Technical Report
(GTA/PEE/COPPE – UFRJ)

Impact of Relative Speed on Node Vicinity Dynamics in VANETs

Dianne S. V. Medeiros · Dayro A. B. Hernandez · Miguel Elias M. Campista · Aloysio de Castro P. Pedroza

August 2017

Abstract Communication protocols generally rely on the existence of very long multihop paths to reach distant nodes. They disregard, however, how often such paths indeed occur, and how long they persist, especially in highly dynamic mobile networks. In this direction, this paper evaluates quantitatively the influence of node relative speed on path establishment and maintenance, using real and synthetic vehicular network traces. We propose a methodology for vehicular network analysis where both relative speeds and hop distances are used as parameters to characterize node vicinity. Results show that contact opportunities highly depend on the relative speed and the hop distance between nodes. In sparser scenarios, the number of contacts between nodes separated by more than 3 hops or even between neighbors with relative speed above 40 km/h is negligible. This confirms the intuition that contacts at lower relative speeds and at few hop distances happen more often. In addition, contacts last longer as the number of hops between nodes decreases. Nevertheless, we can still find multihop paths able to transmit messages at high relative speeds, even though less often. We also demonstrate that relative speeds reduce the number of useful contacts more severely when compared to the hop distance. For last, we show that it is possible to increase the number of successful packet transmissions by simply applying the outcomes of this work, without any sophisticated model, avoiding the waste of resources, such as energy and bandwidth.

Dianne S. V. Medeiros · Dayro A. B. Hernandez · Miguel Elias M. Campista · Aloysio de Castro P. Pedroza
GTA / PEE-COPPE / DEL-Poli
Universidade Federal do Rio de Janeiro
Rio de Janeiro / RJ – Brazil
E-mail: {dianne, dayro, miguel, aloysio}@gta.ufrj.br

Keywords Relative speeds · VANETs · Vicinity analysis.

1 Introduction

Mobile Ad Hoc Networks (MANETs) are still in spotlight even after almost two decades of extensive investigation. Characterization of mobility patterns, e.g., remains an open research issue especially in challenged networks, where contact information is not known a priori and there is no infrastructure to provide connectivity. A good example of such scenario are Vehicular Ad Hoc Networks (VANETs) [1], especially if they rely only on vehicle-to-vehicle (V2V) communication. VANETs are highly dynamic and the intense node mobility contributes to the intermittent connectivity and lack of end-to-end-paths, which can hinder communications during contact opportunities, increasing the difficulty to achieve efficient data transfer [2]. Several studies concerning mobility patterns and connectivity in MANETs already exist and many of them are compiled in several surveys [3–8]. These works are valuable to provide better insights to develop routing protocols to disseminate messages in challenged networks.

Regarding ordinary networks, routing protocols are designed considering that lack of end-to-end paths are transitory. In challenged networks, however, this is not a valid assumption, as the probability of not having a connected network is very high. Therefore, routing protocols must consider that end-to-end paths are rare and nodes need to decide at each encounter whether they will forward a message to the newly encountered node or carry it a little farther. To make this decision, nodes need to know how to explore contact opportunities in the best way they are able to. Some works analyze such

opportunities through the study of node vicinity [9–11]. Typically, they consider that two nodes are in *contact* if they are within mutual radio range, limiting node vicinity to directly reachable neighbors. This definition incurs frequent vicinity changes and a limited view of contact opportunities. As an alternative to circumvent this restriction, some works propose prediction mechanisms to anticipate contact availability [12], disruption [13], or even stabilization [14–16]. A totally different approach is proposed by Phe-neau et al. [17,18], when they extend the concept of *contact* to also include nodes reachable via multiple hops. In this case, the resultant extended vicinity incorporates nodes even if they are out of mutual radio range and, as a consequence, nodes find more contact opportunities. Hoque et al. [19] use this idea to develop an algorithm to analyze multihop connectivity and network partitioning in VANETs. Nevertheless, independent whether the extended vicinity is considered, we did not find any work that analyzes quantitatively the influence of relative speeds and the number of hops on the connectivity of very dynamic networks.

This paper aims at analyzing contact opportunities, taking into account the relationship between the extended vicinity of nodes and their relative speeds. The goal is to quantitatively evaluate the expected notion of "better connectivity at lower relative speeds" to shed more light into multihop communications in typical vehicular scenarios. To this end, we propose a methodology to group nodes according to their relative speeds, i.e., we consider links only between nodes at a certain interval of relative speeds. This definition permits identifying conditions for multihop communications, which in a broader sense, depend on whether opportunistic contacts appear for long enough to be considered useful.

In the case scenarios studied in this work, results considering different radio ranges show that a significant number of useful contacts can happen even between nodes at high relative speeds, separated by multihop distances. Even in such conditions, we show that nodes can transfer MB-size messages according to the contact duration. Besides more general results, we also observe that contacts with longer duration become less often for relative speeds higher than 40 km/h and most likely happen between nodes less than 3 hops away, in sparser scenarios. On the other hand, even considering lower relative speeds, results show that contacts between nodes separated by more than 6 hops are not frequent. We also note that high relative speeds can potentially degrade the number of useful contacts more severely than the hop distance. Finally, we highlight the importance of this work by proposing three different forwarding strategies, based on our results, and comparing them with the `OLSR`. We show that the best strategy is the one the simultaneously considers the reachability of nodes according to their relative speed and, further, restricts the construction of node vicinity using the relative speed with the neighbor nodes. This strategy can increase the average packet delivery ratio while reducing the waste of resources. We believe then that the vicinity of a node must also include relative speeds both in more theoretical evaluations and in practical settings [6] and also that the results obtained herein can be used as a step forward to more sophisticated message dissemination schemes in vehicular networks. In summary, our main contributions are:

- We identify the importance of quantifying the influence of relative speeds on the network connectivity.
- We propose an extended definition of node vicinity to include both nodes at multihop distances from an ego node and their relative speeds, instead of only focusing on the node adjacent vicinity.
- We apply our definition to three different scenarios in order to analyze the behavior of the node vicinity under distinct conditions, such as varied node density, without being attached to any mobility model.
- We demonstrate through simulations that we can potentially reduce network resource consumption, and even improve packet delivery ratio, using the correlation between relative speeds, hop distance, and contact duration to make forwarding decisions.

This paper is organized as follows. Section 2 presents node vicinity definitions used in our analysis. Section 3 introduces real and synthetic datasets used in this work and provides a preliminary evaluation of them. Section 4 discusses the problem tackled in this work. In Section 5 we propose a methodology for node vicinity analysis. We present the analysis results in Section 6. In Section 7 we apply our results to a network simulation and discuss some other possible applications. Finally, Section 8 concludes this work and discusses future directions.

2 Definitions

The analysis proposed herein requires the formalization of node vicinity. To this end, we generalize the definitions in Pheneau et al. [18] to include node relative speeds.

We model a vehicular network as a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of mobile vehicular nodes, and $\mathcal{E}$ is the set of wireless links connecting adjacent nodes in $\mathcal{V}$. Each node v_i moves at speed $\mathbf{s_i}$, where $\{|\mathbf{s_i}| \in \mathbb{R}^+ \mid s_{min} \leq |\mathbf{s_i}| < s_{max}\}$, and s_{min} and s_{max} are the minimum and the maximum absolute speeds allowed,

respectively. The relative speed of nodes v_i and v_j is, thus, given by $\mathbf{r_{ij}} = \mathbf{s_i} - \mathbf{s_j}$, where $|\mathbf{r_{ij}}| = |\mathbf{s_i} - \mathbf{s_j}|$ and $|\mathbf{r_{ij}}| \in [0, 2 \times s_{max}[$. For the sake of simplicity, we use the notation s_i and r_{ij} to represent, respectively, $|\mathbf{s_i}|$ and $|\mathbf{r_{ij}}|$, whenever possible.

Dividing the set of all relative speeds into m consecutive subsets, we have that $r_{ij} \in \bigcup_{\kappa=0}^{m-1} R_\kappa$, $\forall v_i, v_j \in \mathcal{V}$, where $R_\kappa = [\kappa \times s_\delta, (\kappa+1) \times s_\delta[$ and $s_\delta = \frac{2 \times s_{max}}{m}$. In this case, we can group all pairs of nodes v_i, v_j with $r_{ij} \in R_\kappa$ in a subset of nodes $\mathcal{V}(R_\kappa) \subseteq \mathcal{V}$. Consequently, we can obtain the subgraph $\mathcal{G}_{R_\kappa}(\mathcal{V}(R_\kappa), \mathcal{E}(R_\kappa)) \subseteq \mathcal{G}(\mathcal{V}, \mathcal{E})$, where $\mathcal{E}(R_\kappa)$ is the set of existing links connecting adjacent nodes in $\mathcal{G}_{R_\kappa}$ with $r_{ij} \in R_\kappa$. Thus, although two adjacent nodes v_i, v_j may be included in $\mathcal{V}(R_\kappa)$ due to their relative speed, a link connecting them will only exist in $\mathcal{E}(R_\kappa)$ if $r_{ij} \in R_\kappa$ and they are within mutual radio range. Otherwise, v_i, v_j may still be mutually reachable if they are interconnected by *a sequence of adjacent links* between pairs of nodes also in $\mathcal{V}(R_\kappa)$. Hence, according to our definition, if v_i, v_k, v_w, v_j are in $\mathcal{V}(R_\kappa)$, and if there is a link between v_i, v_k and another between v_w, v_j, then $\{r_{ik}, r_{wj}\} \in R_\kappa$. A path from v_i to v_j will exist in $\mathcal{G}_{R_\kappa}$ only if $r_{kw} \in R_\kappa$.

As a corollary, if $m = 1$, all relative speeds are within the same subset $R_0 = [0, s_\delta[= [0, 2 \times s_{max}[$. Analogously to relative speeds, we divide the set of absolute speeds into consecutive subsets, represented by S_κ.

Definition 1 (κ, ρ)**-vicinity**: *The* (κ, ρ)*-vicinity of a node* $v_i \in \mathcal{G}_{R_\kappa}$ *is the set of all nodes also in* $\mathcal{G}_{R_\kappa}$ *for which the shortest path from* v_i *is* ρ *hops at most.*

The vicinity of a node $v_i \in \mathcal{G}_{R_\kappa}$ can be characterized only by parameters κ and ρ, where ρ defines the maximum number of hops from v_i, while κ defines the range of relative speeds considered and, consequently, which subgraph must be used. Hence, nodes in $\mathcal{G}_{R_\kappa}$ may not belong to the same (κ, ρ)-vicinity of v_i, according to the number of hops (ρ) of the shortest path interconnecting them.

Figure 1a depicts a network of nodes moving at speeds within $[0, 45$ km/h[(arrows starting at nodes indicate their absolute speed). In this figure, nodes v_i, v_j, and v_k move, respectively, at absolute speeds s_i, s_j, and s_k, where $s_i \in [0, 15$ km/h[, $s_j \in [15, 30$ km/h[, and $s_k \in [30, 45$ km/h[. If $m = 1$, the relative speed between all pairs of nodes in the network lies within $R_0 = [0, 90$ km/h[. Thus, all nodes are within the $(0, \rho)$-vicinity of v_i. Figure 1b shows v_i's $(0,1)$- and $(0,2)$-vicinities, and the subgraph $\mathcal{G}_{R_0}$ obtained from the subset $\mathcal{V}(R_0)$. Note that with $m = 1$, $\mathcal{V}(R_0) = \mathcal{V}$. Hence, all the links connecting nodes in the network do exist and can be used for path computing. Consequently, $\mathcal{G}_{R_0} = \mathcal{G}$ and the definition of node v_i's vicinity does not change according to the different relative speeds, similarly to [18].

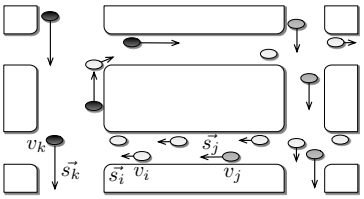

(a) Vehicles v_i, v_j, and v_k moving at speeds, s_i, s_j, and s_k, respectively.

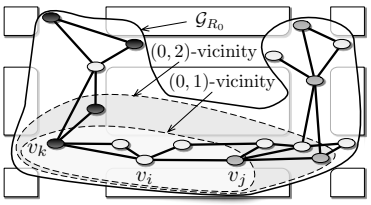

(b) $\mathcal{G}_{R_0}$, $(0,1)$- and $(0,2)$-vicinity of node v_i.

Fig. 1 Definition of v_i's (κ, ρ)-vicinity, disregarding vehicles relative speeds ($m = 1$). Subgraph $\mathcal{G}_{R_0}$ coincides with $\mathcal{G}$, including all nodes in the network.

If we consider $m = 3$, we have three different subsets of relative speeds: $R_0 = [0, 30$ km/h[, $R_1 = [30, 60$ km/h[, and $R_2 = [60, 90$ km/h[. In this case, we can separate the pairs of nodes within $\mathcal{V}$ in subsets, according to their relative speeds: $\mathcal{V}(R_0)$, $\mathcal{V}(R_1)$, and $\mathcal{V}(R_2)$. From these subsets we obtain the subgraphs illustrated in Figure 2, $\mathcal{G}_{R_0}$, $\mathcal{G}_{R_1}$, and $\mathcal{G}_{R_2}$. Note that, we can compute the shortest paths to obtain the (κ, ρ)-vicinity of node v_i only after finding $\mathcal{G}_{R_\kappa}$.

Figure 2a shows $\mathcal{G}_{R_0}$ and the $(0,1)$- and $(0,2)$-vicinity of node v_i. We observe that although the $(0,2)$-vicinity includes all nodes in the $(0,1)$-vicinity, it does not include all nodes in $\mathcal{G}_{R_0}$. Therefore, node v_i requires more than two hops to reach a node which is not in its $(0,2)$-vicinity. In the worst case, no paths connecting v_i to these nodes exist in $\mathcal{G}_{R_0}$, which means that $\rho \to \infty$ for the subset $\mathcal{V}(R_0)$. Figures 2b and 2c show, respectively, the $(1,1)$- and $(1,2)$-vicinity, and the $(2,1)$- and $(2,2)$-vicinity of node v_i, as well as $\mathcal{G}_{R_1}$ and $\mathcal{G}_{R_2}$.

Definition 2 State: *The State* ρ *stores the distance, in number of hops, between a pair of nodes in* $\mathcal{G}_{R_\kappa}$.

We model the vicinity dynamics of a node pair as a Markovian process. Hence, for a given pair of nodes v_i, v_j, there is a random variable X_{ij}^e representing the hop distance between them in a given epoch e. Thus, if nodes $v_i, v_j \in \mathcal{G}_{R_\kappa}$ are in direct contact, they are in State $\rho = 1$. If v_i can only reach v_j through an additional node $v_k \in \mathcal{G}_{R_\kappa}$, linked to both v_i and v_j, the pair v_i, v_j is in State 2, and so on. If no path exists between

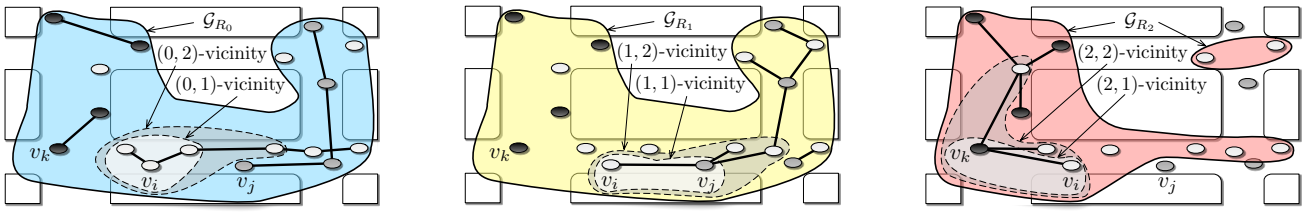

(a) $\mathcal{G}_{R_0}$, and $(0,1)$- and $(0,2)$-vicinity. (b) $\mathcal{G}_{R_1}$, and $(1,1)$- and $(1,2)$-vicinity. (c) $\mathcal{G}_{R_2}$, and $(2,1)$- and $(2,2)$-vicinity.

Fig. 2 Example of $\mathcal{G}_{R_\kappa}$ for nodes with relative speed in R_κ, the links connecting them, and the (κ, ρ)-vicinity of node v_i.

v_i, v_j in $\mathcal{G}_{R_\kappa}$, this pair of nodes is in State ∞, which only represents the absence of intermediate nodes in the same $\mathcal{G}_{R_\kappa}$ to set up a path between v_i and v_j. This does not necessarily mean that v_i and v_j are out of reach. In addition, each pair v_i, v_j may change its state only once per epoch e and the number of states is equal to the maximum number of hops interconnecting a pair of nodes plus the State ∞.

We consider that our model follows the Markov property and, therefore, the next movement depends only on the current state and not on historical behavior. Consequently, if v_i and v_j are n-hops distant in e, there is a non-negative probability that the hop distance between them will be m in $e+1$. Hence, we have that $p_{ab} = \mathbb{P}(X_{ij}^{e+1} : \rho = b \mid X_{ij}^{e} : \rho = a) \geq 0$. In addition, we also consider that each movement is independent of the epoch and that the duration on each state follows an exponential distribution. These properties are convenient since we are most interested in capturing the vicinity changes of nodes at different relative speeds. Thus, we maintain our analysis agnostic to duration in a given state to be independent of the time sampling frequency of events.

Definition 3 **κ-vicinity timeline**: *The κ-vicinity timeline is the progression of the shortest hop distance between a given pair of nodes in $\mathcal{G}_{R_\kappa}$ over time.*

The κ-vicinity timeline of a pair of nodes in $\mathcal{G}_{R_\kappa}$ is a sequence of shortest distances in number of hops (ρ) between these two nodes over time. Each entry in the timeline is an event represented by a tuple $\langle t_i, t_f, v_i, v_j, \rho, r_{ij} \rangle$, where t_i and t_f are the initial and final instants of time of the event, v_i, v_j is the pair of nodes, ρ is the shortest hop distance between them, and r_{ij} is their relative speed. State ∞ is represented by $\rho = 0$ in the κ-vicinity timeline. Time intervals are atomic, i.e., there is no other event in the whole timeline starting or finishing at an instant of time t, where $t_i < t < t_f$. This is important to better understand concurrent events. Further, state transitions in the node (κ, ρ)-vicinity, stored in timelines, allow determination of state transition probabilities, detailing how nodes move relative to each other.

3 Datasets

We analyze three traces in this work, the Mobility Dataset [20] ("Taxi"), Ad Hoc City Dataset [21] ("Bus"), and TAPASCologne Dataset [22] ("Synthetic").

3.1 Taxi trace

In the Mobility Dataset [20] the movement over time of 536 taxis in San Francisco, California – USA is recorded through GPS each 10 seconds. We analyze one day among the 30 provided.

The set of images in Figure 3 shows an approximation of the normalized density of vehicles per hour for each trace. Most results lie within $[0, 0.1$ vehicles per square meter per hour]. For better visualization, we considered as denser areas those where the result is greater than 0.15. In this scenario, although vehicles do not have predefined routes or time schedules, it is common to observe higher density of vehicles over time in the city center, as shown in Figure 3a.

The analysis of 1-day trace provides a distribution of absolute speeds where we observe that approximately 27% of the absolute speeds are within $S_0 = [0, 2$ km/h[. Such percentage of very low absolute speeds can be a consequence of waiting for passengers at taxi stands, in addition to stops due to traffic lights and street intersections. Traffic jam also contributes to this percentage and its presence in this scenario is highly plausible, as approximately 90% of the taxis move at low absolute speeds within $[0, 40$ km/h[and the flow of vehicles through the city center is intense.

3.2 Bus trace

The Ad Hoc City Dataset [21] registers the mobility of 1,200 city buses in Seattle, Washington – USA,

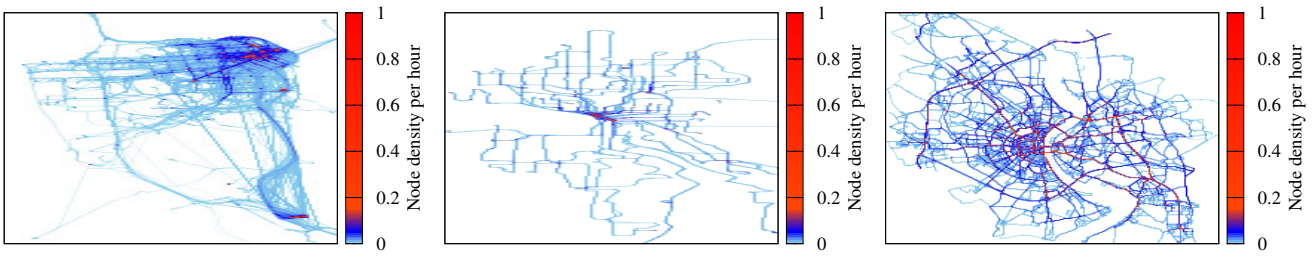

(a) Taxi trace. (b) Bus trace. (c) Synthetic trace.

Fig. 3 Density of vehicles per hour in each scenario.

using GPS. Movement is recorded 536 times per day during 30 days, among which we analyze 1 day. Comparing to the Taxi scenario, we expect to find less vehicles distributed throughout the city. We further expect that buses have predefined routes and time schedules. Figure 3b confirms that buses are distributed sparsely throughout the city and only a small area of high concentration can be found, which is probably the city center. The 1-day trace analyzed shows that approximately 35% of the absolute speeds in this scenario lie within $S_0 = [0, 2$ km/h[, indicating that buses stop more than taxis. Similarly to the Taxi scenario, most buses move at speeds within $[0, 40$ km/h[. Nevertheless, bus absolute speed in the interval $[2, 40$ km/h[is distributed more uniformly.

3.3 Synthetic trace

The TAPASCologne Dataset [22] models the vehicular traffic in the city of Cologne, Germany. The dataset is a synthetic model, built with a set of tools to simulate vehicular mobility. The complete dataset provides the movement of 121,400 vehicles during 2 hours, sampled each 1 second. Routes and time schedules in this trace are not predefined, although they usually follow a pattern for each person, and very high absolute speeds are registered due to the presence of highways crossing the city. We analyze a 10-minute subset of this dataset with almost 9,000 vehicles. The flow of vehicles in the scenario is intense and it is more distributed throughout the city, as Figure 3c shows. Absolute speeds in the 10-minute subset trace are more distributed compared to the other traces, and they are concentrated within a few intervals that sum up almost half of absolute speeds in the scenario. This is probably a consequence of the coexistence in the same scenario of different roads and streets, with different achievable speeds.

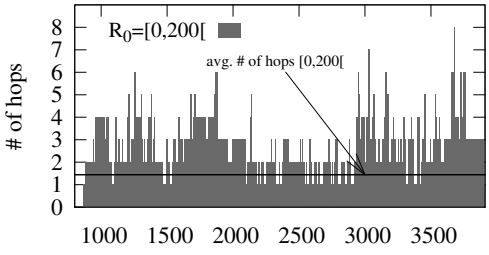

(a) $m = 1$.

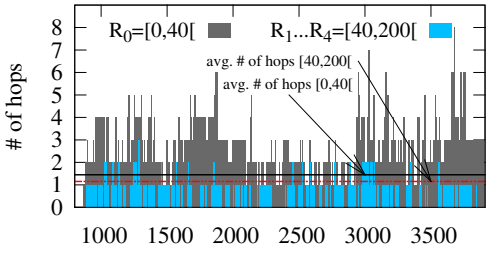

(b) $m = 5$.

Fig. 4 Example of κ-vicinity timelines, considering different values for m to show the influence of absolute relative speeds using the Bus trace.

4 Problem Statement

The investigation of mobility patterns provides information about suitable contact opportunities. This is essential for VANETs, because nodes usually experiment short contacts. Hence, the analysis of the vicinity behavior, considering the influence of relative speed and its effects on both state transitions and achievable states, yields better knowledge of what a node can expect from its vicinity.

Figure 4 shows the κ-vicinity timeline, as well as the average number of hops, for all pairs of nodes in the Bus scenario. We consider that buses move in a

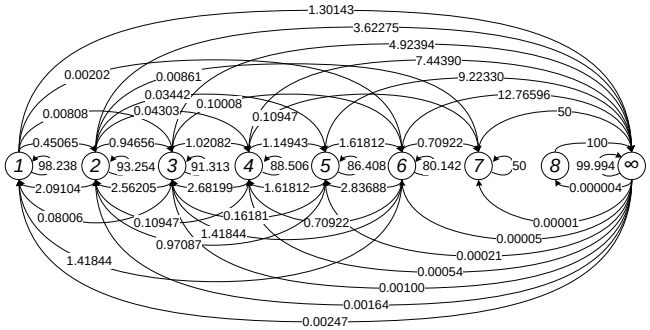

(a) $m = 1$ and $\mathcal{V}(R_0)$.

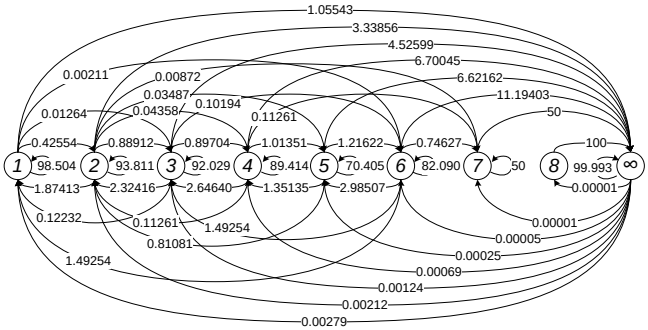

(b) $m = 5$ and $\mathcal{V}(R_0)$.

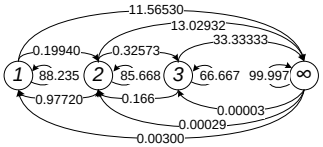

(c) $m = 5$ and $\bigcup_{\kappa=1}^{4} \mathcal{V}(S_\kappa)$.

Fig. 5 State transition probabilities (p_{nm}) in percentage, considering $m = 1$ and $m = 5$.

lossless medium and, thus, all nodes within 100 m radio range are able to communicate with each other. We set $m = 1$ and $m = 5$ to demonstrate the influence of relative speeds on the vicinity timeline. If $m = 1$, $r_{ij} \in R_0 = [0, 200$ km/h[, $\forall v_i, v_j$; whereas for $m = 5$, we have five intervals of 40 km/h. Figure 4a shows that $\rho_{max} = 8$ is the maximum number of hops between pairs of nodes with $r_{ij} \in [0, 200$ km/h[. In addition, the average number of hops is approximately equal to 1.5, meaning that most contacts happen at 1- or 2-hop distance.

Figure 4b shows the results for pairs of nodes with $r_{ij} \in [0, 40$ km/h[and $r_{ij} \in [40, 200$ km/h[. For the first range, we quickly note the significant similarity with Figure 4a. It is clear that both ρ_{max} and the average number of hops remain quite the same. For $r_{ij} \in [40, 200$ km/h[, graph bars become sparser and ρ_{max} drops to 3, consequently reducing the average number of hops. Hence, most contacts in the Bus scenario happen for relative speeds lower than 40 km/h and contacts at higher relative speeds happen mainly at a few hops. As a consequence, we should not consider longer hop distances for communications if nodes are moving at higher relative speeds. All scenarios are further investigated in Section 6.

Figure 5 shows the state transition probabilities considering all nodes in the Bus scenario. Note that values are given in percentage (%). In this scenario, contacts are rare and tend to happen for low relative speeds and states. In Figure 5a we consider a single range of relative speeds, whereas in Figures 5b and 5c we use $m = 5$ to split this range. Figure 5b represents the subset of nodes with relative speeds within [0, 40 km/h], while Figure 5c shows the transition probabilities considering all nodes at relative speeds within the other 4 subsets ([40, 200 km/h[). We observe that Figures 5a and 5b are very similar, maintaining the same ρ_{max} and the great majority of state transitions. On the other hand, Figure 5c shows that the maximum number of hops found for nodes within [40, 200 km/h[is significantly lower, confirming our observations from Figure 4, and most state transitions disappeared, especially for $\rho \geq 4$. This indicates that the highest contribution for multi-hop contacts in this scenario is concentrated at speeds lower than 40 km/h.

We observe in Figure 5 that, independently of the relative speed interval, transitions from one state to itself are most likely to happen. This probability usually decreases as ρ increases, indicating that longer hop contacts are more difficult to maintain for a long time. The highest probability is found for the transition starting and finishing at infinity, in addition, the lowest probabilities are found for the transitions starting at infinity and finishing at any other state. Yet, the higher the state, the higher the probability of going to infinity. We also observe that any pair of nodes in State ρ presents higher probability to go from its current state to State $\rho \pm 1$ rather than to State $\rho \pm n$, with $n \geq 2$. Moreover, nodes are more inclined to return to State $\rho - 1$ than to go forward to State $\rho + 1$. Thus, we can suppose that

State infinity is the most stable state and the lower the state, the more stable it is. This means that, once disconnected, nodes will much likely remain disconnected, but if such nodes fortunately come into the radio range again, they will be at 1- or 2-hop distance from their peers.

Comparing Figures 5b and 5c we further observe that the probability of going to state infinity from any other state is much higher for higher relative speeds, reinforcing the instability of multihop contacts and, moreover, the harmful influence of high relative speeds on these contacts. In fact, the instability increases even for direct contacts (State 1). For instance, a pair of nodes at relative speed within $[0, 40$ km/h$[$ has 1.055% probability of disconnecting, against a 11.565% probability if they were at higher relative speeds. These results can be easily obtained for other datasets and we generalize them for the other scenarios investigated in this work.

5 Analysis Methodology

As discussed in Section 4, nodes relative speeds can, indeed, influence the behavior of contacts. Therefore, we should consider this feature when developing protocols and applications for mobile wireless networks, especially vehicular networks. To further evaluate the impact of such speeds on multihop communications we analyze the (κ, ρ)-vicinity, which includes nodes relative speeds. To perform such analysis, we apply the following procedure:

1. Mobility trace parsing;
2. Discovery of $(0, 1)$-vicinity;
3. Computation of (κ, ρ)-vicinity;
4. Generation of κ-vicinity timeline.

5.1 Mobility trace parsing

We first parse an input mobility trace described in any format to generate an output file as required by the next step. The output file contains information such as position of nodes (x, y) at each time t, the node identification (v_i), and its absolute (s_i) and vector (s_{ix}, s_{iy}) speeds. We calculate s_i, when not provided, and inconsistencies are removed. Duplicated data is ignored, while incoherent absolute speeds are fixed through linear interpolation and use of a threshold, based on what is intuitively expected for the scenario. If the interpolated value remains above the threshold, it is ignored. This results in less than 0.3% of data being ignored due to inconsistencies, which does not affect our final conclusions. Each trace samples events at a different rate. Thus, after obtaining absolute speeds, we format the results with a uniform granularity to facilitate the computation of the $(0, 1)$-vicinity. Nodes position is updated considering that they maintain a constant speed between two adjacent points.

5.2 Discovery of $(0, 1)$-vicinity

In this step, we compute the $(0, 1)$-vicinity for every pair of nodes in the network, considering that they are all within $\mathcal{G}_{R_0}$. This means that nodes within a given radio range are considered to be in contact no matter the relative speed between them. Contacts in such situation are always at 1-hop distance. Hence, at a first moment, we do not take relative speeds into account and we use different fixed radio ranges to observe the effect of the medium attenuation. This step generates an output file containing the time t at which the pair of nodes identified by v_i, v_j exists in the trace simultaneously, the absolute speeds (s_i, s_j), the vector $(s_{xi}, s_{yi}, s_{xj}, s_{yj})$ of speeds of the same pair of nodes, the absolute relative speed between them (r_{ij}), and a flag (n_flag) that indicates whether the pair of nodes is 1-hop away.

5.3 Computation of (κ, ρ)-vicinity

Next, we obtain all (κ, ρ)-vicinities for every node in the scenario, through the computation of the shortest paths between all nodes in the $(0, 1)$-vicinity that are within the same atomic interval. Each event is recorded in an output file containing the identification of the nodes (v_i, v_j), the number of hops between them (ρ), the absolute relative speed between them (r_{ij}) and the time t at which v_i, v_j are at ρ-hops distance moving at r_{ij} relative speed. Note that at this point it is possible to obtain any specific (κ, ρ)-vicinity of any node v_i. To this end, it is only necessary to define the interval of relative speeds, i.e., R_κ, and the number of hops.

5.4 Generation of κ-vicinity timeline

We then proceed to the generation of the κ-vicinity timeline of a pair of nodes using as input the evolution of the (κ, ρ)-vicinities of all nodes in the scenario. Such vicinities provide information about the initial, t_i, and final, t_f, instants of time of the contact at ρ-hops distance and r_{ij} relative speed. This information is needed to compute the entries of the κ-vicinity timeline. Hence, the events stored in the (κ, ρ)-vicinity include the components of the κ-vicinity timeline, which is generated by

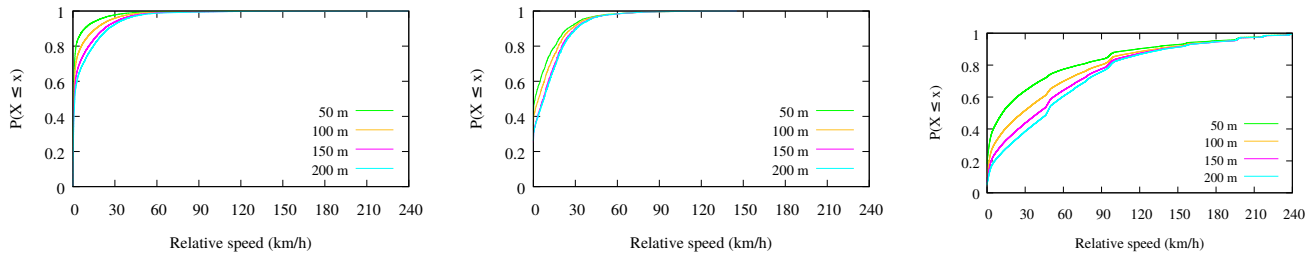

(a) Taxi scenario. (b) Bus scenario. (c) Synthetic scenario.

Fig. 6 Cumulative distribution function of relative speeds for contacts at 1-hop distance.

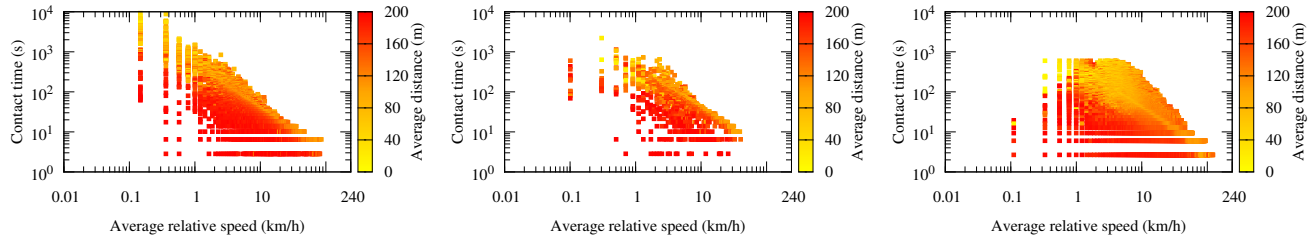

(a) Taxis, $c = 200$ m. (b) Buses, $c = 200$ m. (c) Synthetic, $c = 200$ m.

Fig. 7 Contact duration for 1-hop contacts as a function of average relative speed for each scenario.

storing the state evolution over time of a chosen pair of nodes in the (κ, ρ)-vicinity.

6 Vicinity Analysis

We analyze node vicinities using the method in Section 5, for the radio ranges in $\mathcal{C} = \{50, 100, 150, 200\}$ meters. We investigate the following features.

1. Behavior of node relative speeds;
2. Influence of relative speeds on 1-hop contacts duration;
3. Behavior of number and duration of contacts per (κ, ρ)-vicinity;
4. Average time spent by nodes in each State ρ;
5. Number of useful contacts according to the (κ, ρ)-vicinity.

6.1 Behavior of node relative speeds

We first characterize the scenarios through the investigation of node relative speeds distribution in each one for all radio ranges $c \in \mathcal{C}$. The goal is to find at which relative speed contacts are more common, considering that all nodes within c are in contact with each other at 1-hop distance. Figure 6 shows the resultant cumulative distribution function of relative speeds for each trace.

The Bus and Taxi scenarios present a similar distribution of relative speeds, as depicted in Figure 6. Most vehicles in contact moves at low relative speeds, but this percentage decreases for longer radio ranges. For instance, 90% of nodes in contact move at relative speeds lower than 27 km/h for $c = 200$ m in the Taxi scenario, while for $c = 50$ m this value increases to 99%. In the Bus scenario, we also have the majority of nodes in contact moving at speeds lower than 30 km/h for $c = 200$ m, as shown in Figure 6b. For $c = 50$ m the percentage increases from 90% to 92%. The Synthetic scenario shows a different behavior, with nodes in contact being able to move at very high relative speeds. For instance, in Figure 6c, 90% of vehicles in contact move at relative speeds up to 130 km/h for $c = 200$ m. The presence of such high speeds can be a consequence of the highways crossing the city, in the Synthetic scenario.

Another remarkable characteristic is shared by all scenarios. In Figure 6, we observe that the distributions do not significantly change for $150 \leq c \leq 200$ m, even though longer coverage ranges may enclose more vehicles. Hence, we can conclude that, from a certain radio range on, the additional nodes enclosed will not necessarily have such different relative speeds able to change the distribution.

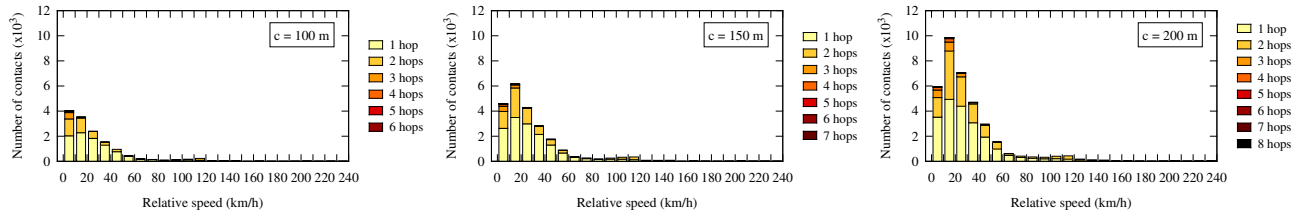

(a) Taxi scenario.

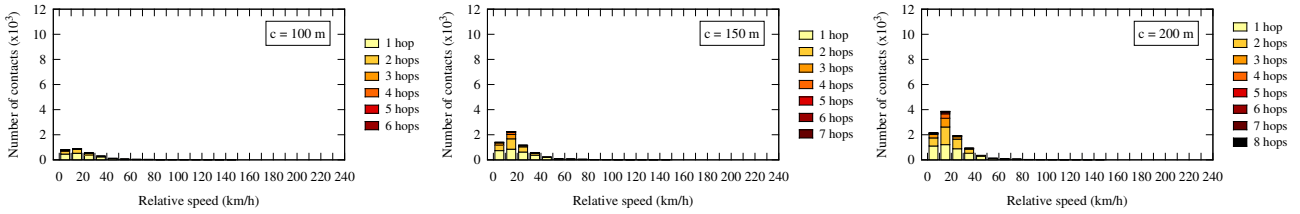

(b) Bus scenario.

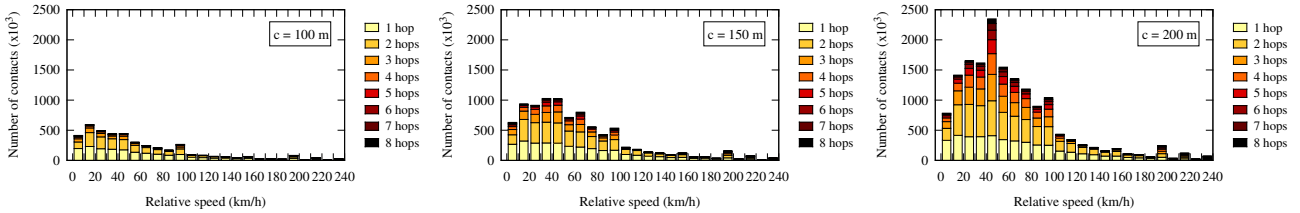

(c) Synthetic scenario.

Fig. 8 Total number of contacts as a function of the (κ, ρ)-vicinity for each scenario, using $100 \leq c \leq 200$ m.

6.2 Influence of relative speeds on 1-hop contacts duration

We investigate the duration of contacts at 1-hop distance to evaluate if long or short contacts are more common when nodes meet directly. We also investigate the relation with the physical distance between nodes. Note that when referring to distance in hops we use "hop distance", whereas for physical distance we use only "distance". From our results, we observe that (i) even though most vehicles in contact at 1-hop distance moves at low relative speeds, we can also find some vehicles in contact at higher relative speeds, and (ii) contacts at 1-hop distance are *usually* short, even considering all scenarios. Obviously, the exact duration of 1-hop contacts depends on the scenario. We noticed that 1-hop contacts established at longer distances (> 120 m) last for less than 200 s in all scenarios. A small percentage can surpass 1,000 s. We aim to confirm these observations investigating the behavior of the contact duration as a function of relative speeds. Figure 7 shows these results for $c = 200$ m, using log scale.

Each point in Figure 7 represents a contact between a pair of nodes. The duration of the contact is represented by the Y-axis. Both the relative speed and the distance between nodes during the contact can change, but we expect the variation to be small. Therefore, the X-axis represents the average relative speed and each point is colored according to the average distance during the contact. The distance is upper-limited by the coverage range c used. The Synthetic scenario shows the greatest number of points, meaning that contacts occur more often than in the other scenarios. Note that the Y-axis in this scenario is upper-limited by the duration of the 10-minute dataset sample (600 s).

We observe in Figure 7 a clear correlation between the contact duration and the average distance during the contact. Most contacts longer than 200 s tend to happen at small distances (< 120 m) between nodes, whereas nodes farther away usually establish shorter contacts. Regarding the relation between the relative speed and the average distance during the contact, we observe that long distance contacts can happen at low or high relative speeds, as shown by the reddish points between $]0, 240[$. Short distance contacts, in turn, only happen for low relative speeds. Focusing on the relation between the contact duration and the relative speed, it is clear that nodes are not able to establish long contacts at high relative speeds. For instance, contacts longer than 200 s happen only for relative speeds lower than 6 km/h in the Bus and Taxi scenarios, and for

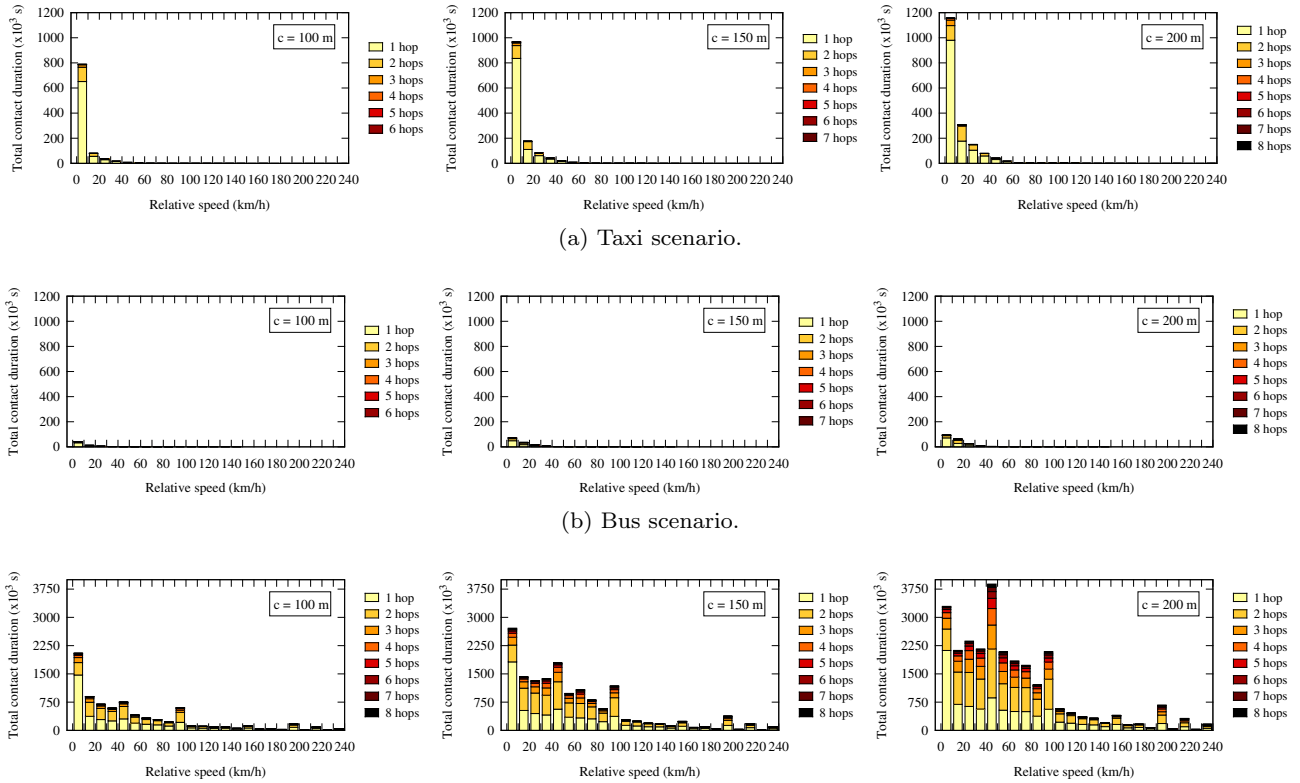

(a) Taxi scenario.

(b) Bus scenario.

(c) Synthetic scenario.

Fig. 9 Total duration of contacts as a function of the (κ, ρ)-vicinity for each scenario, using $100 \leq c \leq 200$ m.

relative speeds lower than 11 km/h in the Synthetic scenario. Very short contacts, on the other hand, can happen at any relative speed, in all scenarios. This can be observed by the presence of several points under 20 s within $]0, 240[$ range.

Summarizing our observations: nodes at high relative speeds never participate in long contacts; nodes at low relative speeds can participate in both long or short contacts, depending on the distance between them; hence, longer contacts can only happen at small distances and low relative speeds; and shorter contacts can happen at long distances and any relative speed.

The analyses carried out indicate that the best opportunities for exploiting the (κ, ρ)-vicinity happen at low relative speeds and short distances, for which contacts last longer. Note that only the relative speed is not enough to define the duration of the contact, because short contacts can happen at low or high relative speeds. Therefore, on the one hand, if we know that nodes are moving at high relative speeds, we can be sure that contacts are very likely short. On the other hand, if they are moving at low relative speeds, we cannot infer anything about contact duration. We can only infer whether the contact will be long (short distance) or short (long distance), if we also know the distance between them.

6.3 Behavior of number and duration of contacts per (κ, ρ)-vicinity

We characterize the (κ, ρ)-vicinity for each scenario regarding the total number and duration of contacts. We analyze the vicinity for ρ up to 8, because the contribution of longer hop distances to the number of contacts is insignificant, even for $c = 200$ m. The goal is to find out how far relative speeds influence multihop vicinity. This information can be used by routing protocols from VANETs to better adjust the maximum expected number of hops a message should be forwarded and prioritize sending messages to neighbors that most likely would provide longer contacts.

We divide the set of relative speeds using $s_\delta = 10$ km/h. For radio ranges c shorter than 50 m we obtain few contacts and ρ_{max} is often smaller than 8. In the Bus scenario, e.g., for $c = 50$ m we obtain $\rho_{max} = 4$. These graphs are omitted herein due to space limitations. Figure 8 shows the results for the number of contacts, for $100 \leq c \leq 200$ m. We observe that, in spite of the increasing number of contacts from one scenario to the other, all of them have similar behaviors. Indeed, we note that more contacts are established at lower relative speeds and, in all scenarios, the number of contacts

increase for each R_κ when we increase the radio range. We further observe that the growth rate of the number of contacts for R_κ with $\kappa \geq 1$ tends to be greater than for $\kappa = 0$ when c increases. This is especially true for R_1 in both Bus and Taxi scenarios. In the Syntethic scenario, the most affected relative speed range by the radio range is R_4. We believe this is due to the presence of the unlimited speed highways present in the scenario, which gives room to the presence of higher relative speeds in the scenario. In the Bus and Taxi scenarios, the contribution on the number of contacts for $\rho \leq 3$ is more significant than for $\rho \geq 4$, regardless of which $\mathcal{G}_{R_\kappa}$ we analyze. Concerning the Synthetic scenario, we can find significant number of contacts up to 6 hops. We attribute this to the higher density of the scenario compared to the others.

Figure 9 presents the results for the total duration of contacts for $100 \leq c \leq 200$ m. We quickly observe longer total duration of contacts in the Synthetic scenario, compared to the other scenarios. Again, we observe a very similar behavior for the Taxi and Bus scenarios, where contacts last longer for $R_\kappa = 0$, whereas the Synthetic scenario behaves differently. In this scenario, the total duration of contacts is higher for $R_\kappa \leq 4$. In any case, we can conclude that lower relative speeds contribute more to longer contacts. These results corroborate Figure 7. We also observe that the most significant contribution to the total contact duration is obtained for $\rho \leq 3$ in the Bus and Taxi scenarios, reaching up to 6 hops in the Synthetic scenario. Comparing Figures 8 and 9, we observe that although the highest number of contacts are found for $R_\kappa = 1$ (Bus and Taxi) or $R_\kappa = 4$ (Synthetic), the total duration of contacts is longer for $R_\kappa = 0$, except for $c = 200$ m in the Synthetic scenario, where the range of relative speeds with longest total duration corresponds to the one that has the highest number of contacts. This indicates that contacts at extremely low speeds ($0 \leq r_{ij} \leq 10$ km/h) have the longest duration. Note that these nodes do not need necessarily to be moving at low absolute speeds.

6.4 Average time spent by nodes in each State ρ

We further investigate the average time spent by nodes in each State ρ, regarding all relative speed intervals. This is important to evaluate how we can adjust protocols to exploit contact opportunities in VANETs more efficiently.

Table 1 shows the average time spent in each state, in seconds for each scenario, using different radio ranges (c). We observe that the average time in contact is always longer for State 1 in all scenarios. In the Bus and Taxi scenarios, the contribution of 8-hop contacts is barely significant. In the Synthetic scenario, the average time in contact is quite similar for all states with multihop connectivity ($\rho > 1$).

Table 1 Average time spent in each state, in seconds.

Traces	c (m)	State ρ								
		∞	1	2	3	4	5	6	7	8
Taxi	50	58.17	73.97	63.36	66.99	36.28	16.91	7.17	–	–
	100	58.00	79.41	38.53	33.25	24.24	27.32	36.60	–	–
	150	57.88	72.48	29.22	30.65	30.83	30.44	17.53	8.00	–
	200	57.77	65.77	27.06	31.87	27.80	24.96	11.92	6.33	5.00
Bus	50	29.62	23.31	10.56	4.60	7.00	–	–	–	–
	100	29.59	26.05	13.88	15.73	13.44	11.44	16	–	–
	150	29.56	27.80	18.68	14.28	8.20	4.55	2.36	1.00	–
	200	29.54	29.18	17.38	9.25	6.73	5.58	4.07	2.79	2.5
Synthetic	50	2.07	1.61	1.41	1.21	1.25	1.24	1.27	1.30	1.28
	100	2.08	2.17	1.73	1.39	1.33	1.32	1.34	1.36	1.36
	150	2.11	2.20	1.84	1.40	1.31	1.26	1.26	1.28	1.32
	200	2.08	2.06	1.82	1.37	1.26	1.19	1.17	1.13	1.13

We observe in Table 1 that, contrarily to expected, sometimes the average time in State ∞ increases. For instance, in the Synthetic scenario, when we increase the radio range from 50 m to 100 m and from 100 m to 150 m, the average time out of contact increases by a small amount. Even though increasing c one would expect higher network connectivity, depending on the scenario, we can also have more transitions from vehicles that were never connected. In the Synthetic scenario, this phenomenon has as a consequence a small increase on the average time spent in State ∞. This occurs because the number of vehicles transitioning, but still only connected for a short-time period, is higher than in the Taxi and Bus scenarios. The level of connectivity of these last scenarios does not change as much as the Synthetic one. Note that every time we change the radio range, a new graph appears with new transitions lasting for different amounts of time. As a consequence, non-linearities can happen on the average time spent on all states and not only on the State ∞. It is worth mentioning that no matter the radio range, the procedure proposed can be reproduced to calculate state transitions, even including the idea of different G_{R_κ}.

6.5 Number of useful contacts according to the (κ, ρ)-vicinity

Finally, we investigate the number of useful contacts according to the (κ, ρ)-vicinity, regarding all hop distances. We consider that a contact is useful if it can transmit a bundle using the lowest rate allowed by the IEEE 802.11p standard (USA). In this work we use bundles of 1 MB and 5 MB to exemplify, which need approximately 2.67 and 13.33 seconds, respectively, to be transmitted. We do not use single packets because the shortest contact in our datasets is equal to 1 second,

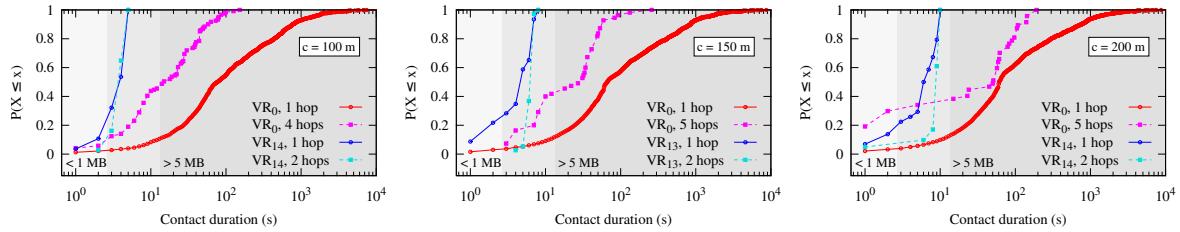

(a) Taxi scenario.

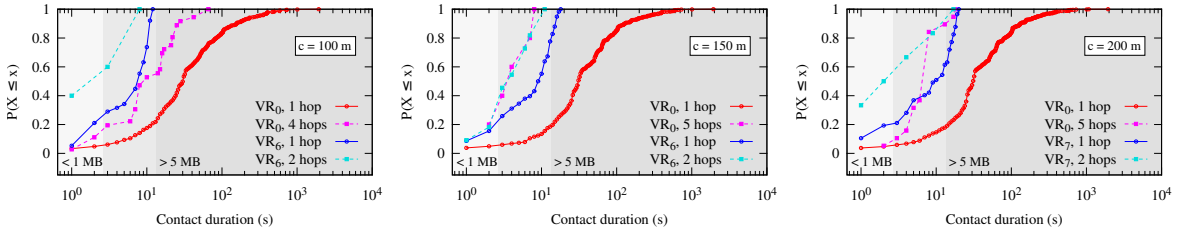

(b) Bus scenario.

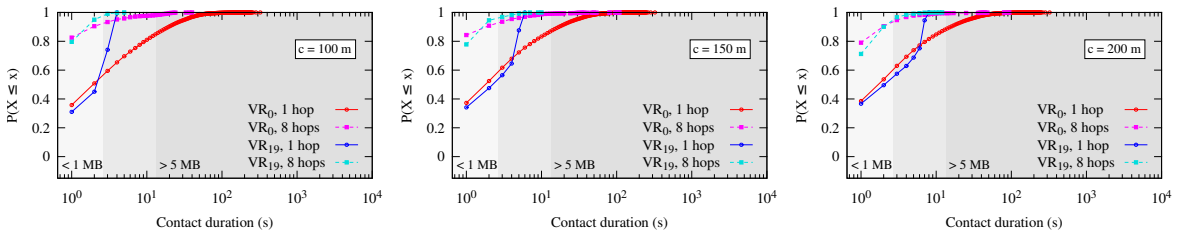

(c) Synthetic scenario.

Fig. 10 Cumulative distribution function of the duration of contacts for the first vicinity and the last vicinity able to successfully communicate.

which is enough to transmit up to 384 KB at the lowest rate of 3 Mbps, 6 times more than IPs usual 64 KB packets and at least 1966 times more than packets used in VANET safety applications [23].

We investigate all (κ, ρ)-vicinities separately and we exclude from our analysis all those that do not have enough number of contacts. Due to the lack of space, we only show the most significant results: the first vicinity, $\mathcal{V}(R_0)$, where relative speeds are within the range $[0, 10$ km/h$[$, and the vicinity with the highest possible relative speeds and number of hops where communications are still observed. Figure 10 plots the Cumulative Distribution Function (CDF) of contact duration. This figure is subdivided in 3 regions, marked by the different background colors. All contacts in the third area are able to transmit bundles greater than 5 MB. In the second area, they can transmit bundles greater than 1 MB, and in the first area, it is possible to transmit only small bundles with less than 1 MB. In all scenarios, we found that the greater the radio range, the more useful contacts happen for the transmission of both 1 and 5 MB bundle sizes, as expected. We clearly observe that the number of contacts that can transmit these bundles is a function of both number of hops and relative speed: it is more numerous for shorter hop distances at lower relative speeds and become less often for higher relative speeds or number of hops. An interesting finding, however, is that relative speeds reduce more severely the number of useful contacts when compared to the hop distance. Further, we observe that this influence is more harsh to the 5 MB bundle, for which the number of useful contacts decreases much more quickly, than to the 1 MB bundle. This behavior is repeated in all scenarios evaluated, although the specific number of useful contacts changes from one scenario to the other. Despite only confirming the notion that lower relative speeds provide better contact opportunities, we also observe that a significant number of useful contacts between nodes at very high relative speeds communicating at multihop distances also exist, even though less often. For instance, 4% of the nodes can successfully communicate at 8-hop distance within $[190, 200$ km/h$[$ at 200 m radio range, in the synthetic scenario. Considering the same radio range, we observe that, much more frequently, nodes can communicate at 2-hop distance

within [60, 70 km/h[in the Taxi scenario, and within [140, 150 km/h[, in the Bus scenario.

7 Discussion

The main goal of this work is to demonstrate the impact of nodes' relative speed on multihop communications. Moreover, we aim at deriving recommendations taking into account the correlation between hop distance and relative speed so as to improve mobile communications. Although absolute speeds are more easily obtained in vehicular networks, we consider that relative speeds are more important since it determines the contact duration. The relative speed between nodes within a predetermined radio range can be estimated even in real time, as shown by Wang et al. [24]. In this paper, our results show that:

1. nodes tend to remain in the same state at low relative speeds and, at high relative speeds, they are usually out of contact or in contact at few-hop distances;
2. direct contacts between nodes within mutual radio range last longer when their relative speed is low and their distance is less than 120 m. If the distance is greater than 120 m, the contact gets shorter, independent of the relative speed. If, however, the relative speed is high, the contact does not last for long;
3. contacts are more frequent at low relative speeds in less dense scenarios (Bus and Taxi), and they tend to happen at most at 3-hop distance. In high density scenarios (Synthetic), in turn, contacts at high relative speeds are found and they can happen even at longer hop distances. A significant number of such contacts, however, cannot be used to transfer data due to its short duration.
4. For last, the number of useful contacts decreases when the relative speed or the number of hops increases.

We observed in each evaluation that the Synthetic scenario shows a different behavior compared to the other two GPS-based scenarios. Surprisingly, despite being the most dense scenario, with the highest number of contacts, even achieving longer hop distances (6 hops), we found that it has the least number of useful contacts. We attribute this to the singularities of the scenario, such as the presence of very different types of roads, including unlimited-speed highways, and we do not generalize these findings to all dense scenarios. Following, we describe some possible applications of our findings and a proof of concept to highlight the importance of our observations.

7.1 Applications

Let us consider the purest and simplest idea of routing in challenging networks, such as VANETs. We do not have information, a priori, about contacts and there is no infrastructure to provide connectivity. In such environments, it is important to know when contact opportunities happen, how long they last, how large is the available bandwidth during contact, and how message priority could be defined. Based on the knowledge of nodes, during an encounter, they have to decide if the message will be forwarded or if it will be carried farther. In order to make better decisions, several works use information about node mobility [25–27]. This increases the probability of successfully forwarding a message to the destination. Nodes for which this probability is high are said to have high utility [28]. The relative speed could be used to compute this metric. For instance, if the relative speed between v_i and v_j is too high, node v_i should not try to forward a message through v_j. Alternatively, node v_i could forward the message but with a modified maximum expected number of hops, that would be set according to the relative speed, reducing the overhead on the network due to forward useless packets.

Shelly et al. [27] use the relative speed of nodes to predict the residual lifetime of a link, using the node attached to the longest-lasting link to forward the message. This reduces the probability of link breakage during the communication. In the best case, 90% of the predictions have an inaccuracy of less than 30% for a small Hello interval. Several packets can be dropped because of the failed prediction. The accuracy could be improved, for instance, if authors include in their model the relation between the hop distance to the destination and the relative speed (Figure 8). As such, if the destination is too far, the current node could give up sending the message.

If we take into account some kind of infrastructure, we can use a communication model based on clusters, where a special node, the cluster head, is responsible for disseminating messages to nodes in the cluster [29]. The cluster head must be carefully chosen in order to increase message delivery while minimizing the overhead. To this end, the relative speed could be a suitable parameter to determine the nodes in the cluster, including the cluster head. For instance, the cluster should be created considering only nodes for which relative speeds are very low, so that the cluster remains unchangeable for longer time intervals. He et al. [26] propose a minimum delay routing algorithm considering that messages are disseminated using clusters. In their scenario, the destination point is fixed and vehicles in

the same road travel with the same speed, with zero relative speed as a consequence. The cluster formation could benefit from our findings if authors considered that relative speeds can assume values different from zero. According to our results, protocols based on clusters or communities could use up to 3-hop communications at low relative speeds in sparser scenarios, i.e., a cluster could be composed by the (κ, ρ)-vicinities of the cluster head, with $0 \leq \kappa \leq 3$ and $1 \leq \rho \leq 3$. The algorithm proposed by He et al. neither takes into account the relative speed nor the contact time between clusters crossing each other paths. These quantities could be used to determine if the message can be fully forwarded to the next cluster. For instance, considering a sparse scenario and a source node that needs to send a 5 MB file to the destination, if the relative speed between the cluster and the next cluster is higher than 70 km/h, there is a high probability that the link between two clusters would break even before the transfer completes (results from Figure 10b). One must question if nodes at high relative speeds are always considered useless contacts. We know that if two vehicles are moving at high relative speed, at least one of them must be moving at high absolute speed, which means that it will likely encounter a significant number of nodes. This fast node could be used as a data mule that could disseminate small messages among clusters encountered along its trip.

The outcome of this work could be used to improve prediction models, and consequently, reduce the amount of resources wasted to forward packets that will not likely arrive at the destination. It could be useful also to improve packet delivery. Bazzi et al. [25] propose two routing algorithms focusing on cellular networks offloading in the VANET context. The hop-count-based algorithm considers that if a path to the nearest Roadside Unit (RSU) does not exist, the vehicle must send the packet using the cellular network. Bazzi et al. do not focus on packet delivery ratio improvement, nor on saving vehicular network resources. They would be able to improve their metric by considering that packets sent to a given RSU will never surely arrive, even when a path exists, due to vehicle mobility. Considering a dense scenario, like the Synthetic one, if the RSU is 8-hops away from the current node carrying the message and the relative speed with the next hop is higher than 10 km/h, a 1 MB message will not probably arrive at the RSU (Figure 10c). Therefore, the routing algorithm could determine that this packet should be sent to the cellular network, consequently increasing the packet delivery ratio (at the cost of increasing the cellular network usage). Hence, the information gathered in this work can be used by VANET routing protocols to better adjust the expected lifetime of a message in the network, or to choose the next hop, improving prediction algorithms. As a consequence, we can reduce the waste of resources due to useless packets, or even increase the packet delivery ratio if we rely on cellular networks.

We also studied the behavior of state transitions and we found singular patterns for both low and high relative speeds. The model can be used, for instance, to artificially extend a dataset to include information of node mobility beyond the duration of the provided dataset. The resultant mobility pattern can be used in simulations without being limited by the dataset duration.

7.2 Relative-speed-aware packet forwarding:

Our discussions in Subsection 7.1 raised the possibility of taking better decisions concerning packet forwarding, if relative speeds between nodes were considered. In this section, we evaluate this notion by showing that a local forwarding decision made by the source node can have already significant impact on the network performance. More precisely, we can save network resources by avoiding useless transmissions. Using the Network Simulator 3 (NS-3), we then simulate three different forwarding strategies that take into account information regarding relative speed. These three strategies are proposed herein and individually applied to OLSR. We evaluate their efficiency, and furthermore our analysis, using packet forwarding decisions based on relative speed awareness as a use case.

We used a subset of the 10-minute TAPASCologne dataset as the mobility model. This subset consists of an area with approximately 4.8 km^2 containing 688 nodes. OLSR was used as the baseline routing protocol and we randomly installed 100 UDP client-server applications on the available nodes. A single node can host several applications simultaneously. The only restriction is that a client node cannot be its own server, and vice-versa. Each application sends packets of 1,500 bytes at a rate of 1 packet per second. The effect of the propagation medium is addressed by the combination of two propagation models included on NS-3: 3-Log-Distance, which allows different attenuation factors for each distance range between the transmitter and the receiver; and Nakagami-m, which models fast fading.

In our simulations, the `OLSR` is used without any modification to play the role of baseline routing protocol. In our first forwarding strategy, we impose a relative speed restriction according to the hop distance to the destination (`RelSpeedR` – Relative Speed Restricted). Hence, when the source node has a packet

to a destination node, it must first check the hop distance to the destination, i.e., path length to the destination, and the next hop. The source node then will only forward the packet to the next hop, if the neighbor is within a relative speed range that still allow reaching the destination at that hop distance. In the second strategy, we use the (κ, ρ)-vicinity to modify how the OLSR chooses its vicinity (`VicR` – Vicinity Restricted). Two nodes are neighbors only if they are within a relative speed that allows 1-hop communications. In the third strategy, besides the new computation of the immediate vicinity, we also use the relative speed restriction to multihop communications, employed in the first strategy (`RelSpeedVicR` – Relative Speed and Vicinity Restricted). The relative speed thresholds are set as follows: nodes at 140 km/h can communicate up to 1-hop distance; at 120 km/h, up to 2-hop distance; at 100 km/h, up to 5-hop distance; at 80 km/h, up to 7-hop distance; and at 70 km/h, up to 8-hop distance. If nodes are moving at a relative speed higher than 140 km/h or if they are more than 8-hops away, the source node automatically discards the message. These values were obtained from Figure 10c, where we considered the relative speed that would reduce the number of contacts by only 20%.

We analyze the packet delivery ratio at the application level and the forwarding ratio of dropped packets provided by each forwarding strategy. The delivery ratio is defined as the total number of packets correctly delivered at the destination divided by the total number of packets sent. The second metric, the forwarding ratio of dropped packets, is used to investigate the number of times packets not delivered are forwarded before being dropped. With this metric, we aim to have an idea of the amount of resources wasted with useless transmissions. The forwarding ratio of dropped packets is then defined as the number of used hops before being dropped divided by the number of necessary hops to reach the destination. Hence, the highest the forwarding ratio of dropped packets, the greater the amount of resources wasted with useless transmissions. We emphasize that we do not include delivered packets in this metric. Hence, we have that `#used_hops`/`#necessary_hops`$\leq$ 1, and the unitary upper bound happens when the last hop before the destination transmits the packet but the destination does not receive it.

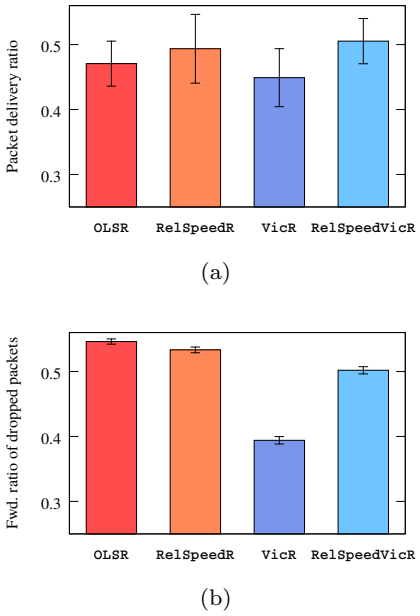

Fig. 11 (a) Packet delivery ratio and (b) forwarding ratio of dropped packets for the OLSR and each relative-speed-aware forwarding strategy based on the OLSR.

Figure 11a shows the packet delivery ratio for each scenario, using a 95% confidence interval. In average, the delivery ratio is higher for `RelSpeedR` and `RelSpeedVicR`, which are the two strategies using information regarding the number of hops reachable at specific relative speeds. This figure also shows that vicinity restrictions alone (`VicR`) are not enough to improve the packet delivery ratio. `VicR` uses with higher probability neighbors or even paths not able to deliver the packet to the final destination. If we compare, however, `RelSpeedR` with `RelSpeedVicR`, we note that `RelSpeedVicR` presents the best delivery ratio, achieving 50.5%, against 49.4% for the former. This happens because the `RelSpeedVicR` combines both restrictions, i.e., relative speed and vicinity.

Figure 11b shows the results for the forwarding ratio of dropped packets, also using a 95% confidence interval. Comparing the forwarding strategies, `VicR` wastes less network resources with useless transmissions than the others, reaching a forwarding ratio of dropped packets equal to 39.4%. The `RelSpeedVicR` is the second best forwarding strategy in terms of wasting less resources with packets. This strategy uses, in average, 50.2% of the necessary number of hops before dropping the packet, whereas `RelSpeedR` and `OLSR` achieve 53.3% and 54.6%, respectively. Therefore, we can conclude that, in the TAPASCologne scenario, submitted to the simulation parameters we described earlier, there is a clear tradeoff between delivery ratio and resource utilization. Considering both average packet delivery ratio and average forwarding ratio of dropped packets, the best strategy is `RelSpeedVicR`, as it combines number of hops and relative speed knowledge. This strategy is able to improve the average delivery ratio and simultaneously reduce the amount of useless transmissions compared with `OLSR`. Nevertheless, `VicR` is the most resource-efficient, at the cost of also reducing the

packet delivery ratio by a small amount, even compared with `OLSR`. Depending on the application, e.g., if saving network resource is a major concern, one can decide to use `VicR` instead of `RelSpeedVicR`.

8 Conclusions

This work analyzed how relative speeds can impact node vicinity. To this end, we quantitatively confirmed the intuition behind "contacts happen more often at few hops and low relative speeds" and, moreover, we demonstrate that useful contacts can also happen between nodes at higher relative speeds, separated by longer hop distances, even though less often. The relationship between number of hops and relative speeds of vehicles has never been quantified, as far as we know. To accomplish that, we used a methodology based on the idea of extended node vicinity, which can be reproduced for different scenarios. In a broader sense, results from three typical vehicular scenarios revealed similar characteristics that can be used to enhance data dissemination.

More specific results were also obtained. In this paper, we observed that whenever in contact, nodes at relative speeds higher than 40 km/h can often communicate using at most 3-hop paths; whereas at lower relative speeds, they use at most 6-hop paths. Consequently, we can infer that networks in which nodes move at higher relative speeds do not benefit much from long multihop paths, even though we can still find some useful contacts. Nevertheless, nodes at lower relative speeds can only double the number of hops from 3 to 6. We conclude then that the number of useful contacts is highly influenced by both hop distance and relative speed. These results show the importance of also considering the relative speed for path establishment when developing an application for vehicular networks to avoid unnecessary transmissions, especially when transmitting bigger messages. The same idea could also be applied to adjust the number of hops a message could be forwarded. This would be a function of the relative speed in comparison to encountered nodes. To demonstrate the impact of our results on mobile communications, we proposed three forwarding strategies based on the outcomes of this work and we ran simulations using each one. We showed that we can indeed reduce the waste of network resources with useless transmissions, while maintaining a satisfactory level of successful delivered packets. Hence, the outcomes of this work can also be used to improve existing forwarding schemes. As future work we plan to extend this analysis to other real datasets using data mining tools. In addition, we intend to propose a routing protocol that could benefit from the knowledge obtained with node relative speeds.

Acknowledgements The authors would like to thank CAPES, CNPq, and FAPERJ for their support.

References

1. S. Olariu, G. Yan, S. Salleh. A probabilistic routing protocol in VANET. IJMCMC **2**(4), 21 (2010)
2. N. Belblidia, M. Sammarco, L.H.M.K. Costa, M.D. de Amorim. EPICS: Fair opportunistic multi-content dissemination. Trans. on Mobile Computing **14**(9), 1847 (2015)
3. C. Sommer, F. Dressler. Progressing toward realistic mobility models in VANET simulations. IEEE Comm. Mag. **46**(11), 132 (2008)
4. J. Harri, F. Filali, C. Bonnet. Mobility models for vehicular ad hoc networks: a survey and taxonomy. IEEE Comm. Surveys and Tutorials **11**(4), 19 (2009)
5. D.S. Gaikwad, M. Zaveri, *VANET Routing Protocols and Mobility Models: A Survey* (2011)
6. E. Spaho, L. Barolli, G. Mino, F. Xhafa, V. Kolici. VANET simulators: A survey on mobility and routing protocols. in *Proc. of BWCCA '11*, pp. 1–10 (2011) (2011)
7. S. Zeadally, R. Hunt, Y.S. Chen, A. Irwin, A. Hassan. Vehicular ad hoc networks (VANETS): status, results, and challenges. Telecommunication Systems **50**(4), 217 (2012)
8. S. Madi, H. Al-Qamzi. A survey on realistic mobility models for vehicular ad hoc networks (VANETs). in *Proc. of ICNSC '13*, pp. 333–339 (2013) (2013)
9. V. Conan, J. Leguay, T. Friedman. Characterizing pairwise inter-contact patterns in delay tolerant networks. in *Proc. of Autonomics '07*, pp. 19:1–19:9 (2007) (2007)
10. M.C. Gonzalez, C.A. Hidalgo, A.L. Barabasi. Understanding individual human mobility patterns. Nature **453**(7196), 779 (2008). DOI 10.1038/nature06958
11. A. Passarella, M. Conti. Characterising aggregate inter-contact times in heterogeneous opportunistic networks. in *Proc. of NETWORKING'11*, pp. 301–313 (2011) (2011)
12. C.G. Rezende, R.W. Pazzi, A. Boukerche. An efficient neighborhood prediction protocol to estimate link availability in VANETs. in *Proc. of MobiWAC '09*, pp. 83–90 (2009) (2009)
13. T. Taleb, E. Sakhaee, A. Jamalipour, K. Hashimoto, N. Kato, Y. Nemoto. A stable routing protocol to support ITS services in VANET networks. Trans. on Vehicular Technology **56**(6), 3337 (2007)
14. M. Gerharz, C. de Waal, M. Frank, P. Martini. Link stability in mobile wireless ad hoc networks. in *Proc. of LCN '02*, pp. 30–39 (2002) (2002)
15. H. Menouar, M. Lenardi, F. Filali. Movement prediction-based routing (MOPR) concept for position-based routing in vehicular networks. in *Proc. of VTC '07*, pp. 2101–2105 (2007) (2007)
16. S. Barghi, A. Benslimane, C. Assi. A lifetime-based routing protocol for connecting VANETs to the internet. in *Proc. of WoWMoM '09*, pp. 1–9 (2009) (2009)
17. T. Phe-Neau, M. Dias de Amorim, V. Conan. Vicinity-based DTN characterization. in *Proc. of MobiOpp '12*, pp. 37–44 (2012) (2012)

18. T. Phe-Neau, M. Dias de Amorim, M.E.M. Campista, V. Conan. Examining vicinity dynamics in opportunistic networks. in *Proc. of PM2HW2N '13*, pp. 153–160 (2013) (2013)
19. M.A. Hoque, X. Hong, B. Dixon. Efficient multi-hop connectivity analysis in urban vehicular networks. Vehicular Communications **1**(2), 78 (2014)
20. M. Piorkowski, N. Sarafijanovic-Djukic, M. Grossglauser. CRAWDAD data set epfl/mobility (v. 2009-02-24). [Online]. Available jan/2017 in http://crawdad.org/epfl/mobility/ (2009)
21. J.G. Jetcheva, Y.C. Hu, S. PalChaudhuri, A.K. Saha, D.B. Johnson. CRAWDAD data set rice/ad_hoc_city (v. 2003-09-11). [Online]. Available jan/2017 in http://crawdad.org/rice/ad_hoc_city/ (2003)
22. S. Uppoor, M. Fiore. Large-scale urban vehicular mobility for networking research. in *Proc. of VNC '11*, pp. 62–69 (2011) (2011)
23. H. Hartenstein, K. Laberteaux, I. Ebrary. *VANET: vehicular applications and inter-networking technologies* (Wiley Online Library, 2010) (2010)
24. X. Wang, C. Wang, G. Cui, Q. Yang. Practical link duration prediction model in Vehicular Ad Hoc Networks. International Journal of Distributed Sensor Networks **11**(3) (2015)
25. A. Bazzi, B.M. Masini, A. Zanella, G. Pasolini. {IEEE} 802.11p for cellular offloading in vehicular sensor networks. Computer Communications **60** (2015)
26. J. He, L. Cai, J. Pan, P. Cheng. Delay analysis and routing for two-dimensional vanets using carry-and-forward mechanism. IEEE Transactions on Mobile Computing **16**(7), 1830 (2017)
27. S. Shelly, A.V. Babu. Link residual lifetime-based next hop selection scheme for vehicular ad hoc networks. EURASIP Journal on Wireless Communications and Networking **2017**(1), 23:1 (2017)
28. T. Spyropoulos, R.N.B. Rais, T. Turletti, K. Obraczka, A. Vasilakos. Routing for disruption tolerant networks: taxonomy and design. Wireless Networks **16**(8), 2349 (2010)
29. M. Fathian, A.R. Jafarian-Moghaddam. New clustering algorithms for vehicular ad-hoc network in a highway communication environment. Wireless Networks **21**(8), 2765 (2015)