

Received 23 February 2023, accepted 24 March 2023, date of publication 5 April 2023, date of current version 14 April 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3264668

RESEARCH ARTICLE

Methodology and Performance Assessment of Three-Dimensional Vehicular Ad-Hoc Network Simulation

ALEXANDER BRUMMER¹, REINHARD GERMAN, AND ANATOLI DJANATLIEV¹

Computer Networks and Communication Systems, Friedrich-Alexander-Universität (FAU) Erlangen-Nürnberg, 91058 Erlangen, Germany

Corresponding author: Alexander Brummer (alexander.brummer@fau.de)

ABSTRACT Packet-based simulation is a key tool for the research and development of Vehicular Ad-hoc Networks (VANETs). Over the last decade, many models throughout the communication stack have been presented, which have increased the degree of realism that can be achieved with popular simulation frameworks. Nevertheless, the three-dimensional aspects of many real-world traffic situations barely find consideration. In this paper, we present a holistic approach to simulate large-scale three-dimensional VANET scenarios. We briefly summarize our previously presented models covering different aspects of communication in 3D scenarios, including an environmental diffraction model, an n-ray ground interference model, and the consideration of multi-floor communication. We then describe the principle of a model selector, which applies the appropriate models depending on the environment of the currently transmitted packet. Subsequently, we use the outlined methodology implemented in our Veins 3D framework to simulate a large urban reference scenario. The results differ significantly from comparable 2D simulations, demonstrating the necessity of three-dimensional considerations. However, they also show strongly increased execution times. Therefore, we further suggest different approaches to improve the simulation performance. Based on these optimizations, simulation durations in the same order of magnitude as a comparable 2D simulation can be achieved.

INDEX TERMS Vehicular ad-hoc network, V2X, packet-based simulation, three-dimensional, propagation model, simulation performance.

I. INTRODUCTION

Vehicle-to-Everything (V2X) communication is a promising technology with the goal of connecting vehicles to other road participants and their environment. The exchange of relevant information can lead to improved road safety and efficiency in conventional traffic, and can also be seen as an important factor for future autonomous cars. Various applications and protocols building upon cellular radio technologies like C-V2X [1] or the adapted WLAN standards IEEE 802.11p [2] and IEEE 802.11bd [3] have been researched in the recent past.

Simulation solutions are a crucial element in the research and development process of Intelligent Transportation Sys-

tems (ITS's) and V2X applications. They allow for repeatable investigations of newly developed aspects at varying levels of abstraction. The spectrum ranges from highly accurate ray-tracing simulations at link level (e.g., [4]) to system-level simulation approaches, which facilitate the investigation of large-scale Vehicular Ad-Hoc Networks (VANETs) [5]. In general, simulation models always constitute a simplification of the real-world system and its environment. The degree of realism is strongly interconnected with the computational complexity and should be chosen based on the goal of the investigation. As a consequence, analyzing vehicular networks consisting of hundreds of nodes at the highest level of detail is usually not feasible, and most likely also not reasonable. Instead, packet-based simulators such as the Veins framework [6] are employed, abstracting away from symbol-accurate considerations.

The associate editor coordinating the review of this manuscript and approving it for publication was Jie Gao¹.

Nevertheless, the VANET research of the last decade has brought forth accurate models of different mechanisms throughout the network layers. Examples are models for realistic Cooperative Awareness Message (CAM) generation [8], detailed aspects of IEEE 802.11p's and ETSI ITS G5's medium access [9], [10], mechanisms representing C-V2X sidelink communication [11], or 5G's Quality of Service (QoS) model [12]. Moreover, various propagation models have been proposed. They are of high importance in packet-based simulators, as the decision of whether a packet is decodable at all is made based on the Signal-to-Interference-plus-Noise Ratio (SINR), which in turn depends on the packet's received power [5]. Therefore, appropriate propagation models lay the foundation for credible simulation results, while trying to preserve the scalability of simulations. As examples, one could mention the two-ray ground interference model [13], a computationally efficient approach to capture shadowing caused by buildings [14], or more sophisticated models also considering diffraction and reflections of buildings and vehicles [15].

In spite of these steps forward, a frequent abstraction remaining is the assumption of a (mostly) two-dimensional environment. This means that the roads and potentially irregular terrain surrounding the scene appear to be perfectly planar. Also, the impact of vehicles as obstacles is still neglected in many studies, effectively rendering them flat as well. Overpasses, parking garages, or stacked road layers in general can usually not be represented at all. These simplifications consequently lead to limitations towards scenarios that can be investigated, as real-world traffic is rarely perfectly flat (for an example see Fig. 1). Furthermore, the application of 2D models in non-planar scenarios may lead to substantial errors. Packets which appear to be decodable when using the conventional simulation approach may turn out to be heavily attenuated after taking the third dimension into consideration.

To deal with these issues, we have recently proposed approaches to enable the investigation of different characteristic aspects of 3D VANET scenarios. These include

- the consideration of 3D antenna patterns, and an environmental diffraction model which considers the impact of both the terrain and vehicles between sender and receiver [16],
- a floor attenuation model to enable the investigation of communication across multiple road layers [17], including the special case of multi-story parking garages [18],
- and an n-ray ground interference model which can deal with arbitrarily shaped terrain, and thus with an arbitrary number of ground reflections [19].

In our previous studies, we could demonstrate that the application of each of these 3D models can cause significant deviations in the simulation results. Moreover, we have validated them by performing several field tests [18], [20], [21].

In the present work, we seek to bring the different models together to allow for the investigation of holistic three-dimensional scenarios. Depending on the characteristics of



FIGURE 1. Example of a three-dimensional traffic situation, which cannot be analyzed properly with conventional vehicular network simulators (photo by Superchilum, cropped, used under CC BY-SA 3.0 [7]).

the respective sender, receiver, and environment, the suitable models have to be selected and applied automatically. Obviously, the increased level of detail leads to longer computation times. In order to be able to simulate large-scale scenarios, we have further examined potential countermeasures to deal with the increased complexity.

The contributions of this article can thus be summarized as follows:

- We briefly review the proposed models covering different aspects of three-dimensional vehicular networking scenarios.
- We present a methodology to seamlessly apply these models, ensuring the selection of the relevant models for each communication link.
- In order to test this approach, we describe the setup of a three-dimensional reference scenario and analyze the utilization of the different models.
- Finally, we assess the impact on the computational performance and present approaches to accelerate the execution.

The individual models as well as the combined methodology have been implemented in our publicly available open-source framework Veins 3D,¹ which is based on the widely used Veins simulator [6].

II. RELATED WORK

Since the mid-2000s, various vehicular network simulators have been presented [22]. Well-known examples (which are still maintained and/or extended) are Eclipse MOSAIC (formerly VSimRTI) [23], ezCar2X [24], or the already-mentioned Veins framework [6]. Most solutions bring together existing mobility (e.g., SUMO [25] or VISSIM [26]) and network simulators (e.g., ns-3 [27] or OMNeT++ [28]). In terms of propagation models, these frameworks usually rely on the models offered by the employed network simulators, partially extended by custom implementations. This includes simple models covering free-space path loss and two-ray ground interference, as well as statistical models for log-distance path loss or Rayleigh fading.

¹<https://github.com/cs7org/veins3d>

Apart from such simple (statistical) models, there also exist geometry-based approaches to model the vehicular communication channel. These take into account the exact locations and shapes of surrounding objects, making it possible to accurately compute diffractions and reflections in 3D space. Despite the high degree of realism, ray-tracing methods such as [4] require enormous modeling efforts, and lead to a high computational complexity, effectively inhibiting the investigation of large vehicular networks.

Therefore, propagation models for VANET simulators going beyond the simple statistical approaches often require a compromise between realism and scalability. An example is the obstacle-shadowing model proposed by Sommer et al. [14]. By taking into account the outlines of buildings, their implementation determines the number of intersections with walls and the distance the direct signal travels through the buildings. Based on measurements, they weight and add those two values, yielding a certain attenuation representing the large-scale fading. Boban et al. present a more sophisticated set of models [15], which also include the potential impact of vehicles and foliage. Instead of assuming the direct signal path only, they consider single-interaction reflections and diffraction effects.

Although mostly neglected in vehicular networks, the influence of irregular terrain has already been addressed in the context of Mobile Ad-Hoc Networks (MANETs). As an example, the connectivity of wireless sensor network links may be constrained by terrain characteristics such as hills. The authors in [29], [30], and [31] present models which take Digital Elevation Models (DEMs) into account. They are all based on Durkin's model [32], which calculates the knife-edge diffraction loss in case of an obstructed Line-of-Sight (LOS). Another research area dealing with three-dimensionality by definition is the field of connected Unmanned Aerial Vehicles (UAVs). For the simulation of UAVs communicating with cars, Hadiwardoyo et al. also make use of elevation data in combination with a multiple knife-edge model [33].

In the context of VANET simulation, there has not been a particular focus on the three-dimensional character of real-world traffic situations so far. Still, the relevance of 3D-aware modeling opportunities can be shown by looking at applied research in different directions.

Message forwarding in VANETs is often based on position-based routing protocols. Lin et al. outline the problem that conventional position-based routing protocols assume all vehicles to be located in a plane [34]. In contrast, they describe a Three-Dimensional Scenario Oriented Routing (TDR) protocol which is also aware of different road layers. This way, the delivery ratio can be increased by roughly 40%. Boban et al. explicitly leverage the height differences of vehicles for forwarding messages [35]. They argue that considerable attenuation is caused by obstructing vehicles even in case of sparse traffic densities. To increase transmission ranges, they propose Tall Vehicle Relaying

(TVR) to use the increased elevation of the antennas on large vehicles.

Other examples requiring 3D considerations are applications in multi-story parking garages. For instance, there are many projects which investigate automated valet parking [36]. Besides Vehicle-to-Vehicle (V2V) connections also Vehicle-to-Infrastructure (V2I) links for the interaction with a potential global controller module need to be considered. Appropriate models are necessary to evaluate the communication aspects in this special environment.

In addition to such use cases explicitly demanding 3D models, it has been unclear to what extent the neglect of a three-dimensional environment might induce errors in simulation results. In our previous investigations, we could show that the deviations can be of considerable extent [16], [17], [19], [20], [21]. With this work, we seek to provide a unified way of incorporating various aspects of 3D scenarios in vehicular network simulation. In the next section, we outline a set of models which lay the foundation for that purpose, before we describe how they can be combined accordingly in Section IV.

III. PROPOSED MODELS FOR 3D VANET SIMULATIONS

An initial requirement for the simulation of three-dimensional VANET scenarios is the integration of information about the altitude of roads and the surrounding terrain. To this end, we make use of DEMs, i.e., datasets providing elevation data for a certain region of geographic coordinates. All subsequently implemented models rely on this knowledge. Furthermore, the road network used by the road traffic simulator has to be extended by adding z-coordinates to all junctions and intermediate nodes. In our case, the SUMO mobility simulator offers the tool *netconvert*, which can be used to import DEMs in Shapefile or GeoTiff format. This way, the simulated vehicles obtain a position and orientation in 3D space.

A. 3D ANTENNA PATTERNS

Antennas are crucial components for wireless communications in general. Depending on the angle of incidence, they can cause a significant attenuation or gain, leading to characteristic radiation patterns. In case of vehicular communication, these are further influenced by the vehicles themselves as shown by various studies (e.g., [37], [38], [39]). Especially in the vertical directions, strongly negative gains have to be expected, which should not be neglected in VANET simulation.

In [16], we have thus described how 3D antenna patterns can be incorporated. These are usually available in the form of the two principal planes (see the example in Fig. 2a), which are sampled equidistantly and stored in an XML file. In order to determine the antenna gain for a given transmission, we first calculate the azimuth and elevation angles, ϕ and θ . With the LOS vector \vec{v}_{LOS} between sender and receiver, and the orientation vector of the vehicle \vec{v}_{orient} , the angles are

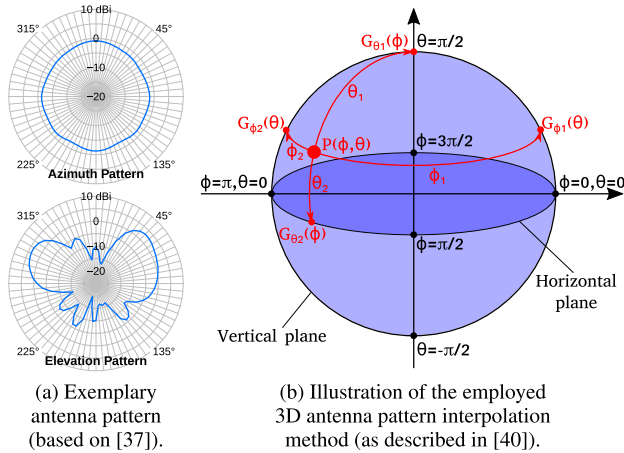


FIGURE 2. Consideration of 3D antenna patterns as described in [40] (exemplary pattern based on [37]). (a) Exemplary antenna pattern. (b) Illustration of the employed 3D antenna pattern interpolation method.

computed as [16]

$$\phi = \arccos \left(\frac{\vec{v}_{orient} \circ \vec{v}_{LOS}}{|\vec{v}_{orient}| \cdot |\vec{v}_{LOS}|} \right), \quad (1)$$

$$\left. \begin{aligned} \theta_{LOS} &= \arcsin \left(\frac{v_{LOS}^z}{|\vec{v}_{LOS}|} \right) \\ \theta_{orient} &= \arcsin \left(\frac{v_{orient}^z}{|\vec{v}_{orient}|} \right) \end{aligned} \right\} \Rightarrow \theta = \theta_{LOS} - \theta_{orient}. \quad (2)$$

We then employ the antenna pattern interpolation method presented by Leonor et al. [40] to estimate the 3D antenna gain in the direction of $P(\phi, \theta)$. As illustrated in Fig. 2b, the two angles are used to query the horizontal and vertical radiation patterns, yielding four gain values $G_{\phi_1}(\theta)$, $G_{\phi_2}(\theta)$, $G_{\theta_1}(\phi)$, and $G_{\theta_2}(\phi)$. Next, the four gains are weighted based on ϕ and θ to obtain a horizontal and vertical interpolation value. Finally, these two values are combined based on a third weight function, resulting in the antenna gain estimate for the given direction (see [16] for more details).

B. ENVIRONMENTAL DIFFRACTION MODEL

If a general three-dimensional environment is considered, buildings are not the only objects that can obscure the transmitted signal. On the one hand, the impact of other vehicles has to be taken into account, as already described in [41]. On the other hand, the terrain shape may block the direct signal path. In [16], we have presented an approach to model environmental diffraction effects considering both vehicles as well as terrain characteristics as potential obstacles.

Diffraction effects are often analyzed by treating the obstacles as knife edges. This approach is valid if the signal’s wavelength is small in comparison to the obstacles, which is fulfilled for signals in the ITS 5.9 GHz band. In general, a knife edge can be characterized by the dimensionless parameter ν , which depends on the wavelength, the distances to transmit and receive antenna, and the height of the obstacle above the direct signal path.

For our purpose, we need to create a set of knife edges representing the terrain as well as potential vehicles between




FIGURE 3. Exemplary set of knife edges representing the terrain shape and vehicles in the LOS; the terrain point spacing equals 5 m, not drawn to scale (based on [16]).

sender and receiver, as illustrated in Fig. 3. This involves two steps:

- 1) The DEM is queried at equidistant locations along the LOS. The spacing can be adjusted and should be chosen with respect to the DEM’s horizontal resolution. For each point, the altitude and corresponding distance to the sender are added to a map.
- 2) We check whether there are vehicles blocking the LOS. Therefore, the four segments representing a vehicle’s outline are checked for a potential intersection with the LOS. If that is the case, the elevation plus the vehicle height are added to the map. This procedure is repeated for every vehicle in the simulation.

Based on the obtained set of knife edges, we estimate the diffraction loss by applying the cascaded knife-edge approach described in ITU Recommendation ITU-R P.526-11 [42], which is an adapted version of the Deygout method [43] extended by correction terms. The first step consists of determining the knife edge with the highest impact. To this end, the diffraction parameter ν_n for the n-th profile point in the generated set is computed as

$$\nu_n = h \cdot \sqrt{\frac{2 \cdot d_{sr}}{\lambda \cdot d_{sn} \cdot d_{nr}}}, \quad (3)$$

where d_{xy} refers to the distances between sender s, receiver r, and the respective profile point n. λ denotes the signal’s wavelength and h is the knife edge’s height above the LOS [42]. The point with the maximum value for ν_n is considered as principal edge ν_p .

A value of $\nu_p > -0.78$ means that the direct signal path is affected (i.e., we cannot assume free-space propagation). In this case, the procedure is repeated for the subsets between sender and principal edge, as well as between receiver and principal edge, resulting in two further values ν_s and ν_r . Finally, the determined parameters are combined to estimate the overall diffraction loss [42]:

$$L_{diff} = J(\nu_p) + T \cdot (J(\nu_s) + J(\nu_r) + C). \quad (4)$$

Here, T and C are correction terms, and $J(\nu)$ is a function to determine the attenuation for a given knife edge in dB based on an approximation of the Fresnel integrals (see [16] and [42] for more details).

In [16], we could show that the application of antenna and diffraction models can already lead to significantly differing simulation results. For example, the number of reachable neighbors in a city-wide scenario decreased from 34 in the 2D setup to eleven when the extensions were included. The diffraction model's applicability was also evaluated through a field test in a hilly environment [20].

C. MULTI-FLOOR COMMUNICATION

Another type of three-dimensional propagation occurs when cars are located on different road layers. In a conventional VANET simulator, such a scenario cannot be represented at all. For example, a highway interchange with various levels would simply appear to be a large, extraordinary intersection with respect to the communication aspects. In reality, such an environment may lead to unusual antenna elevation angles (and thus, negative gains) as well as to significant additional attenuation due to the floor slabs obstructing the direct signal path. While the former aspect has already been addressed in Section III-A, the latter requires an additional floor attenuation model, which we initially introduced in [17].

In preparation of such scenarios, the height values of the elevated road structures need to be set accordingly. For large parts, this can be achieved by using a special type of DEMs, namely Digital Surface Models (DSMs). As they represent the top layer, they also include the height of overpasses, for instance. However, special attention is required in regions where multiple road layers are directly located above each other, since the lower layers may have been assigned the (wrong) elevation of the top-most layer. In this case, the user has to correct them by interpolating the two closest undistorted height values of the respective roads.

When the simulation is started, a set of polygons which represent the floor surfaces is constructed. This is done by querying the road shapes from the traffic simulator (in our case using the Traffic Control Interface (TraCI) to query SUMO). The constructed polygons are stored in a three-dimensional R-tree, an efficient, hierarchical spatial data structure [44]. In order to determine the floor attenuation for a transmission, this R-tree is then queried and returns all polygons (i.e., floor slabs) that are located in the bounding box spanned by the sender's and receiver's antennas (see Fig. 4). Afterwards, we perform polygon-segment intersection checks on these floors, which yields the number of floors intersecting with the LOS.

Next, we apply the Floor Attenuation Factor (FAF) model originally presented in [45] and [46]. The simple idea is to add a predetermined attenuation factor that changes depending on the number of penetrated floors. The additional attenuation is thus given by

$$L_{FAF}(n) = FAF(n) + X, \quad (5)$$

where $FAF(n)$ is the attenuation factor for n floors, and X is a normally distributed random variable.

To test the impact of obstructing floor slabs, we applied the model to a multi-level highway interchange scenario [17],




FIGURE 4. Illustration of the determination of floor slabs intersecting with the LOS. The R-tree yields all road layers within the signal path's bounding box (here three). These are checked for actual intersections (here two). From [17], © 2019 IEEE.

which also led to a decrease of reachable neighbors of roughly 40%. Furthermore, we used the model for the special case of multi-story parking garages, which showed good agreement with real-world measurements [18].

D. N-RAY GROUND INTERFERENCE

Even if the direct signal path is free of obstructions, a generally non-planar environment also demands special attention. A well-known model often applied under LOS conditions is the two-ray interference model [13]. It takes into consideration that the transmitted signal does not only reach the receiver via the direct path, but also via reflection off the ground between sender and receiver. Due to the superposition of the two versions of the signal, the received power undergoes characteristic fluctuations. If a flat terrain is assumed, there is exactly one reflection. In the case of a 3D environment, this assumption cannot be guaranteed anymore. Depending on the actual shape of the ground along the LOS, there can be an arbitrary number of reflections, which cannot be predicted in a general form.

To solve this issue, we have presented the n-ray ground interference model [19]. First, a height profile has to be generated by querying the DEM equidistantly (similar to the first step of the environmental diffraction model). This results in a set of segments spanned by consecutive profile points. As depicted in Fig. 5, each segment is checked for a potential reflection. After that, the corresponding parameters for each reflection have to be computed. This includes the path length of the i -th reflected ray $d_{ref,i}$, the resulting phase




FIGURE 5. Illustration of the n-ray ground reflection model. In this example, all four ground segments cause reflections that interfere at the receiver. From [21], © 2020 IEEE.

difference $\Delta\phi_i$, and the reflection coefficient $\Gamma_{\perp,i}$ (assuming horizontal polarization). $\Gamma_{\perp,i}$ itself depends on the incidence angle θ , and on the relative permittivity of the ground ϵ_r . Superimposing the direct ray (with path length d_{LOS}) with the n reflected rays yields the received power

$$P_r = P_t \cdot \left(\frac{\lambda}{4\pi}\right)^2 \cdot \left| \frac{1}{d_{LOS}} + \sum_{i=1}^n \frac{\Gamma_{\perp,i} \cdot e^{j\Delta\phi_i}}{d_{ref,i}} \right|^2 \quad (6)$$

based on transmit power P_t . The detailed geometry is explained in [19].


Similar to the previous models, we performed measurements on various test tracks to evaluate the applicability of the n -ray ground interference model [21]. In all cases, the course of the measured received power could be reproduced well. Furthermore, significant deviations from the two-ray interference model (up to 20 dB, differing from one test site to another) could be demonstrated.

IV. SITUATION-DEPENDENT COMBINATION OF APPROPRIATE MODELS

Each of the different models outlined in the previous section has been tested and evaluated, demonstrating the relevance to the respective use cases. In order to make them applicable to large-scale scenarios involving diverse environments, a further step is necessary.

Most wireless network simulators require the selection and assignment of propagation models to be used by network nodes before the simulation is started. For example, Veins expects so-called `AnalogueModels` representing propagation models to be specified in an XML file. These are then assigned to the physical layer of each node, and applied to all packet transmissions (see Fig. 6a). However, a general three-dimensional simulation scenario requires a more fine-grained selection and combination of the available models depending on the individual surroundings of each transmission.

Therefore, we have introduced a new `ModelSelector` class, which is owned by the physical layer (see Fig. 6b). This model selector is now responsible for all the path loss models instead of the physical layer itself. It also considers the models specified in the already mentioned XML file, but does not apply all of them to all transmissions. Instead,



(a) Original case, all models to be used are directly assigned to the `BasePhyLayer`.

it differentiates various cases with the goal of adding only the attenuation values of the currently suitable models.

A. NON-LINE-OF-SIGHT VERSUS LINE-OF-SIGHT

The first question to be answered is whether the direct signal path is influenced by any kind of object. In this case, we would have Non-Line-of-Sight (NLOS) conditions, which can require several models depending on the type of influencing or even obstructing objects. In fact, we also use these models to see which case applies in the first place:

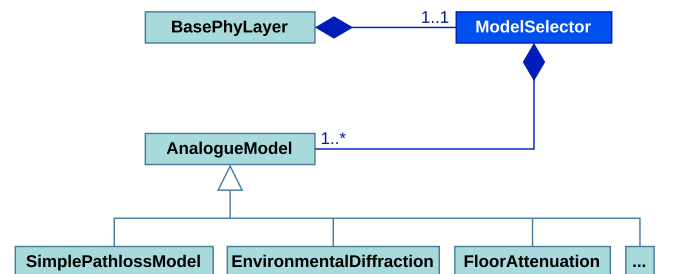
- To begin with, we check whether multi-floor communication takes place. If this is the case, the floor attenuation factor corresponding to the number of penetrated floor slabs will be determined (see Section III-C).
- Otherwise, there can still be other objects on the same level disturbing the signal, i.e., vehicles or the terrain itself (covered by the environmental diffraction model, see Section III-B),
- or buildings, which we consider by applying the simple obstacle shadowing model by Sommer et al. [14].

If all of the NLOS models yield an attenuation of 0 dB, we can infer that the direct signal path is clear. In this case, the model selector proceeds with the LOS models. Other than that, the path loss resulting from the NLOS models is added to the signal along with distance-dependent path loss, and the model selector returns.

B. SELECTION OF A SUITABLE LOS MODEL

In the case of an unobstructed LOS, the simplest option is to assume free-space propagation, i.e., adding distance-dependent path loss. In many simulators, the user can also choose to apply the popular two-ray interference model instead. As was discussed in Section III-D, a three-dimensional environment requires a more general consideration of ground reflections. In any case, it does not always make sense to use ground reflection models.

On the one hand, if there is a floor located along the LOS (e.g., an overpass between sender and receiver), it may influence the ground reflections. The n -ray model, which takes the terrain shape as an input, is unaware of the floors and



(b) Adaptation consisting of a `ModelSelector` responsible for selecting the right models for each transmission.

FIGURE 6. Class diagrams showing the relationship between the physical layer and the path loss models to be applied (as realized in Veins and Veins 3D).




FIGURE 7. Screenshot of an exemplary SUMO road network. The blue polygons represent the estimated street widths.

would thus produce wrong results. Therefore, falling back on the free-space path loss is a more reasonable choice.

On the other hand, as Erlacher et al. have investigated, the applicability of the (in their case two-ray) ground interference model is influenced by the width of the street on which the cars are currently located [47]. The street width actually refers to the distance between the buildings on either side of the road. If the street is very narrow, there are many reflections off the nearby buildings which may reach the receiver with a relatively large signal power. As a consequence, a smaller street width leads to a more complex multipath propagation environment, and the impact of the ground reflections declines. Erlacher and colleagues performed measurements on roads with different street widths. They could show that the two-ray interference model works very well for wide streets, and that it still provides a good fit for a medium street width of 28 m. For a street as narrow as 18 m, the two-ray model was not applicable anymore [47].

These insights should also find consideration during the selection of the right LOS model. This requires knowledge of the street widths of all the roads in the investigated scenario. Therefore, we have implemented a preprocessing script in Python, which expects the road network and building outlines as inputs, and estimates the respective street widths. The estimation for each edge is based on calculating the distances of surrounding buildings to the respective edge. These values are averaged, resulting in the mean distance to a building. Doubling this value yields the street width. An exemplary result is shown in Fig. 7. The estimated values are then attached to the road network's edges as additional attributes.

When the simulation is executed, the street width can be queried via TraCI. If this value is below 18 m, the n-ray ground interference model cannot be used, and we employ only distance-dependent path loss. Based on the results in [47], we figure that the medium range around 28 m can be seen as a transitional zone in which the ground reflection model can be applied. However, the impact of the ground reflections should be adapted by scaling the relative permittivity ϵ_r : The closer the street width is to the lower bound of 18 m, the closer $\epsilon_{r, \text{effective}}$ should be to 1, as a relative permittivity of 1 effectively equals free-space propagation. Above 38 m, we consider the street to be wide

enough to apply the n-ray ground interference model without scaling.

C. SPECIAL ENVIRONMENTS

For the simulation of large-scale scenarios, the model selector also has to be able to deal with certain environments that require special treatment. So far, we have identified and considered two types.

First, parking garages constitute an extraordinary environment. As we have examined in [18], the communication is characterized by potential multi-floor transmissions, shadowing and diffraction, as well as small-scale fading that should not be neglected. To this end, we have proposed an appropriate set of models dealing with these characteristics, which includes the floor attenuation model, diffraction and obstacle shadowing models, and a module for the generation of Rayleigh/Rician fading.

The second type of traffic situation that cannot be represented by the conventional models are tunnels. Although they are not the focus of our work, we also had to find a suitable representation, since our reference scenario (see Section V) includes a small number of them. Therefore, we have implemented the following simplistic model:

- If the sender is located inside the tunnel and the receiver is outside (or vice versa) with tunnel walls in between, we imply that the signal cannot travel through the walls, and set the received power to 0.
- On the other hand, if both vehicles are driving in the same tunnel, we assume good signal propagation, since the tunnel walls can have a waveguiding effect [48]. Thus, only distance-dependent attenuation is added.

In order to identify such special locations during simulation, one option is to define them as regions of interest based on coordinates, which can then be checked by the model selector. Another way is tagging the relevant roads in the road network, and querying whether this tag is set on the roads that sender and receiver are currently located on.

D. OVERALL MODEL SELECTION PROCEDURE

The complete working principle of the model selector is depicted in Fig. 8. Based on the locations of sender and receiver, we first check whether they are currently located in a special environment. Otherwise, a normal traffic situation can be assumed. Thus, the NLOS (floor, environmental diffraction, and obstacle shadowing) models are evaluated. If this step reveals a clear LOS, we examine whether the n-ray ground interference model is applicable. If this is not the case, the free-space path loss is computed instead.

It should be noted that all parts of this procedure, especially the applied models, can easily be replaced or adapted, making the model selector flexible and extensible. Furthermore, its usage can be disabled based on a single flag in the configuration file, so that users can also choose the conventional way of applying path loss models.

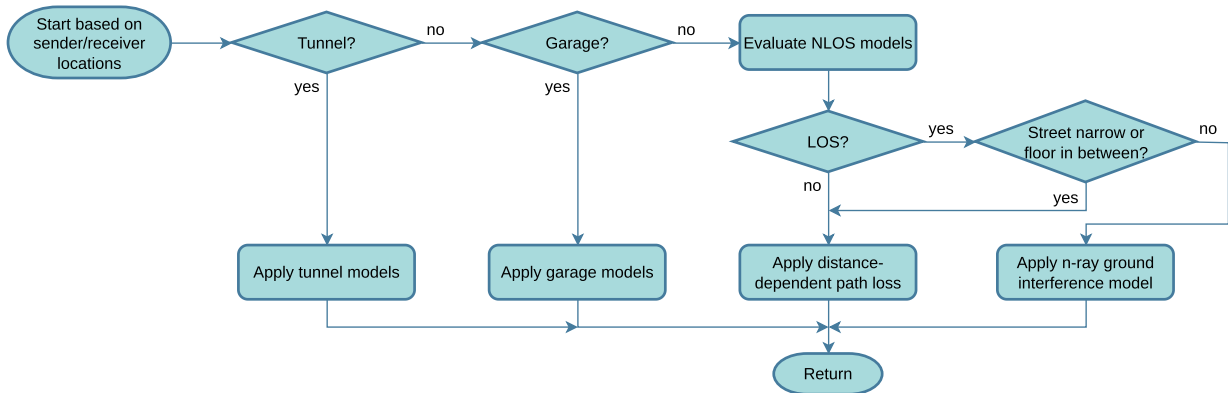


FIGURE 8. Flowchart illustrating the simplified working principle of the model selector.

V. REFERENCE SCENARIO AND PERFORMANCE ASSESSMENT

To test the interplay of the proposed models, we have used the model selector to simulate an exemplary large-scale scenario. Based on that, it is also possible to assess how the simulation runtime is affected by the increased level of detail.

A. SCENARIO DESCRIPTION

Our reference scenario has to consist of different environments to make sure that all of the 3D models presented in Section III come into use. Therefore, we have chosen the well-known Luxembourg SUMO Traffic (LuST) scenario [49] as a starting point. It provides realistic urban and highway traffic in the city of Luxembourg for 24 hours with varying traffic densities over time. We have extended the base scenario by several 3D aspects.

- First of all, altitude data were imported into the road network from an official DEM offered by the Luxembourgish Administration of Cadastre and Topography.² We processed the LiDAR-based dataset to have a horizontal resolution of 1 m and cropped it to the city of Luxembourg. These elevation data are also used by the environmental diffraction and n-ray ground interference models.
- We identified and tagged floors and tunnels in the road network.
- Furthermore, the street widths were estimated and added to the road network with the help of the script already mentioned in Section IV-B.

All vehicles were configured to broadcast periodic Basic Safety Messages (BSMs) at an interval of 10 s, which is sufficient to get a general impression of the connectivity properties of the vehicular network. Moreover, the monopole antenna pattern depicted in Fig. 2a was assigned. Further simulation parameters are given in Table 1.

For comparison, we also simulated the scenario with the conventional 2D setup. In this case, the two-ray interference

²<https://data.public.lu/fr/datasets/lidar-2019-modele-numerique-de-terrain-mnt/>

TABLE 1. General simulation parameters.

Parameter	Value
Carrier frequency	5.89 GHz
Transmit power	20 mW
Sensitivity threshold	-98 dBm
Thermal noise	-98 dBm
Maximum interference distance	2600 m
Beacon (BSM) interval	10 s
Bitrate	6 Mbit/s
Obstacle shadowing [14]	9 dB/wall, 0.4 dB/meter
Diffraction model DEM sampling interval	connectivity evaluation: 1 m duration measurements: 5 m

model as well as the obstacle shadowing model were applied to all transmissions, and only the horizontal antenna pattern was used. All simulations were performed using our extended Veins 3D framework (which is based on Veins v. 4.7.1), OMNeT++ v. 5.2, as well as SUMO v. 1.8.0.

B. CONNECTIVITY EVALUATION

Although the focus of the present work is not on demonstrating the necessity of 3D simulations (see our previous studies for that), we first have a brief look at the differences in the VANET’s connectivity resulting from the two different setups, i.e., Veins 4.7.1 with conventional propagation models, and our extended framework Veins 3D including the model selector. For this purpose, we simulated one minute of the described scenario, including ten repetitions with different random seeds. To see the impact of varying traffic densities, four different starting times were chosen: At 4 am, there are roughly 350 simulated vehicles, at 5 am, 6 am, and 7 am around 600, 1500, and 3200, respectively. The mean vehicle speed ranges between 50 km h⁻¹ (7 am) and 80 km h⁻¹ (4 am).

Fig. 9 shows the average number of reachable neighbors for the different configurations. This metric represents how many vehicles could receive a single transmitted beacon on average. Not surprisingly, this value increases from 4 am to 7 am, as more and more vehicles are on the roads as the morning rush hour is approached. More importantly, we can see

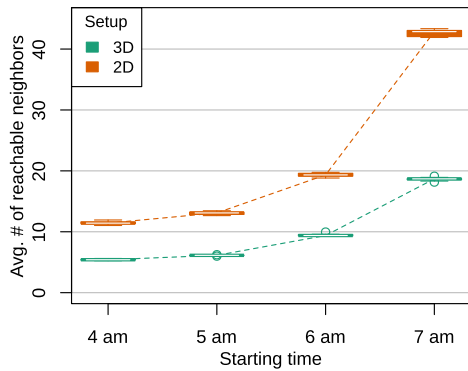


FIGURE 9. Average number of reachable neighbors for different starting times, based on 2D and 3D models.

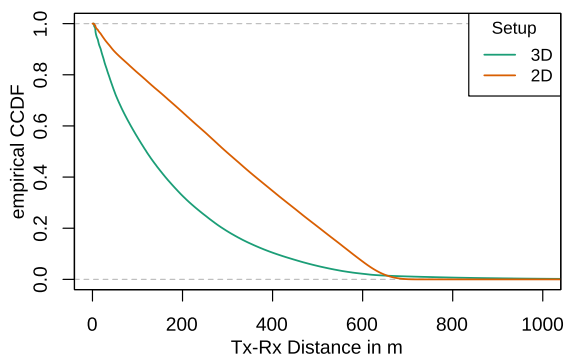


FIGURE 10. Complementary cumulative distribution function (CCDF) of the sender-receiver distance of successful packet receptions for starting time 7 am, based on 2D and 3D models.

that the incorporation of the three-dimensional considerations leads to a significant decrease in the number of reachable neighbors. For all four starting times, the numbers are more than halved.

We further looked at the sender-receiver distances of successful packet transmissions at 7 am. In Fig. 10, the empirical complementary cumulative distribution functions of the distance for both setups are depicted. As can be seen, the 3D setup curve drops significantly faster. While roughly 50 % of the packets are received at distances > 300 m in the conventional setup, this percentage is less than 20 % with our three-dimensional extensions. As we have already discussed in our previous studies, it becomes clear that these deviations may lead to vastly different conclusions if actual applications and protocols are investigated.

C. MODEL SELECTOR STATISTICS

Next, we examine how often and where the different models were used by evaluating a single simulation run at 7 am as an example. Only 10 s (instead of 60 s) were simulated to obtain a clear, comprehensible illustration. As the pie chart in Fig. 11 demonstrates, most of the evaluated transmissions took place under NLOS conditions. Thus, the floor attenuation, environmental diffraction, or obstacle shadowing model

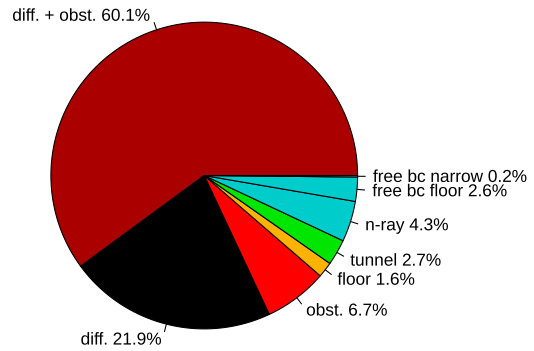


FIGURE 11. Pie chart representing the share of models applied by the model selector (10 simulated seconds at 7 am with a maximum interference distance of 1000 m, corresponding to Fig. 12).

(as well as a combination thereof) was applied in the majority of cases. The LOS models were used in about 7 % of all evaluated packet transmissions; the remainder corresponds to the special case of tunnels. Note that the maximum interference distance for this study was reduced to 1000 m, since larger-distance transmissions are of less interest here, as they almost only increase the NLOS share.

Fig. 12a illustrates where which types of models were applied. Similar to the pie chart, this figure is dominated by the colors red and black (or a mixture of both), which represent the obstacle shadowing and diffraction models. This is especially the case for potential long-distance transmissions. Furthermore, we can see that the LOS models (cyan) find use especially along the roads. The floor attenuation and tunnel models are limited to the locations of overpasses/highway interchanges, and tunnels respectively. In contrast to the previous illustration, Fig. 12b only shows the transmissions that were actually successful (around 1.2 % of all evaluated transmissions). As can be seen, these are mainly connections under LOS conditions, as well as transmissions among vehicles driving in the same tunnel. Most of the red and black connections do not show up anymore.

We would like to emphasize that increasing the interference distance mainly increases the percentage of the diffraction and obstacle shadowing models, as the probability of terrain, other vehicles, or buildings obstructing the LOS rises with increasing sender-receiver separation. For the subsequent investigations, it was reset to the initial value of 2600 m. Furthermore, the shown results are obviously specific to the given reference scenario. Depending on the terrain shape, distribution of buildings, number of overpasses, etc., the shares of the used models may look quite different.

D. UNOPTIMIZED SIMULATION PERFORMANCE

Finally, the impact of our three-dimensional extensions on the simulation’s execution times was assessed. The configuration settings are almost equal to the ones used in Section V-B. Only the diffraction model’s DEM sampling interval was set to 5 m, and the simulated time period was reduced from 60 s to 10 s, since this amount of time is sufficient to observe the

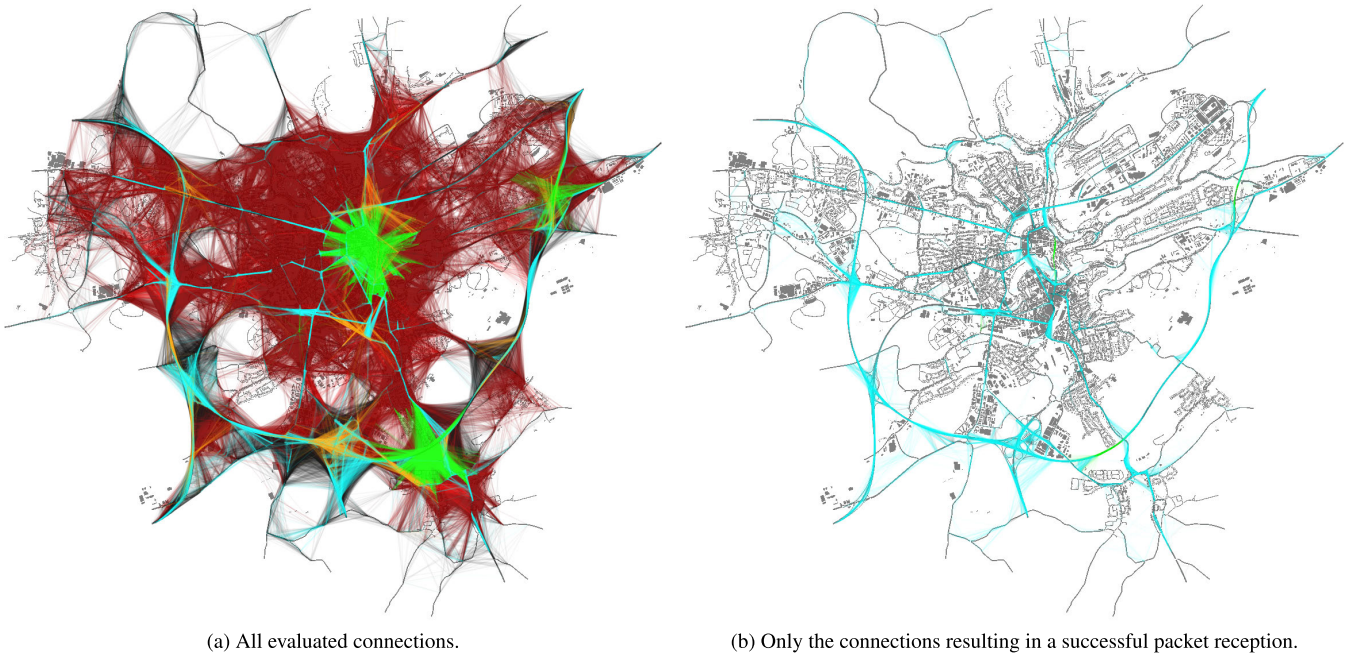


FIGURE 12. Map of the simulated scenario overlaid with the potential packet transmissions handled by the model selector (10 simulated seconds at 7 am with a maximum interference distance of 1000 m). The different colors represent the models that have been applied to each connection (see Fig. 11).




FIGURE 13. Unoptimized simulation durations resulting from running 10 s of the reference scenario.

differences in the resulting simulation durations. To obtain comparable, representative numbers, all of the performance measurements were executed on a normal Desktop computer with an Intel® Core™ i5-4590 CPU (clock speed 3.3 GHz) and 16 GB of memory. The results are depicted in Fig. 13.

Similar to the number of reachable neighbors, we can see that the simulation duration increases for later starting times. As already mentioned, the number of vehicles running simultaneously develops from about 350 to 3200. With the chosen maximum interference distance of 2600 m, the number of potential receivers increases drastically, and thus the attenuation models are evaluated much more often. This also explains why the absolute difference between the 2D and 3D setups is rather small at 4 am. At 5 am, 6 am, and

TABLE 2. Excerpt of the *perf* profiling results.

Samples	Function
69.33 %	ModelSelector::filterSignal
68.67 %	↳ ModelSelector::applyNLOSModels
34.66 %	↳ EnvironmentalDiffraction::calcAttenuation
33.26 %	↳ SimpleObstacleShadowing::calcAttenuation

7 am, however, we can see that the higher level of detail gained through the 3D considerations comes at the cost of significantly prolonged runtimes. While the unmodified Veins framework requires around 25 minutes to simulate the 7 am scenario on average, the unoptimized 3D extensions lead to execution times of roughly 20 hours.

To find the parts of the simulation that consume the most time, we employed *perf*,³ a sampling profiler that periodically samples a program to query the currently active function. The relevant results are summarized in Table 2. Almost 70 % of the samples refer to the model selector applying the respective path loss models. In line with what Fig. 11 already suggested, the profiling result in fact reveals that the majority of CPU time is spent applying the NLOS models, with the environmental diffraction and obstacle shadowing models having approximately equal shares.

The high time consumption of the environmental diffraction model is not only caused by the fact that it is evaluated in most of the cases (see Figs. 8 and 11), but also due to the various calculations the model has to perform: The elevation model has to be queried, other vehicles have to be checked for

³https://perf.wiki.kernel.org/index.php/Main_Page

potentially blocking the LOS, the resulting set of knife edges has to be analyzed, and the corresponding diffraction loss has to be computed. The smaller the DEM sampling interval and the greater the number of vehicles in the simulation, the more checks and calculations have to be performed.

Again, we would like to note that the execution times shown here refer to our reference scenario, and may change depending on the composition of other scenarios.

VI. PERFORMANCE IMPROVEMENT APPROACHES

As the previous section indicates, the three-dimensional models can increase the degree of realism leading to significantly differing results, but are hardly usable for large-scale scenarios in this form. Therefore, the simulation performance needs to be improved. There already exist general approaches to speed up the computation of the received power. Bronner et al. describe a “thresholding” approach [50], which stops the evaluation of further models if the signal power is already too low to be detected. However, all models might still have to be evaluated when the impact of interfering signals (to obtain the SINR) needs to be determined.

Instead, we propose four performance improvement approaches tailored to the introduced 3D extensions. The first two constitute optimizations of the environmental diffraction model (which caused the largest slowdown in the previous section), the other two refer to the model selector in general. This way, we seek to accelerate the simulations, whilst maintaining the increased level of detail.

A. MANAGEMENT OF VEHICLE POSITIONS

As we have already mentioned in the previous section, one of the computationally expensive tasks of the environmental diffraction model consists of testing whether other vehicles obstruct the direct signal path of the current packet transmission (see the second step in Section III-B). To do that, the simulator has to iterate through all currently simulated vehicles, generate the four segments corresponding to a car’s outline based on its position and orientation, and determine whether there is an intersection with the segment representing the LOS. The greater the number of simulated vehicles, the more intersection checks need to be performed. Even vehicles which are located far away from the current transmission are taken into consideration because only after checking for an intersection, the model knows about it.

To reduce this workload, the number of potentially obstructing vehicles has to be decreased. For this purpose, we overlay the investigated scenario with a grid. Based on the latitude and longitude (or x - and y -coordinates) of the vehicles, they are assigned to the corresponding grid cells. Instead of checking all vehicles for intersections with the LOS, only those assigned to grid cells that overlap with the bounding box spanned by the sender and receiver have to be considered. Especially for scenarios with a large geographical extent, this can be very beneficial. The only disadvantage is that after every update received from SUMO, the vehicles which have crossed cell boundaries need to be reassigned.




FIGURE 14. Exemplary scene illustrating how a grid can reduce the number of vehicles that have to be taken into consideration as potential obstacles. Instead of all vehicles, only the ones within the grid cells highlighted in blue are checked (based on screenshot from SUMO, vehicles not drawn to scale).

Fig. 14 illustrates this methodology. Originally, all vehicles in the given scenario have to be checked to find the ones which obstruct the transmission between the cars circled in red (here, 67). By only querying the nodes that are assigned to the grid cells spanned by sender and receiver (in this example represented by the blue rectangle), the computations can be reduced significantly (here, only 23 vehicles need to be examined). Of course, the reductions are maximized if the communicating nodes are close to each other; if they are far away from each other, many cells (potentially containing many vehicles) have to be assessed.

A question remaining is what might be a good value for the size of the grid cells. If the cells are too large, the effect of this approach is small, since there are always many cars to check. On the other hand, if the cells are too small, many vehicles have to be assigned to other cells after each location update from SUMO. We have resimulated the reference scenario with the same settings as in Section V-D including this adaptation and varied the grid cell size. The resulting simulation durations are depicted in Fig. 15. We can see that all tested cell sizes perform almost similarly with a simulation duration of around 17 h, which equals a speed-up of 15%. Only values greater than 1000 m cause slightly longer execution times, since too many irrelevant vehicles have to be checked.

As we have mentioned in Section II, considering vehicles as obstacles based on multiple knife-edge models has already been discussed by other researchers. Boban et al. manage the vehicle positions by storing them in a two-dimensional r -tree [41]. In order to compare its performance to our simplistic grid approach, we also implemented the use of an r -tree, and resimulated the reference performance scenario. The rightmost boxplot in Fig. 15 shows that this approach achieves roughly the same reduction in execution time.

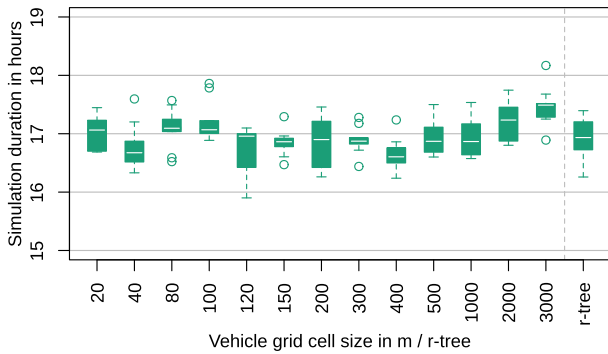


FIGURE 15. Simulation durations resulting from running 10 s of the reference scenario at 7 am with different variants of managing the vehicle positions: Based on uniform grids with varying cell sizes, as well as based on an r-tree.




FIGURE 16. Exemplary scene demonstrating the caching of elevation values. If the vehicles highlighted in red want to communicate with each other, the values cached in the blue grid cells can be reused (based on screenshot from SUMO).

B. CACHING OF ELEVATION VALUES

The second time-consuming aspect of the environmental diffraction model is the consideration of the terrain shape. For the generation of the height profile along the LOS, the DEM needs to be queried at equidistant locations. We employ the same class used by SUMO to generate the 3D road network, which interpolates the wanted elevation from the three closest DEM raster points: The plane spanned by those three points is determined and its z-coordinate corresponding to the given location is computed. This querying procedure is a time-consuming, and recurring task. Thus, it would be beneficial to reuse already queried values.

For this purpose, we overlay the investigated scenario with another grid structure, where each cell is used to cache the corresponding altitude. Whenever a height profile has to be generated, the environmental diffraction model first checks whether the cache grid cell at the given location already contains a value. If this is the case, this elevation is used. Otherwise, the DEM is queried as usual, the obtained value is used for the height profile, and it is stored in the cache grid. This principle is illustrated in Fig. 16. Suppose the left car highlighted in red transmits a beacon. The decodability of this packet will be checked by all surrounding vehicles. Looking at the two highlighted cars on the right-hand side, we can see that roughly the same elevation values are required to generate the LOS’s height profile. After the DEM has been queried by the first vehicle, the other one can make use of the cached values.

Similar to the vehicle grid, the cell size for the DEM cache can be adjusted by the user. The larger the cells are, the smaller the number of required DEM queries. However, the selected cell size obviously affects the accuracy, as all locations within a grid cell will be mapped to the cached elevation. Thus, the cache’s cell size should not be set to overly large values. To investigate both the performance gains and the loss in accuracy we simulated the reference scenario using the DEM cache and varied the grid cell size.

For the execution time measurements, the diffraction model’s DEM sampling interval (i.e., the distance between

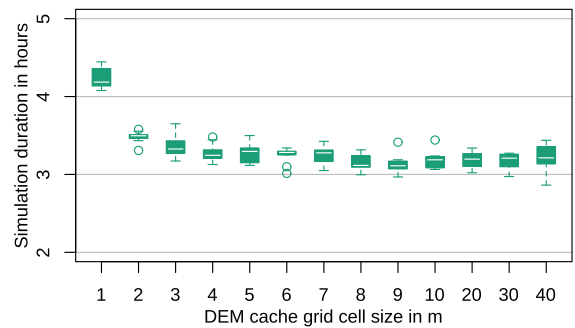


FIGURE 17. Simulation durations resulting from running 10 s of the reference scenario at 7 am with varying DEM cache grid cell sizes.

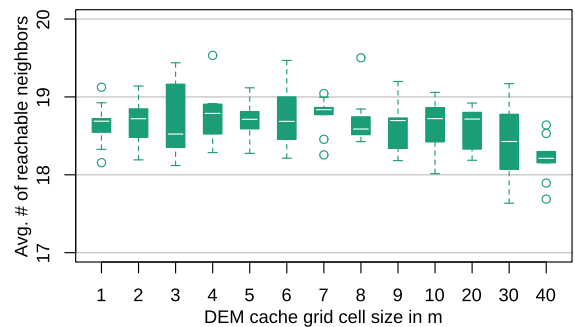


FIGURE 18. Average number of reachable neighbors for varying DEM cache grid cell sizes.

height profile points) was set to 5 m, similar to the unoptimized performance measurements. The resulting durations are given in Fig. 17. As we can see, even the use of a relatively dense grid with a cell size of 1 m accelerates the simulation significantly, resulting in a duration of under 4.5 h. Cell sizes above 3 m only cause small differences, leading to execution times of slightly more than 3 h (i.e., a speed-up of around 84 %).

To see how the cache cell size impacts the accuracy, the diffraction model’s DEM sampling interval was set to be equal to the respective cell size. The obtained numbers of

reachable neighbors are shown in Fig. 18. The results indicate that a DEM cache cell size greater than 20 m leads to inaccuracies, which results in a reduced number of neighbors in this scenario. Below 20 m the cell size does not have a big impact on the conclusions drawn on the overall connectivity. However, one should keep in mind that the average number of reachable neighbors is an aggregated metric. Individual transmissions can indeed be evaluated differently if the generated height profiles vary too much. Thus, for more detailed investigations of smaller areas (e.g., individual intersections), the DEM cache grid should not be configured too coarse.

C. PARALLELIZATION

A widely used technique to speed up the computation of (sub)tasks is parallel execution. More precisely, we do not mean the (generally possible) parallel execution of multiple simulation runs (potentially distributed across various hosts in a computer cluster or in the cloud), but the parallelization of the simulation program itself. Especially nowadays, as multicore and multiprocessor architectures are available at relatively low cost, this is a viable option. Network simulators are usually based on Discrete Event Simulation (DES). Although there are approaches concerning parallel simulation [51], most of them operate in a single-threaded manner, processing one event at a time in chronological order.

In wireless network simulation, packet receptions are also represented by events. As we have seen, these usually involve considerable computations: In order to determine the received power and the SINR, the path loss models need to be evaluated. An example of how this can be parallelized is given in [52]. At the sending time, signal copies corresponding to the potential receivers of a transmitted packet are grouped. The computations required to evaluate these signal copies, including the application of the attenuation models, are performed by asynchronous background threads. Such approaches can also be leveraged to accelerate the evaluation of the models presented here.

However, said technique requires various adaptations to the physical layer and related classes. Instead, we experimented with a simpler approach which only affects our model selector. Looking at the profiler result reveals that by far the most time is spent on the environmental diffraction and obstacle shadowing models. Thus, evaluating these two models in parallel promises great potential, and can be implemented directly in the model selector. For this purpose, we make use of OpenMP,⁴ a well-known multi-platform API for parallel programming. The two models are usually applied sequentially one after the other. By embedding them into OpenMP *parallel sections*, a *team of threads* is instantiated so that both models can be run simultaneously. Special synchronization or protection measures are not required because both models are completely independent of each other and only return numbers which represent the respective attenuation. Once

both models have been evaluated, the simulation can continue as usual.

We resimulated the reference scenario with the same settings as in Section V-D, including the described parallelization. As illustrated in Fig. 19a, this adaptation results in a 17 % reduction in execution time (median: 16.6 h vs. 20.0 h).

D. IDENTIFICATION OF 3D REGIONS

Since the 3D models lead to an increased computational complexity, it would make sense to only apply them in regions where they are actually needed. If the environment is approximately flat, it is sufficient to use the two-ray interference model - applying the n-ray ground interference model including the required generation of a height profile would produce a similar attenuation value. Furthermore, the environmental diffraction model can be limited to the consideration of vehicles as potential knife edges, as an evaluation of the terrain shape is also not necessary in this case.

The important question is how the 3D regions of interest can be determined. On the one hand, the user can explicitly state the corner points of such areas. When the simulation is run, the model selector can check if the signal travels through such regions, and select the relevant models accordingly.

On the other hand, it would be convenient to automate the identification of three-dimensional regions. Note that a constant slope does not require 3D considerations because it still resembles a flat surface. Instead, a 3D region is characterized by *changes* in terrain inclination. Therefore, we implemented a preprocessing script in Python based on edge detection, which generates a set of rectangles representing three-dimensional areas. The approach can be summarized in four steps:

- 1) Read the DEM file, resulting in a raster of height values (similar to the pixels of an image).
- 2) Perform the edge detection by applying a Gaussian filter followed by a Laplacian filter (also known as Laplacian of Gaussian, or LoG filter [53]). Return the pixels whose value after filtering exceeds a certain threshold.
- 3) Group the identified points into rectangles.
- 4) Write the coordinates of the rectangles' corner points to a file.

Once created, the generated file can be used by the model selector. To manage the potentially large number of 3D areas, they are stored in another r-tree. For a given set of sender and receiver coordinates, the model selector queries this r-tree for the existence of 3D regions in between.

As shown in Fig. 19a, this approach only leads to a small reduction in execution time of roughly 2.2 % for the given reference scenario (median of 19.56 h vs. 20.00 h). This is caused by the fact that the script determines many 3D regions for the elevation model of Luxembourg. The number of identified 3D areas depends on the parameterization of the filtering process. A more relaxed configuration results in fewer 3D rectangles and faster execution, but also negatively affects the simulation accuracy.

⁴<https://www.openmp.org/>

E. COMPARISON OF PERFORMANCE IMPROVEMENT APPROACHES

To compare the effectiveness of the approaches presented in the previous subsections, we have selected the best parameterization for each case. The respective execution times are plotted in Fig. 19a. As can be seen, especially the use of the DEM cache brings the simulation duration closer to the execution time of a conventional 2D simulation. It should also be mentioned that the speed-up resulting from limiting the 3D models to actual three-dimensional regions strongly depends on the composition of the scenario under investigation.

Of course, it is further possible to combine the different approaches. Using both the vehicle grid and the DEM cache grid leads to a median duration of 57 minutes. Including the parallelization of diffraction and obstacle shadowing models, the performance reference simulation finishes after 39 minutes. Although these simulation setups still take longer than the 2D simulation, they come relatively close - while also providing the higher level of detail resulting from the 3D considerations.

In particular, the two approaches of managing the vehicles in a grid and parallelizing the model evaluation do not influence the level of detail at all, as the data provided to the models are unchanged. This is not necessarily the case for the other two presented techniques. The limitation to 3D regions may cause deviations if the preprocessing omits areas where the 3D models would in fact lead to additional attenuation. As already mentioned, this strongly depends on the parameterization of the applied filters. The DEM caching can also lead to a loss of accuracy. However, as we have shown in Fig. 18, this mainly has an observable effect if grid cell sizes over 20 m are used. Thus, in the case of our reference scenario, all of the 3D setups (including those executing fastest) lead to a similar average number of neighbors, in contrast to the conventional 2D setup (see Fig. 19b).

VII. CONCLUSION AND FUTURE WORK

In this paper, we have described a holistic approach to simulate three-dimensional vehicular ad-hoc networks. We have summarized our previously presented models, which cover different aspects of communication in 3D traffic situations. In order to use them in combination, and dynamically apply them to large-scale scenarios, a new model selector has been introduced. Depending on the environment of the current transmission, the appropriate models are applied to determine the packet’s received power. Afterwards, the outlined methodology has been used to simulate an urban reference scenario. Similar to our previous studies, the results showed significant deviations between our approach and the conventional 2D simulation setup, motivating the necessity of three-dimensional considerations.

However, the increased level of detail also leads to considerably extended simulation durations. To address this problem, we have described four possible performance improvement approaches. We showed how the environmental diffraction model can be optimized to reduce the number of necessary computations. Moreover, it is possible to parallelize the evaluation of frequently used models and to limit some 3D considerations to regions where the investigated scenario actually shows 3D characteristics. This way, acceptable execution times as well as a high degree of realism can be achieved.

Yet, both aspects can be further improved in future work. On the one hand, a further reduction in execution times is desirable. Apart from model-specific optimizations, especially the identification of 3D regions requires further investigations to parameterize the edge detection filters based on the DEM’s resolution. On the other hand, the set of available 3D propagation models can be optimized. The special tunnel environment requires a more detailed treatment, whereby the application of a two-slope path loss model might be a good choice [48]. Moreover, it would be interesting to test alternatives to our proposed models (e.g., other diffraction models) and to include further factors not considered so far. A summary of the aspects that Veins 3D already supports is given in Table 3, alongside a (non-exhaustive) list of points we have not included by now, such as models to account for the influence of vegetation. In this context, embedding other

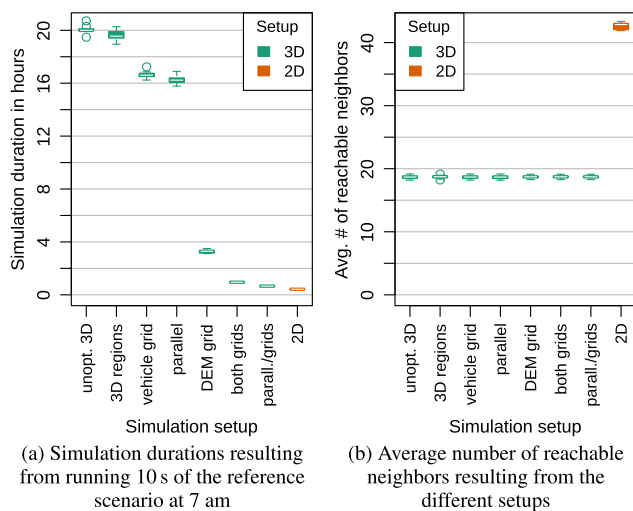


FIGURE 19. Comparison of the unoptimized 3D setup, the conventional 2D setup, and the different performance improvement approaches (if used, the vehicle grid cell size was set to 400 m and the DEM cache cell size to 5 m).

TABLE 3. Summary of the (environmental) aspects considered / not considered in Veins 3D including the model selector.

Considered	Yet to be explored
<ul style="list-style-type: none"> shadowing by buildings diffraction by terrain diffraction by vehicles arbitrary ground reflections (depending on the street width) penetration of floor slabs parking garages tunnels (simplified) 3D antenna patterns 	<ul style="list-style-type: none"> impact of vegetation diffraction and reflections off buildings reflections off vehicles

sets of models such as GEMV² [15] in our 3D framework might be promising.

REFERENCES

- [1] R. Molina-Masegosa and J. Gozávez, "LTE-V for sidelink 5G V2X vehicular communications: A new 5G technology for short-range vehicle-to-everything communications," *IEEE Veh. Technol. Mag.*, vol. 12, no. 4, pp. 30–39, Dec. 2017.
- [2] J. B. Kenney, "Dedicated short-range communications (DSRC) standards in the United States," *Proc. IEEE*, vol. 99, no. 7, pp. 1162–1182, Jul. 2011.
- [3] G. Naik, B. Choudhury, and J.-M. Park, "IEEE 802.11bd & 5G NR V2X: Evolution of radio access technologies for V2X communications," *IEEE Access*, vol. 7, pp. 70169–70184, 2019.
- [4] N. Dreyer and T. Kürner, "An analytical raytracer for efficient D2D path loss predictions," in *Proc. 13th Eur. Conf. Antennas Propag. (EuCAP)*, Krakow, Poland, 2019, pp. 1–5.
- [5] C. Sommer, J. Häri, F. Hrizi, B. Schünemann, and F. Dressler, "Simulation tools and techniques for vehicular communications and applications," in *Vehicular Ad Hoc Networks*, C. Campolo, A. Molinaro, and R. Scopigno, Eds. Berlin, Germany: Springer, May 2015, pp. 365–392.
- [6] C. Sommer, R. German, and F. Dressler, "Bidirectionally coupled network and road traffic simulation for improved IVC analysis," *IEEE Trans. Mobile Comput.*, vol. 10, no. 1, pp. 3–15, Jan. 2011.
- [7] Superchilum. (2009). *Street in San Francisco*. [Online]. Available: <https://commons.wikimedia.org/wiki/File> and <https://creativecommons.org/licenses/by-sa/3.0/legalcode>
- [8] R. Molina-Masegosa, M. Sepulcre, J. Gozávez, F. Berens, and V. Martínez, "Empirical models for the realistic generation of cooperative awareness messages in vehicular networks," *IEEE Trans. Veh. Technol.*, vol. 69, no. 5, pp. 5713–5717, May 2020.
- [9] D. Eckhoff and C. Sommer, "A multi-channel IEEE 1609.4 and 802.11p EDCA model for the veins framework," in *Proc. 5th ACM/ICST Int. Conf. Simul. Tools Techn. Commun., Netw. Syst.*, Desenzano, Italy, Mar. 2012, pp. 1–2.
- [10] D. Eckhoff, N. Sofra, and R. German, "A performance study of cooperative awareness in ETSI ITS G5 and IEEE WAVE," in *Proc. 10th Annu. Conf. Wireless Demand Netw. Syst. (WONS)*, Banff, AB, Canada, Mar. 2013, pp. 196–200.
- [11] A. Hegde and A. Festag, "Artery-C: An OMNeT++ based discrete event simulation framework for cellular V2X," in *Proc. 23rd Int. ACM Conf. Model., Anal. Simul. Wireless Mobile Syst.*, M. Aguilar-Igartua, P. Bellavista, A. A. F. Loureiro, and C. Giannelli, Eds. Nov. 2020, pp. 47–51.
- [12] T. Deinlein, M. Roshdi, T. Nan, T. Heyn, A. Djanatliev, and R. German, "On the impact of priority-based MAC layer scheduling in 5G V2N multi-application scenarios," in *Proc. 13th IFIP Wireless Mobile Netw. Conf. (WMNC)*, Oct. 2021, pp. 63–70.
- [13] C. Sommer, S. Joerer, and F. Dressler, "On the applicability of two-ray path loss models for vehicular network simulation," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Seoul, (South) Korea, Nov. 2012, pp. 64–69.
- [14] C. Sommer, D. Eckhoff, R. German, and F. Dressler, "A computationally inexpensive empirical model of IEEE 802.11p radio shadowing in urban environments," in *Proc. WONS*, Bardonecchia, Italy, Jan. 2011, pp. 84–90.
- [15] M. Boban, J. Barros, and O. K. Tonguz, "Geometry-based vehicle-to-vehicle channel modeling for large-scale simulation," *IEEE Trans. Veh. Technol.*, vol. 63, no. 9, pp. 4146–4164, Nov. 2014.
- [16] A. Brummer, R. German, and A. Djanatliev, "On the necessity of three-dimensional considerations in vehicular network simulation," in *Proc. 14th Annu. Conf. Wireless Demand Netw. Syst. Services (WONS)*, Feb. 2018, pp. 75–82.
- [17] A. Brummer, R. German, and A. Djanatliev, "Modeling V2X communications across multiple road levels," in *Proc. IEEE 2nd Connected Automated Vehicles Symp. (CAVS)*, Honolulu, HI, USA, Sep. 2019, pp. 1–5.
- [18] A. Brummer, T. Deinlein, and A. Djanatliev, "On the simulation of vehicular networking scenarios in multi-story parking garages," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Ulm, Germany, Nov. 2021, pp. 87–94.
- [19] A. Brummer, L. Ammon, and A. Djanatliev, "N-ray ground interference: Extending the two-ray interference model for 3D terrain shapes," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Los Angeles, CA, USA, Dec. 2019, pp. 1–4.
- [20] A. Brummer, T. Deinlein, K.-S. Hielscher, R. German, and A. Djanatliev, "Measurement-based evaluation of environmental diffraction modeling for 3D vehicle-to-X simulation," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Taipei, Taiwan, Dec. 2018, pp. 1–8.
- [21] A. Brummer, M. Gütlein, M. Schäfer, R. German, and A. Djanatliev, "Experimental evaluation of the N-ray ground interference model," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Ulm, Germany, Dec. 2020, pp. 1–8.
- [22] S. Babu and A. R. Kumar P, "A comprehensive survey on simulators, emulators, and testbeds for VANETs," *Int. J. Commun. Syst.*, vol. 35, no. 8, pp. 1–32, May 2022.
- [23] Eclipse MOSAIC Core Team. *Eclipse MOSAIC: A Multi-Domain and Multi-Scale Simulation Framework for Connected and Automated Mobility*. Accessed: Dec. 8, 2022. [Online]. Available: <https://eclipse.org/mosaic>
- [24] K. Roscher, S. Bittl, A. Gonzalez, M. Myrtus, and J. Jiru, "ezCar2X: Rapid-prototyping of communication technologies and cooperative its applications on real targets and inside simulation environments," in *Proc. 11th Conf. Wireless Commun. Inf.*, 2014, pp. 51–62.
- [25] P. A. Lopez, E. Wiessner, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flotterod, R. Hilbrich, L. Lucken, J. Rummel, and P. Wagner, "Microscopic traffic simulation using SUMO," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 2575–2582.
- [26] PTV. *PTV VISSIM: Multimodal Traffic Simulation Software*. Accessed: Dec. 8, 2022. [Online]. Available: <https://www.myptv.com/en/mobility-software/ptv-vissim>
- [27] G. F. Riley and T. R. Henderson, *The NS-3 Network Simulator*. Berlin, Germany: Springer, 2010, pp. 15–34.
- [28] A. Varga and R. Hornig, "An overview of the OMNeT++ simulation environment," in *Proc. 1st Int. ICST Conf. Simul. Tools Techn. Commun. Netw. Syst.*, Marseille, France, 2008, pp. 1–10.
- [29] S. Filiposka and D. Trajanov, "Terrain-aware three-dimensional radio-propagation model extension for NS-2," *Simulation*, vol. 87, nos. 1–2, pp. 7–23, Jan. 2011.
- [30] H. H. Nguyen, S. Krug, and J. Seitz, "Simulation of 3D signal propagation based on real world terrains for ad-hoc network evaluation," in *Proc. 9th IFIP Wireless Mobile Netw. Conf. (WMNC)*, Colmar, France, Jul. 2016, pp. 131–137.
- [31] S. Mansfield, K. Veenstra, and K. Obraczka, "Modeling communication over terrain for realistic simulation of outdoor sensor network deployments," *ACM Trans. Model. Perform. Eval. Comput. Syst.*, vol. 6, no. 4, pp. 1–22, Apr. 2022.
- [32] R. Edwards and J. Durkin, "Computer prediction of service areas for VHF mobile radio networks," *Proc. Inst. Electr. Eng.*, vol. 116, pp. 1493–1500, Sep. 1969.
- [33] S. A. Hadiwardoyo, C. T. Calafate, J.-C. Cano, Y. Ji, E. Hernandez-Orallo, and P. Manzoni, "3D simulation modeling of UAV-to-car communications," *IEEE Access*, vol. 7, pp. 8808–8823, 2019.
- [34] Q. Lin, C. Li, X. Wang, and L. Zhu, "A three-dimensional scenario oriented routing protocol in vehicular ad hoc networks," in *Proc. IEEE 77th Veh. Technol. Conf. (VTC Spring)*, Dresden, Germany, Jun. 2013, pp. 1–5.
- [35] M. Boban, R. Meireles, J. Barros, P. Steenkiste, and O. K. Tonguz, "TVR—Tall vehicle relaying in vehicular networks," *IEEE Trans. Mobile Comput.*, vol. 13, no. 5, pp. 1118–1131, May 2014.
- [36] H. Banzhaf, D. Nienhüser, S. Knoop, and J. M. Zöllner, "The future of parking: A survey on automated valet parking with an outlook on high density parking," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2017, pp. 1827–1834.
- [37] D. Kornek, M. Schack, E. Slotke, O. Klemp, I. Rolfes, and T. Kürner, "Effects of antenna characteristics and placements on a vehicle-to-vehicle channel scenario," in *Proc. Int. Conf. Commun. Workshops*, Capetown, (South) Africa, May 2010, pp. 1–5.
- [38] A. Kwoczek, Z. Raida, J. Láčik, M. Pokorny, J. Puskely, and P. Vágner, "Influence of car panorama glass roofs on Car2Car communication," in *Proc. 3rd IEEE Veh. Netw. Conf.*, Amsterdam, The Netherlands, Nov. 2011, pp. 246–251.
- [39] D. Eckhoff, A. Brummer, and C. Sommer, "On the impact of antenna patterns on VANET simulation," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Columbus, OH, USA, Dec. 2016, pp. 17–20.
- [40] N. R. Leonor, R. F. S. Caldeirinha, M. G. Sánchez, and T. R. Fernandes, "A three-dimensional directive antenna pattern interpolation method," *IEEE Antennas Wireless Propag. Lett.*, vol. 15, pp. 881–884, 2016.
- [41] M. Boban, T. T. V. Vinhoza, M. Ferreira, J. Barros, and O. K. Tonguz, "Impact of vehicles as obstacles in vehicular ad hoc networks," *IEEE J. Sel. Areas Commun.*, vol. 29, no. 1, pp. 15–28, Jan. 2011.
- [42] *Propagation by Diffraction*, document ITU-R P.526-11, International Telecommunication Union, Geneva, Switzerland, Oct. 2009.
- [43] J. Deygout, "Multiple knife-edge diffraction of microwaves," *IEEE Trans. Antennas Propag.*, vol. AP-14, no. 4, pp. 480–489, Jul. 1966.

- [44] A. Guttman, "R-trees: A dynamic index structure for spatial searching," in *Proc. ACM SIGMOD Int. Conf. Manag. Data*, Boston, MA, USA, 1984, pp. 47–57.
- [45] J. M. Keenan and A. J. Motley, "Radio coverage in buildings," *Bell Syst. Tech. J.*, vol. 8, pp. 19–24, Jan. 1990.
- [46] S. Y. Seidel and T. S. Rappaport, "914 MHz path loss prediction models for indoor wireless communications in multifloored buildings," *IEEE Trans. Antennas Propag.*, vol. 40, no. 2, pp. 207–217, Feb. 1992.
- [47] F. Erlacher, F. Klingler, C. Sommer, and F. Dressler, "On the impact of street width on 5.9 GHz radio signal propagation in vehicular networks," in *Proc. 11th Annu. Conf. Wireless Demand Netw. Syst. Services (WONS)*, Obergurgl, Austria, Apr. 2014, pp. 143–146.
- [48] A. Hrovat, G. Kandus, and T. Javornik, "A survey of radio propagation modeling for tunnels," *IEEE Commun. Surveys Tuts.*, vol. 16, no. 2, pp. 658–669, 2nd Quart., 2014.
- [49] L. Codeca, R. Frank, and T. Engel, "LuST: A 24-hour scenario of Luxembourg city for SUMO traffic simulations," in *Proc. SUMO User Conf.*, Berlin, Germany, May 2015, pp. 1–10.
- [50] F. Bronner and C. Sommer, "Efficient multi-channel simulation of wireless communications," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Taipei, Taiwan, Dec. 2018, pp. 1–8.
- [51] R. M. Fujimoto, "Research challenges in parallel and distributed simulation," *ACM Trans. Model. Comput. Simul.*, vol. 26, no. 4, pp. 1–29, May 2016.
- [52] D. S. Buse, G. Echterling, and F. Dressler, "Accelerating the simulation of wireless communication protocols using asynchronous parallelism," in *Proc. 24th Int. ACM Conf. Model., Anal. Simul. Wireless Mobile Syst.*, Alicante, Spain, Nov. 2021, pp. 57–66.
- [53] D. Marr and E. Hildreth, "Theory of edge detection," *Proc. Roy. Soc. London, Ser. B, Biol. Sci.*, vol. 207, pp. 187–217, Feb. 1980.



ALEXANDER BRUMMER received the B.Sc. and M.Sc. degrees in information and communication technology from the Faculty of Engineering, Friedrich-Alexander-Universität (FAU) Erlangen-Nürnberg, in 2014 and 2017, respectively. He is currently pursuing the Ph.D. degree with the Computer Networks and Communication Systems Lab, FAU Erlangen-Nürnberg. He is currently a Doctoral Research Assistant with FAU Erlangen-Nürnberg. His research interests include wireless communication systems and the modeling and simulation of vehicular ad-hoc networks, with a focus on three-dimensional aspects.



REINHARD GERMAN received the master's degree in computer science and the Ph.D. degree from the Computer Science Department, Technical University of Berlin, Germany, in 1991 and 1994, respectively. He is a Full Professor with the Computer Networks Lab, Department of Computer Science, Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany. He is also an Adjunct Professor with the Faculty of Information Technology, Monash University, Melbourne, Australia. His research interests include the performance and dependability analysis of interconnected systems based on numerical analysis, network calculus, discrete-event simulation, measurements, and testing. Vehicular communications and smart energy constitute major application domains.



ANATOLI DJANATLIEV received the M.Sc. (Dipl.-Inf. Univ.) and Ph.D. (Dr.-Ing.) degrees in computer science from the Faculty of Engineering, Friedrich-Alexander-Universität (FAU) Erlangen-Nürnberg, in 2008 and 2015, respectively. He is the Head of the Connected Mobility Group at the Computer Networks and Communication Systems Lab, FAU Erlangen-Nürnberg. His current research interests include various topics on simulation and modeling. His major application areas are the simulation of vehicular networks, innovative aspects of connected mobility, and future mobility services.

• • •