

This is a postprint version of the following published document:

Glacet, Christian; Fiore, Marco; Gramaglia, Marco.
(2015). Temporal connectivity of vehicular networks:
the power of store-carry-and-forward. *2015 IEEE
Vehicular Networking Conference (VNC), Kyoto,
Japan, 16-18 December 2015: [proceedings]*. [8] p.

DOI: [10.1109/VNC.2015.7385546](https://doi.org/10.1109/VNC.2015.7385546)

©2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Temporal Connectivity of Vehicular Networks: The Power of Store-Carry-and-Forward

Christian Glacet*, Marco Fiore*, Marco Gramaglia†

*CNR-IEIIT, Torino, Italy

{christian.glacet,marco.fiore}@ieiit.cnr.it

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

mgramagl@it.uc3m.es

Abstract—Store-carry-and-forward is extensively used in vehicular environments for many and varied purposes, including routing, disseminating, downloading, uploading, or offloading delay-tolerant content. The performance gain of store-carry-and-forward over traditional connected forwarding is primarily determined by the fact that it grants a much improved network connectivity. Indeed, by letting vehicles physically carry data, store-carry-and-forward adds a temporal dimension to the (typically fragmented) instantaneous network topology that is employed by connected forwarding. Temporal connectivity has thus an important role in the operation of a wide range of vehicular network protocols. Still, our understanding of the dynamics of the temporal connectivity of vehicular networks is extremely limited. In this paper, we shed light on this underrated aspect of vehicular networking, by exploring a vast space of scenarios through an evolving graph-theoretical approach. Our results show that using store-carry-and-forward greatly increases connectivity, especially in very sparse networks. Moreover, using store-carry-and-forward mechanisms to share content within a geographically-bounded area can be very efficient, i.e., new entering vehicles can be reached rapidly.

I. INTRODUCTION

Connectivity promises to bring disruptive innovations in the automotive industry. Already today, connected vehicles equipped with cellular interfaces enable services such as road traffic estimation and smart navigation [1], remote vehicle monitoring [2], or insurance data recording [3]. Cellular connectivity will be soon complemented by vehicle-to-vehicle communications. The technology is standardized [4]–[6], and may become mandatory on all new cars sold in the US as early as 2017 [7].

Direct communication among vehicles is primarily expected to enable road safety services that have stringent latency requirements, as well as to provide decentralized support to traffic efficiency applications [8]. However, vehicle-to-vehicle communications will also warrant a wide range of infotainment services to drivers and passengers.

In the light of the interest towards direct vehicular communications, the networking research community has made a significant effort in proposing architectures and protocols

This work was done while Marco Gramaglia was at CNR-IEIIT. The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Unions Seventh Framework Programme (FP7/2007-2013) under REA grant agreement n.630211 ReFlex. The work of Christian Glacet was carried out during the tenure of an ERCIM “Alain Bensoussan” Fellowship Programme.

for vehicular environments. This has led to a vast literature, thoroughly reviewed by comprehensive surveys [9]–[11].

A subject that has attracted significant attention within the networking community is that of multi-hop communication among moving vehicles. This data transfer paradigm is especially challenging, since it turns short-range radio-equipped vehicles into a fully distributed network of highly dynamic nodes. In such an environment, leveraging traditional connected forwarding (as done in wired networks or wireless networks of static nodes) proved exceedingly hard. Initial attempts at performing connected forwarding were gradually turned down, in favor of store-carry-and-forward approaches. The latter limit the support of the multi-hop vehicular network to delay-tolerant services only, but with a much higher reliability. They have been leveraged for multiple purposes, including unicast routing [12], [13], multicast routing [14], dissemination [15], content downloading [16], cellular network offloading [17], or floating car data collection [18].

The fundamental reasons behind the limited viability of connected forwarding in spontaneous networks of vehicles were neatly exposed by topological analyses of the instantaneous connectivity of the network. Specifically, studies carried out on large-scale road traffic datasets unveiled an inadequate presence of connected components, often unstable and hard to navigate [19]–[21]. These works are in fact the tip of a fairly large literature on the connectivity analysis of vehicular networks, including studies on analytical models [19], [22]–[27], synthetic mobility traces [20], [21], [27]–[30] and real-world vehicle movement datasets [31], in either highway [22], [26], [27], [29], [30] or urban [19]–[21], [23]–[25], [27], [28], [31] environments.

Despite the significant efforts made in studying the instantaneous topology of vehicular networks, there is no equivalent investigation on the connectivity of store-carry-and-forward networks. Some works have considered pairwise temporal metrics, such as the duration of links between vehicles [19], [28], or the duration of the interval during which two vehicles remain (dis)connected in a multi-hop fashion [19]. However, these metrics do not provide a global view of the overall network connectivity in presence of store-carry-and-forward transmissions. The only work to date to address that aspect is in [32], where preliminary results are drawn in a single non-validated mobility scenario, considering unit-disc propagation

and a low (i.e., 10%) penetration rate of communication-enabled vehicles.

In this paper, we leverage an evolving graph-theoretical approach to characterize the temporal connectivity of vehicular networks, and better understand its dynamics. Our results confirm the preliminary findings in [32], and extend them. Namely, we (i) consider validated road traffic scenarios, (ii) adopt a realistic radio-frequency signal propagation model, (iii) generalize to multiple urban and suburban areas in different cities, (iv) evaluate critical system parameters such as the impact of the traffic conditions and of the technology penetration rate.

Overall, we provide a first comprehensive analysis of the temporal connectivity of vehicular networks. The outcome of our study allows answering questions that are currently open, such as “can store-carry-and-forward achieve complete network connectivity?”, “how long does it take to do so?”, or “how does the number of nodes that can be reached from a source evolves during that time interval?”. The responses to these issues allow for a clear understanding of the temporal connectivity granted by very popular delay-tolerant network-ing techniques employed in vehicular environments.

II. EVOLVING GRAPH AND REACHABILITY GRAPH MODELS

In our experiments we will use a graph model, either called *reachability graph* [33] or *transitive closure* [34], to represent temporal connectivity. Constructing such a graph is done in two steps: firstly an *evolving graph* is derived from the trace; secondly, based on the existence of some temporal paths (journeys) in this evolving graph, the transitive closure is calculated.

A. From input trace to evolving graph

An evolving graph \mathcal{G} is a set of static graphs, or *snapshots*, $\mathcal{G} = \{G_1, G_2, \dots, G_\ell\}$ that will represent, in our case, the direct communication connection between pairs of vehicles at a given time. Every $G_i = (V_i, E_i) \in \mathcal{G}$ is a static graph with node (vehicle) set V_i and edge (communication links) set E_i . We will denote the time indexes at which these snapshots are taken from the mobility trace by t_1, t_2, \dots, t_ℓ , where the integer ℓ denotes the last time step of the vehicular mobility under study. We implicitly consider that G_i is the snapshot of \mathcal{G} at time t_i . At time t_i , every two vehicles a and b such that $\{a, b\} \in E_i$ are considered to be within communication range in the mobility trace at time t_i . The approach we adopt in order to determine whether two given vehicles are in communication range is defined in Section III-C.

B. Evolving graphs and journeys

A journey is a set of hops (from node to node) in time. It is important to define the journey model as there are, in the literature, distinctive classes of journeys. They distinguish *strict* and *non-strict*, as well as *direct* and *indirect* journeys. Informally, a journey is said to be strict if only one edge can be traversed between two consecutive time steps, otherwise the journey is said to be non-strict. Also, a journey is said to be

indirect if it is allowed to stay put at some node over time; if instead the journey has to perform one hop at every time step, it is called direct. In our study we will only consider strict and indirect journeys, therefore we will only define formally this type of journeys and omit the qualifying term ‘strict-indirect’:

Definition 1 (Itinerary and Journey). *For a given evolving graph $\mathcal{G} = (G_1, G_2, \dots, G_\ell)$ an itinerary is an ordered sequence of tuples where every tuple is made of a pair of nodes and a time stamp:*

$$I = \left((\{a = a_s, b_s\}, t_s), (\{a_{s+1}, b_{s+1}\}, t_{s+1}), \dots, (\{a_f, b_f = b\}, t_f) \right), \text{ such that } \forall i \in [s, f], b_i = a_{i+1}$$

If the itinerary I is feasible in \mathcal{G} , then it is referred as a journey, i.e., if $1 \leq s \leq f \leq \ell$ and $\forall i \in [s, f], \{a_i, b_i\} \in E_i \cup \{\{u, u\} \mid u \in V\}$.

C. Reachability graph (Transitive closure)

In our case the evolving graph are extracted from various mobility traces with the sole purpose of computing the corresponding transitive closure, which is why the transitive closure will only be considered from time t_1 (i.e., $s = 1$). For any evolving graph $\mathcal{G} = \{G_i = (V_i, E_i) \mid i \in [1, \ell]\}$ and any time-interval $[t_1, t_k]$ with $k \leq \ell$, we define $\mathcal{C}_k = (\mathcal{V}_k, \mathcal{J}_k)$ to be the *transitive closure* of \mathcal{G} at time k . Note that when referring to the closure of the whole graph sequence \mathcal{G} we will omit the subscript and simply write \mathcal{C} ($= \mathcal{C}_\ell$). In the literature, \mathcal{C} is also called *reachability graph* of \mathcal{G} . The set of nodes \mathcal{V}_k is a subset of $\{V_i \mid i \in [1, \ell]\}$ and represents the set of vehicles that were sampled in the given time-interval. The set of arcs $\mathcal{J}_k \subseteq \mathcal{V}_k \times \mathcal{V}_k$ represents every available journey in \mathcal{G} within t_k . In other words, if $(u, v) \in \mathcal{J}_k$, then it exists a temporal path from u to v in \mathcal{G} such that the arrival time is smaller than k . Which intuitively means that vehicle u is able to communicate with vehicle v in the trace represented by the evolving graph $\{G_i \mid i \in [1, k]\}$. Notice that the relation “can communicate with” is not symmetric: considering three vehicles A, B and C, if A meets B and B later meets C then A can send a message to C via B but not the other way around.

We compute the transitive closure using an algorithm inspired by [34], which is the temporal equivalent of a Breadth-First Search. At every time step t_i , the algorithm adds the new available journeys to the closure $\mathcal{C}_{i-1} = (\mathcal{V}_{i-1}, \mathcal{J}_{i-1})$. Thus, at time i , for any given node $u \in \mathcal{V}_{i-1}$ a new journey from u to another node v exists in \mathcal{C}_i either if it exists some node $w \in V_i$ such that $(u, w) \in \mathcal{J}_{i-1}$ and $\{w, v\} \in E_i$, or if $\{u, v\} \in E_i$. The closure of \mathcal{G} over the time interval $[1, k]$ is computed recursively using this aggregation of journeys. The algorithm we use to compute the closure from an input evolving graph \mathcal{G} is provided in Listing 1.

III. REFERENCE SCENARIOS

A. Road traffic

We consider different road traffic scenarios in our analysis of temporal connectivity in order to generalize our findings as much as possible. As we aim at assessing the performance of

```

1 def compute_closure(dynamic_graph G):
2   let C be stable graph
3   for every snapshot G_i in G:
4     closure_increment(C, G_i)
5   return C
6
7 def closure_increment(graph C, dynamic_graph G_i):
8   add every new node from G_i to C
9   for every node u in C:
10    # direct journeys
11    let neighbors of u in G_i be neighbors of u in C
12    # long journeys w. several hops in time
13    for every node v in neighborhood of u:
14      let neighbors of v in G_i be neighbors of u in C

```

Listing 1: Pseudocode for the incremental construction of the closure C given an input evolving graph \mathcal{G} .



Fig. 1. Bologna Ringway scenario. Road layout, darker colors denote more trafficked road segments.

store-carry-and-forward in presence of significant technology penetration rates, we need a complete description of the vehicular traffic. This rules out real-world mobility traces of, e.g., taxis [31], since they only provide a partial view of the overall circulating traffic.

Instead, we rely on synthetic mobility scenarios. Since these are issued from simulations, and realism becomes a major question, we only consider properly validated road traffic traces, presented next.

1) *Bologna Ringway dataset*: The first scenario covers a 20 km² area in Bologna, Italy, comprising downtown, the arterial ringway around it, as well as part of the periphery. The dataset describes road traffic during the morning, in various traffic conditions, ranging from the rush hour to mild traffic.

The mobility in the region was generated using OpenStreetMap (OSM) road infrastructure data, the well-known SUMO simulator [35], and a realistic origin-destination (O-D) matrix of road traffic flows compiled within the collaborative research project iTetris [36]. A representative map of the area, including road traffic intensity levels on each road segment, is depicted in Fig. 1.

The synthetic mobility was validated through an original methodology that leverages publicly accessible data provided by free navigation services and concerning routing and travel time estimation functions. Full details on the dataset, its generation process, and the validation procedure are available in [37]. In the following we will refer to this scenario as *bo-ringway*.

2) *Cologne dataset*: The second scenario describes road traffic in the conurbation of Cologne, Germany. It encompasses a wide region of 400 km², and includes various

traffic conditions, from sparse overnight vehicular mobility to heavy congestion during peak traffic hours. The dataset was generated the same way as the Bologna one. Road layout information was obtained from OSM, and the microscopic behavior of individual drivers was modeled through SUMO. This time, the O-D matrix was derived using the TAPAS methodology on census data and large-scale surveys of the local population [38]. The mobility description was validated against live traffic service data, as well as against real-world road traffic counts. We refer the interested reader to [39], [40] for a detailed description of the dataset and of its validation.

We extract from the Cologne dataset six scenarios, each encompassing a 25-km² region, and thus comparable to the Bologna one. Areas have distinctive features, and range from downtown to suburban and per-urban regions. The six scenarios are depicted in Fig. 2. They will be referred to as *co-A* to *co-F* in the remainder of the paper.

B. Dissemination procedure

Datasets connectivity properties are evaluated under different system parameters. The first one is the road traffic density. The *bo-ringway* dataset is hence also scaled to 70, 80, 90 and 100% of the original feed to mimic diverse traffic levels [37]. In *co-[A-F]*, the traffic intensity in the different areas of a the city are already provided at different times of the day mapping to heterogeneous traffic densities.

The second parameter is the percentage of equipped vehicles, or penetration rate. In real-world applications, it is very unlikely that all vehicles deployed in the road network are equipped with a wireless interface, especially during the early adoption phase. Therefore, we study the dataset considering different penetration rates p , ranging from 1% to 100%.

C. Signal propagation

A correct modeling of the radio-frequency signal propagation is critical to the dependability of vehicular network simulations. Indeed, vehicle-to-vehicle communication channels feature fairly specific properties that need dedicated modeling [41].

In our evaluation, we resort to one of the state-of-the-art models proposed in the literature, and complement it with packet error rate information, as detailed next.

1) *Received power*: The radio-frequency signal propagation model we consider is based on that proposed in [42]. The model includes obstacle information, has been experimentally validated, yields a good trade-off between complexity and accuracy, and is implemented in popular tools for the simulation of vehicular networks, such as Veins [43].

The model works as follows. For any pair of vehicles (a, b) we define \mathcal{P} to be the shortest euclidean path between a and b in the input mobility trace. Let $d(a, b)$, $D_b(a, b)$ and $N_b(a, b)$ be the length of \mathcal{P} (in meters), the total distance within \mathcal{P} that crosses building on the city-map (also in meters), and the number of buildings crossed by \mathcal{P} , respectively. The two vehicles a, b are in communication range if and only if the

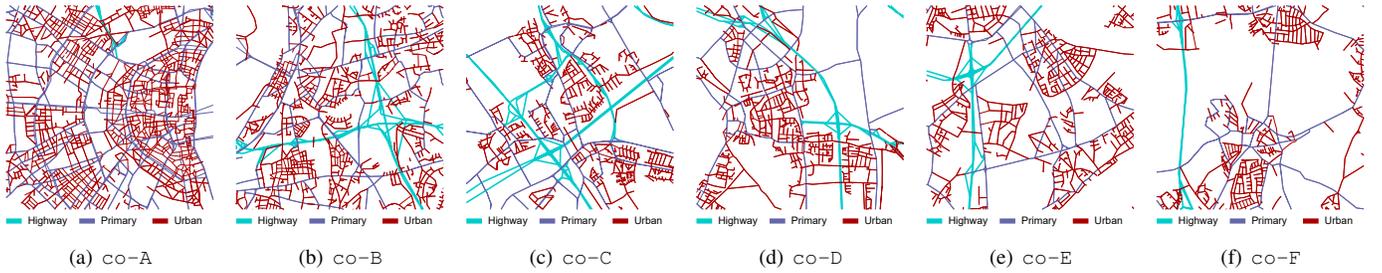


Fig. 2. Cologne scenarios areas: (A) city downtown, (B) industrial/transit, (C) suburban highways, (D) residential outskirts, (E) peri-urban and (F) rural highways. Different colors denote highway, primary and urban roads, respectively.

following inequation is satisfied:

$$P_r(a, b) = P_t - 10 \log_{10} \left(\frac{f}{d(a, b)^\alpha} \right) - \beta \cdot N_b(a, b) - \gamma \cdot D_b(a, b) > \text{RSS}_{min}. \quad (1)$$

The left-hand side in (1) represents the received power P_r for transmissions between the two vehicles, when the source emits at power P_t . The right-hand side is the received power threshold RSS_{min} above which the communication can occur. The parameters are set according to [42]: $P_t = 20$ dBm, $\alpha = 2.2$, $\beta = 9.2$ dBm, $\gamma = 0.32$ dBm, $\text{RSS}_{min} = -90$ dB, and $f = \frac{G_t \cdot G_r \cdot W^2}{16\pi^2}$ with $G_t = 7$, $G_r = 7$, and $W = 0.050812$.

2) *Packet loss ratio*: The model in (1) provides a simple binary connectivity decision. Network simulators typically use the actual received power information to determine the bit error rate (BER); from that, the erroneous bits can be computed on each individual packet, and the latter can be correctly received or discarded. As we rely on an evolving graph-theoretical approach rather than on packet-level simulation, we need to integrate the model above with a direct representation of the packet loss probability.

To that end, we replace the fixed threshold RSS_{min} with a probability function accounting for the fact that the weaker the received signal is, the more likely it is to have packet losses. In order to tune the function, we rely on experimental results on vehicle-to-vehicle communication that analyze the relation between the received power and the packet loss ratio [44]. Specifically, the packet loss probability $\text{loss}(P_r)$ for a received message with signal strength P_r (in dBm) is:

$$\text{loss}(P_r) = \begin{cases} 0, & \text{if } P_r > P_r^{max} \\ \min \left\{ 1, \left(\frac{P_r^{max} - P_r}{P_r^{max} - P_r^{min}} \right)^\tau \right\}, & \text{otherwise} \end{cases} \quad (2)$$

with $P_r^{max} = -78$ dBm, $P_r^{min} = -91$ dBm, and $\tau = 3.6$.

Applying (2) to our evolving graph model, we obtain a probabilistic graph where, at any moment in time, two vehicles that establish a communication link with received power P_r are connected with probability $1 - \text{loss}(P_r)$.

Fig. 3 outlines some statistical properties of the links resulting from the propagation model described above. The scatterplot clearly shows the logarithmic relationship that ties the distance and received power in line-of-sight conditions (apparent solid line). The presence of buildings introduces a negative offset in the received power, which depends on the number of interposing obstacles (clouds of points towards

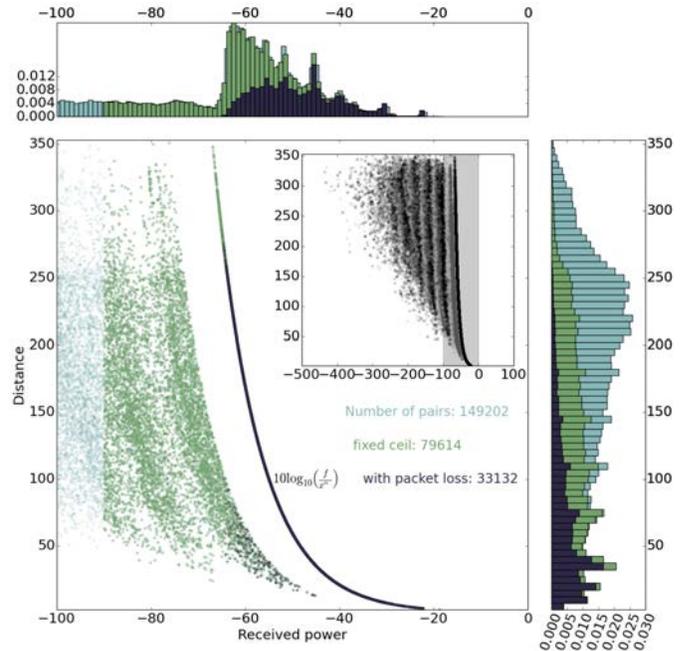


Fig. 3. Signal propagation model. The plot portrays, for every pair of vehicles $a, b \in \mathcal{V}^2$ in the bo-ringway road traffic scenario, the distance $d(a, b)$ of the vehicles as a function of received power $P_r(a, b)$. Plots on the top and right represent probability density distributions on each axis. Colors tell apart all links (light-blue), links that satisfy (1) (green), and links in the probabilistic graph obtained from (2) (dark-blue). Figure best viewed in colors.

more negative received powers, more evident in the zoomed-out inset plot). When applying the binary threshold in (1), only links whose signal strength is above -91 dBm are retained (green dots). However, looking at the distributions at the top and right of the main plot, this retains numerous links with low signal strength and large distance. By introducing the probabilistic model of packet loss ratio, we limit the number of links to realistic ones (dark-blue dots): we stress that the resulting distances, typically around 50-100 meters but reaching up to 250 m, are consistent with those of experimental assessments of direct vehicular communication in urban environments [44].

IV. SIMULATION SETUP AND RESULTS

Let us first study the bo-ringway scenario with full traffic density, all vehicles equipped with radio interfaces and considering the strict-indirect journey model. The first metric we observe is the (*temporal*) *reachability ratio*:

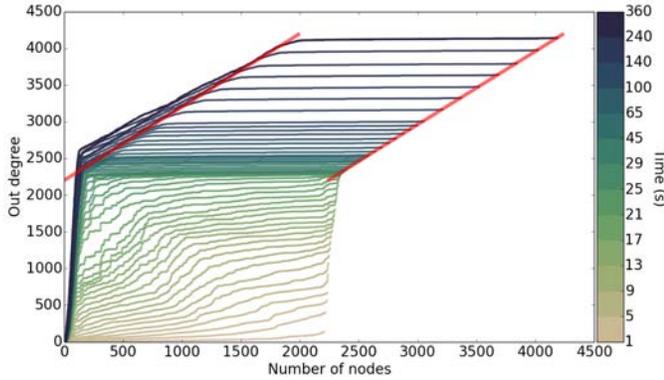


Fig. 4. `bo-ringway` scenario. Every plot represents the outgoing degree distribution of the closure at time t , i.e., \mathcal{C}_t . The x-axis is the set of node ordered by degree and the y-axis corresponds to the degree of these nodes.

Definition 2 (Reachability ratio). *The expected ratio of vehicles that a given vehicle can communicate with after k units of time. Formally, for an evolving graph $\mathcal{G} = \{G_i \mid i \in [1, \ell]\}$, the reachability ratio at time $k \leq \ell$ is $\frac{|\mathcal{J}_k|}{|\mathcal{V}_k|(|\mathcal{V}_k|-1)}$. Note that this ratio is exactly equal to the density of the transitive closure at time k .*

We will show that this ratio never reaches one. This is due to the fact that some vehicles enter the observed area after others have left it, and it is true even in the case of 100% penetration rate. The effect is observed in Fig. 4, where a clear cut in the out-degree distribution appears. This cut starts around the 30-s mark and has approximately the same slope as the total number of vehicles (parallel red lines help to visualize this). Remark that if a vehicle leaves the observed area at time t_l , all the vehicles entering at time $t > t_l$ will never be able to reach it. Therefore, for any given vehicle, its outgoing degree is upper bounded by the total number of vehicles that are or will be inside the observed area. A similar effect appears in the distribution of ingoing degrees (not shown for the sake of brevity). A further confirmation is in Fig. 5, where the closure average degree never reaches the total number of vehicles¹, a gap separating the two curves even after long observation intervals. These observations are an artifact induced by the geographical boundaries of the considered scenario, rather than an intrinsic property of vehicular networks connectivity. Nevertheless, in some application cases, observing geographically delimited areas may be pertinent: e.g., in the case of location-based services where contents are meaningful only within well-defined spatial regions.

Fig. 5 yields additional information, as it shows, for every time step i of the input evolving graph: (i) the total number of vehicles observed, $|\mathcal{V}_i|$; (ii) the number of vehicles present in the observed area, $|\mathcal{V}_i|$; (iii) the size of the largest connected component of \mathcal{G}_i ; (iv) the average degree of nodes in the closure, i.e., $2 \cdot |\mathcal{J}_i|/|\mathcal{V}_i|$. This figure shows a clear step transition around the 30-s mark when the closure graph has converged to its maximum reachability ratio. Afterwards, the

¹If the total number of vehicles is equal to the average degree in the closure, then, by definition, the reachability ratio is equal to one.

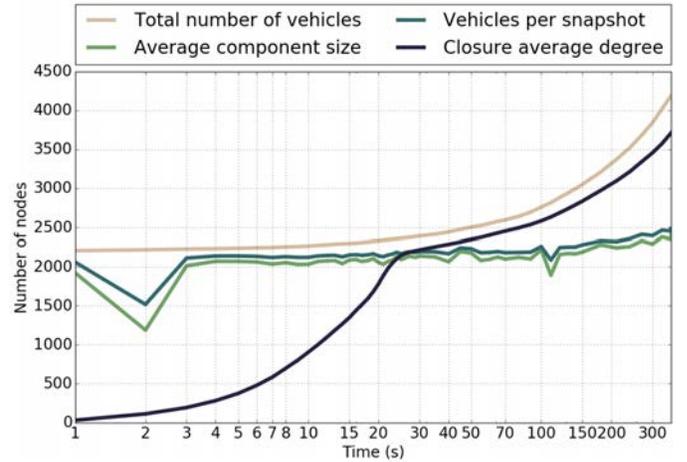


Fig. 5. `bo-ringway` scenario. Evolution of the closure connectivity function of time together with the evolving graph connectivity.

total number of vehicles provides the upper bound on the nodes average degree. In other words, a representative majority of the new arcs in the transitive closure from this particular moment in time come from (or go to) new vehicles (i.e., vehicles that appeared after the 30-s mark) and a large amount of vehicles that enter the simulation are involved in the creation of new arcs. This observation could be also derived from the slope of the cut in Fig. 5, and is partly due to the fact that vehicles moving outside the observed area can still increase their outgoing degree in the closure as long as they are connected (within the closure) to vehicles that are still inside the area. In practice, our results show that if a message is to be sent to all vehicles entering an area, after a certain amount of time (in this particular experiment, 30 s) the majority of new arriving vehicles receive it with a very low delay.

a) *Impact of vehicle traffic density:* The variation of traffic density has a marginal impact on the reachability ratio, as shown in Fig. 6. However, reachability is dependent on the graph connectivity, and the `bo-ringway` scenario shows a unusually good vehicular network connectivity. To show this effect, we measure the *expected instantaneous connectivity*:

Definition 3 (Expected instantaneous connectivity). *This connectivity metric is equal to the expected component size of a random node at a random time step in the evolving graph.*

We also use a normalized version of this metric (i.e., divided by the number of nodes in the snapshot of the selected node). The latter is equal to the proportion of vehicles that a randomly picked vehicle could send a message to, when no store-carry-and-forward is allowed and only connected forwarding is employed. For vehicular traffic densities of 100%, 90%, 80% and 70% in the `bo-ringway` scenario, the normalized expected instantaneous connectivity score very high values at 0.925, 0.869, 0.852 and 0.584, respectively: in other words a node can reach almost 60% of the vehicles through connected forwarding, even in the sparsest road traffic conditions. In this excellent instantaneous connectivity situation, the convergence time to a maximally connected closure is driven by

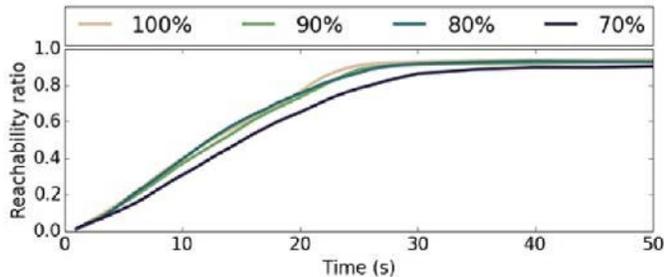


Fig. 6. bo-ringway scenario. Evolution of density (reachability ratio) function of time for different traffic densities.

TABLE I

bo-ringway SCENARIO. EVOLVING GRAPH MEASUREMENTS VERSUS PENETRATION RATES, WITH 100% ROAD TRAFFIC DENSITY.

Penetration rate (%)	5	10	20	30	50	80	100
Expected instant connectivity (%)	6.7	10.0	14.2	21.7	45.5	81.5	93.0
Largest component diameter	5.6	9.7	13.0	19.1	32.2	44.1	42.3

the diameter of the snapshots². Our experiments for traffic ratios of 100%, 90%, 80% and 70% in the bo-ringway scenario show that the average diameter of the snapshots are respectively 42.3, 43.6, 47.4 and 41.4 whereas the convergence times are respectively around 25, 28, 30 and 37 seconds.

To sum up, under all traffic conditions in the bo-ringway scenario, a randomly picked vehicle is able to send a message to at least 91% of the vehicles, as shown in Fig. 6. This occurs in at most 37 communication hops (and less than 37 seconds), thanks to the excellent instantaneous connectivity.

b) *Impact of the penetration rate:* The previous results assume a 100% penetration rate. We now investigate the impact of lower penetration rates, leading to sparser evolving graphs. Tab.I shows how the penetration rate influences the instantaneous connectivity, whereas Fig. 7 and Fig. 8 show the impact on the convergence to a maximally connected closure. The reduced network density best highlights the power of store-carry-and-forward. In very sparse cases where only 10% of vehicles are communication-enabled, the instantaneous connectivity is limited to 10%, yet the reachability ratio is at almost 70%, with a convergence time of around four minutes. Nevertheless, the performance degradation seems to scale in a superlinear way, showing the importance of the penetration rate. Specifically, Fig. 7 and Fig. 8 also show two phases: the first, up to a 40% penetration rate, characterized by a linear decrease of the convergence time as the fraction of participating vehicles grows; the second, from a 40% penetration rate onwards, where the convergence is always fast and efficient.

However, even for small penetration rates such as 10%, store-carry-and-forward mechanisms allow, for example, to increase the reachability ratio from 10% (in the connected forwarding case, see Tab.I) to 20% after 60 seconds and 60% after five minutes (see Fig. 7). After two minutes and 30 seconds the number of vehicles that are reached by a message

²This is due to the slowly changing nature of connectivity graphs: in most cases, given two consecutive snapshots, the distance (number of hops in the graph) between two vehicles is often the same. Instead, the convergence time is less than the evolving graph size if every snapshot is composed of a single connected component.

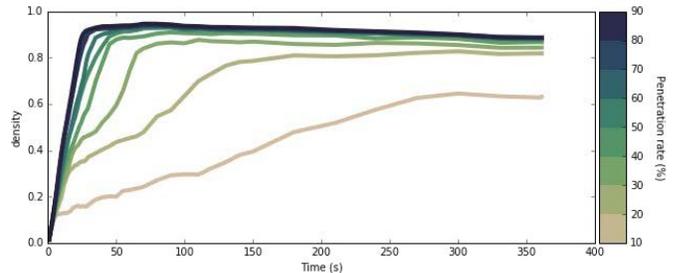


Fig. 7. bo-ringway scenario. Temporal evolution of the average reachability ratio for different penetration rates. This allows appreciating the convergence time to maximal density.

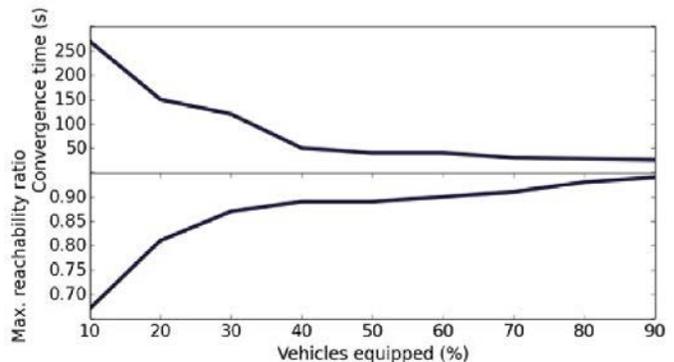


Fig. 8. bo-ringway scenario. Convergence time to maximal reachability ratio for different penetration rates.

is four times bigger than the number of vehicles reachable without using store-carry-and-forward.

c) *Generalization to other urban areas:* Tab. II shows how the vehicular network connectivity varies under the heterogeneous road traffic scenarios provided by the co-[A-F] datasets. The table reports results for different daytimes (top part) and urban regions (bottom part). It concerns both the instantaneous connectivity (in terms of normalized expected connectivity and largest component diameter) and the transitive closure (in terms of the average degree of vehicles in the evolving graph). Results are ordered by closure average degree. The table let us appreciate the very strong variability that the daytime and urban region induce in both instantaneous and temporal connectivity: e.g., the expected connectivity can range from little more than 5% to almost 50%, and the closure degree from a hundred to more than a thousand neighbors.

A complete summary of temporal connectivity, for all combinations of daytime and urban region, is provided in Fig. 9 and Fig. 10. There, we observe how higher closure average degrees map to transitive closures that converge faster and to higher reachability ratios. In typical cases, the gain brought by store-carry-and-forward is strongly correlated to the average instantaneous connectivity itself. The gain is higher if the connectivity of the evolving graph is low, as shown in Fig. 11. In other words, it is in sparse vehicular networks that store-carry-and-forward shows its superiority over connected forwarding: for example, in co-D at 8am, the average connected component size is 8.5%, after 30 s the temporal reachability ratio is 27% and goes up to 46% after one minute. Yet, some cases may not obey this tendency: e.g.,

TABLE II
EVOLVING GRAPH MEASUREMENTS FOR VARIOUS PERIODS OF THE DAY
AND THEN FOR DIFFERENT AREAS OF COLOGNE.

Day time	7am	5pm	8am	1pm	11am	9pm
Expected connectivity (%)	33.1	26.6	27.4	13.4	5.8	5.3
Largest component diameter	30.0	24.3	19.0	15.7	9.6	6.8
Closure average degree	880.7	726.2	539.5	376.8	254.4	198.2
Zone identifier	A	B	D	C	F	E
Expected connectivity (%)	49.0	12.6	9.5	7.0	18.0	15.7
Largest component diameter	36.0	22.3	13.7	13.7	15.9	3.6
Closure average degree	1327.6	708.0	307.4	282.5	237.6	112.6

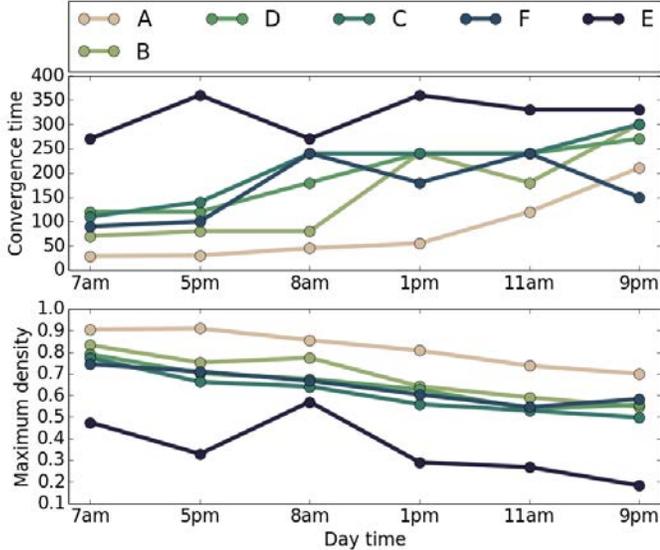


Fig. 9. Temporal connectivity with different urban regions and daytimes.

this is the case for zone E at 8pm. However, this is explained by an unusually high instantaneous connectivity at 51.7%, leading to an average degree of 227, or 1.9% of the average 12205 vehicles per snapshot.

At this point, we would like to unveil the physical reason behind the variability induced by scenario settings such as the daytime or urban region, so as to derive a unifying theory of the temporal connectivity of vehicular networks. To that end, we explore the relationship between the transitive closure and the road traffic density, which has already been demonstrated to play a major role in the instantaneous connectivity of urban vehicular networks [20]. In Fig. 12, we observe a fairly neat correlation of the convergence time and maximum reachability ratio in the transitive closure, independently of the urban zone and daytime. Similar results were obtained for other cases as well, but are omitted for the sake of brevity. Moreover, the resulting trends are very similar to the ones obtained for the bo-ringway dataset under different penetration ratios (i.e., road traffic density), in Fig. 8, which further confirms that the trend has general validity. The average Pearson’s correlation coefficients over all samples are very high, at 0.88.

These results unveil how the temporal connectivity (and thus the performance of store-carry-and-forward) are driven solely by the number of vehicles traveling in an area, and is instead only marginally affected by, e.g., the road layout or the travel

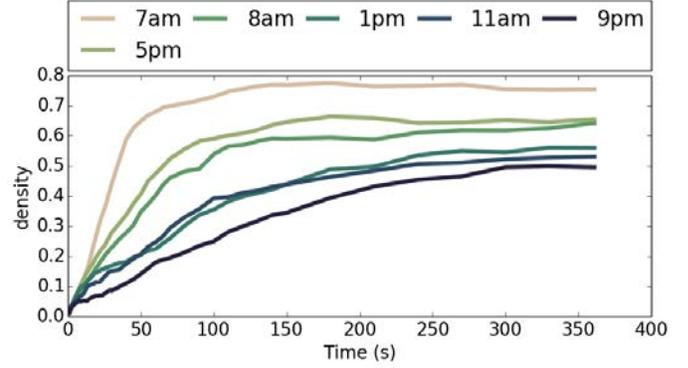


Fig. 10. Reachability ratio in c-o-C for different periods of the day.

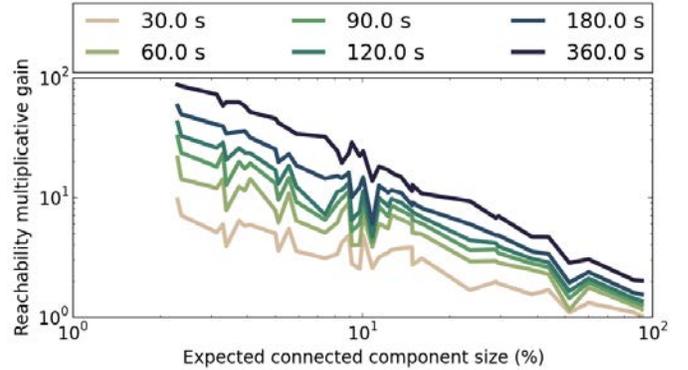


Fig. 11. Reachability ratio versus the expected instantaneous connectivity.

demand flows. Moreover, Fig. 12 proves how the dependence is sub-linear. Thus, high store-carry-and-forward performance gains are expected when increasing vehicles participation at low penetration rate regimes or in presence of sparse traffic; less so, once a critical density of communication-enabled vehicles is achieved. We remark that the critical density is very low in all our tests, at around 8 veh/km².

V. CONCLUSION AND FUTURE WORKS

We showed that, even though it cannot not allow full network connectivity, store-carry-and-forward improves significantly the instantaneous connectivity attained through simple connected forwarding. In dense networks, the convergence to a maximal connectivity is fast (a few tens of seconds in a region of a few km²). In sparse networks, store-carry-and-forward increases the reachability (by a factor up to 90 in our experiments) with convergence times of the order of few minutes. Moreover, once connectivity has converged to its maximum, new arriving vehicles in the area get new messages very quickly: a useful feature for local data dissemination.

Our work paves the way to additional research on the temporal connectivity of vehicular networks. We analyzed two different cities, Bologna and Cologne, however further scenarios may let additional behaviors emerge. Also, other store-carry-and-forward models can be taken into account, e.g., hybrid models in between strict and non-strict, or direct and indirect, journeys; this would allow modelling cases where a vehicle is allowed to multi-hop information at each snapshot, or to carry information for a limited amount of time.

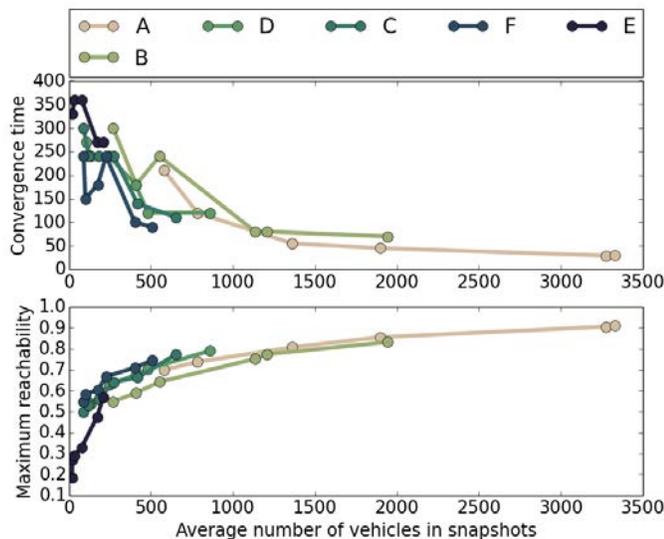


Fig. 12. Temporal connectivity versus the average road traffic density.

REFERENCES

- [1] TomTom, "Travel Time Measurements using GSM and GPS Probe Data", *White Paper*, 2010.
- [2] Uconnect Systems, <http://www.driveuconnect.com/>.
- [3] Octo Telematics, <http://www.octotelematics.com/en/>.
- [4] IEEE, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", *IEEE 802.11-2012*, 2012.
- [5] OSI, "Communications access for land mobiles (ITS-CALM-M5)", *OSI 21215 Standard*.
- [6] ETSI, "Intelligent Transportation Systems (ITS)", *EN 302 665*, ETSI Standard.
- [7] Reuters, "Obama Backs Highway Fund Fix, Touts 'Talking' Cars", *The New York Times*, 2014.
- [8] ETSI, "Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Definitions," *TR 102 638 V1.1.1*, ETSI Technical Report.
- [9] H. Hartenstein, K.P. Laberteaux, "A tutorial survey on vehicular ad hoc networks", *IEEE Communications Magazine*, 46(6):164–171, 2008.
- [10] G. Karagiannis, O. Altintas, E. Ekici, G. Heijnen, B. Jarupan, K. Lin, T. Weil, "Vehicular networking: a survey and tutorial on requirements, architectures, challenges, standards and solutions", *IEEE Communications Surveys and Tutorials*, 13(4):584–616, 2011.
- [11] X. Wu, S. Subramanian, R. Guha, R.G. White, J. Li, K.W. Lu, A. Bucceri, T. Zhang, "Vehicular communications using DSRC: challenges, enhancements, and evolution", *IEEE Journal on Selected Areas in Communications*, 31(9):399–408, 2013.
- [12] J. Burgess, B. Gallagher, D. Jensen, B.N. Levine, "MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks", *IEEE INFOCOM*, 2006.
- [13] H.-Y. Huang, P.-E. Luo, M. Li, D. Li, X. Li, W. Shu, M.-Y. Wu, "Performance Evaluation of SUVnet With Real-Time Traffic Data", *IEEE Transactions on Vehicular Technology*, 56(6):3381–3396, 2007.
- [14] R. Jiang, Y. Zhu, X. Wang, L.M. Ni, "TMC: Exploiting Trajectories for Multicast in Sparse Vehicular Networks", *IEEE Transactions on Parallel and Distributed Systems*, 26(1):262–271, 2015.
- [15] H. Wu, R. Fujimoto, R. Guensler, M. Hunter, "MDDV: a mobility-centric data dissemination algorithm for vehicular networks", *ACM VANET*, 2004.
- [16] F. Malandrino, C. Casetti, C.-F. Chiasserini, M. Fiore, "Optimal Content Downloading in Vehicular Networks", *IEEE Transactions on Mobile Computing*, 12(7):1377–1391, 2013.
- [17] J. Whitbeck, Y. Lopez, J. Leguay, V. Conan, M. Dias De Amorim, "Push-and-track: Saving infrastructure bandwidth through opportunistic forwarding", *Pervasive and Mobile Computing*, 8(5):682–697, 2012.
- [18] R. Stanica, M. Fiore, F. Malandrino, "Offloading Floating Car Data", *IEEE WoWMoM*, 2013.
- [19] W. Viriyasitavat, F. Bai, O.K. Tonguz, "Dynamics of Network Connectivity in Urban Vehicular Networks", *IEEE Journal on Selected Areas in Communications*, 29(3):515–533, 2011.
- [20] D. Naboulsi, M. Fiore, "On the Instantaneous Topology of a Large-scale Urban Vehicular Network: the Cologne case", *ACM MobiHoc*, 2013.
- [21] A. Gryzbek, G. Danoy, M. Seredynski, P. Bouvry, "Evaluation of Dynamic Communities in Large-Scale Vehicular Networks," *ACM DI-VANet*, 2013.
- [22] M. Khabazian, M.K. Ali, "A Performance Modeling of Connectivity in Vehicular Ad Hoc Networks", *IEEE Transactions on Vehicular Technology*, 57(4):2440–2450, 2008.
- [23] M. Kafsi, P. Papadimitratos, O. Dousse, T. Alpcan, J.-P. Hubaux, "VANET Connectivity Analysis", *IEEE AutoNet*, 2008.
- [24] M. Fiore, J. Härri "The Networking Shape of Vehicular Mobility", *ACM MobiHoc*, 2008.
- [25] G.H. Mohimani, F. Ashtiani, A. Javanmard, M. Hamdi, "Mobility Modeling, Spatial Traffic Distribution, and Probability of Connectivity for Sparse and Dense Vehicular Ad Hoc Networks", *IEEE Transactions on Vehicular Technology*, 58(4):1998–2006, 2009.
- [26] S. Durrani, X. Zhou, A. Chandra, "Effect of Vehicle Mobility on Connectivity of Vehicular Ad Hoc Networks", *IEEE VTC-Fall*, Ottawa, Canada, Sep. 2010.
- [27] R. Monteiro, S. Sargento, W. Viriyasitavat, O.K. Tonguz, "Improving VANET protocols via network science," *IEEE VNC*, 2012.
- [28] G. Pallis, D. Katsaros, M. D. Dikaiakos, N. Loulloudes, L. Tassioulas, "On the Structure and Evolution of Vehicular Networks", *IEEE/ACM MASCOTS*, 2009.
- [29] M. Gramaglia, O. Trullols-Cruces, D. Naboulsi, M. Fiore, M. Calderon, "Vehicular Networks on Two Madrid Highways", *IEEE SECON*, 2014.
- [30] 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*, 64(1):248–262, 2015.
- [31] M.A. Hoque, X. Hong, B. Dixon, "Efficient multi-hop connectivity analysis in urban vehicular networks", *Vehicular Communications* 1, 78–90, 2014.
- [32] H. Conceição, M. Ferreira, J. Barros, "On the Urban Connectivity of Vehicular Sensor Networks", *DCOSS*, 2008.
- [33] J. Whitbeck, M. Dias de Amorim, V. Conan, J.-L. Guillaume, "Temporal reachability graphs," *ACM MobiCom*, 2012.
- [34] M. Barjon, A. Casteigts, S. Chaumette, C. Johnen, Y.M. Neggaz, "Testing temporal connectivity in sparse dynamic graphs," arXiv:1404.7634 [cs.DS], 2014.
- [35] 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*, 5(3):128–138, 2012.
- [36] M. Rondinone *et al.*, "ITETRIS: a modular simulation platform for the large scale evaluation of cooperative ITS applications", *Simulation Modelling Practice and Theory*, 34:99–125, 2013.
- [37] L. Bedogni, M. Gramaglia, A. Vesco, M. Fiore, J. Hrri, F. Ferrero, "The Bologna Ringway dataset: improving road network conversion in SUMO and validating urban mobility via navigation services", *IEEE Transactions on Vehicular Technology*, to appear.
- [38] G. Hertkorn, P. Wagner, "The application of microscopic activity based travel demand modeling in large scale simulations," *World Conference on Transport Research*, 2004.
- [39] 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*, 13(5), 2014.
- [40] S. Uppoor, M. Fiore, "Characterizing pervasive vehicular access to the cellular RAN infrastructure: an urban case study," *IEEE Transactions on Vehicular Technology*, 64(6), 2015.
- [41] W. Viriyasitavat, M. Boban, H.-M. Tsai, A. Vasilakos, "Vehicular Communications: Survey and Challenges of Channel and Propagation Models," *IEEE Vehicular Technology Magazine*, 10(2):55–66, 2015.
- [42] C. Sommer, D. Eckhoff, R. German, F. Dressler, "A Computationally Inexpensive Empirical Model of IEEE 802.11p Radio Shadowing in Urban Environments," *IEEE/IFIP WONS*, 2011.
- [43] C. Sommer, R. German, F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, 10(1):3–15, 2011.
- [44] P. Alexander, D. Haley, A. Grant, "Cooperative intelligent transport systems: 5.9-ghz field trials," *Proceedings of the IEEE*, 99(7):1213–1235, 2011.