

On the Level of Detail of Synthetic Highway Traffic Necessary to Vehicular Networking Studies

Marco Gramaglia*, Marco Fiore†

*University Carlos III of Madrid, IMDEA Networks Institute, Madrid, Spain

†CNR-IEIIT, Torino, Italy

Abstract—The proper modeling of road traffic is paramount to the dependability of studies on vehicular networking solutions intended for highway environments. Yet, it is not clear which is the actual level of detail in the mobility representation that is sufficient and necessary to such studies. This uncertainty results into a variety of approaches being adopted in the literature, and ultimately undermines the reliability and reproducibility of research outcomes. We explore the space of possible mobility models and performance metrics, and pinpoint the level of detail needed for different types of vehicular networking studies.

I. INTRODUCTION

Communication-enabled vehicles are nowadays a reality. A growing number of car models feature multi-technology radio interfaces, including GPRS, HSPA, LTE, Bluetooth, and, soon enough, DSRC [1]. Vehicles are turning into communication hubs supporting a full range of services related to road safety, traffic management, and passenger entertainment.

The huge potential of vehicular communications has attracted significant attention from both the academic and the industrial research communities over the past decade. Countless proposals for network architectures, protocols and algorithms have been formulated that cover the whole network stack and support a full range of application use cases [2], [3].

In the vast majority of vehicular networking studies, the design and performance evaluation of solutions partially or fully relies on synthetic models of road traffic. Indeed, experimental evaluations would require large-scale testbeds comprising hundreds of vehicles, with overwhelming costs and complexity.

The need for reliable simulation environments has fostered the development of increasingly realistic representations of road traffic in both urban [4], [5] and highway [6], [7] environments. Notwithstanding the public availability of fairly complete mobility traces, a significant portion of the research community is still reluctant to adopt them in their studies [8]. This is a critical aspect that limits the dependability and reproducibility of research in the vehicular networking domain.

We argue that part of the problem stems from the fact that we currently lack a clear understanding of which level of realism in road traffic modeling is actually *sufficient* and *necessary* to the simulation of vehicular networking solutions. Some efforts have been made in urban environments, where it was shown that simplistic constant-speed or stochastic models of drivers' behavior can bias the results of the simulation of vehicular networks, and that more complex car-following and lane-changing models are necessary [4], [9].

However, the existing literature does not provide an answer in the case of highway environments. Some works propose to use aggregate statistics to describe vehicle inflows [7], [10]–[14], whereas others employ fine-grained, per-vehicle traffic count data [6]. Some works employ stochastic models of drivers' behavior [10]–[13], whereas others leverage complex microscopic models [6], [7], [14]. Many works neglect the presence of entry and exit ramps [6], [10]–[13], whereas others consider them [7], [14]. Overall, which of the settings above can be deemed correct for the evaluation of vehicular networking solutions for highway environments – and which are not – remains unclear.

In this paper, we aim at filling the gap outlined above, and provide an ultimate definition of the level of road traffic realism *necessary* to the simulation of vehicular networks in highway environments. To that end, we proceed as follows. First, we identify a set of performance metrics that is restricted, yet meaningful to a large number of networking studies, in Sec. II. Then, we introduce the different combinations of road traffic models we consider, selected so as to cover the full space of possible simulation settings, in Sec. III. Finally, the results of our comparative evaluation are shown and commented in Sec. IV. Sec. V concludes the paper.

II. PERFORMANCE METRICS

We aim at investigating the impact of mobility models on the simulation of vehicular networks and their supported services. To that end, we must make a choice of performance metrics that is necessarily circumscribed, and yet covers a vast portion of the many and varied vehicular networking use cases. Moreover, for the sake of generality, it is desirable that the metrics are not specific to any network architecture or protocol.

We propose, in Tab. I, a taxonomy of relevant case studies in vehicular network environments. We leverage such a taxonomy to show how the performance metrics we will consider in this work allow for a decent coverage of the research topic space.

The top rows of the table outline our classification of services and network operations related to vehicular networking. The *underlying function* row denotes inter-dependencies between the aforementioned services and network operations, whereas each subsequent row refers to one (set of) performance metrics we will consider in our work. Finally, colored cells in the table indicate which metrics are relevant to each service or network operation: dark blue evidences metrics that are needed for the assessment of a service or operation, and light red marks metrics relevant to the underlying functions required by the service and operation. An asterisk (*) symbol denotes situations where the metric is still pertinent to the case study, but not in mainstream approaches.

The research leading to these results was carried out while Marco Gramaglia was at CNR-IEIIT, and has received funding from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement n.630211 ReFlex.

TABLE I: Categories of connectivity metrics, and their relevance to vehicular networking studies. Colors and symbols read as follows. Relevant. Relevant to underlying function. * Seldom considered. Irrelevant.

	Application-level services									Network-level services					Network architectures & protocols			
	Road safety		Traffic efficiency		Infotainment					Beaconing	Upload	Download	Dissemination	Fusion	Cellular RAN	RSU	Routing	MAC
	Cooperative awareness	Hazard warning	Cooperative speed control	Cooperative navigation	Insurance	Fleet management	Content access	Massive notification	Sensing									
Underlying functions	(1)		(1)	(2) (3) or (1) (4) (5)	(2)	(2)	(3)	(4)	(5)	(d)	(a) (b)	(a) (b)				(d)		
I Network connectivity	*	 	*	 	 	 	 	 	 	 	 	 	 	 	 	 	 	
II Component availability	 	 	 	*	*	*	*	 	 	 	*	*	 	 	 	 	*	
III Component stability	 	 	 	*	*	*	*	 	 	 	*	*	 	 	 	 	*	
IV Single-hop connectivity	 	 	 	 	 	 	 	 	 	 	 	 	 	 	 	 	 	
V Contact duration	 	 	 	*	 	 	 	 	 	 	 	 	 	 	 	 	 	
VI Cellular timings	 	 	 	 	 	 	 	 	 	 	 	 	 	 	 	 	 	

A. Services and network operations

We identify three main classes of case studies in the vehicular networking literature, which we detail separately.

Application-level services. These are the actual utilities that sit on top of the vehicular network system, and implement intelligent transport system (ITS) solutions related to road safety and traffic efficiency, as well as inboard infotainment. We tell apart application the use cases below, in agreement with ETSI definitions [16].

Concerning road safety, *cooperative awareness* includes all safety services that require periodic broadcasting by vehicles, e.g., emergency and slow vehicle warning, or lane change assistance; *hazard warning* involves instead on-event messaging for, e.g., accident, collision risk, or roadwork warning. All these services rely on a V2V communication that is reliable and fulfills precise latency requirements. Cooperative awareness also builds on periodic broadcasting, for which beaconing is a fundamental network-level service.

Traffic efficiency includes services that aim at reducing travel times and at increasing road capacity utilization. They primarily target *cooperative speed control* in presence of road congestion or traffic lights, and *cooperative navigation* allowing for real-time traffic notification and route calculation. The first service fully relies on periodic broadcasting, while the second requires, depending on the implementation, upload of floating car data and download of processed information (centralized approach) or periodic broadcasting and in-network fusion and dissemination (decentralized approach).

Infotainment use cases are in the first place commercial, such as those associated to vehicle *insurance* and *fleet management*. These services only require periodic upload of information to some Internet-based server. A more general task is *content access* from the Internet, typically for delay-bounded applications such as web browsing or media streaming. Here, services basically rely on download functionalities of the network. Additional use cases are *massive notification* to a large set of vehicles, often concerning small-sized data (containing, e.g., software or map updates, advertisements, or breaking news), and *sensing* of some physical phenomenon using inboard sensors (e.g., pollution or acoustic noise). In both situations, communication occurs only within the vehicular network, and is thus based on dissemination and fusion of

information, respectively. Involvement of the infrastructure for the retrieval of seed information or the final upload of sensed data may be also envisioned, yet it is marginal to the process.

Network-level services. These are high-level network facilities that perform complex telecommunication tasks involving connected vehicles.

Periodic *beaconing* is a foremost operation in vehicular networking, allowing each communication-enabled vehicle to broadcast its current status at a high frequency (typically, several times per second). Since it is a one-hop broadcast communication, it heavily relies on MAC-layer network functionalities. Unlike beaconing, *upload* and *download* network-level services involve data transfer to and from the Internet, and thus depend on some kind of access infrastructure, be it the traditional cellular radio access network (RAN) or a dedicated roadside unit (RSU) deployment. A decentralized, distributed networking paradigm is behind *dissemination* and *fusion* of data. These in-network operations fully rely on multi-hop V2V connectivity in order to distribute (one-to-many) or gather (many-to-one) some data, respectively.

Network architecture and protocols. These are the basic network functionalities allowing vehicles to communicate with each other and with the infrastructure.

The categories are fairly self-explanatory: *cellular RAN* and *RSU* concern the planning, dimensioning, and management of cellular and RSU access infrastructures; *routing* concerns the multi-hop point-to-point or point-to-multipoint transfer of messages within the vehicular network, typically in a fully distributed fashion; *MAC* relates to all medium access control operations, including data rate adaptation and power control.

B. Description of metrics

As previously mentioned, our choice of performance metrics is driven by three guidelines: (i) the set of metrics should not be too large, for the sake of readability of results; (ii) the metrics should cover well the space of case studies in the vehicular networking literature; (iii) the metrics should not be specific to any protocol or system architecture.

In order to fulfill these requirements, we pick metrics that allow assessing those properties of vehicular networks that are fundamental and transversal to the different case studies, i.e., the topological features of the network. We thus

primarily aim at investigating the structure of a network built by communication-enabled vehicles (rows I to V in Tab.I). However, this leaves out those studies that concern interactions of connected vehicles with the cellular network, and thus we also include metrics that relate to those type of analysis (row VI in Tab.I). As they are quite different in nature, we discuss the two types of metrics separately in the following.

Connectivity metrics. Most of the metrics we consider describe the instantaneous connectivity of vehicular networks. Formally, at each time instant t , we represent the network as an undirected graph $G(\mathbb{V}(t), \mathbb{E}(t))$. Each vertex in the set $\mathbb{V}(t) = \{v_i(t)\}$ maps to a vehicle i in the network at time t , and each edge in the set $\mathbb{E}(t) = \{e_{ij}(t)\}$ connects $v_i(t)$ and $v_j(t)$ if a V2V communication link exists between vehicles i and j at time t . We also denote as $\mathcal{N}(t) = \|\mathbb{V}(t)\|$ the number of vertices in the graph (i.e., the number of vehicles in the road scenario) at time t .

The *network connectivity* metrics describe the global structure of the vehicular network, evaluating its overall level of connectivity or fragmentation. Formally, let us define a component $C_m(t) = G(\mathbb{V}_m(t), \mathbb{E}_m(t))$ as a subgraph of $G(\mathbb{V}(t), \mathbb{E}(t))$, where $\mathbb{V}_m(t) \subset \mathbb{V}(t)$ includes all and only the vertices corresponding to vehicles that can reach each other via direct or multi-hop communication at time t . Equivalently, $\mathbb{E}_m(t) = \{e_{ij}(t) \mid v_i(t), v_j(t) \in \mathbb{V}_m(t)\} \subseteq \mathbb{E}(t)$. We denote as $\mathcal{S}_m(t) = \|\mathbb{V}_m(t)\|$ the size of the component $C_m(t)$. Since components are disjoint by definition, $\mathcal{C}(t) = \|\{C_m(t)\}\|$ is the number of components appearing in the network at time t . These number and size of components in the network at each time instant will be our network connectivity metrics.

The *component availability* and *component stability* metrics study large connected components emerging in the network, which are especially interesting as they allow for significant multi-hop communication opportunities. In particular the two metrics focus on (i) the presence and (ii) the temporal fluctuations of such large components. Formally, we refer to the largest component appearing in the network at time t as $C_{max}(t) = G(\mathbb{V}_{max}(t), \mathbb{E}_{max}(t)) = C_m(t) \mid m = \arg_n \max \mathcal{S}_n(t)$. Then, $\mathcal{S}_{max}(t) = \|\mathbb{V}_{max}(t)\|$ is the size of the largest component at the same time instant. The normalized value of $\frac{\mathcal{S}_{max}(t)}{\mathcal{N}}$ at each instant will be our reference metric for the study of the component availability, whereas its temporal variations will be leveraged to analyze the component stability.

The component stability is assessed through the correlograms of $\mathcal{S}_{max}(t)$. Correlograms are derived by dividing $\mathcal{S}_{max}(t)$ time series into 10-minute windows¹ and computing the temporal autocorrelation at different lags, for each window.

The *single-hop connectivity* metrics focus on the direct neighborhood observed by each vehicle in the network. Formally, we denote as $k_i(t) = \|\{v_j(t) \text{ s.t. } \exists e_{ij}(t)\}\|$ the vertex degree, i.e., the number of vertices directly connected to a given vertex $v_i(t)$ at time t . The degree of vertex $v_i(t)$ thus maps to the number of direct vehicle-to-vehicle (V2V) communication neighbors of vehicle i , and we will employ it to investigate single-hop connectivity of vehicular networks.

¹Time windows need to be long enough not to truncate any significant autocorrelation. In all our scenarios, autocorrelations tend to disappear after a few tens of seconds, thus 10 minutes are largely sufficient to the study.

The *contact time* of a single-hop V2V link is instead defined as $\tau_{ij}(t) = \tau_{ij}^e(t) - \tau_{ij}^s(t) = \arg \max_t \{t \mid e_{ij}(t) \in \mathbb{E}(t)\} - \arg \min_t \{t \mid e_{ij}(t) \in \mathbb{E}(t)\} \iff e_{ij}(t) \in \mathbb{E}(t) \forall t \in [t_s \dots t_e]$. In other words, it is the amount of consecutive time instants that include time t and feature a valid communication link between vehicles i and j .

In the remainder of the paper, we will drop the time notation for the sake of brevity, and refer all metrics to a generic time instant. We will thus use \mathcal{N} to indicate the number of vertices in the network, \mathcal{C} for the number of components, and \mathcal{S}_{max} the largest component size. Similarly, k is to be intended as the node degree of a generic vertex, and τ as the duration of a generic V2V link.

Cellular network metrics. User mobility has a significant impact on multiple facets of cellular networks, including deployment, handover management, and resource allocation. Two measures are especially relevant to all of the above aspects, and we include them in our choice of metrics. The first measure is the *cell inter-arrival*, i.e., the time elapsed between two subsequent arrivals of users at a cell. Inter-arrivals can be mapped to handover rates, and their characterization has significant implications especially in terms of signalization overhead in cellular networks. The second measure is the *cell residence time*, also referred to as sojourn time, i.e., the temporal interval spent by users within a cell. The residence time is paramount to estimate the access load and dimension network capacity.

C. Coherence of use cases and metrics

Tab.I provides an intuitive visualization of the good coverage of vehicular networking case studies provided by the performance metrics presented above.

Road safety services related to cooperative awareness and hazard warning mainly rely on single-hop communication, and their feasibility and performance depend on the volume of vehicles within reach of the broadcaster. Thus, single-hop connectivity metrics are highly relevant to these applications. Also, cooperative awareness is based on periodic beaconing, making additional metrics pertinent to its study, although in an indirect way. Finally, some hazard notification use cases, such as decentralized floating car data [16], leverage multi-hop communication among vehicles, making the impact of global network connectivity important in this type of services.

Cooperative traffic efficiency and infotainment services entail more complex operations than one-hop broadcasting, and thus all rely on some network-level services. This makes the different metrics associated to the network-level services also important for the overlay application-level services.

Concerning the aforementioned network-level services, beaconing is basically one-hop broadcasting; thus metrics defining single-hop connectivity should capture any impact of mobility on beaconing. In some, less frequent, cases, multi-hop beaconing is also considered, making network connectivity metrics also meaningful to the performance of this network-level service. Visibly, the performance of MAC-layer protocols are critical to beaconing, and so are the associated metrics.

Upload and download services heavily rely on the access infrastructure. However, a significant amount of work has focused on complementing cellular and RSU data transfers

with in-network data transfer, so as to offload the former from part of their load. In this case, the global level of connectivity in the network drives the offloading capabilities of the system, and the associated metrics become relevant to studies on vehicular upload and download. In particular, the instantaneous connectivity is especially important in presence of the delay-bounded services that typically rely on upload and download functionalities. When it comes to offloading in the downlink, we also deem contact duration a critical aspect, since the size of contents can be large, and the amount of transferable data upon each V2V contact becomes important.

Dissemination and fusion of data are visibly in-network functionalities that leverage to a maximum the (multi-hop) transmission opportunities offered by V2V communication. As such, not only the overall level of connectivity of the vehicular network, but also the availability and stability of large connected components are essential to the system performance. The contact duration is another important metric, since dissemination may concern large-sized contents.

The architecture- and protocol- independent nature of the metrics we consider let us associate them to multiple basic networking operations. From the access infrastructure viewpoint, cellular metrics are relevant, rather obviously, to the cellular architecture. In addition, RSU deployment strategies take often into account the structure of the V2V network, so as to improve its connectivity and robustness through fixed units; less frequently, the internal connectivity of large components is also accounted for in RSU planning studies.

Connected multi-hop routing is instead completely built on V2V communication. Since it aims at allowing in-network unicast or multicast transfers, both network and component connectivity metrics are good proxies of the support provided to this functionality by connected vehicles. It is interesting to remark how no higher-level service among those listed in the table actually relies on routing: we consider this an indicator of the lack of practical applications for this communication model in vehicular environments.

Finally, MAC-layer operations only concern the immediate communication surroundings of a vehicle. Thus, the important parameters here are the number of vehicles within range, which determines the channel congestion level, and the duration of links, which drives responsiveness requirements for, e.g., contention window and data rate adaptation.

Overall, our proposed metrics cover well the wide variety of application- and network-level use cases we identified. We will thus study the impact that different mobility models have on the metrics, knowing that variations in the metrics will then reflect on the performance of services or network architectures and solutions as shown in Tab. I.

III. ROAD TRAFFIC MODEL SPACE

We identify the following components required for the simulation of highway vehicular networks.

- **Traffic input feed.** The information concerning the inflow of vehicles in the simulated highway segment.
- **Mobility model.** The representation of the driving behavior of vehicles that travel on the simulated segment.
- **Propagation model.** The model of the radio-frequency signal propagation. It is especially relevant to the V2V

communication use case, as it determines whether vehicles are within range.

- **Network simulator.** The representation of the protocol stack for the packet-level simulation of the network.

Clearly, only the first two items pertain to the road traffic modeling we are interested in. Thus, they represent the main focus of our study. We also need a propagation model, since several of the performance metrics we introduced in Sec. II are based on a graph representation of V2V communications. Instead, as our metrics are protocol-independent, we will not need to run a packet-level network simulation, and the last item will not be part of our study. Next, we detail how we model the relevant items in our analysis.

A. Traffic input feed

There are two endpoints categories for the input fed to vehicular simulations in the literature.

The *real* traffic input feed is straightforward: vehicles are inserted into the simulation according to a fine-grained real-world traffic count source. Such traffic count sources provide information on the transit of each vehicle, and include data such as the lane, time-stamp, speed, and possibly length of the vehicle. These high-precision sources are not easily collected, yet they have been employed in the literature [6].

The *synthetic* traffic input feed relies instead on probability distributions of the inter-arrival or inter-spacing between subsequent vehicles to generate the inflow into the simulated highway segment. The distributions employed in the literature to that end vary, and include deterministic [7], [14], exponential [12] and log-normal [21] arrivals, up to generative models for mixture distributions [15]. Also the initial speed is derived from a probability distribution, typically uniform [10].

In some cases, synthetic traffic input feeds can be trained on real-world traffic counts. In this case, traffic counts are typically aggregated by fitting a theoretical distribution on the inter-arrival times of all the vehicles flowing within a time window w [10], [13]. Clearly, the shorter is the time window w the more accurate is the input feed.

We opted for a set of five different input feeds: *real* indicates real traffic input feed, whereas four versions of synthetic traffic counts are denoted by *synthetic- w* . In the *real* dataset, vehicles are inserted into the simulation using their real initial time and speed. The parameter w used for the generation of *synthetic- w* feeds is the selected time window over which the traffic count data is aggregated: 5 minutes, 10 minutes, 30 minutes or one hour. The inter-arrival times for the feed *synthetic- w* are exponentially distributed as $f_w(t) = \lambda_w e^{-\lambda_w t}$ where $\lambda_w = \frac{N_w}{w}$ is the average number of vehicles per unit of time [10]. The initial speed is distributed according to a uniform random variable $f_w(s) = \mathcal{U}(S_w^{min}, S_w^{max})$, where $S_w^{min} = 0.9 \bar{S}_w$ and $S_w^{max} = 1.1 \bar{S}_w$ and \bar{S}_w is the average inflow speed observed during time window w [10]. Starting lanes are randomly selected in the *synthetic- w* cases.

In our experiments, we use or train distributions on real-world traffic counts kindly provided by the Madrid City Council, Spain. The data spans 8 hours, from 4 a.m. to 12 p.m., of a typical working day, and describe transits on M30, a 4-lane highway in proximity of the city of Madrid. Despite

their limited timespan, these real-world traffic counts yield very heterogeneous vehicular densities, covering all of the typical conditions encountered in everyday traffic. Specifically, the data is representative of very sparse overnight traffic (until 7.30 a.m. approximately), rush hour congestion with speed breakdown (until 8.30 a.m.) and free-flow moderate-density traffic (until the end of the feed). We simulated a 4-Km road, i.e., a distance separating consecutive ramps and long enough to allow for a wide range of network studies.

B. Mobility model

The literature on highway vehicular simulation features heterogeneous mobility models. Some works use low complexity representations [10], [11], whereas others rely on dedicated simulators [7], [14]. In some cases, researchers developed proprietary tools tailored to their needs [6]. We tested all of these approaches, as follows.

The `unstructured` approach simply assign a speed to each vehicle entering the simulated highway segment, and lets each vehicle travel at that constant velocity along the whole segment. Clearly, this model completely neglects interactions among vehicles, and possibly lets them overlap. It is however a computationally inexpensive approach that has been largely adopted in vehicular networking research.

The `SUMO` approach leverages the `SUMO` tool, i.e., the *de-facto* standard open-source software for the simulation of vehicular mobility [17]. `SUMO` implements microscopic car-following and lane-changing models. The former is Krauss' model [18], and regulates each vehicle acceleration as a function of the distance to the leading one, the current speed, the safety distance, or the acceleration and deceleration profiles. The latter is Krajzewicz's model [19], and lets vehicles take overtaking and lane-change decisions, considering the position and speed of nearby vehicles on different lanes. These models provide a much more complex and realistic representation of the movement of each vehicle within the traffic flow.

The `IDM` approach is an example of own-coded software². Similarly to `SUMO`, the tool implements well-known microscopic models. The Intelligent Driver Model (`IDM`) [20] is a popular car-following model similar in nature to Krauss', with lane-changing functionalities based on a game-theoretical approach. Unlike `SUMO`, this is not a general-purpose tool that can be used to simulate very different road environments: instead, it is a simple software for pure highway mobility.

All of the mobility models above are fed with the data presented in Sec. III-A. For the `unstructured` approach, the initial speed is derived from a probability distribution fitted on the real-world traffic counts collected in Madrid. This approach is consistent with those used in the literature [10]. For `SUMO` and `IDM` the calibration of the models is more complex. In the both cases, the target speed and minimum gap between subsequent vehicles are calculated as described in [6], whereas all other parameters in Krauss', Krajzewicz's, and `IDM` models are left to their default.

C. Propagation model

Our focus is on road traffic modeling and not on the way the propagation of radio-frequency signals is represented. This

notwithstanding, the choice of the propagation model is a very important aspect when performing vehicular network simulations [22]. It is thus critical that a single, dependable model is used throughout our performance evaluation.

A plethora of wireless propagation models are available nowadays, each one taking into account different aspects of V2V channels, such as the presence of interfering vehicles or all kinds of obstacles. Since buildings are not an issue in highway vehicular simulations, for the sake of scalability we discard complex deterministic propagation models based on e.g., ray tracing. Instead, we extracted the coverage distance from a state-of-the-art propagation model [23], considering a transmission power is set to 20 dBm, a received signal strength threshold of -91 dBm, and a reliability of .99. Shadowing effects due to nearby vehicles are considered as well, as an additional pathloss when the latter obstruct the line-of-sight.

IV. RESULTS

We discuss the impact of road traffic models on our set of network metrics by separating V2V and cellular connectivity.

A. Connectivity metrics

We first assess the impact of mobility modeling on the global network connectivity, expressed as the number of components \mathcal{C} . Since they display equivalent trends, we show at the same time the results on component availability, i.e., the ratio between \mathcal{S}_{max} and \mathcal{N} . Fig. 1 and Fig. 2 portray smoothed scatter plots that refer to each mobility description, for the two metrics separately. All plots show the metrics as functions of the road traffic density, in vehicles per Km. When comparing the plots within each figure, remarkable differences emerge.

First, we observe that the parameter w (in minutes) strongly influences the connectivity and availability metrics. While Fig. 1a, 1b and 1c show a comparable and realistic behavior; using synthetic traffic with $w = 60$, in Fig. 1d, yields an abrupt transition between the disconnected ($\mathcal{C} \sim 10$) and fully connected ($\mathcal{C} \sim 1$) phases. Equivalent considerations hold when synthetic traffic is combined with microscopic mobility, see Fig. 1g and Fig. 1j, as well as in the case of availability, in Fig. 2. We conclude that *an exceedingly coarse inflow granularity risks to completely lose the state transitions that occur in real-world traffic, as well as the associated connectivity and availability states*. Unfortunately, w is often a non-configurable parameter decided by the data providers, who are typically only interested in rough aggregates of the inflow traffic for statistical purposes.

Second, the use of `SUMO` appears to cause issues with the observed metrics. All plots where `SUMO` is used to model the vehicular mobility show that the mobility generator is just unable to insert all the vehicles in the simulation. This is clear when looking at Fig. 1f and Fig. 1g-Fig. 1j for the network connectivity, but the same considerations also hold for the component availability, in Fig. 2. While `unstructured` and `IDM` attain almost 100 vehicles per Km over the four highway lanes, the traffic peak in `SUMO` is 25% lower. This is a parametrization issue: the default settings of the Krauss' model do not allow accommodating high inflows observed in the real world, which forces `SUMO` to delay the insertion of a vehicle until Krauss' model safety requirements are fulfilled. In turn, this affects network connectivity and availability.

²Available at <http://www.it.uc3m.es/madrid-traces>.

Fig. 1: \mathcal{C} versus the vehicular density, the red line denotes the average.

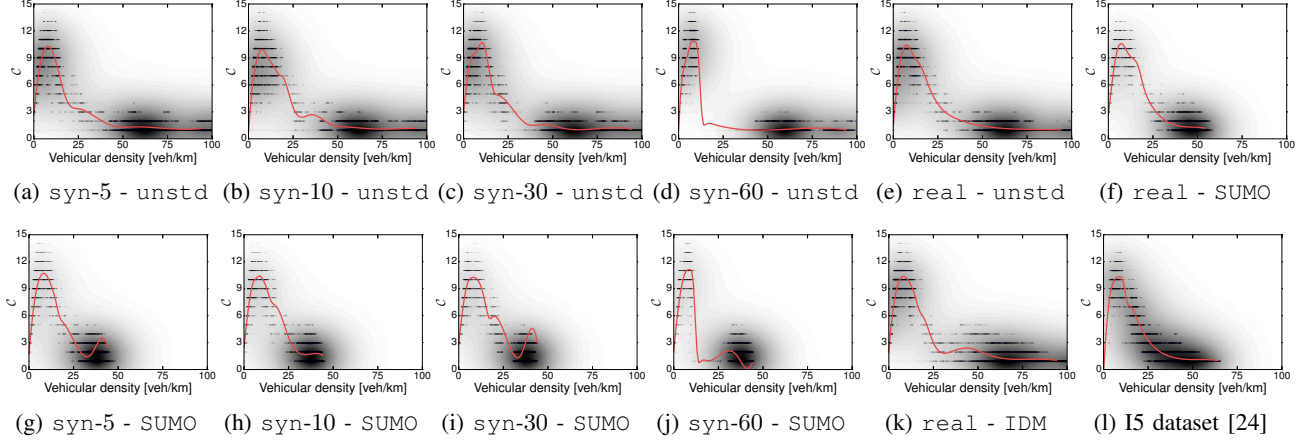



Fig. 2: The relative availability versus the vehicular density, the red line denotes the average.



These results prove that *using a validated microscopic model of vehicular mobility is not sufficient to obtain a realistic representation of road traffic: the parametrization of the model is extremely important, and a careless setting can lead to biased simulation outcomes*. Clearly, this is not a problem of Krauss' model per-se. In order to prove it, we also show the connectivity and availability metrics obtained using the mobility dataset generated by Akhtar *et al.* [24], which was generated using SUMO with customized (but undisclosed) parametrization. Fig. 2l and Fig. 1l show similar trends to those obtained with unstructured and IDM.

Third, a very interesting observation is that *a very simple constant-speed simulator using synthetic (but sufficiently detailed) traffic input feed results in a network connectivity and availability comparable to those attained by much more complex models*. Fig. 1a, 1e, and 1k show precisely this effect.

Fourth, we stress that the highway road traffic trace by Akhtar *et al.* [24] describes traffic in a different scenario, i.e., Interstate highway 5 (I5) in CA, USA. Still, the connectivity and availability scatter-plots and mean curves are identical to those of our reference scenario, i.e., M30 in Spain. This result let us speculate on the general validity of our findings, which could apply to different highway environments.

The correlograms of S_{max} in Fig. 3 display the temporal variation of the largest connected component in the network: they map to the component stability metric. Here, we only

display a subset of the results, for the sake of brevity and since w did not appear to influence the component stability. Again, SUMO, in Fig. 3b and 3d, shows a very different trend due to the maximum density issue we already discussed. However, the important result here is that the unstructured mobility model starts showing limitations. Fig. 3a and 3c prove how the lack of interaction among vehicles in these models results in correlograms that differ from that obtained with IDM, in Fig. 3e. In the latter model, drivers are forced to adjust their speed according to the surrounding road traffic conditions, which leads to well-known phenomena, such as synchronized traffic: in turn, the global reduction of speed and queuing of vehicles noticeably improve connected component lifetime. We conclude that *a simplistic representation of microscopic mobility does not impact network-wide metrics, but leads to connected components that may be significantly less stable in time than what would occur in the real world*.

As far as single-hop connectivity is considered, we report the Cumulative Distribution Function (CDF) of the relevant metrics in Fig. 4 and Fig. 5. Also in this case, the lower vehicular densities enforced by SUMO result in distinctive and biased V2V link durations (much shorter than normal) and node degrees (sensibly lower than those obtained under the other models). These are artifacts of the fact that vehicles are injected more slowly in SUMO, and experience less congested traffic conditions.




Fig. 3: $\mathcal{S}_{max}/\mathcal{N}$ correlograms.




Fig. 4: Contact time τ .




Fig. 5: Node degree k .

Feeding an unstructured simulator with synthetic data leads to V2V contacts, in Fig. 4, that are longer than those recorded with SUMO or IDM. The reason is that, under the former model, the speed is constant and extracted from a limited value range (see Sec. III-A). So it often happens that nearby vehicles have similar speed and travel a long way within each other communication range. We argue that *the speed modeling becomes paramount when using very simple mobility models, and a simple (and widely adopted) approach such as that in Sec. III-A clearly biases the simulation results.*

In addition, we remark that the lack of a target speed in the unstructured approach significantly shifts the k distribution with respect to those obtained under more realistic mobility, as depicted in Fig. 5. In the unstructured case, vehicles do not speed up: hence, during the congested traffic period, vehicles enter the simulation at high inflow rate and low speed, generating unrealistically large one-hop neighborhoods. Instead, IDM allows vehicles to accelerate and overtake, which eliminates the long tail of very high vertex degrees.

B. Cellular metrics

We next analyze the effect of the different simulation approaches on the noteworthy metrics for cellular networks. The first considered metric is the inter-arrival time, computed on a




Fig. 6: Inter-arrival time.




Fig. 7: Residence time.

radio cell boundary placed halfway of the highway segment. From the CDFs shown Fig. 6, SUMO induces longer inter-arrival times due to its lower vehicular density. Instead, IDM and unstructured (using both real and synthetic data) perform quite similarly. Interestingly, unstructured provides good coherence with the actual vehicular spacing even though vehicles are not changing lane, but overlapping. As a corollary, the inter-arrival time distribution is never purely exponential, as all the simulation techniques failed the KS test. Thus, we stress that *the common practice of modeling inter-arrivals as a Poisson process may bias networking results.*

The CDF of the vehicle residence time in a reference cell covering 2 Km of the highway is depicted in Fig. 7. SUMO and IDM perform similarly. Conversely, the unstructured mobility model generates a heavy tail of long residence times: again, this is due to constant speeds that remain unrealistically low during congested periods. We conclude that *simplicistic mobility models that neglect drivers' desired speeds may not be fit to studies involving cellular networks.* We vary the cell coverage in the violin plot of Fig. 8, so as to account for different deployments and technologies (e.g., DSRC RSUs usually cover just a few hundred meters). Results are consistent with the discussion above: heavy tails are observed in all unstructured distributions.

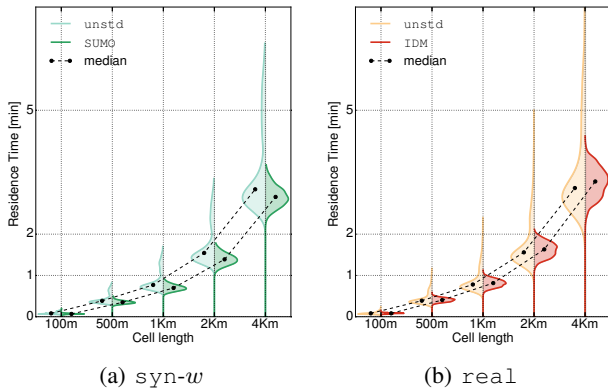


Fig. 8: Residence time versus cell size.

TABLE II: Summary of results

Feed	Synthetic		Real		
Simulator	unstd	SUMO	unstd	SUMO	IDM
Network Connectivity	✓(w)	×	✓	×	✓
Component Availability	×	×	×	×	✓
Component Stability	×	×	×	×	✓
Single Hop Connectivity	×	×	×	×	✓
Contact Duration	×	×	✓	×	✓
Cell Inter-arrivals	✓	×	✓	×	✓
Cell Residence	×	×	×	×	✓
Services	massive notification, sensing, dissemination, fusion	none	massive notification, sensing, dissemination, fusion	none	all

V. DISCUSSION AND CONCLUSIONS

We investigated the kind of vehicular mobility representation that is necessary for a correct assessment of the performance of specific applications or network services. Tab. II summarizes our results.

A specialized highway simulator like IDM, fine-tuned on a real data feed provides the best accuracy, however it requires fine-grained real-world traffic counts and a complex parametrization process such as that described in [6].

In our tests, a state-of-the-art mobility simulator such as SUMO never yielded good results: the default parameters of the microscopic mobility models could not accommodate congested road traffic situations. The problem may be solved by a dedicated configuration of the tool such as that used, but not disclosed, in [24].

As an alternative, an unstructured simulator where vehicles travel at constant speed may be a correct choice, depending on the networking metrics we consider. This simple approach works for high-level aggregate metrics, e.g., the network-wide connectivity, the availability of large connected components, or the inter-arrivals at RAN or RSU cells. Instead, it tends to fail when more precise metrics are considered.

We also observe that `synthetic` data can be used to feed simulators, if not aggregated over too large temporal windows w . The value of w should be small enough to capture the state transitions in real-world traffic.

By cross-referencing the results in Tab. II and the associations between metrics and services in Tab. I, we can draw some conclusions on how different road traffic representations fit vehicular networking studies. The bottom row of Tab. II evidences how, based on our results, the IDM approach can be employed with any type of services, whereas SUMO with

default settings is deprecated. A simple unstructured simulation with real or sufficiently detailed synthetic may be used (at the cost of some approximation, since component stability is not properly modeled in this case) for a limited subset of services, including massive notification, sensing, data dissemination and fusion.

REFERENCES

- [1] Reuters, "Obama Backs Highway Fund Fix, Touts 'Talking' Cars," The New York Times, 2014.
- [2] H. Hartenstein, K.P. Laberteaux, "A tutorial survey on vehicular ad hoc networks," IEEE Communications Magazine, 46(6):164–171, 2008.
- [3] K.-H. Chang, Wireless communications for vehicular safety, IEEE Wireless Communications, 22(1):6–7, 2015.
- [4] S. Uppoor, O. Trullols-Cruces, M. Fiore, J. Barcelo-Ordinas, "Generation and analysis of a large-scale urban vehicular mobility dataset," IEEE Transactions on Mobile Computing, 2014.
- [5] Y. Pigné, G. Danoy, P. Bouvry, "A vehicular mobility model based on real traffic counting data," Nets4Cars/Nets4Trains, 2011.
- [6] M. Gramaglia, O. Trullols-Cruces, D. Naboulsi, M. Fiore, M. Calderon, "Vehicular Networks on Two Madrid Highways," IEEE SECON, 2014.
- [7] N. Akhtar, S. Coleri Ergen, O. Ozkasap, "Vehicle Mobility and Communication Channel Models for Realistic and Efficient Highway VANET Simulation," IEEE Transactions on Vehicular Technology, 2015.
- [8] S. Joerer, C. Sommer, F. Dressler, "Toward reproducibility and comparability of IVC simulation studies: a literature survey," IEEE Communications Magazine, 2012.
- [9] M. Fiore, J. Härrri, "The networking shape of vehicular mobility," ACM MobiHoc, 2008.
- [10] F. Bai, B. Krishnamachari, "Spatio-temporal variations of vehicle traffic in VANETs: facts and implications," ACM VANET, 2009.
- [11] S. Yousefi, E. Altman, R. El-Azouzi, M. Fathy, "Analytical model for connectivity in vehicular ad hoc networks," IEEE Transactions on Vehicular Technology, 2008.
- [12] M. Khabazian, M.K. Ali, "A Performance Modeling of Connectivity in Vehicular Ad Hoc Networks," IEEE Transactions on Vehicular Technology, 2008.
- [13] R. Monteiro, S. Sargento, W. Viriyasitavat, O.K. Tonguz, "Improving VANET protocols via network science," IEEE VNC, 2012.
- [14] M. De Felice, A. Baiocchi, F. Cuomo, G. Fusco, C. Colombaroni, "Traffic monitoring and incident detection through VANETs," IEEE WONS, 2014.
- [15] M. Gramaglia, M. Fiore, M. Calderon, "Measurement-based modeling of interarrivals for the simulation of highway vehicular networks," IEEE Communications Letters, 2014.
- [16] European Telecommunications Standards Institute, "Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Definitions," ETSI TR 102 638 V1.1.1, 2009.
- [17] D. Krajzewicz, J. Erdmann, M. Behrisch, L. Bieker, "Recent Development and Applications of SUMO - Simulation of Urban MObility," International Journal On Advances in Systems and Measurements, 2012.
- [18] S. Krauß, P. Wagner, C. Gawron, "Metastable States in a Microscopic Model of Traffic Flow," Physical Review E, 1997.
- [19] D. Krajzewicz, "Kombination von taktischen und strategischen Einflüssen in einer mikroskopischen Verkehrsflusssimulation", in T. Jürgensohn, H. Kolrep (editors), Fahrermodellierung in Wissenschaft und Wirtschaft, VDI-Verlag, 104–115, Berlin-Adlershof, Germany, 2009.
- [20] M. Treiber, A. Hennecke, D. Helbing, "Congested Traffic States in Empirical Observations and Microscopic Simulations," Physics Review E, 2000.
- [21] N. Wisitpongphan, F. Bai, P. Mudalige, V. Sadekar, O. Tonguz, "Routing in sparse vehicular ad hoc wireless networks," IEEE Journal Selected Areas in Communications, 2007.
- [22] W. Viriyasitavat, M. Boban, H. Tsai and A. Vasilakos, "Vehicular Communications: Survey and Challenges of Channel and Propagation Models," IEEE Vehicular Technology Magazine, 2015.
- [23] T. Abbas, K. Sjöberg, J. Karedal, F. Tufvesson, "A Measurement Based Shadow Fading Model for Vehicle-to-Vehicle Network Simulations," International Journal of Antennas and Propagation, 2015.
- [24] N. Akhtar, S. C. Ergen, and O. Ozkasap, "Vehicle mobility and communication channel models for realistic and efficient highway VANET simulation," IEEE Transactions on Vehicular Technology, 2015.