MilliCar – An ns-3 Module for mmWave NR V2X Networks

Matteo Drago, Tommaso Zugno, Michele Polese, Marco Giordani, Michele Zorzi

Department of Information Engineering, University of Padova, Padova, Italy

e-mail: {name.surname}@dei.unipd.it

ABSTRACT

Vehicle-to-vehicle (V2V) communications have opened the way towards cooperative automated driving as a means to guarantee improved road safety and traffic efficiency. The use of the millimeter wave (mmWave) spectrum for V2V, in particular, holds great promise since the large bandwidth available offers the possibility of realizing high-data-rate connections. However, this potential is hindered by the significant path and penetration loss experienced at these frequencies. It then becomes fundamental to practically evaluate the feasibility of installing mmWave-based technologies in the vehicular scenario, in view of the strict latency and throughput requirements of future automotive applications. To do so, in this paper we present MilliCar, the first ns-3 module for V2V mmWave networks, which features a detailed implementation of the sidelink Physical (PHY) and Medium Access Control (MAC) layers based on the latest NR V2X specifications, the 3GPP standard for next-generation vehicular systems. Our module is open-source and enables researchers to compare possible design options and their relative performance through an end-to-end full-stack approach, thereby stimulating further research on this topic.

CCS CONCEPTS

• **Networks** → **Network simulations**; **Mobile networks**;

KEYWORDS

ns-3, NR V2X, 3GPP, mmWave, vehicular

ACM Reference Format:

Matteo Drago, Tommaso Zugno, Michele Polese, Marco Giordani, and Michele Zorzi. 2020. MilliCar – An ns-3 Module for mmWave NR V2X Networks. In *2020 Workshop on ns-3 (WNS3 2020)*, June 17–18, 2020, Gaithersburg, MD, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3389400.3389402>

1 INTRODUCTION

The 5th generation (5G) of mobile networks will support new market verticals and applications, thanks to a flexible design of the network architecture and of the protocol stack [2]. One of the fundamental novelties of NR (the 3rd Generation Partnership Project (3GPP) Radio Access Network (RAN) for 5G) is the communication at millimeter wave (mmWave) frequencies. 3GPP NR, indeed, supports a carrier frequency up to 52.6 GHz in Release 15, with possible extensions to the even higher spectrum planned for Releases 16

and 17. The usage of mmWaves enables unprecedented data rates (in the order of multi gigabit per second) in the access network [7], thanks to the availability of large chunks of free spectrum (a single carrier in NR can operate in a bandwidth of up to 400 MHz).

In particular, the mmWave bands support the most bandwidth-hungry applications in a vehicular context, where sensors (e.g., radars, LiDARs, cameras) and infotainment systems are expected to generate data at a rate of Terabytes per driving hour [9]. Standardization bodies have already started considering the integration of mmWaves in the specifications for next-generation vehicular networks, such as IEEE 802.11bd [14] and 3GPP NR V2X [6].

The communication at mmWaves in a vehicular scenario, however, introduces numerous challenges that need to be carefully addressed to guarantee a high quality of service for the end users and make such systems reliable and robust to the vehicle mobility [12]. First, the severe propagation loss experienced at high frequencies prevents long-range communications. This can be partially compensated using large antenna arrays, that can focus the transmitted power in sharp beams and increase the link budget. The usage of directional transmissions, however, implies a coordination between the two communication endpoints, which have to continuously align their beams to experience a sufficient beamforming gain [16]. Moreover, mmWave signals can be easily blocked by common obstacles such as, for example, vehicles, road signs, and pedestrians. Additionally, the metal of the vehicle bodies acts as a reflector for mmWaves, creating strong and bursty interference [21]. The interaction of the mmWave channel with the vehicular environment also creates an extremely dynamic channel, with a fluctuating capacity and link availability, that may have a serious impact on the whole protocol stack [13].

These challenges become even harder in a Vehicle-to-Vehicle (V2V)¹ context, which, in principle, should be designed to work also in the absence of the fixed network infrastructure that can be exploited, instead, in a Vehicle-to-Infrastructure (V2I) use case. The research community has started proposing solutions to improve the performance of V2V communications at mmWaves [17, 20]. As of today, the lack of testbeds for mmWave V2V scenarios makes simulation the preferred means for the performance evaluation of novel networking designs. However, to the best of our knowledge, an open-source, publicly available network simulator that integrates mmWaves and V2V scenarios is not currently available: simulation tools for mmWaves, indeed, only support fixed infrastructure scenarios [18, 19], while simulators for ad-hoc communications (e.g., Device-to-Device (D2D)) only model sub-6 GHz frequencies [22].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.
WNS3 2020, June 17–18, 2020, Gaithersburg, MD, USA

© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-7537-5/20/06...\$15.00
<https://doi.org/10.1145/3389400.3389402>

¹While 3GPP refers to NR-V2X, i.e., the incorporation of different paradigms, including Vehicle-to-(Roadway) Infrastructure (V2I), Vehicle-to-Network (V2N), Vehicle-to-Vehicle (V2V), Vehicle-to-Pedestrian (V2P) and Vehicle-to-Device (V2D), our effort is focused specifically on the design of V2V protocols. From now on, when we refer to the 3GPP standard, we will use NR-V2X, whereas when we describe our module's operations we will use V2V terminology.

Figure 1: General Overview of the MilliCar Module, with the Features Implemented at each Layer of the Stack

To fill this gap, in this paper we introduce MilliCar, an open-source ns-3 module for V2V mmWave networks². The module introduces a characterization of sidelink Physical (PHY) and Medium Access Control (MAC) layers that follow the 3GPP numerologies for NR V2X [6], and enables the study and development of beamforming, link adaptation and medium access techniques for mmWave V2V in end-to-end, full-stack simulations. Additionally, the module features the 3GPP channel model introduced in [5], which has been designed for vehicular simulations in urban and highway scenarios. MilliCar integrates the Long Term Evolution (LTE) Service Access Point (SAP) to connect the MAC layer to Radio Link Control (RLC) and Packet Data Convergence Protocol (PDCP), and implements a new ns-3 NetDevice (i.e., MmWaveVehicularNetDevice) to take care of the integration with the TCP/IP stack of ns-3. Finally, the module incorporates (i) a helper, that can be used to easily set up simulations; (ii) unit tests, to guarantee that the module behaves as expected even when adding new features; and (iii) several examples, to simulate scenarios with a varying number of vehicles and different deployments.

MilliCar was developed as a standalone module (e.g., with respect to the other ns-3 modules that support mmWaves [18, 19]), to separate the sidelink implementation from that of scheduled cellular protocol stacks. One of the design goals of this module, indeed, is a lean implementation, and the possibility to extend it with new features without having to deal with the complexity of protocol stacks that have not been designed from the start to support a sidelink. Nonetheless, the MilliCar module can still rely on the higher layers from the LTE module through SAPs, to run end-to-end simulations with an NR-V2X-like protocol stack. We believe that this constitutes a good tradeoff between integration with ns-3 and flexibility to develop new components.

We also report a preliminary performance evaluation, to validate the main functionalities of the module, with throughput, latency and Signal to Interference plus Noise Ratio (SINR) results in different scenarios, varying system parameters such as the bandwidth, source rate and mobility of vehicles.

The remainder of this paper is organized as follows. In Section 2, we provide an overview on the research and standardization activities related to vehicular networks. The module is described in Section 3, where we comment on the channel model and the sidelink physical and MAC layers. We discuss the structure of tests

and helpers in Section 4, while the example scenarios are introduced in Section 5. Finally, in Section 6 we conclude the paper and provide suggestions for future extensions.

2 NEXT-GENERATION VEHICULAR NETWORKS

The recent evolution of hardware, software, and communication technologies in the automotive sector has paved the way towards Connected Intelligent Transportation Systems (C-ITSs) as a means to support traffic safety, traffic efficiency and infotainment services. More specifically, C-ITSs promise to reduce the number of road accidents and carbon emission by more than 90% and 60% respectively, as well as to improve drivers' productivity during commutes [10]. Overall, the market potential of C-ITSs is estimated in 1.3 trillion USD annually in the U.S. alone, thereby stimulating further research in the field of next-generation vehicular networks.

When fully deployed, C-ITSs will address the demands and business trends of the 2030 society, and support the new use cases highlighted in Figure 1. They include (but are not limited to):

- *platooning*, where vehicles can travel in close proximity to one another at very high speeds;
- *cooperative perception*, where vehicles broadcast sensors' observations to increase the perception range of their on-board instrumentation;
- *semi- or fully-automated driving*, where vehicles can sense the surrounding environment and move with minimal or no human input;
- *infotainment*, i.e., a set of services that deliver a combination of information and entertainment.

These applications have very strict demands in terms of data rate (in the order of Terabytes per driving hour, according to some estimates), latency (from approximately 3 ms up to 100 ms depending on the degrees of automation) and reliability (up to 99.999% for the most safety-critical services) [3]. In order for these heterogeneous requirements to be satisfied, C-ITSs will enable wireless communications to and from roadside infrastructures and among vehicles, a concept that is referred to as Vehicle-To-Everything (V2X) connectivity. While current V2X standard technologies, i.e., IEEE 802.11p and 3GPP Cellular-V2X (C-V2X), may lack the level of reliability and availability requested by future vehicular applications, the IEEE and 3GPP are promoting standardization efforts, i.e., 802.11bd [14] and NR V2X [6] respectively, to overcome current technology limitations [27], as described in the following paragraphs.

²Available at <https://github.com/signetlabdei/millicar>

Figure 2: Simplified Unified Modeling Language (UML) Diagram of the MilliCar Classes

IEEE 802.11bd. The 802.11bd standard will enhance IEEE 802.11p, targeting future V2X application requirements. Although technical details have not yet been thoroughly discussed, the 802.11bd evolution is positioned to support reduced latency and twice the throughput and the communication range of 802.11p through (i) new transmission mechanisms, (ii) dual carrier modulation with midambles to provide better channel estimation in fast-fading channels [15], and (iii) re-designed PHY and MAC features, including a flexible sub-carrier spacing with up to 40 MHz channel bandwidth.

3GPP NR V2X. The NR V2X specifications will amend 3GPP C-V2X as part of the ongoing 5G effort. New developments will include (i) more sophisticated channel models, (ii) sidelink and network architecture improvements, (iii) support of mini-slot scheduling for latency-critical services, (iv) a flexible numerology, along the lines of the 3GPP Rel. 15 specifications, and (v) new resource allocation schemes, where sidelink resources are scheduled either by the base station (mode 1) or autonomously by the vehicles (mode 2).

In particular, both 3GPP NR V2X and IEEE 802.11bd standards will support operations at mmWave frequencies, i.e., in the 24.25 – 52.6 GHz and 57 – 71 GHz ranges, respectively. As discussed in Section 1, the large spectrum availability at these frequencies promotes high-capacity low-latency communication, however working at mmWave raises several challenges from a communication point of view. For this reason, it is important to demonstrate the practical feasibility of designing protocols to support vehicular operations at such high frequencies.

While testbed validation is impractical due to the prohibitive costs of the mmWave hardware and due to the difficulties associated

to the design and implementation of a representative and not-too-simple application scenario, end-to-end simulation is the preferred solution for performance evaluation. Currently, V2I communications can be simulated using the ns-3 mmWave module designed by NYU and the University of Padova (UNIPD) in 2015 [18] or 5G Lena [19], which implement a complete 3GPP-like protocol stack, but focus on cellular-like, infrastructure-based deployments. For the V2V scenario, instead, there are currently no open-source simulators capable of modeling the mmWave channel as well as a full TCP/IP protocol stack and the mobility of vehicles. Specifically, traditional vehicular simulators, e.g., Veins [24] or V2X Simulation Runtime Infrastructure (VSimRTI) [23], support development, training, and validation of autonomous urban driving systems, but do not support mmWave communications.

In this paper, we fill this gap by providing the first open-source ns-3 module for V2V networks operating at mmWaves, so as to enable realistic full-stack simulations in the V2V environment, as described in the next sections.

3 A MODULE FOR NR V2X

In this section we describe the main characteristics of the MilliCar module for V2V networking. The general features of each component of the module are depicted in Figure 1, while Figure 2 provides a simplified UML diagram.

In Section 3.1 we present the channel model proposed by the 3GPP to characterize the V2V propagation at mmWaves. Then, in Secs. 3.2 and 3.3 we present the key elements of the MilliCar PHY and MAC layers, respectively. They provide functionalities for packet transmission and reception over an NR-V2X-compliant frame structure, and a proper scheduling of the radio resources.

As mentioned in Section 2, the study item in [6] considers both in-coverage (mode 1) and out-of-coverage (mode 2) options for resource allocation. The protocol stack of MilliCar natively supports mode 2, which is the more likely to be implemented in an early deployment of NR V2X, given that mode 1 would require an update of base stations following standard specifications [6, 27]. Also, mode 2 is of particular interest for researchers since it poses several challenges that remain to be addressed.

Finally, in Section 3.4 we describe how the module integrates with the higher layers of the ns-3 protocol stack. In particular, the module currently relies on the implementation of the RLC and PDCP layers and of the traffic flow filters of the lte module.

Notice that MilliCar is integrated with ns-3, and reuses a number of data structures and classes (for example - to hold the configuration parameters of the frame structure) from the NYU/UNIPD mmwave module for cellular communications. Finally, we plan to complete the integration with the new features of the spectrum module, described in [28], which would allow us to separate the code base of the channel model from that of the protocol stack.

3.1 Vehicular Channel Model Implementation

The accurate characterization of the channel behavior is paramount to obtain reliable simulation results. Therefore, MilliCar implements the propagation and fading models that the 3GPP suggests for V2V communications at mmWaves [5, 11], by means of the APIs provided by the ns-3 Spectrum module. In particular, devices communicating through the same wireless channel are attached to a single instance of SpectrumChannel which accounts for the modeling of the propagation phenomena using the interfaces PropagationLossModel and SpectrumPropagationLossModel.

The pathloss has been implemented in MmWaveVehicularPropagationLossModel, which extends the PropagationLossModel interface and adopts the structure described in [26]. Following the standard's guidelines [5], we defined two scenarios, V2V-Highway and V2V-Urban, where a pair of vehicles can be in one of the following states:

- Line-of-Sight (LOS), if the vehicles are in the same street or highway and the path is free from obstacles;
- Non-Line-of-Sight (NLOS), if the path is blocked by static objects, e.g., buildings;
- Vehicle Non-Line-of-Sight (NLOSv), if the vehicles are in the same street but the path is blocked by other vehicles.

This state can be fixed during the whole simulation by setting the ChannelCondition attribute. By default, if not specified, the LOS state is updated when the overall channel characterization is updated.

Notice that the model in [5] determines the LOS state in a multi-step process. The first is deterministic, based on the simulation scenario. If two vehicles are on different streets or highways, the channel condition is set to NLOS and the procedure ends. Otherwise, for vehicles that are in the same street or highway, the 3GPP proposes a stochastic model to distinguish between LOS and NLOSv conditions. For the second step in the V2V-Highway scenario, the LOS probability is given by [5, Section 6.2]

$$P_{\text{LOS}} = \begin{cases} \min \{1, a d^2 + b d + c\} & d \leq 475 \text{ m} \\ \max \{0, 0.54 - 0.001(d - 475)\} & d > 475 \text{ m}, \end{cases} \quad (1)$$

where $a = 2.1013 \cdot 10^{-6}$, $b = -0.002$, $c = 1.0193$, and d is the distance between the vehicles, measured in meters. For the V2V-Urban scenario, instead, the LOS probability is evaluated as

$$P_{\text{LOS}} = \min \left\{ 1, 1.05 e^{-0.0114d} \right\}. \quad (2)$$

In both scenarios, $P_{\text{NLOSv}} = 1 - P_{\text{LOS}}$.

To also account for a probabilistic characterization of the NLOS state, the MilliCar module defines two additional scenarios, namely Extended-V2V-Highway and Extended-V2V-Urban, that implement the state probabilities detailed in [8]. For the Extended-V2V-Highway scenario, the LOS and NLOS probabilities are given by

$$P_{\text{LOS}} = \min \left\{ 1, \max \{0, 2.7 \cdot 10^{-6} d^2 - 0.0025 d + 1\} \right\}, \quad (3)$$

$$P_{\text{NLOS}} = \min \left\{ 1, \max \{0, -3.7 \cdot 10^{-7} d^2 + 0.00061 d + 0.015\} \right\}. \quad (4)$$

$$P_{\text{NLOSv}} = 1 - P_{\text{LOS}} - P_{\text{NLOS}}. \quad (5)$$

For the Extended-V2V-Urban scenario, instead, the LOS and NLOS probabilities are equal to

$$P_{\text{LOS}} = \min \{1, \max \{0, 0.8372 e^{-0.0114d}\} \}, \quad (6)$$

$$P_{\text{NLOSv}} = \min \left\{ 1, \max \left\{ 0, \frac{1}{0.0312d} e^{\frac{(-\ln(d)-5.0063))^2}{2.4544}} \right\} \right\}, \quad (7)$$

$$P_{\text{NLOS}} = 1 - P_{\text{LOS}} - P_{\text{NLOSv}}. \quad (8)$$

All these terms represent a vehicular traffic condition of medium density (i.e., 120 vehicles/km² in the urban scenario, and 1500 vehicles/h/dir in the highway scenario).

For each scenario, the link pathloss is implemented accurately following the 3GPP specifications [5]. For the highway scenarios (i.e., V2V-Highway and Extended-V2V-Highway), the LOS pathloss is

$$PL_{\text{LOS}} = 32.4 + 20 \log_{10} d + 20 \log_{10} f_c \text{ [dB]}, \quad (9)$$

where the distance d is in meters and the center frequency f_c is in GHz. For the urban scenarios (i.e., V2V-Urban and Extended-V2V-Urban), instead, the LOS pathloss is given by

$$PL_{\text{LOS}} = 38.77 + 16.7 \log_{10} d + 18.2 \log_{10} f_c \text{ [dB]}. \quad (10)$$

In case of NLOS, the model does not distinguish between urban and highway propagation, so that the pathloss equation becomes

$$PL_{\text{NLOS}} = 36.85 + 30 \log_{10} d + 18.9 \log_{10} f_c \text{ [dB]}. \quad (11)$$

When the link is obstructed by a vehicle (i.e., in NLOSv state), an incremental shadowing term, modeled as a log-normal random variable, is added to the PL_{LOS} equation. This term depends on the blocker height, which is randomly selected among three possible vehicle types (to be configured in the simulation), as detailed in [5].

The operations associated to fast fading and beamforming gain computation have been implemented in the class MmWaveVehicularSpectrumPropagationLossModel which extends the interface of SpectrumPropagationLossModel to guarantee consistency among different channel models provided by ns-3. The procedure to generate the channel is based on the specifications in [4], while the dual mobility (to compute the Doppler effect) and the parameters of the equations that define the channel are modeled as described in [5]. Also, to take into account the non-isotropic behavior of real antennas, we extended the class AntennaArrayModel of the mmwave module, including the radiation model specified in [5]. This

class handles also the computation of the beamforming vectors using a Discrete Fourier Transform (DFT) based approach [25].

3.2 Physical Layer

The MilliCar physical layer is composed of two classes, namely, `MmWaveSidelinkSpectrumPhy` and `MmWaveSidelinkPhy`. `MmWaveSidelinkSpectrumPhy` extends the abstract class `SpectrumPhy` and acts as an interface between `MmWaveVehicularNetDevice` and `SpectrumChannel`. In fact, it handles the transmission and reception operations through the methods `StartTxDataFrames` and `StartRx`. The method `StartTxDataFrames` generates the signal to be transmitted over the channel, represented by the structure `MmWaveSidelinkSpectrumSignalParameters`. Then, it forwards it to the `SpectrumChannel` instance by calling the method `SpectrumChannel::StartTx`. Conversely, when a signal is received from `SpectrumChannel`, the method `StartRx` checks whether or not it can be decoded by applying an error model and, if so, forwards it to the upper layer. The error model that is currently supported by our module is based on the one described in [18], which derives the error probability taking as input the received SINR and the Modulation and Coding Scheme (MCS) used to encode the signal. To compute the SINR, `MmWaveSidelinkSpectrumPhy` relies on the classes `mmWaveInterference` and `mmWaveChunkProcessor`, of the `mmwave` module [18]. Moreover, `MmWaveSidelinkSpectrumPhy` takes care of the periodic generation of the Channel State Information (CSI) reported to the upper layers.

The class `MmWaveSidelinkPhy` is in charge of maintaining the system synchronization, and of managing the physical channel used for the transmission and reception of the transport blocks (our module currently supports the modeling of Physical Sidelink Shared Channel (PSSCH) only). The frame structure used by MilliCar is compliant with NR specifications, i.e., a frame of 10 ms is divided in 10 subframes, each containing a variable number of slots. Each slot is composed of 14 Orthogonal Frequency Division Multiplexing (OFDM) symbols, whose duration depends on the selected numerology configuration [1]. Following the proposal in [6], our module currently supports NR numerologies 2 and 3, i.e., with 4 and 8 slots per subframe, respectively, corresponding to a Sub-carrier Spacing (SCS) of 60 kHz or 120 kHz. Figure 3 shows the implemented frame structure, where different colors represent different possible allocation patterns. A transmission buffer is used to store the transport blocks to be sent during the first available slot, together with information regarding the MCS to use and the allocated OFDM symbols.

The method `StartSlot` marks the beginning of each slot and takes care of transmitting the transport blocks stored in the buffer, by scheduling multiple calls to the method `StartTxDataFrames` of `MmWaveSidelinkSpectrumPhy`. Moreover, `MmWaveSidelinkPhy` takes care of forwarding the received transport blocks to the upper layer and managing the beamforming operations to properly point the beam towards the other end device (at this stage, perfect beam alignment is assumed, and further refinements are left for future work).

3.3 MAC Layer

The MAC layer functionalities are implemented in the `MmWaveSidelinkMac` class, which includes: (i) the management of the medium

Figure 3: Frame Structures Supported by MilliCar

access, (ii) the scheduling of the available resources, (iii) the support of transmission and reception over multiple logical channels, and (iv) the link adaptation.

`MmWaveSidelinkMac` supports a Time Division Multiple Access (TDMA)-based access scheme, where different vehicles transmit in different slots, as generally assumed for directional mmWave operations [18]. Similarly to mode 2c defined in [6], the MAC layer is pre-configured with a fixed scheduling pattern, which determines how the slots are assigned to the vehicles on a per-subframe basis. By default, each vehicle can use a single slot per subframe, but this pattern can be customized using the `SetSfAllocationInfo` method.

At the beginning of each slot, the MAC layer retrieves the scheduling pattern and executes `DoSlotIndication` to decide whether to perform the transmit or receive operation. In case the slot is intended for transmission, the method `ScheduleResources` divides the available resources among the active logical channels³ and notifies the scheduling decision to the upper layers. Then, it builds the transport block using the Service Data Units (SDUs) received from the higher layers, which is then forwarded to the PHY layer by calling the method `AddTransportBlock`. To avoid the allocation of unnecessary resources, the buffers at the upper layers are monitored through a periodic buffer status reporting procedure. Such reports are used to decide the amount of resources to be reserved for each logical channel. Conversely, if the slot is dedicated to another device, the PHY layer is informed about a possible incoming reception by the MAC, which then performs de-multiplexing operations to map the received packets onto the proper logical channels.

Moreover, automatic link adaptation functionalities are provided based on CSI reports received from the PHY. This mechanism is handled by the `MmWaveAmc` class, which uses the last received CSI report to determine the optimal modulation and coding scheme to be used for the transmission.

3.4 Integration with the Higher Layers

MilliCar also provides full integration with the higher layers of the protocol stack ensuring, by means of SAP, high flexibility for future improvements of the stack design. We attach to each `MmWaveSidelinkMac` object multiple instances of `LteRlc`, each connected to an

³A logical channel represents an end-to-end connection at the MAC and physical layers.

LtePdcP object. Closing the gap, a specific class implementing the SAP is used to connect the PDCP object to our ad hoc MmWaveVehicularNetDevice, envisioning in this way a full bottom-up and top-down integration. The instances of these layers for each end-to-end connection are managed inside the MmWaveVehicularNetDevice class, which extends NetDevice and implements all the virtual classes commonly used to set up the communication to and from the TCP/IP stack.

In order for two nodes to communicate, a radio bearer must be set up. Once a MmWaveVehicularNetDevice is associated to each node, the method `MmWaveVehicularNetDevice::ActivateBearer` is executed on both communication endpoints. This function accepts as input an integer number representing the bearerID, the Radio Network Temporary Identifier (RNTI) of the destination (an integer number that differentiates distinct nodes in the network) and the IP address of the pairing node. In particular, each bearerID must unequivocally identify a radio bearer, and cannot be shared among different pairs of devices. The consistency of this assignment (along with that of the RNTI) among different nodes is guaranteed in the helper's configuration method, which will be described in Section 4. At the MAC layer, a bearer is mapped to a logical channel identifier, as defined by the 3GPP standard. However, at this stage of development, logical channels and radio bearers have a one-to-one correspondence.

The operations carried out by the `ActivateBearer` method are:

- the creation of a rule to classify packets generated from different sources, using `EpcTft::PacketFilter`;
- the instantiation of an `LteRlc` object, which can be identified by the RNTI of the destination node and the logical channel identifier. The RLC object is then linked to the MAC layer instance associated to the node;
- the creation of an `LtePdcP` object, which has to be connected to the `MmWaveVehicularNetDevice` and the RLC object created in the previous step.

After these steps, the RLC and PDCP objects are stored in a dedicated structure, i.e., `SidelinkRadioBearerInfo`, which is then identified with the univocal bearerID and saved in the `m_bearerToInfoMap` variable. Currently, the version of the RLC supported by this module is `LteRlcUm`, which provides segmentation and concatenation but no retransmissions.

Once `MmWaveVehicularNetDevice::Send` receives a packet from the IP layer, it accesses the `m_tftClassifier` variable to retrieve the bearerID that associates the RNTI and the logical channel identifier to the packet, and stores them in the `TransmitPdcPParams` struct of `LtePdcPProvider`. This is then forwarded to the PDCP. Conversely, in the reception phase, a packet is simply sent from PDCP to the NetDevice, and from the NetDevice to the upper layers.

4 HELPERS AND TEST FRAMEWORK

The mmWave vehicular module is also equipped with *helpers* (Section 4.1) to allow the users to easily set up the simulation, and *unit tests* (Section 4.2) to check basic functionalities of the module and facilitate future class developments.

4.1 Helpers

The main helper is `MmWaveVehicularHelper` which (i) creates and configures the objects for the channel computation; (ii) computes the parameters for the frame structure, according to the selected 3GPP numerology; (iii) installs the networking stack on the vehicles; and (iv) connects groups of vehicles that will communicate together. The first operation is performed during initialization, and relies on three `StringValue` attributes that configure the propagation loss model (`PropagationLossModel`), the fading model (`SpectrumPropagationLossModel`), and the propagation delay model (`PropagationDelayModel`). A typical configuration would include the propagation and fading classes described in Section 3.1, without a delay model, as this is included in the spectrum model. However, the user can change and select different options (e.g., a simple Friis propagation loss) and combine also a delay model. The `Numerology` attribute, which is linked to the `SetNumerology` method, accepts 2 or 3 as integer value, to select among the two different numerologies currently foreseen for NR V2X.

The method `InstallMmWaveVehicularNetDevices` accepts a container of `Node` objects, and returns the `NetDeviceContainer` with the `MmWaveVehicularNetDevice` objects. Additionally, for each vehicle, this method sets up the instances of the PHY and MAC layers, configures the antenna at the vehicle and connects it to the channel.

The `MmWaveVehicularHelper` also configures and connects to each device another helper, called `MmWaveVehicularTracesHelper`, which, at runtime, generates a trace of the SINR and MCS for each transmitted packet.

Finally, `PairDevices` configures a bearer and connects, pair by pair, all the devices in a container of `NetDevices` passed as input argument. This makes it possible to create multiple groups of vehicles, with independent scheduling patterns, that generate interference with concurrent transmissions. The vehicles in the same group are all logically connected, thus packets could (in principle) be exchanged among any pair of nodes.

4.2 Unit Test

The unit test suite contains four tests. `MmWaveVehicularSidelinkSpectrumPhyTestSuite` has a test case that checks if the Signal-to-Noise-Ratio (SNR) computed by `MmWaveSidelinkSpectrumPhy` for a single transmission is in line with the expected SNR (considering isotropic antennas and an ideal channel). Similarly, `MmWaveVehicularInterferenceTestSuite` tests if the interference among two groups of vehicles is correctly computed. `MmWaveVehicularRateTestCase`, instead, features vehicles equipped with a full protocol stack (with User Datagram Protocol (UDP) at the transport layer), and tests full buffer transmissions for different values of the MCS (i.e., from 0 to 28), checking if there is a one-to-one correspondence between transmitted and received packets.

5 EXAMPLE SCENARIOS

We provide four examples of scenarios in the `examples` folder. They represent different simple use cases for the module, and can be used as a reference by a user of the module to understand which are the necessary steps to generate a vehicular scenario at mmWave

(a) Throughput

(b) Latency

Figure 4: Results of vehicular-simple-one.cc, using V2V-Highway Scenario, Bandwidth BW and Distance between Vehicles d

frequencies. In particular, these are the common operations to be implemented:

- (1) Configure relevant parameters for the simulation scenario, using the ns-3 Config attribute system. For example, relevant parameters are related to the carrier frequency, the system bandwidth, whether to use Adaptive Modulation and Coding (AMC) or not, the type of RLC and the size of its buffer.
- (2) Create nodes of the scenario, setting their positions and mobility models.
- (3) Initialize the MmWaveHelper object, for which it is possible to set the numerology and the channel models, as discussed in Section 4.
- (4) Install the MmWaveVehicularNetDevice objects in the nodes.
- (5) Configure the TCP/IP stack and assign IP addresses to the vehicles.
- (6) Configure and install the applications.

The simplest scenario is vehicular-simple-one.cc. In this example, there are two devices in the same lane, one behind the other at fixed distance and speed. It is possible to select the speed and the simulation duration from the command line, as well as the MCS value or if AMC should be used. The applications are an echo client and server using UDP as transport. The vehicular-simple-two.cc example, instead, features two groups with two vehicles each, moving in two different lanes at opposite speeds. Each vehicle exchanges UDP packets with the other vehicle in the same lane. vehicular-simple-three.cc also has four vehicles, grouped in pairs of two, but moving as a platoon, on the same lane, at a constant speed and at a safety distance of 20 m. Finally, vehicular-simple-four.cc is a scenario with three vehicles, with two of them moving at a constant distance and speed, and the third moving away from the first two. One of the vehicles acts as a server, which echoes the UDP packets sent by the others.

As an example, we report in this section some metrics that it is possible to obtain by running the available examples. Figures 4a and 4b have been obtained choosing different configurations of parameters in vehicular-simple-one.cc, such as the initial distance d between the vehicles, the bandwidth (100 MHz or 400 MHz, which is the maximum for an NR carrier), and the source rate at the application layer. In this simulation campaign the channel was configured to transmit using a carrier frequency $f_c = 28$ GHz, with the LOS condition computed using the equations outlined in

Section 3.1 for the reference scenario V2V-Highway. Moreover, the LteRlcUm buffer size was set to 500 packets. Buffering is necessary to guarantee high throughput, as (i) the scheduling at the MAC layer happens on a slot basis (i.e., 250 μ s with numerology 2), but the application may generate packets more regularly, with interpacket intervals as short as 10 μ s; (ii) the default allocation pattern is used, i.e., each vehicle is assigned a single slot per subframe; and (iii) the fluctuations of the channel may temporarily degrade the capacity available at the physical layer. On the other hand, an excessively large buffer could lead to an increase in the experienced delay of the communication, due to packet buffering in the queue.

Figure 4a compares the throughput (averaged over the simulation time of 2 s, and over 100 independent simulation runs) for different values of the source rate. The results obtained for source rates up to 80 Mbps are similar when using a bandwidth of 100 MHz or 400 MHz, with small differences at different distances. The gap between the two configurations becomes more marked as we increase the application source rate, up to 800 Mbps. Figure 4b, instead, reports the average delay in the same conditions. It has to be highlighted that the end-to-end delay obtained when the channel bandwidth is set to 400 MHz, the initial distance is 50 m and the source rate is ~ 8 Mbps, is in the order of microseconds (~ 1 OFDM symbol). In this case, the traffic injected in the link is negligible with respect to the available resources, and the packet only experiences transmission delay. Future extension of this work will study how to model the delay introduced by the processing of packets at the different layers of the protocol stack.

Figure 5 shows the trend of the instantaneous SINR over time, obtained with a single run of vehicle-simple-two.cc, to validate how the interference is modeled by MilliCar. The plots show a decrease in the received SINR, with a minimum at ~ 1 s, and a consequent increase. This reduction is caused by a separate pair of vehicles interfering with the transmission that the nodes considered for the SINR metric are carrying out. This setup has been evaluated in both scenarios defined by 3GPP, i.e., V2V-Highway (Figure 5a) and V2V-Urban (Figure 5b), and with different dimensions of the antenna array. Notice that, while different architectures have a significant impact on the received power and, thus, on the SINR, the two scenarios differ only by 1.2 dB (for the configuration with 4 antennas at the transmitter and 4 at the receiver).

(a) V2V-Highway scenario

(b) V2V-Urban scenario

Figure 5: SINR for a Pair of Vehicles in vehicular-simple-two.cc. 4×4 (8×8) Indicates that Both Endpoints Have Antenna Arrays with 4 (8) Elements

6 CONCLUSIONS

The wireless networking standardization bodies have started focusing on new market verticals to find new use cases and applications for 5G and beyond. For example, the 3GPP is considering vehicle-to-vehicle communications in NR V2X, which will support vehicular networking in the mmWave frequency bands. In this paper, we introduced MilliCar, the first implementation of an open-source ns-3 module for the simulation of NR-V2X networks at mmWaves. The module enables end-to-end, full-stack simulations of vehicular networks with a 3GPP channel model for V2V propagation and fading at mmWaves, physical and MAC layers redesigned for NR V2X, and integration with the higher layers of the protocol stack (e.g., RLC) from ns-3. We believe that this contribution, being an open-source tool, easily extensible and available to the overall wireless community, will help foster more research on the design and performance evaluation of mmWave vehicular networks.

As future works, we plan to improve the functionalities of MilliCar by adding a number of new features, including a more robust and realistic beam management framework and a multiple medium access scheme, and by keeping the module updated to the latest proposals in the 3GPP standardization process.

ACKNOWLEDGMENTS

This work was partially supported by NIST through Awards No. 70NANB17H166 and 70NANB18H273, and by MIUR (Italian Ministry for Education and Research) under the initiative "Departments of Excellence" (Law 232/2016).

REFERENCES

- [1] 3GPP. 2018. NR - Physical Channels and Modulation. TS 38.211 (Release 15). (2018).
- [2] 3GPP. 2018. NR and NG-RAN Overall Description. TS 38.300 (Release 15). (2018).
- [3] 3GPP. 2018. Service Requirements for Enhanced V2X Scenarios. TS 22.186 (Release 15). (2018).
- [4] 3GPP. 2018. Study on Channel Model for Frequencies from 0.5 to 100 GHz (Release 14). TR 38.901. (2018).
- [5] 3GPP. 2019. Study on Evaluation Methodology of New Vehicle-to-Everything (V2X) Use Cases for LTE and NR. TR 37.885 (Release 15). (June 2019).
- [6] 3GPP. 2019. Study on NR Vehicle-to-Everything (V2X). TR 38.885 (Release 16). (March 2019).
- [7] S. Akoum, O. El Ayach, and R. W. Heath. 2012. Coverage and Capacity in mmWave Cellular Systems. In *46th Asilomar Conference on Signals, Systems and Computers*. IEEE, Pacific Grove, CA, USA, pp. 688–692.
- [8] M. Boban, X. Gong, and W. Xu. 2016. Modeling the Evolution of Line-of-Sight Blockage for V2V Channels. In *IEEE 84th Vehicular Technology Conference (VTC-Fall)*. Montreal, Canada.
- [9] J. Choi, V. Va, N. Gonzalez-Prelcic, R. Daniels, C. R. Bhat, and R. W. Heath. 2016. Millimeter-Wave Vehicular Communication to Support Massive Automotive Sensing. *IEEE Communications Magazine* 54, 12 (December 2016), pp. 160–167.
- [10] L. M. Clements and K. M. Kockelman. 2017. Economic Effects of Automated Vehicles. *Transportation Research Record* 2606 (2017), pp. 106–114.
- [11] M. Giordani, T. Shimizu, A. Zanella, T. Higuchi, O. Altintas, and M. Zorzi. 2019. Path Loss Models for V2V mmWave Communication: Performance Evaluation and Open Challenges. In *IEEE 2nd Connected and Automated Vehicles Symposium*. Honolulu, Hawaii, USA.
- [12] M. Giordani, A. Zanella, and M. Zorzi. 2017. Millimeter Wave Communication in Vehicular Networks: Challenges and Opportunities. In *6th International Conference on Modern Circuits and Systems Technologies (MOCAST)*. Thessaloniki, Greece.
- [13] M. Giordani, A. Zanella, and M. Zorzi. 2019. LTE and Millimeter Waves for V2I Communications: An End-to-End Performance Comparison. In *IEEE 89th Vehicular Technology Conference (VTC2019-Spring)*. Kuala Lumpur, Malaysia.
- [14] IEEE. 2019. 802.11 NGV Proposed PAR. Study Group on 802.11bd (TGbd) 802.11-18/0861r8.
- [15] IEEE. 2019. PHY Numerology Discussions. Study Group on 802.11bd (TGbd) 802.11-19/0686.
- [16] S. Kuttty and D. Sen. 2016. Beamforming for Millimeter Wave Communications: An Inclusive Survey. *IEEE Communications Surveys & Tutorials* 18, 2 (Second quarter 2016), pp. 949–973.
- [17] I. Mavromatis, A. Tassi, R. J. Piechocki, and A. Nix. 2018. Efficient V2V Communication Scheme for 5G mmWave Hyper-Connected CAVs. In *IEEE International Conference on Communications Workshops*. Kansas City, MO, USA.
- [18] M. Mezzavilla, M. Zhang, M. Polese, R. Ford, S. Dutta, S. Rangan, and M. Zorzi. 2018. End-to-End Simulation of 5G mmWave Networks. *IEEE Communications Surveys & Tutorials* 20, 3 (Third quarter 2018), pp. 2237–2263.
- [19] N. Patriciello, S. Lagen, B. Bojovic, and L. Giupponi. 2019. An E2E Simulator for 5G NR Networks. *Simulation Modelling Practice and Theory* 96, 101933 (2019).
- [20] C. Perfecto, J. Del Ser, and M. Bennis. 2017. Millimeter-Wave V2V Communications: Distributed Association and Beam Alignment. *IEEE Journal on Selected Areas in Communications* 35, 9 (Sep. 2017), pp. 2148–2162.
- [21] V. Petrov, J. Kokkonen, D. Moltchanov, J. Lehtomäki, M. Juntti, and Y. Koucheryavy. 2018. The Impact of Interference From the Side Lanes on mmWave/THz Band V2V Communication Systems With Directional Antennas. *IEEE Transactions on Vehicular Technology* 67, 6 (June 2018), pp. 5028–5041.
- [22] R. Rouil, F. J. Cintrón, A. Ben Mosbah, and S. Gamboa. 2017. Implementation and Validation of an LTE D2D Model for ns-3. In *Proceedings of the Workshop on ns-3 (WNS3 '17)*. Association for Computing Machinery, Porto, Portugal, pp. 55–62.
- [23] B. Schunemann, K. Massow, and I. Radusch. 2008. A Novel Approach for Realistic Emulation of Vehicle-2-X Communication Applications. In *IEEE Vehicular Technology Conference*. Calgary, Alberta, pp. 2709–2713.
- [24] C. Sommer, R. German, and F. Dressler. 2011. Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis. *IEEE Transactions on Mobile Computing* 10, 1 (January 2011), pp. 3–15.
- [25] D. Yang, L. Yang, and L. Hanzo. 2010. DFT-Based Beamforming Weight-Vector Codebook Design for Spatially Correlated Channels in the Unitary Precoding Aided Multiuser Downlink. In *IEEE International Conference on Communications*. Cape Town, South Africa.
- [26] M. Zhang, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi. 2017. ns-3 Implementation of the 3GPP MIMO Channel Model for Frequency Spectrum above 6 GHz. In *Proceedings of the Workshop on ns-3 (WNS3 '17)*. Association for Computing Machinery, Porto, Portugal, pp. 71–78.
- [27] T. Zugno, M. Drago, M. Giordani, M. Polese, and M. Zorzi. 2020. Towards Standardization of Millimeter Wave Vehicle-to-Vehicle Networks: Open Challenges and Performance Evaluation. *Submitted to IEEE Communications Magazine* (2020). <https://arxiv.org/abs/1910.00300>
- [28] T. Zugno, M. Polese, N. Patriciello, B. Bojovic, S. Lagen, and M. Zorzi. 2020. Implementation of A Spatial Channel Model for ns-3. In *Proceedings of the Workshop on ns-3 (WNS3 '20)*. Association for Computing Machinery, Gaithersburg, MD, USA. <https://arxiv.org/abs/2002.09341>