LTE and Millimeter Waves for V2I Communications: an End-to-End Performance Comparison

Marco Giordani, Andrea Zanella, Michele Zorzi Department of Information Engineering, University of Padova, Italy Email:{giordani, zanella, zorzi}@dei.unipd.it

Abstract—The Long Term Evolution (LTE) standard enables, besides cellular connectivity, basic automotive services to promote road safety through vehicle-to-infrastructure (V2I) communications. Nevertheless, stakeholders and research institutions, driven by the ambitious technological advances expected from fully autonomous and intelligent transportation systems, have recently investigated new radio technologies as a means to support vehicular applications. In particular, the millimeter wave (mmWave) spectrum holds great promise because of the large available bandwidth that may provide the required link capacity. Communications at high frequencies, however, suffer from severe propagation and absorption loss, which may cause communication disconnections especially considering high mobility scenarios. It is therefore important to validate, through simulations, the actual feasibility of establishing V2I communications in the above-6 GHz bands. Following this rationale, in this paper we provide the first comparative end-to-end evaluation of the performance of the LTE and mmWave technologies in a vehicular scenario. The simulation framework includes detailed measurement-based channel models as well as the full details of MAC, RLC and transport protocols. Our results show that, although LTE still represents a promising access solution to guarantee robust and fair connections, mmWaves satisfy the foreseen extreme throughput demands of most emerging automotive applications.

Index Terms—Vehicular communications; LTE; millimeter waves (mmWaves); end-to-end performance; ns-3.

I. INTRODUCTION

In recent years, the automotive industry has rapidly evolved towards the development of advanced automotive services offering safer traveling, improved traffic management, and support to infotainment applications. A key enabler of this evolution is Vehicle-to-Infrastructure (V2I) communication, which allows vehicles to communicate with road-side infrastructures and the Internet. The Long Term Evolution (LTE) standard presently represents the principal wireless interface offering V2I transmission services [2]. However, future Connected-Intelligent Transportation Systems (C-ITSs) will need to satisfy unprecedentedly stringent demands in terms of latency and throughput (i.e., in the order of terabytes per driving hour according to some estimates [3]) which may saturate the capacity of traditional technologies for vehicular communications. In this perspective, industry players have devoted efforts into specifying new communication solutions as enablers of the performance requirements of next-generation automotive networks. The millimeter wave (mmWave) spectrum - roughly above 10 GHz [4] - currently holds great promise because of the large available bandwidth that may guarantee data rates in the order of multi-gigabit-per-second.

Although the application of mmWaves in a vehicular context is not new (automotive radars operating in the 77 GHz band are already in the market [5]), the severe isotropic path loss and blockage absorption of mmWave signals, as well as the increased Doppler effect in high mobility scenarios, make the design of wireless systems in the above-6 GHz bands very challenging [6]. Before unleashing the potential of this technology into a V2I deployment, it is therefore fundamental to validate the practical feasibility of designing mmWaveaware strategies in view of the strict requirements and the specific features of future transportation systems.

Motivated by the above introduction, our paper targets the following objectives. First, we provide the first comprehensive end-to-end performance evaluation of the mmWave and the LTE paradigms in a vehicular environment. In particular, we characterize the system throughput and latency when varying the density of network infrastructures, the target application's demands and the channel model. We also consider both urban and highway scenarios, to characterize different mobility and propagation regimes. Unlike traditional performance analyses, e.g., [7]-[9], which rely on Physical (PHY) or Medium Access Control (MAC) layer quality metrics (e.g., achievable transmission range or packet transmission probability), our work investigates the impact of the upper layers on the network behavior, thereby guaranteeing more accurate system-level analyses. Moreover, unlike analytical evaluations, e.g. [10], [11], which typically adopt conservative assumptions on the signal propagation, our paper considers full-stack simulations, which allow to estimate the system performance accounting for detailed protocol implementations. Second, we evaluate through numerical comparisons whether the mmWave technology might represent a more promising solution in creating a safer and more efficient driving ecosystem than its LTE counterpart. Third, based on our extensive simulation results, we provide guidelines to identify the most appropriate network interface that satisfies V2I service requirements while establishing high-capacity channels, a research task that, to date, has not been throughly investigated yet.

Our simulation campaign proves that, although LTE still represents a promising access solution to maintain robust communications, mmWave technology emerges as an enabler of the foreseen extreme throughput demands of future automotive applications. Moreover, we demonstrate that, while densification is beneficial to urban mmWave deployments for both throughput and latency, it may have a negative impact on the performance of LTE systems and in highway scenarios. Our study reveals also several important findings on the interaction between the transport layer mechanisms and the underlying physical propagation.

V2N Application	Communication Scenario	End-to-end latency [ms]	Reliability [%]	Data Rate [Mbps]
Vehicle Platooning	Cooperative driving (low degree of automation) Cooperative driving (high degree of automation) Information Sharing	25 10 20	90 99.99 N.D. $^{(*)}$	$\begin{array}{c} 0.096\\1\\65\end{array}$
Advanced Driving	Cooperative collision avoidance Information sharing (high degree of automation)	10 100	99.99 ${ m N.D.}^{(*)}$	10 50
Extended Sensors	Sensor data sharing (medium degree of automation) Sensor data sharing (high degree of automation) Video sharing	25 10 50	99.99 99.99 90	250 1000 10
Remote Driving	Information exchange	5	99.999	UL:25 DL: 1

TABLE I: Performance requirements of different use cases specified by the 3GPP in [1].

^(*) This requirement is still under discussion (or, in some cases, not yet discussed) in 3GPP.

The rest of the paper is organized as follows. Sec. II overviews the characteristics of the LTE and the mmWave radios as enabling technologies for V2I communications. Sec. III describes our simulation setup and Sec. IV presents our main findings and comparative results. Finally, Sec. V concludes the paper and discusses possible research extensions.

II. V2I RADIO TECHNOLOGIES: AN OVERVIEW

Connected and Autonomous Vehicles (CAVs), when fully commercialized, will address the societal and business trends of 2020 and beyond, and will have ever more stringent regulations in terms of road safety and traffic efficiency [12]. In this regard, the 3rd Generation Partnership Project (3GPP), in its Release 15, defines new use cases specific to future vehicular services whose requirements, although not yet fully specified, have already been outlined in [1], as summarized in Table I.

- *Vehicles Platooning*. It refers to the set of services that enable the vehicles to cooperatively travel in close proximity to one another at highway speeds. The data rate ranges from a few Kbps up to 65 Mbps depending on whether sensor sharing is required, while the latency ranges from 10 ms to 500 ms depending on the inter-vehicle distance. Vehicle platooning poses also very strict requirements in terms of connection reliability.
- *Advanced Driving.* It enables semi- or fully-automated driving through persistent dissemination of perception data. While the required data rate is relatively low (i.e., less than 50 Mbps), the latency must be very small (i.e., less than 100 ms for high degree of automation) to ensure prompt reactions to unpredictable events.
- *Extended Sensors.* It enables the exchange of raw or processed data gathered through local sensors, thereby enhancing the perception range of the vehicles beyond the capabilities of their on-board instrumentation. The data rate demands are proportional to the resolution of the acquired sensory data and range from around 10 Mbps for a 300-beam 32-bit LIDAR up to approximately 1 Gbps for high-quality uncompressed camera images [13]. Due to the sensitive nature of the exchanged information, the maximum tolerable latency varies from approximately 3 ms up to 100 ms for lower degrees of automation.

• *Remote Driving.* It enables remote control of a vehicle by either a human operator or cloud computing, to support coordination between vehicles in dangerous conditions. Remote driving services require high uplink throughput connections (i.e., around 25 Mbps) and an end-to-end latency lower than 5 ms for fast vehicle teleoperations. Moreover, ultra-high reliability (i.e., 99.999% or higher) shall be guaranteed to avoid application malfunctions.

Given the variety of automotive services and the heterogeneity of their requirements, it is unlikely that V2I communications will be supported by a single radio solution, rather the orchestration of multiple technologies is recommended. In this section we therefore overview the characteristics of candidate radio interfaces currently being considered for V2I communications, i.e., the LTE and the mmWave technologies.

A. Long Term Evolution (LTE)

Since its inception, the LTE cellular technology, operating in the sub-6 GHz spectrum, has represented an ideal candidate to support V2I operations [14]. First, LTE relies on a capillary deployment of eNodeBs (eNBs) offering wide area coverage and long-lived connectivity. Second, resource allocation is centrally managed by an eNB at every transmission opportunity, thereby satisfying service quality constraints while managing priorities in case of V2I applications competing for resources [12]. Third, LTE operates through omnidirectional transmissions and therefore supports broadcast data distribution [2]. Fourth, the LTE interface may guarantee transfer latencies in the radio access theoretically lower than 100 ms, which is particularly beneficial for delay-sensitive vehicular applications.

Nevertheless, LTE was originally designed for mobile broadband traffic and its capability to support V2I communications is still an open issue. The main concern comes from LTE's architecture, that is configured to keep non-active terminals in idle mode: transitions to connected mode may require several seconds [15], which is intolerable for vehicular services. The access and transmission latency also increases with the number of users in the cell, thus raising issues. Moreover, despite the almost ubiquitous coverage of LTE, still the connection may not be always available (e.g., in underground areas). Finally, LTE offers limited downlink capacity, which might not be enough to satisfy the requirements of some V2I applications.

B. Millimeter Waves (mmWaves)

Recently, the mmWave band has been investigated as a means to enhance automated driving and address the stringent throughput and latency demands of emerging vehicular applications. These frequencies, combined with high-order modulation and Multiple Input Multiple Output (MIMO) techniques, offer orders of magnitude higher bitrates than legacy vehicular technologies [4]. Moreover, unlike in LTE, mmWave systems operate through highly directional communications which tend to isolate the users and deliver reduced interference. Inherent security and privacy is also improved because of the short-range transmissions which are typically established [16].

Although mmWave-assisted V2I operations are very attractive from the throughput perspective, they still pose significant challenges [6], [17]. Signals propagating in the mmWave spectrum suffer from severe path loss and susceptibility to shadowing, thereby preventing long-range transmissions (assuming isotropic propagation). Furthermore, directionality requires precise beam alignment of the transmitter and the receiver. In high density and/or high mobility scenarios, the corresponding peer may change frequently, thus implying increased control overhead and communication disconnections. Additionally, mmWave links are highly sensitive to blockage and have ever more stringent requirements on electronic components, size, and power consumption. Given that the challenging radio conditions derived from the mobility of vehicles are further exacerbated considering the dynamic topology of the vehicular networks, the direct applicability of mmWave technology to a V2I deployment is still not clear and has become a research focus in the area of intelligent autonomous systems [18].

III. EVALUATION METHODOLOGY AND SETUP

In this section we give an overview of the methodology we use to assess the performance of the V2I deployment. In detail, in Sec. III-A we briefly describe the architecture of the LTE and the mmWave modules in ns-3, while in Secs. III-B and III-C we introduce our system-level simulation parameters and performance metrics, respectively.

A. The ns-3 Architecture

Our performance evaluation is conducted using ns-3 [19], an open source software which allows the simulation of complex networks with a very high level of detail. The ns-3 simulator features both an LTE and a mmWave full protocol stack, as described in the following paragraphs.

LTE Module. The LTE ns-3 module, designed by Centre Tecnològic Telecomunicacions Catalunya (CTTC) in 2011, provides a basic implementation of LTE devices, including propagation models, PHY and MAC layers. A complete description of the LTE module is presented in [20]: it features (i) a basic implementation of both the User Equipment (UE) and the eNB devices, (ii) Radio Resource Control (RRC) entities for both the UE and the eNB, (iii) handover mechanisms for UE mobility management, (iv) Radio Resource Management (RRM) of the data radio bearers, the MAC queues and the Radio Link Control (RLC) instances, (v) support for both

uplink and downlik packet scheduling, and (vi) a PHY layer model with Resource Block (RB) level granularity.

The path loss is based on pure geometric considerations which deterministically evaluate whether the V2I link is blocked or not by buildings. In this paper, we further extend the LTE module introducing a probabilistic model for the characterization of the channel between the UE and the eNB devices as a function of the distance d for both Line of Sight (LOS) and Non Line of Sight (NLOS) propagation [21]. In case of urban (UMi) scenario, a vehicle is in LOS with probability

$$P_{\text{LOS}}^{\text{UMi}}(d) = \min\left(\frac{0.018}{d}, 1\right) \left[1 - \exp\left(\frac{-d}{0.063}\right)\right] + \exp\left(\frac{-d}{0.063}\right)$$

and in NLOS with probability $P_{\rm NLOS}(d) = 1 - P_{\rm LOS}^{\rm UMi}(d)$. The path loss, for both LOS and NLOS cases, is implemented in the new Lte3gppPropagationLossModel class following the model in [21]. In case of highway (RMa) scenario, the channel characterization follows the Friis free-space model. In addition, we consider a fast Rayleigh fading, which is modeled as a stochastic gain with unit power (in linear scale).

mmWave Module. The mmWave ns-3 module, designed by NYU and the University of Padova in 2015, builds upon the LTE module and implements a complete 3GPP-like cellular protocol stack. A complete description of the mmWave module is presented in [22]: it features (i) a custom PHY/MAC layer implementation for both UE and eNB devices, (ii) support for directional transmissions through analog beamforming, (iii) a dynamic Time Division Duplexing (TDD) scheme designed for low latency communications, (iv) RLC, Packet Data Convergence Protocol (PDCP) and RRC layers, and (v) a complete TCP/IP protocol suite.

The propagation model is based on the 3GPP channel model for frequencies above 6 GHz [23], which characterizes the time correlation among the channel impulse responses to account for spatial consistency. The LOS probability for both UMi and RMa scenarios, in case the propagation is free of obstructions, is given by

$$\int P_{\text{LOS}}^{\text{RMa}}(d) = \exp\left(-\frac{d-10}{1000}\right) \text{ for } 10 \text{ m} < d;$$
(2)

$$P_{\text{LOS}}^{\text{UMi}}(d) = \frac{18}{d} + \exp\left(-\frac{d}{36}\right) \left(1 - \frac{18}{d}\right) \text{ for } 18 \text{ m} < d. \quad (3)$$

The path loss, for both LOS and NLOS cases, is finally implemented in the MmWave3gppPropagationLossModel class, as described in [23, Sec. 7.4]. Moreover, since the effects of high mobility result in rapidly time-varying channels at mmWaves, ns-3 implements a very detailed fading model in the MmWave3gppChannel class. In particular, the model characterizes spatial clusters, subpaths, angular beamspreads and the Doppler shift, which is a function of the total angular dispersion, carrier frequency and mobile velocity.

B. Simulation Setup

The simulation parameters are based on realistic system design considerations and are summarized in Table II. The

TABLE II: Main simulation parameters.

Parameter	Value	
mmWave bandwidth $W_{\rm mmW}$	1 GHz	
mmWave carrier frequency $f_{\rm c,mmW}$	28 GHz	
eNB array size $M_{\rm UPA,eNB}$	8×8	
Vehicle array size $M_{\rm UPA,V}$	4×4	
LTE carrier frequency $f_{c,LTE}$	2 GHz	
LTE bandwidth $W_{\rm LTE}$	20 MHz	
eNB density λ_{eNB}	$\{480\}/km^2$	
TX power P_{TX}	30 dBm	
Packet size D	1400 B	
Noise figure NF	5 dB	
Application rate R	$\{224, 11, 1\}$ Mbps	
RLC buffer size $B_{\rm RLC}$	10 MB	
Vehicles per $eNBM_V$	10	
RLC AM reordering timer $\tau_{\rm RLC}$	1 ms	
Number of simulation runs $N_{\rm sim}$	100	

mmWave and LTE eNBs are deployed over an area of 500×500 meters according to a Poisson Point Processes (PPP) of density λ_{eNB} , with λ_{eNB} varying from 4 to 80 eNB/km² (the trade-off involves signal coverage and deployment cost). We also deploy an average of $M_V = 10$ vehicles per eNB, as foreseen in [24] for a dense environment. We consider urban and highway scenarios, i.e., UMi-Street-Canyon and RMa, according to the 3GPP terminology, to characterize different mobility and propagation regimes, as described in Sec. III-A.

At the PHY layer, LTE eNBs operate in the 2 GHz band, with 20 MHz of bandwidth and omnidirectional transmissions. Conversely, mmWave eNBs operate at 28 GHz with 1 GHz of bandwidth and are equipped with Uniform Planar Arrays (UPAs) of 8×8 elements to establish directional communications through beamforming. Vehicles are also equipped with 4×4 UPAs.¹ For both LTE and mmWave systems, the transmission power and noise figure are set to $P_{\rm TX} = 30$ dBm and NF = 5 dB, respectively.

The MAC layer performs Hybrid Automatic Repeat reQuest (HARQ) to enable fast retransmissions in case of corrupted receptions, and the RLC layer, whose buffer is $B_{\rm RLC} = 10$ MB, uses Acknowledged Mode (AM) to offer additional reliability.

User Datagram Protocol (UDP) is used at the transport layer. Each vehicular application generates packets of D = 1400bytes at a constant interarrival rate ranging from $\tau_{\min} = 50 \ \mu s$ $\tau_{\max} = 10000 \ \mu s$, corresponding to application rates ranging from $R_{\max} \simeq 224$ Mbps to $R_{\min} \simeq 1$ Mbps, to test the performance of LTE and mmWaves in relation with different service requirements. According to Table I, high-rate transmissions are compatible with V2I applications offering extended sensor sharing services, while lower source rates are typical of platooning systems having very stringent requirements in terms of communication delay and reliability but for which the size of the exchanged messages is reasonably small.

C. Performance Metrics

The statistical results are derived through a Monte Carlo approach, where 100 independent simulations are repeated to get different quantities of interest. In particular, we analyze the following end-to-end performance metrics.



Fig. 1: Average UDP throughput vs. $\lambda_{e\rm NB}$ and the application rate in urban scenarios. Narrow (wide) bars refer to a mmWave (LTE) system.

- Average UDP throughput, which is computed as the total number of received bytes divided by the total simulation time, averaged over the $N_{\rm sim}$ simulations.
- *Total UDP throughput*, the sum of the throughput experienced by all vehicles within the coverage of a given eNB.
- 5th (and 10th) percentile UDP throughput, the average throughput relative to the worst 5% (10%) of the vehicles (which represents the performance of cell-edge nodes, the most resource-constrained network entities).
- Average UDP latency, which is measured for each packet, from the time it is generated at the application layer to when it is successfully received assuming perfect beam alignment (it is therefore the latency of only the correctly received packets).
- Jain's fairness index, which is used to determine whether vehicles are receiving a fair share of the cell resources. This index is defined as

$$J = \frac{\left(\sum_{i=1}^{M_V} S_i\right)^2}{M_V \sum_{i=1}^{M_V} S_i^2},$$
(4)

where M_V is the number of users in the cell and S_i is the throughput experienced by the i-th vehicle. The result ranges from $1/M_V$ (most unfair) to 1 (most fair).

IV. END-TO-END PERFORMANCE EVALUATION

In this section we provide some numerical results to evaluate the end-to-end performance of the LTE and the mmWave technologies in a V2I scenario.

UDP Throughput. Fig. 1 shows the average experienced UDP throughput for different eNB densities. We observe that, for the low source rate scenario (i.e., R = 1 Mbps), both LTE and mmWave systems deliver comparable values of throughput, which is almost equal to the full rate offered by the application layer. Conversely, higher-rate applications (i.e., R = 11 Mbps and R = 224 Mbps) are not well-supported by LTE connections which are constrained by the limited capacity of the low-bandwidth physical channel. The mmWave spectrum, in turn, offers orders of magnitude higher data rates than lower frequencies even in case of congested channels

¹For a discussion on the impact of the antenna array size on the overall system performance, we refer the interested reader to our previous work [8].



Fig. 2: Average UDP latency vs. $\lambda_{e\rm NB}$ and the application rate in urban scenarios. Narrow (wide) bars refer to a mmWave (LTE) system.

(110 Mbps vs. 7 Mbps for $\lambda_{eNB} = 80 \text{ eNB/km}^2$), although still not satisfying the requirements of the most demanding vehicular services.

Moreover, we see that the throughput generally increases with the eNB density, as a consequence of stronger channels. The effect of densification is particularly evident for mmWave networks (i.e., the throughput increases by more than 50% from 4 to 80 eNB/km² for R = 224 Mbps) since the endpoints are progressively closer thus guaranteeing improved signal quality and higher received power. On the other hand, densification delivers very negligible improvements for the LTE case due to the more serious impact of interference in case of omnidirectional communications. The above discussion exemplifies how, unlike in legacy networks, the harsh propagation characteristics of the above-6 GHz bands advocate for a high-density deployment of eNBs, to guarantee LOS at any given time and decrease the outage probability [25].² We finally highlight that, for low-rate applications, the UDP traffic injected in the system is sufficiently well handled by the buffer, with no overflow, also considering sparsely deployed networks.

UDP Latency. In Fig. 2 we measure the average communication latency as a function of λ_{eNB} for different application rates.³ We observe that, for R = 1 Mbps, both LTE and mmWave guarantee very low latency (i.e., below 20 ms) since the MAC buffers are empty most of the time. For R = 11Mbps, although the two technologies were proven to offer comparable average throughput (7 Mbps vs. 9 Mbps, respectively, for $\lambda_{eNB} = 40$ eNB/km²), mmWave systems guarantee 5 times lower latency than legacy systems, which cannot ensure time critical message dissemination in case of highly saturated channels. For higher application rates, the end-to-end latency increases uncontrollably in all investigated configurations as a consequence of longer queueing at the MAC layer, although the overall average latency for the mmWave deployment (i.e., around 150 ms for $\lambda_{eNB} = 40 \text{ eNB/km}^2$) is still more than 50% lower than its LTE counterpart.



(a) Highway (RMa) vs. urban (UMi) channel model in LTE systems.



(b) Highway (RMa) vs. urban (UMi) channel model in mmWave systems.

Fig. 3: Total UDP throughput for different values of the λ_{eNB} and application rate. Dotted (straight) bars refer to a highway (urban) scenario. The performance of the LTE and the mmWave technologies are compared.

Additionally, Fig. 2 illustrates that increasing the eNB density in mmWave scenarios has beneficial effects in terms of latency reduction (more than 35% as a results of densification from 4 to 80 eNB/km² for R = 224 Mbps) as compared to a reduction of only 1% in case of LTE connections. According to Table I, these latencies would likely satisfy the envisioned requirements for most next-generation V2I applications (e.g., those supporting advanced driving services).

UMi vs. RMa Propagation. In Fig. 3 we plot the total UDP throughput as a function of the eNB density for both highway (RMa) and urban (UMi) scenarios. Considering highly saturated channels (i.e., R = 224 Mbps), RMa generally guarantees throughput improvements (i.e., +40% and +63% for LTE and mmWave systems, respectively) with respect to UMi in case of sparse, thus noise-limited, networks (i.e., $\lambda_{eNB} = 4$ eNB/km²): despite the increased Doppler effect of high mobility scenarios, free-space propagation indeed results in reduced outage probability. On the other hand, when pushing the network into interference-limited regimes, thus for dense and extremely dense deployments, the gain progressively reduces with λ_{eNB} because of the increasing impact of the interference from the surrounding cells. For LTE deployments (Fig. 3a), the

²Overdensification, in turn, might lead to performance degradation if the number of handovers increases uncontrollably, e.g., in high mobility scenarios.

³The above results were derived considering an RLC buffer of 10 MB. However, the buffer size is critical for the performance of the network [26].



(b) The 10th percentile UDP throughput.

Fig. 4: The 5th and 10th percentile UDP throughput vs. $\lambda_{e\rm NB}$ and the application rate in urban scenarios. The performance of the LTE and the mmWave technologies are compared.

total throughput starts decreasing for $\lambda_{eNB} \ge 40 \text{ eNB/km}^2$ as a result of RMa propagation which, while generally ensuring better signal quality, increases interfering signal strength unintentionally due to the transition of a large number of interference paths from NLOS to LOS [27]. For mmWave deployments (Fig. 3b), RMa propagation induces more than 30% throughput decrease for $\lambda_{eNB} = 80 \text{ eNB/km}^2$ compared to UMi propagation. In fact, while in the highway environment the propagating signals attenuate over distance following the square power law, i.e., Friis' law, the waveguide effect resulting from the more likely signal reflections and scattering in dense urban canyons results in reduced attenuation. Moreover, the presence of blockages in the UMi scenario may actually reduce the impact of the interference from neighboring eNBs when the obstructions block the interfering signals [28].

Considering non-congested scenarios (i.e., R = 1 Mbps) instead, Fig. 3 proves that the experienced throughput becomes independent of the eNB density and the propagation environment since both UMi and RMa channels, regardless of their propagation characteristics, can support well the loose requirements typical of low source rate V2I applications.

5th/10th Percentile Throughput. Fig. 4 represents the 5th and 10th percentile throughput for different application



Fig. 5: Jain's fairness index of the UDP throughput vs. $\lambda_{e\rm NB}$ and the application rate in urban scenarios. Narrow (wide) bars refer to a mmWave (LTE) system.

rates. First, we observe that, for sparsely deployed networks, LTE eNBs offer higher throughput to cell-edge vehicles than mmWave eNBs. In this region, most vehicles are in NLOS and, unlike sub-6 GHz propagation, the challenging communication characteristics of high-frequency channels might result in outage to the serving cell. Moreover, as edge vehicles are power-limited, they are unable to fully exploit the potential of the increased spectrum availability at mmWaves [25]. Densification, in turn, increases the LOS probability and avoids the presence of coverage holes, thereby making the LTE and mmWave radio solutions roughly comparable in terms of cell-edge throughput.

Second, Fig. 4 shows that, for LTE deployments, the mutual interference from omnidirectional eNBs eventually impacts on the cell-edge throughput, which decreases for increasing values of λ_{eNB} . Similarly, we see that, although the directional nature of mmWave systems guarantees reduced interference, there are some special cases where interference is not negligible, i.e., when $\lambda_{eNB} > 45$ eNB/km² for R = 224 Mbps.

Third, while for LTE the 5th and 10th percentile rates reported in Fig. 4 compare similarly to the average values measured in Fig. 1, mmWave systems alone cannot provide uniform capacity, with cell-edge users suffering significantly. In particular, the 5th percentile throughput experiences a dramatic 475 fold decrease (from around 100 Mbps to only 200 Kbps for R = 224 Mbps and considering $\lambda_{eNB} = 40$ eNB/km²) with respect to average conditions, demonstrating a significant limitation of mmWaves under NLOS propagation.

Fairness. In Fig. 5 we plot Jain's fairness index, defined in Eq. (4), for the average vehicle throughput considering both LTE and mmWave scenarios. Although fairness is not always required (e.g., some categories of applications, like those supporting time-critical safety operations, deserve prioritization), it still represents a major concern that should be taken into account to guarantee a minimum performance also to the cell-edge users (or, in general, to users experiencing bad channel conditions). We observe that, for LTE systems, Jain's index is very close to 1 for all density configurations, indicating that (i) cell-edge vehicles experience a throughput comparable to that of other vehicles in the cell regardless of the source application rate, and (ii) densification has a negligible impact on the overall network performance. Conversely, mmWave deployments are generally not compatible with fairness. In particular, the effect of a highly saturated network (i.e., R = 224 Mbps) makes Jain's index fall by an impressive 45% (for $\lambda_{eNB} = 40 \text{ eNB/km}^2$) compared to LTE propagation, as a result of the increased time-variability of the mmWave channel due to scattering and reflection from nearby obstructions, and due to higher Doppler spread. However, such effect is partially mitigated considering denser deployments, i.e., as the probability of path loss outage decreases: in this case, the system is able to increase the coverage of celledge users, i.e., the most resource-constrained network entities. and consequently, provide more uniform quality of service throughout the network (for example, J increases by more than 30% when going from 4 to 80 eNB/km²).

V. CONCLUSIONS AND OPEN CHALLENGES

In this paper we provide the first end-to-end performance comparison between the LTE and the mmWave technologies in a V2I deployment. The impact of several automotive-specific parameters (i.e., the eNB density, the vehicular scenario and the application data rate) was investigated in terms of experienced throughput, communication latency and fairness. We concluded that, although LTE delivers a good compromise between fairness and low latency, the combination of massive bandwidth and spatial degrees of freedom has the potential for mmWave systems to meet some of the boldest requirements of next-generation transportation systems, including high peak per user data rate and very low latency, both in urban and highmobility highway scenarios. We also demonstrated that, unlike in legacy V2I networks, densification of mmWave eNBs is beneficial, for urban propagation, to decrease the outage probability and deliver uniform service quality throughout the cell. In this context, the end-to-end communication performance can be improved by using multiple radios in parallel (i.e., hybrid *networking*), to complement the limitations of each type of network and deliver more flexible and resilient transmissions.

This work opens up interesting research directions. In particular, we will consider more realistic traffic models and more complex evaluation scenarios to address dynamic topologies. Moreover, we will design methods to identify the best radio solution as a function of channel characteristics and the environment in which the vehicles are deployed.

REFERENCES

- 3GPP, "Service requirements for enhanced V2X scenarios (Release 15)," TS 22.186, Sept 2018.
- [2] G. Araniti, C. Campolo, M. Condoluci, A. Iera, and A. Molinaro, "LTE for vehicular networking: a survey," *IEEE Communications Magazine*, vol. 51, no. 5, pp. 148–157, May 2013.
- [3] N. Lu, N. Cheng, N. Zhang, X. Shen, and J. W. Mark, "Connected Vehicles: Solutions and Challenges," *IEEE Internet of Things Journal*, vol. 1, no. 4, pp. 289–299, Aug 2014.
- [4] 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, pp. 160–167, December 2016.

- [5] J. Hasch, E. Topak, R. Schnabel, T. Zwick, R. Weigel, and C. Waldschmidt, "Millimeter-Wave Technology for Automotive Radar Sensors in the 77 GHz Frequency Band," *IEEE Transactions on Microwave Theory* and Techniques, vol. 60, no. 3, pp. 845–860, March 2012.
- [6] M. Giordani, A. Zanella, and M. Zorzi, "Millimeter wave communication in vehicular networks: Challenges and opportunities," in 6th International Conference on Modern Circuits and Systems Technologies (MOCAST), May 2017.
- [7] K. Mase, J. Inoue, and M. Kizu, "Performance evaluation of a roadsideto-vehicle communication system using narrow antenna beam switching based on traffic flow model," in *IEEE GLOBECOM Workshops*, 2008.
- [8] M. Giordani, A. Zanella, T. Higuchi, O. Altintas, and M. Zorzi, "Performance Study of LTE and mmWave in Vehicle-to-Network Communications," *IEEE 17th Annual Mediterranean Ad Hoc Networking Workshop* (*Med-Hoc-Net*), 2018.
- [9] V. Va, T. Shimizu, G. Bansal, and R. W. Heath, "Beam design for beam switching based millimeter wave vehicle-to-infrastructure communications," in *IEEE International Conference on Communications (ICC)*, May 2016.
- [10] M. Giordani, M. Rebato, A. Zanella, and M. Zorzi, "Coverage and connectivity analysis of millimeter wave vehicular networks," *Ad Hoc Networks*, vol. 80, pp. 158–171, Aug 2018.
- [11] A. Tassi, M. Egan, R. J. Piechocki, and A. Nix, "Modeling and Design of Millimeter-Wave Networks for Highway Vehicular Communication," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 12, pp. 10676– 10691, Dec 2017.
- [12] M. Boban, A. Kousaridas, K. Manolakis, J. Eichinger, and W. Xu, "Connected Roads of the Future: Use Cases, Requirements, and Design Considerations for Vehicle-to-Everything Communications," *IEEE Vehicular Technology Magazine*, vol. 13, no. 3, pp. 110–123, Sept 2018.
- [13] S. Kim, B. Qin, Z. J. Chong, X. Shen, W. Liu, M. H. Ang, E. Frazzoli, and D. Rus, "Multivehicle Cooperative Driving Using Cooperative Perception: Design and Experimental Validation," *IEEE Transactions* on *Intelligent Transportation Systems*, vol. 16, no. 2, pp. 663–680, April 2015.
- [14] 3GPP, "Study on LTE support for Vehicle to Everything (V2X) services (Release 14)," *TR* 22.885, 2015.
- [15] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN) (Release 8)," *TR* 36.300, 2018.
- [16] T. S. Rappaport, R. W. Heath Jr, R. C. Daniels, and J. N. Murdock, *Millimeter wave wireless communications*. Pearson Education, 2014.
- [17] Z. Pi and F. Khan, "An introduction to millimeter-wave mobile broadband systems," *IEEE Communications Magazine*, vol. 49, no. 6, June 2011.
- [18] 3GPP, "Study on NR beyond 52.6 GHz (Release 16)," RP-181435, 2018.
- [19] T. R. Henderson, M. Lacage, G. F. Riley, C. Dowell, and J. Kopena, "Network simulations with the ns-3 simulator," *SIGCOMM demonstration*, vol. 14, no. 14, p. 527, 2008.
- [20] G. Piro, N. Baldo, and M. Miozzo, "An LTE module for the ns-3 network simulator," in *Proceedings of the 4th International Conference* on Simulation Tools and Techniques, 2011, pp. 415–422.
- [21] 3GPP, "Study on Small Cell enhancements for E-UTRA and E-UTRAN; Higher layer aspects (Release 12)," TR 36.842, 2014.
- [22] M. Mezzavilla, M. Zhang, M. Polese, R. Ford, S. Dutta, S. Rangan, and M. Zorzi, "End-to-End Simulation of 5G mmWave Networks," *IEEE Communications Surveys Tutorials*, vol. 20, no. 3, pp. 2237–2263, thirdquarter 2018.
- [23] 3GPP, "Study on channel model for frequencies from 0.5 to 100 GHz (Release 14)," TR 38.901, 2018.
- [24] 3GPP, "Study on scenarios and requirements for next generation access technologies (Release 14)," *TR* 38.913, 2018.
- [25] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter-wave cellular wireless networks: Potentials and challenges," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, March 2014.
- [26] M. Zhang, M. Polese, M. Mezzavilla, J. Zhu, S. Rangan, S. Panwar, and a. M. Zorzi, "Will TCP Work in mmWave 5G Cellular Networks?" *IEEE Communications Magazine*, vol. 57, no. 1, pp. 65–71, January 2019.
- [27] M. Ding, P. Wang, D. López-Pérez, G. Mao, and Z. Lin, "Performance Impact of LoS and NLoS Transmissions in Dense Cellular Networks," *IEEE Transactions on Wireless Communications*, vol. 15, no. 3, pp. 2365–2380, March 2016.
- [28] J. Kim, J. Park, S. Kim, S. Kim, K. W. Sung, and K. S. Kim, "Millimeter-Wave Interference Avoidance via Building-Aware Associations," *IEEE Access*, vol. 6, pp. 10618–10634, Feb. 2018.