

mmWave in Vehicular Networks: Leveraging Traffic Signals for Beam Design

Original

mmWave in Vehicular Networks: Leveraging Traffic Signals for Beam Design / LIMANI FAZLIU, Z., Malandrino, F., Chiasserini, C.F.. - ELETTRONICO. - (2019). (IEEE WoWMoM-CCNCPS 2019 Washington, DC (USA) 10-12 June 2019) [10.1109/WoWMoM.2019.8793039].

Availability:

This version is available at: 11583/2731901 since: 2019-05-01T10:14:49Z

Publisher:

IEEE

Published

DOI:10.1109/WoWMoM.2019.8793039

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

mmWave in Vehicular Networks: Leveraging Traffic Signals for Beam Design

Zana Limani Fazliu*, Francesco Malandrino^{†,‡,§}, Carla-Fabiana Chiasserini^{‡,†,§}

* University of Prishtina, Kosovo † CNR-IEIIT, Italy ‡ Politecnico di Torino, Italy § CNIT, Italy

Abstract—Vehicle-to-infrastructure millimeter-wave (mmWave) communication represents a potential solution to capacity shortage in mobile networks. However, effective beam alignment between senders and receivers requires knowledge of the position of vehicles, which is often impractical to obtain in real time. We propose to solve this problem by leveraging the *traffic signals*, e.g., semaphores, that regulate the vehicular mobility. As an example, we may coordinate beams with red semaphore lights, as they correspond to higher vehicle densities and lower speeds. In order to evaluate such intuition, we propose a mmWave communication model accounting for both the distance and the speed of vehicles being served, and use such a model to compare several beam design strategies. For increased realism, we consider as our reference scenario a large-scale, real-world vehicular trace depicting the mobility in Luxembourg. Our results show that our approach outperforms static beam design based on road topology alone, and, remarkably, it yields a performance comparable to that of solutions based on real-time mobility information.

I. INTRODUCTION

High-definition maps, their real-time updates, and on-board multimedia systems are just few of the applications that concur to make vehicles, both ordinary and self-driving, prime consumers of network traffic. Indeed, automotive services – safety – as well as entertainment-related – are among the reference use cases for several next-generation network technologies [1], [2], including C-V2X, 802.11p/ITS-G5, and 5G. In spite of the important differences among these technologies, they all share the goal of providing more network capacity to vehicles and their drivers.

Whenever more capacity is needed, millimeter-wave (mmWave) communications are an appealing option [3]. On the negative side, mmWave suffers from harsh propagation conditions, with severe attenuation and high blockage probability. This has led to the design of *directional* antenna systems, where the available power is concentrated on one or more *beams*. It follows that the performance of mmWave networks critically depends on beam design, i.e., deciding the number, direction, and amplitude of the beams to form. Successful beam design requires knowledge about the location of the user(s) to serve, which explains why the earliest and most mature mmWave applications target static or quasi-static scenarios.




Fig. 1. A semaphore-regulated road crossing, with a mmWave gNB deployed at the center. The north-south road has a green light; the east-west one has a red light. Assume that the gNB can generate two beams, and we have to decide their direction: pointing east and west (red beams) allows us to cover many more vehicles than pointing north and south (green beams), thus improving the network performance. Since traffic light programs are known in advance, there is no need to measure the vehicular traffic and adapt to it in real time.

The need to use narrow, concentrated beams in order to establish mmWave links, complicates, in particular, the initial phases in link establishment such as cell discovery and initial access. Due to the need for beam alignment, beams must be searched at both the base station (gNB) and user side to align them at both ends. The authors in [4] evaluate three potential solutions for the initial access procedure: exhaustive brute-force sequential search, two-stage iterative search, and an algorithm that leverages context information such as GPS location, obtained from legacy LTE base station. For mmWave-based vehicular mobile networks too, beam training is often identified as the main challenge to overcome, due to the associated control overhead and delay [3]. Most popular approaches are predicated on hoarding and leveraging as much location information as possible, coming from road-side sensors as well as from the vehicles themselves [3], [5] and transferred through non-mmWave control channels [5]. Collecting the information, processing it, and re-aligning the beams in near-real-time is a very challenging task. As a consequence, some works, e.g., [6], envision dispensing with beam realignment altogether, statically setting the beam orientation using road topology information.

All of the above mentioned works focus on beam aligning for each individual gNB-user equipment (UE) pair, implicitly assuming that each narrow beam is employed to transmit to a single UE only. However, in

ultra dense scenarios, it is highly likely that even a beam as narrow as 5° can cover several UEs simultaneously at any given time, UEs that can be multiplexed within the same beam. Therefore, in this work, instead of focusing on perfectly aligning beams for each gNB-UE pair, we rather look at how to align the beams at the gNB so as to increase the number of covered vehicular UEs with quality links, i.e., links with acceptable values of signal-to-interference-and-noise (SINR) ratio.

Specifically, we leverage the fact that in urban environments it is possible to know a great deal of information about vehicular mobility *without* detecting it in real time. Consider the situation depicted in Fig. 1: red semaphore lights are associated with a higher vehicle density and lower speeds – two factors that can improve the achievable total throughput. This is very valuable information, and is known *a priori*, thus allowing to design the beams without the need to make real-time decisions.

Our high-level purpose is to assess the performance of this approach, i.e., using traffic signal state information to complement and replace real-time mobility data. To this end, we make the following main contributions:

- we adopt a simplified channel and beamforming model, able to capture the interplay between directivity gain and coverage. Importantly, we tailor aspects of our model using empirical results obtained through real-world traces;
- we formulate the beam configuration problem as an optimization problem, aiming at maximizing the quality of coverage of vehicular UEs. In light of the problem complexity, we define four alternative beam design strategies, requiring different types of information and working at different time scales. Among these, we propose a low-complexity scheme, named Traffic Light (TL), which leverages road topology information and traffic light signals to efficiently configure the beams of the gNBs;
- we evaluate the above strategies leveraging a large-scale trace, including the real-world urban topology and realistic vehicular mobility of Luxembourg City, Luxembourg. Our results show that, in spite of its much lower complexity, the TL scheme yields a performance comparable to that of solutions based on real-time mobility information.

The remainder of this paper is organized as follows. We detail our mmWave communication model in Sec. II and the beam design strategies we consider in Sec. III. Then, in Sec. IV, we describe our reference scenario and some present numerical results. Finally, Sec. V concludes the paper.

II. MMWAVE COMMUNICATION

Here, we introduce the mmWave channel characterization and the antenna model. It is worth highlighting

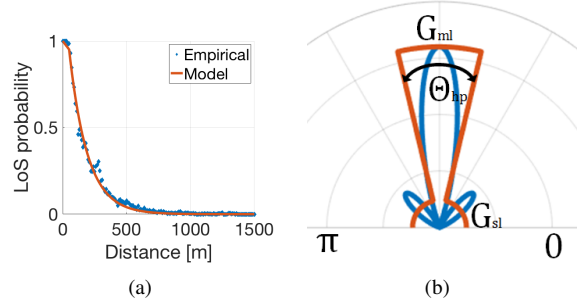


Fig. 2. Empirical vs. tailored LoS probability model (a); sectorized antenna model (b).

that, although our study can be easily extended to uplink communications, in the following we focus on downlink transmissions, from the gNB to the vehicular UEs, in virtue of the traffic asymmetry that still holds.

As previously mentioned, mmWaves experience higher path loss and blockage due to their high frequency, thus accurately modeling line-of-sight (LoS) conditions is crucial. To determine whether the transmitter-receiver pair are in LoS, NLoS, or outage (no communication link possible), we draw on the model in [7] and use the following distance-based probabilities:

$$\begin{aligned}
 p_{out} &= \max(0, 1 - e^{a_{out}d + b_{out}}) \\
 p_{LoS} &= (1 - p_{out})e^{-a_{LoS}d} \\
 p_{nLoS} &= 1 - p_{LoS} - p_{out},
 \end{aligned} \tag{1}$$

where d is the distance and the parameters a_{LoS} , a_{out} , and b_{out} can be tailored to the particular scenario. Consistently with outdoors models in the literature, we do not explicitly perform ray-tracing.

Since the reference scenario under study, described in detail in Sec. IV, is based on real-world data, we adjust the model parameters to fit the particular environment of our scenario. In particular, taking advantage of the knowledge we have about the locations of the buildings, road topology, and vehicle positions and speed, we can derive empirically the LoS condition of each gNB-vehicle pair, at a given time. This allows us to draw an empirical LoS probability curve as a function of the distance between the gNB and the vehicle, depicted in Fig. 2(a) with blue dots. Optimized values for model parameters a_{LoS} , a_{out} , and b_{out} , can then be obtained by applying curve fitting against the empirical probability curve. The optimized model is compared to the empirical curve in Fig. 2(a).

We assume that the gNBs are equipped with directional steerable antenna arrays composed of N_{tx} elements, which can support multiple simultaneous beams. A UE is also equipped with directional steerable antenna arrays composed of N_{rx} elements, but it can support one beam at a time only. We consider that the channel between the transmitter and the receiver

includes a certain number of path clusters, as described in [7].

For analytical tractability, it is common to approximate the actual antenna array beam pattern by a step function with a constant main-lobe over the beamwidth and a constant side-lobe otherwise [8], as shown in Fig. 2(b). This simple model captures the interplay between directivity gain, which ultimately affects the transmission range, i.e., the coverage and half-power beamwidth¹. Specifically, the gain of the main lobe for a single beam pattern is calculated through the approximation provided in [9]: $G_{ml} = \frac{41000}{\Theta_{hp}^\circ \Phi_{hp}^\circ}$ where Θ_{hp}° and Φ_{hp}° are the half-power beamwidths (in degrees), in the azimuth and elevation planes, respectively. For simplicity and consistency with the assumptions made in [9], we set $G_{sl} = 0$.

Finally, we can describe the effective channel that two communicating endpoints experience, including beamforming gains at both ends, $H(t)$, as

$$H(t) = \frac{1}{L} \sum_{k,l} h_{k,l}(t) \sqrt{G_{rx}(\theta_{k,l}^{rx}, \phi_{k,l}^{rx}) G_{tx}(\theta_{k,l}^{tx}, \phi_{k,l}^{tx})}$$

Therein K is the number of clusters and L is the number of paths within a cluster, $h_{k,l}$ are the small-scale fading gains, while G_{tx} and G_{rx} are the gain values of the transmit and receive antennas, respectively, as functions of the azimuth (θ) and elevation (ϕ) arrival and departure angles. The small-scale fading gains are generated as described in [7].

III. GNB SELECTION AND BEAM DESIGN

We now focus on the main aspects of the mmWave communication system we intend to address and on the approach we propose to overcome the existing hurdles. In particular, Sec. III-A introduces an optimization formulation for the beam design problem, formally stating its objective and constraints. Then Sec. III-B presents the heuristic approaches we compare in our performance evaluation, among which our proposed scheme, named Traffic Lights.

A. Optimization formulation

The essential trade-off in mmWave communications is between the directionality gain that can be achieved using beamforming and the spatial coverage that can be offered. Thus, it is clear that both the width and the number of beams used by the gNB should be treated as design considerations during beamforming. Depending on the vehicle distribution and mobility, fewer wider beams may be preferred over multiple narrower beams, and vice-versa.

Based on the above observations, we aim to address the following questions: i) how many beams should

¹The half-power beamwidth refers to the angular width of the main lobe of the antenna pattern between the points that are half the intensity of the maximum beam level (beam peak).

a gNB transmit, and of what beamwidth; ii) which directions should they transmit at; and, finally, iii) which vehicular UEs should be scheduled on which beam. The goal is to find answers to these questions while maximizing coverage and network throughput. For the sake of simplicity, we do not jointly optimize the number, width, and direction of beams; rather, we study several number/width combinations and optimize, for each of them, the direction of the beams.

We consider that the set of gNBs, denoted by \mathcal{B} , and the set of vehicular UEs, denoted by \mathcal{U} , are already mutually aware of the direction they can communicate in. This is a fair assumption, since vehicles routinely broadcast their location and speed [5] in the cooperative awareness message (CAMs), also used for automotive safety-related applications. We further denote the downlink direction between a gNB b and a UE u , as $\tau_{b,u}$. Let n_b be the maximum number of beams a gNB can simultaneously transmit, and Θ_b the width of the beams. We only consider the beam directions in the azimuth plane, i.e., we assume that the elevation angles are fixed towards the optimal direction.

We can then formulate the choice of the directions $\theta_{b,i}$, $i \in [1 \dots n]$ of the beams as an optimization problem, where the decision variables are represented by the θ -values themselves. The objective function to maximize is the (expected) number of served vehicles, i.e., vehicles that are able to exchange information with at least one gNB. For a vehicle u to be served, we need that: (1) there is at least one beam of a gNB covering that vehicle; (2) such a (gNB, vehicle) pair is not in outage. Item 1 depends on the direction $\theta_{b,i}$ of each beam i of gNB b (which is a decision variable) and its width Θ_b (which is a parameter). Specifically, beam i of gNB b covers vehicle u if:

$$|\tau_{b,u} - \theta_{b,i}| \leq \frac{\Theta_b}{2}. \quad (2)$$

Item 2 depends on the blockage model; from the viewpoint of the optimization problem, it is summarized by a parameter $p_{out}(b, u) \in [0, 1]$, expressing the probability that the path from gNB b to vehicle u is blocked, as computed in (1).

Combining condition (2) and parameter $p_{out}(b, u)$, we can write the objective function as:

$$\max_{\theta} \sum_{u \in \mathcal{U}} \max_{b \in \mathcal{B}, i \in [1 \dots n]} \left(1 - p_{out}(b, u) \mathbb{1}_{|\tau_{b,u} - \theta_{b,i}| \leq \frac{\Theta_b}{2}} \right).$$

The above equation can be read from right to left, as follows. The condition within the indicator function is the one in (2); thus, the indicator function takes 1 for those beams that cover vehicle u and 0 otherwise. We then weigh the value of the indicator function by the probability that the path from b to u is not in outage, i.e., $1 - p_{out}$. Finally, for each vehicle, we only consider the beam with the highest value of $1 - p_{out}$, since we do

TABLE I
COMPARISON AMONG THE DIFFERENT STRATEGIES

Strategy	Reconfiguration	Needed traffic information	Clustering
Static	Never	Statistics	Hierarchical [10]
Dynamic	At every step	Real-time	Hierarchical [10]
DBSCAN	At every step	Real-time	DBSCAN [11]
Traffic Light	Periodic	None	No

not allow the same vehicle to be covered by multiple beams. In summary, our objective function expresses the expected number of served vehicles, given the beam directions, the vehicle positions, and the blockage model.

As far as the constraints are concerned, two beams from the same gNB cannot overlap, i.e.,

$$|\theta_{b,i} - \theta_{b,j}| \geq \Theta_b, \quad \forall b \in \mathcal{B}, i, j \in [1 \dots n]: i \neq j.$$

B. Beam design heuristics

Directly solving the above optimization problem, has a very high computational complexity. Indeed, the indicator function and the modulo operators both result in binary variables, which means that the problem falls in the mixed-integer linear programming (MILP) category. Such problems are well-known to be NP-hard (see the reduction from the vertex cover problem in [12]).

In light of this, we cast the problem of maximizing the number of vehicles for which condition (2) holds as a one-dimensional clustering² problem. Indeed, each beam i can be seen as a cluster of directions whose center is $\theta_{b,i}$ and whose maximum amplitude is Θ_b . To solve such a problem, we compare four different *strategies*, detailed next and whose main features are summarized in Tab. I.

Static. Under this strategy, the directions of the beams, i.e., the $\theta_{b,i}$ values, do not change over time. To determine such directions, we formulate a *hierarchical clustering* problem as follows:

- 1) we divide the time in discrete *steps*;
- 2) for each vehicle and time step, we create one observation (i.e., one data point) corresponding to the vehicle's position;
- 3) we compute the pairwise *angular distances* between observations;
- 4) we feed the resulting distance matrix to the Vor Hees Algorithm [10], setting the maximum intra-cluster distance to Θ_b ;
- 5) we consider the n largest clusters, i.e., the clusters including the highest number of observations;
- 6) for each such cluster, we set the direction $\theta_{b,i}$ as the mean between the minimum and maximum angle of the vehicles it includes, i.e., $\theta_{b,i} \leftarrow \frac{1}{2}[\min_u \tau(b, u) + \max_u \tau(b, u)]$.

²Clustering here does not refer to the path clusters dealt with in Sec. II.

Note that, in step 2, we may create multiple observations with the same coordinates, e.g., if two vehicles are observed in the same position at different times. This is intentional and allows us to properly account for the fact that we are more likely to find a vehicle in some locations of the topology than in others.

The angular distance used in step 3 is a metric connected with the cosine similarity. For two generic vectors $\vec{a}, \vec{b} \in \mathbb{R}^n$, we have $d(\vec{a}, \vec{b}) = 1 - \frac{\cos^{-1} s(\vec{a}, \vec{b})}{\pi}$ where $s(\vec{a}, \vec{b}) = \frac{\vec{a} \cdot \vec{b}}{|\vec{a}| |\vec{b}|}$ is the cosine similarity [13]. Such a metric is very well-suited to our scenario, as the beam amplitude Θ_b limits the angular distance between vehicles served by the same beam rather than the Euclidean one.

The Vor Hees algorithm we run in step 4 starts by creating singletons, i.e., clusters with only one observation each. Then, at every iteration, it merges the two closest clusters; the algorithm stops when additional merges would result in clusters with a size larger than a set threshold (Θ_b in our case). Note that, unlike the more popular k-means algorithm, this approach has the advantage that the number of clusters do not have to be set in advance.

The static strategy is arguably the simplest way to leverage aggregate traffic statistics; also, it can be performed offline and requires no reconfiguration of the beams. On the negative side, it cannot account for the time evolution of vehicular mobility, e.g., different traffic patterns at different hours of the day.

Dynamic. It works as the Static one, with the important difference that decisions are re-made *at every time step*, i.e., the clustering procedure described above is repeated at every time step. Implementing the Dynamic strategy would require real-time knowledge of vehicular mobility and almost-instantaneous beam reconfiguration – aspects that render the strategy impractical for real-world implementations. Nonetheless, we use it as a benchmark to compare against.

DBSCAN. It is a variant of the Dynamic strategy, using (as the name suggests) the popular DBSCAN clustering algorithm [11] *in lieu* of the Vor Hees one. DBSCAN is very popular and highly effective: similarly to hierarchical approaches, it returns an undetermined number of clusters with a target maximum inter-cluster distance. Unlike hierarchical approaches, DBSCAN can be seamlessly extended to account for additional factors, e.g., the vehicles distance and speed, hence the data rate they can achieve. This makes DBSCAN a potentially better-performing alternative to the Vor Hees algorithm.

Traffic Lights (TL). Vehicular mobility is constrained not only by the road topology, but also by the state of traffic signals, e.g., semaphores. The TL strategy leverages the available information on semaphore states and points the beams available at each gNB towards the road segments where the semaphore light is red. As




Fig. 3. Real-world scenario: Luxembourg city center. The red circles denote the locations of the traffic lights.




Fig. 4. Total throughput yielded by each strategy, as the number and amplitude of beams changes.




Fig. 5. CDF of the vehicles data rate under each strategy when there are two 5-degree beams.

exemplified in Fig. 1, this increases both the number of vehicles that can be covered and the data rate they can obtain. The TL strategy is more flexible than the Static one, in that beam directions account for the vehicles mobility. At the same time, it is much more practical than the Dynamic and DBSCAN strategies, as it does not require any real-time mobility information and beam reconfigurations are less frequent. In fact, the TL strategy requires *no knowledge* whatsoever of vehicular mobility, hence it can be applied in situations where such information is unavailable or unreliable.

IV. PERFORMANCE EVALUATION

Reference Scenario. We consider the publicly-available Luxembourg scenario [14], which combines:

- the real-world topology of Luxembourg City;
- the location of semaphores and bus stops therein;
- semaphore states (red, yellow, green) at all times;
- the realistic mobility of around 14,000 vehicles over a period of 12 hours, generated with SUMO and based on real-world traffic flows, e.g., commuters traveling to the city center.

We consider a 2×2 km² area of the city center, as depicted in Fig. 3. Throughout such a road topology, we place a total of 51 gNBs, corresponding to traffic lights. Their positions are marked by red dots.

We set the center frequency available for mmWave communication to $f_c = 76$ GHz, as typically assumed for vehicular networks [15], and the available bandwidth to $BW = 1$ GHz. All gNBs are equipped with a 32×32 uniform planar array (UPA) with up to 4 RF chains, and vehicular UEs are equipped with a 8×8 UPA. The parameters used for modeling the mmWave wireless channel are the same as those used in [7], except for the parameters of the blockage model, which are tailored to the Luxembourg scenario. Specifically, we set: $a_{out} = 0.006$, $b_{out} = 0.3$, and $a_{LoS} = 0.001$.

For all strategies, we consider two possible half-power beamwidths: $\Theta_{HP}^\circ = 5^\circ, 10^\circ$, while the number of simultaneous beams vary from 2 to 4. Given the beam

directions, for each gNB-UE pair, we compute the SINR as well as the throughput, using the model in [7]:

$$R = BW \min \left[\log_2 \left(1 + 10^{0.1(SINR - \Delta)} \right), r_{max} \right],$$

where Δ is a loss factor set to 3 dB [7] and r_{max} is the maximum spectral efficiency. Since future networks will support modulation schemes as advanced as 1024 QAM, r_{max} is set to 9.26 bps/Hz as per 3GPP TR 36.213.

Numerical Results. The first aspect we are interested in is the performance of the set-up beams discussed in Sec. III-B, i.e., the total network throughput they can guarantee: these results are summarized in Fig. 4. First, observe that, as expected, having more or wider beams results in better performance, regardless of the strategy. Second, the number of beams (two or four) has a more significant impact than their amplitude (5 or 15 degrees); indeed, in urban scenarios, a narrow beam is often sufficient to illuminate most of the vehicles in a road lane.

Interestingly, TL significantly outperforms Static and yields a throughput that is comparable to that of the Dynamic and DBSCAN strategies. This is a very important result, confirming our intuition that traffic signals – semaphores in this case – can indeed provide valuable information for beam design decisions. Also, recall that, as summarized in Tab. I, Dynamic and DBSCAN require real-time mobility information and are thus not practical to implement. The CDF of the vehicles data rate yielded by different strategies is depicted in Fig. 5, in the case of two 5-degree beams.

Fig. 6(a) is devoted to the second of such major factors, i.e., the number of users each strategy can serve with each beam. We note that DBSCAN serves the highest number of users, a hint that such a clustering algorithm may be more effective than the Voor Hees algorithm used by the other strategies. All other schemes serve a very similar number of users, which means that their difference in performance (see Fig. 4) is almost entirely due to the different SINR (hence, data rate) such users experience, as depicted in Fig. 5.

Fig. 6(b), depicting the number of beams serving each user under different strategies, sheds further light on the

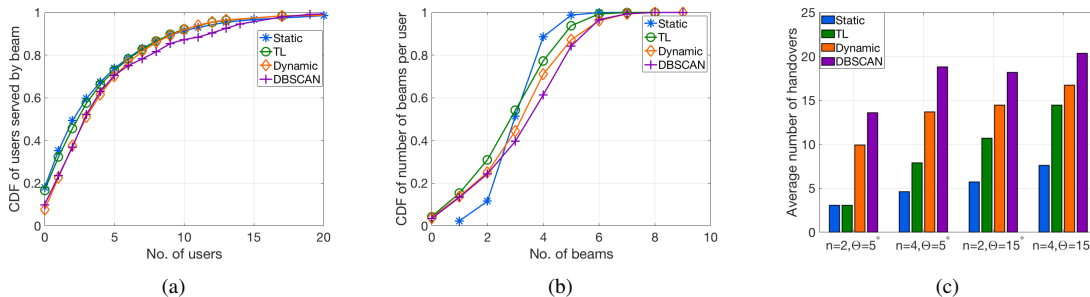


Fig. 6. Number of vehicles served (a) and number of beams serving each vehicle (b) under the different strategies when there are four 15-degree beams; number of handovers under different strategies and configurations (c).

performance of the beam configurations. Under the Dynamic and DBSCAN strategies, beams are much more likely to overlap than under the Static and TL ones. Indeed, under Dynamic and DBSCAN, gNBs direct their beams towards whichever area of the topology is more crowded, regardless of what neighboring gNBs are doing. In TL this effect is less evident, as consecutive semaphores on the same road tend to show different colors. Finally, for Static, overlaps are even less frequent since considering the whole trace duration instead of just one time step makes more likely that neighboring gNBs serve different areas.

A higher number of overlapping beams in Fig. 6(b) corresponds, as shown in Fig. 6(c), to a higher number of handovers between different gNBs. This is a very important result, indeed, as discussed in Sec. II, in our model handovers are ideal and result in no performance penalty, however this is not the case in real-world cellular networks. Typically, a higher number of handovers translates into a significantly higher control overhead and into a higher load on the core network entities (e.g., the Mobility Management Entity (MME) in LTE EPC). From this viewpoint, the TL strategy represents the safest option, guaranteeing a performance comparable to Dynamic and DBSCAN, with a number of handovers comparable to that of Static.

V. CONCLUSION

We have identified mmWave as a promising technology for vehicular networks. However, the performance of mmWave networks depends upon the alignment of beams between gNBs and vehicles, and such an alignment requires knowledge of the vehicles position and speed. Instead of relying on real-time mobility information, in this paper we proposed to rely on traffic signals, e.g., semaphores, which influence the mobility itself. Leveraging traffic semaphore state information for beam design, results in a network performance that exceeds that of baseline approaches (namely, static beam alignment) and is comparable to that of approaches using real-time mobility information. Furthermore, our approach leads to a reduced overlap between beams, hence to a lower number of handover procedures.

REFERENCES

- [1] S. Chen, J. Hu, Y. Shi, Y. Peng, J. Fang, R. Zhao, and L. Zhao, "Vehicle-to-everything (v2x) services supported by lte-based systems and 5g," *IEEE Comm. Standards Mag.*, 2017.
- [2] A. Bazzi, B. M. Masini, A. Zanella, and I. Thibault, "Beaconing from connected vehicles: IEEE 802.11 p vs. LTE-V2V," in *IEEE PIMRC*, 2016.
- [3] 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 Comm. Mag.*, 2016.
- [4] M. Giordani, M. Mezzavilla, and M. Zorzi, "Initial access in 5g mmwave cellular networks," *IEEE Comm. Mag.*, 2016.
- [5] I. Filippini, V. Sciancalepore, F. Devoti, and A. Capone, "Fast Cell Discovery in mm-wave 5G Networks with Context Information," *IEEE Trans. on Mobile Computing*, 2017.
- [6] A. Loch, A. Asadi, G. H. Sim, J. Widmer, and M. Hollick, "mm-Wave on Wheels: Practical 60 GHz Vehicular Communication Without Beam Training," *IEEE COMSNETS*, 2017.
- [7] M. R. Akdeniz, Y. Liu, M. K. Samimi, S. Sun, S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter wave channel modeling and cellular capacity evaluation," *IEEE JSAC*, 2014.
- [8] J. G. Andrews, T. Bai, M. N. Kulkarni, A. Alkhateeb, A. K. Gupta, and R. W. Heath, "Modeling and analyzing millimeter wave cellular systems," *IEEE Trans. on Communications*, 2017.
- [9] J. D. Kraus and R. J. Marhefka, "Antennas: for all applications," *Mc Graw Hill*, 2002.
- [10] E. M. Voorhees, "Implementing agglomerative hierarchic clustering algorithms for use in document retrieval," Cornell University, Tech. Rep., 1986.
- [11] D. Birant and A. Kut, "St-dbscan: An algorithm for clustering spatial-temporal data," *Data & Knowledge Engineering*, 2007.
- [12] A. Schrijver, *Theory of linear and integer programming*. John Wiley & Sons, 1998.
- [13] M. Steinbach, G. Karypis, V. Kumar *et al.*, "A comparison of document clustering techniques," in *ACM KDD*, 2000.
- [14] L. Codeca, R. Frank, S. Faye, and T. Engel, "Luxembourg sumo traffic (LuST) scenario: Traffic demand evaluation," *IEEE Intelligent Transportation Systems Magazine*, 2017.
- [15] B. Antonescu, M. T. Moayyed, and S. Basagni, "mmwave channel propagation modeling for v2x communication systems," in *IEEE PIMRC*, Oct 2017.

ACKNOWLEDGMENT

This work has been performed in the framework of the European Union's Horizon 2020 projects 5G-EVE and 5G-CARMEN, co-funded by the EU under grant agreements (respectively) No. 815074 and No. 825012. The views expressed are those of the authors and do not necessarily represent the projects. The Commission is not liable for any use that may be made of any of the information contained therein.