol. 52, no. 12, 2014.
[5] S. Hur, T. Kim, D. J. Love, J. V. Krogmeier, T. A. Thomas, and A. Ghosh, "Millimeter wave beamforming for wireless backhaul and access in small cell networks," *IEEE Transactions on Communications*, vol. 61, no. 10, 2013.
[6] Z. Marzi, D. Ramasamy, and U. Madhow, "Compressive channel estimation and tracking for large arrays in mm-wave picocells," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 3, 2016.
[7] M. E. Rasekh, Z. Marzi, Y. Zhu, U. Madhow, and H. Zheng, "Noncoherent mmwave path tracking," in *18th International Workshop on Mobile Computing Systems and Applications*, ser. HotMobile '17. ACM, 2017.
[8] A. Aguiar, A. Rodrigues, M. Mattos, B. Maurício, R. Meireles, and P. Steenkiste, "Exploratory Study of 802.11ad Vehicular Links," in *2020 Mediterranean Communication and Computer Networking Conference (MedComNet)*, 2020.
[9] R. Meireles, A. Rodrigues, A. Stanciu, A. Aguiar, and P. Steenkiste, "Exploring Wi-Fi Network Diversity for Vehicle-To-Infrastructure Communication," in *2020 IEEE Vehicular Networking Conference (VNC)*, 2020.
[10] W. Kim, "Experimental Demonstration of MmWave Vehicle-to-Vehicle Communications Using IEEE 802.11ad," *Sensors*, vol. 19, no. 9, 2019.
[11] A. Loch, A. Asadi, G. H. Sim, J. Widmer, and 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.
[12] T. Nitsche, A. Flores, E. Knightly, and J. Widmer, "Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, 2015.
[13] J. Choi, V. Va, N. Gonzalez-Prelcic, R. Daniels, C. R. Bhat, and R. W. Heath, "Millimeter-Wave Vehicular Communication to Support Massive Automotive Sensing," *IEEE Communications Magazine*, vol. 54, no. 12, 2016.
[14] G. R. Muns, K. V. Mishra, C. B. Guerra, Y. C. Eldar, and 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.
[15] D. Steinmetzer, D. Wegemer, and M. Hollick. (2017) Talon tools: The framework for practical ieee 802.11ad research. [Online]. Available: https://seemoo.de/talon-tools (accessed on 2021-03-05).
[16] *Trimble pro series datasheet*, Trimble, 2012.
[17] 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.
[18] C. Leys, C. Ley, O. Klein, P. Bernard, and L. Licata, "Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median," *Journal of Experimental Social Psychology*, vol. 49, no. 4, 2013.