

VANET Connectivity Analysis

Mohamed Kafsi*, Panos Papadimitratos*, Olivier Dousse†, Tansu Alpcan‡, Jean-Pierre Hubaux*

* EPFL, Lausanne, Switzerland

† Nokia Research Center, Lausanne, Switzerland

‡T-Labs, Berlin, Germany

Abstract—Vehicular Ad Hoc Networks (VANETs) are a peculiar subclass of mobile ad hoc networks that raise a number of technical challenges, notably from the point of view of their mobility models. In this paper, we provide a thorough analysis of the connectivity of such networks by leveraging on well-known results of percolation theory. By means of simulations, we study the influence of a number of parameters, including vehicle density, proportion of equipped vehicles, and radio communication range. We also study the influence of traffic lights and roadside units. Our results provide insights on the behavior of connectivity. We believe this paper to be a valuable framework to assess the feasibility and performance of future applications relying on vehicular connectivity in urban scenarios.

I. INTRODUCTION

Vehicular Ad Hoc Networks (VANETs) are a type of mobile ad hoc networks, with mobile nodes being vehicles and Road Side Units (RSUs) as static nodes. All car makers are currently investigating the feasibility and benefits of VANETs. Several prototypes have already been implemented and deployed, but there is consensus on the fact that VANETs still raise formidable design challenges. Predicting the level of VANET connectivity is a notable one, as it is strongly influenced by the peculiarities of VANETs, including the large range of possible node speeds, the mobility constrained by the road network, and the possible presence of traffic lights and RSUs.

In this paper, we study the impact of these unique characteristics on the connectivity of VANETs. More specifically, we assess the connectivity in (two-dimensional) urban scenarios, as one-dimensional, simpler highway scenarios were already investigated. By means of well-known results of percolation theory, we compute the probability for two nodes to be in power range of each other. We then derive the proportion of vehicles in the biggest cluster of the connectivity graph. We corroborate and extend this analysis by a carefully chosen set of simulations, in which we vary the parameters of vehicle density, market penetration of the communication equipment, and radio communication range, as well as several possible combinations of vehicle trips and vehicle flows. We finally study the impact of RSUs and of traffic lights. The overall approach is summarized in Fig. 1.

One of the difficulties with VANETs is that it is nowadays still unclear which applications will prevail, although it is usually admitted that collision avoidance and traffic optimization will be among the most prominent ones. To circumvent this problem, we take a rather *agnostic* approach, assuming it is desirable for vehicles to communicate with each other and with RSUs (in a multi-hop fashion whenever needed), without

worrying about the purpose of that communication. We also do not distinguish between unicast, multicast, broadcast, or geocast communications.

The rest of the paper is organized as follows. Sec. II presents the related work. Sec. III introduces the different models. Sec. IV relates the percolation theory to the connectivity in a square lattice and analyzes the connectivity results in a pure ad-hoc (vehicle-to-vehicle) network. We introduce RSUs in Sec. V and investigate its effect on connectivity. Finally, Sec. VI summarizes the lessons learned from this investigation and concludes the paper.

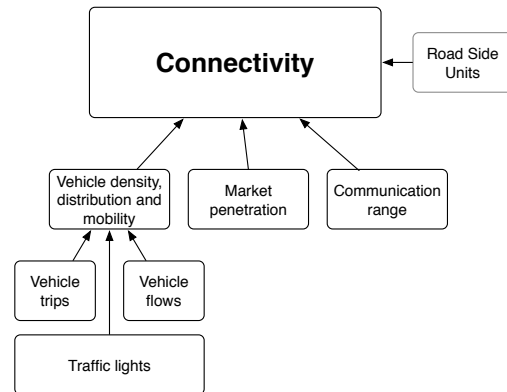


Fig. 1: Connectivity study approach.

II. RELATED WORK

Connectivity in VANETs has been a topic of interest, especially due to the recently increasing research activity. Artimy et al. investigated connectivity in VANETs for both highway and simple road configurations, using the *Roadsim* traffic micro-simulator to generate vehicle movement in multi-lane and unidirectional highways based on a cellular automaton model [15]. They examined factors such as vehicle density, relative velocity, and number of lanes, and their influence on connectivity and in particular the maintenance of an active communication session between a pair of vehicles [2]. Subsequently, in [1], they further investigated the effect of vehicle density and road configuration on the value of minimum transmission range, showing that an increase of vehicle density may not always result in higher connectivity. In a related study, an algorithm was proposed to adapt dynamically the vehicle transmission range based on local density estimation [3].

In [13], Marfia et al. considered a network where all vehicles opportunistically exploit open access points to communicate with other vehicles, looking at how the infrastructure can help inter-vehicular communication and how the stop-and-go behavior of vehicles can increase network congestion and eventually lead to collapse in performance. They observed that a significant impact of the mobility model on performance and significant differences between real data and traffic simulators. In another study [14], mobility scenarios in urban settings are studied and the importance of accurate vehicle distribution modeling on network performance is identified. Ho et al. have studied in [11] the connectivity of a vehicular ad-hoc network consisting of buses, using a hybrid simulator *GrooveNet* [12] developed by Carnegie Mellon University and General Motors, considering the impact of the topology, traffic signals, and vehicle traffic on network connectivity. In a more recent work [9], five analytical mobility models for vehicular networks, developed and used earlier in the literature, are investigated: through simulations, it is assessed how realistic the traffic they generate is and what the resultant connectivity graph properties are.

The connectivity analysis of VANETs is closely related to investigations on the connectivity of ad-hoc and hybrid networks. For example, Bai et al. proposed a framework for analyzing the impact of mobility on performance of routing protocols for ad-hoc networks metrics [4]. Dousse et al. [8] used percolation theory to develop models capturing the behavior of connected clusters as a function of node density. In a related study [5], [6], Cohen et al. investigated the resilience of the Internet to breakdowns and intentional attacks, using percolation theory to derive a lower bound on deactivated nodes that render the network disconnected.

III. MODELS

A. Traffic model

We consider a hash-shaped grid of N vertical and N horizontal roads with N^2 intersections, as depicted in Fig. 2. The distance between two consecutive horizontal (similarly, vertical) roads is L meters. We consider three types of intersections:

- 1) **No Traffic Lights:** Vehicles yield to any vehicle arriving from their right.
- 2) **Synchronized Traffic Lights:** Intersections are equipped with traffic lights. All intersections are synchronized, so that the light is green in the horizontal direction at all intersections at the same time.
- 3) **Alternating Traffic Lights (“Green Wave”):** Intersections are equipped with traffic lights. We adjust the traffic light phase duration to simulate the green wave effect, so that drivers can go through successive intersections in a main direction without stopping.

In our model, vehicles enter the simulation domain with rate f from all $4N$ road end points. The number of horizontal and vertical perpendicular roads range in our investigation from $N = 5$ to 10. Vehicles enter the simulation domain at times




Fig. 2: The simulation domain for $N = 3$. In the random destination scenario, vehicles entering at Point A pick a destination uniformly between Points A, B and C.

chosen uniformly throughout the simulation duration. If the simulation time is sufficiently long, vehicle arrivals can locally seen as a Poisson process of intensity f . Each entering vehicle picks its destination (exit point) in two different ways:

- 1) **Anti-Diametric Destination:** The vehicle source (entry) and destination (exit point) are at the same row or column of the road network.
- 2) **Random Destination:** Vehicles select uniformly their destination among the N road end points at the opposite side of the road network.

After reaching their destination, the vehicles leave the simulation domain. The total number of vehicles at time t is denoted by $n(t)$.

We use SUMO [7], as microscopic, continuous-space and discrete-time traffic simulator to model realistic vehicle behavior. SUMO simulates the movement of every single vehicle implementing *vehicle following model*. The driver behavior depends on the behavior and the speed of the preceding car and a *dynamic user assignment* to find out the route to take. The driver model decides in real time which route to take and her choice is not only dictated by the route length but also by the average speed of other vehicles along the same route. This leads to realistic behavior, as drivers change routes to avoid congestion, preferring, for example a longer but less congested route. Interestingly, as our experiments reveal, for anti-diametric destinations, chosen routes do not include turns in most cases, because the vehicle traffic is mostly uniform across the road network. In the *stationary regime*, i.e., long after the beginning of the simulation), the rate at which vehicles leave the domain is equal to f after stabilization of the simulation, unless the throughput of vehicles cannot be supported by the roads. In this case, *saturation* occurs, i.e., entry road segments fill up with vehicles and the simulator fails to add

Name	Notation [Unit]
Vehicle arrival rate	$f[h^{-1}]$
Market penetration	$\rho[-]$
Connectivity range	$r[m]$
Road segment length	$L[m]$
Isolated vehicles (%)	$\phi[-]$
Vehicles in largest component (%)	$\theta[-]$

TABLE I: Key model parameters.

new vehicles. Finally, we introduce $0 < \rho \leq 1$, the fraction of equipped vehicles (essentially market penetration for vehicular communications). The average number of equipped vehicles is $\rho n(t)$.

We present results for the three types of road traffic control differentiated by the existence and operation of traffic lights: “No Traffic Lights” (NL), “Synchronized Traffic Lights” (SL), and “Green Wave” (GW). We combine those settings with different road destination selection settings: “Anti-diametric” or “Random” destination road end points. For example, SL-R signifies synchronized traffic lights with random destinations. Note that the green wave regime can be meaningful only when vehicles can keep a constant speed. This is not the case for the random destination model, where vehicles are more likely to turn at intersections, decelerate or even stop before doing so, and thus cause significant velocity fluctuations; this is why we do not investigate the GW-R case.

B. Connectivity Model

We assume that two vehicles communicate directly through a wireless link if their distance does not exceed a certain *connectivity range* r . As a result, we obtain at each point in time a connectivity graph, the main object of this study, whose vertices correspond to vehicles equipped with an on-board vehicular communication unit. We use below the terms “vehicle” and “nodes” interchangeably, unless noted otherwise. The connectedness of the graph is measured with two (observed) values: (i) the fraction of (equipped) vehicles that are not connected with any other vehicle, denoted by $\phi(t)$, and (ii) the fraction of vehicles that belong to the largest connected component of the graph, $\theta(t)$. The parameters of our model are summarized in Table I.

IV. VEHICLE-TO-VEHICLE CONNECTIVITY

This section presents our findings on vehicle-to-vehicle (V2V) connectivity, based on analysis and extensive experiments with microscopic vehicle mobility simulation. At first, we seek to understand how vehicles are distributed across the road grid, as a function of their entrance rate. Our findings, detailed in Sec. IV-A, show that traffic has a significant impact on the density and spatial distribution of vehicles. The differing vehicle placement conditions, as a result of the differing rates of entrance (and exit) from the simulated road grid, influence the overall connectivity and most notably the proportion of vehicles in the largest cluster. But, as it will be explained in Sec. IV-B, traffic lights significantly influence the size of largest cluster. On a different dimension

Simulation interval	[0,999s]
Number of lane	2
Grid dimension	5, 10
Traffic lights regimes	Synchronized, Green Wave
Road segment length L [m]	400
f (Vehicles/h/lane)	100..1500
Types of destination	Anti-Diametric, Random
Connectivity range r [m]	100

TABLE II: Simulation parameters

of our investigation, we consider the *critical range*, that is, the minimum communication range that enables all vehicles to be part of a single cluster. As shown in Sec. IV-C, traffic lights heavily influence the value of the critical range. We shift gears in Sec. IV-D, where we are concerned with the presence of vehicles that are not equipped with radios but still affect the spatial distribution of the equipped vehicles and thus the vehicular network connectivity. The simulation parameters are summarized in Table II.

A. Vehicle Density

As the simulator allows us to control the in-flow rate for each source (entry) point, we utilize this as the independent parameter in our experiments. But it is the resultant placement of vehicles that is of real interest. We investigate how the entrance rate (for a given N dimension of the road network grid) relates to the average vehicle density. We note nonetheless that vehicle density and spatial distribution more generally do not depend only on the flow of vehicles entering the network but also on the mobility regime and type of vehicle trips. We differentiate *peripheral road segments* as those $4N$ segments between the vehicle entry (exit) points and the first- (last-) encountered intersection. We term all the remaining road segments in the grid as *central road segments*. This is motivated by the observation that, as explained below, there is a variation of the distribution across the road network.

We first look at the scenarios with anti-diametric destinations in Fig. 3: as the vehicle in-flow rate increases, the vehicle density increases in a similar manner for peripheral and central road segments. But this is so only up to a certain value of f , e.g., $f = 800$ vehicles/h/lane (in Fig. 3.(a)), after which the density in central roads increases very slowly and then it does not increase any further. While, the vehicle density in peripheral roads continues increasing until it also reaches a plateau for a larger in-flow rate, e.g., $f = 1000$ vehicles/h/lane for Fig. 3.(a). This phenomenon occurs consistently for all scenarios, but with different transition f values. Basically, the central road segment λ is lower for GW compared to SL and NL scenarios. At the same time, the differentiation between peripheral and central segments occurs for lower f values for GW, compared to SL and NL.

These variations are due to the role of traffic lights. Looking, for example, at Fig. 3.(c), central road density flattens at $f = 500$ vehicles/h/lane: traffic lights at the first intersection encountered by vehicles, that is, at the limit of each peripheral road segment, control the in-flow to the central road segments.

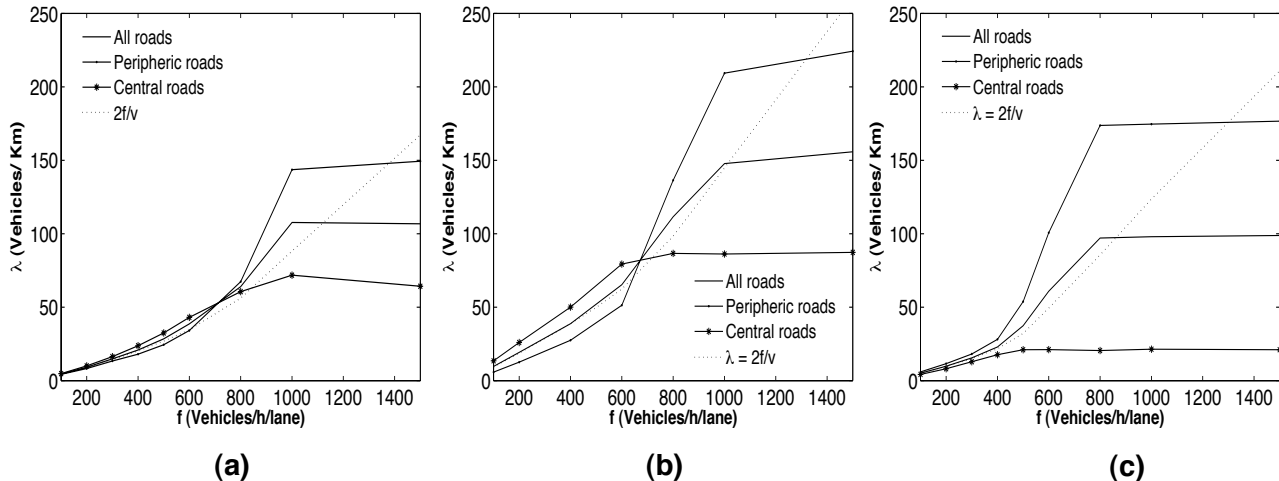


Fig. 3: Vehicle density vs vehicle flow in the scenarios (a) NL-A (b) SL-A and (c) GW-A.

Essentially, they act as a bottleneck. At this point, traffic jams start building up at the peripheral roads, up to the point the density reaches, at $f = 800$ vehicles/h/lane, a plateau: the saturated peripheral roads do not allow additional vehicles to the simulation domain. In contrast, without traffic lights, stop signs at intersections start acting as a bottleneck only when the amount of traffic is such that long queues build up at all sides of the intersection. This essentially “shifts” the saturation point for λ to occur for $f > 1000$ vehicles/h/lane. But it is interesting to see that for NL, the plateau occurs for peripheral and central roads for that same f . Again, peripheral roads reach higher densities, up to the point they get saturated.

For scenarios with random destination(s), in Fig. 4, we also observe the saturation phenomenon and the differentiation between peripheral and central road segments. However, central roads get relatively more congested than peripheral roads. This peculiarity stems from the trip type: with random destinations vehicle routes are longer while vehicles are more likely to turn at intersections to reach their destination. As a consequence, vehicles stay longer in the (simulated) road network and especially in central road segments.

B. Vehicle Clustering

We study the clustering of vehicles both analytically and through simulations. For scenarios with simple intersections, we show that accurate predictions of the network connectivity can be made using percolation theory. For the scenarios involving traffic lights, we study connectivity primarily through simulations.

1) *Analytical study*: For low vehicle arrival rates, road traffic has little impact on the spatial distribution of vehicles. In the stationary regime, we can model this with a spatial Poisson process. We neglect the width of the streets and associate them with lines and we obtain a one-dimensional process along each road; its intensity, λ , is related to the vehicle arrival rate

through the formula:

$$\lambda = \frac{2f}{\bar{v}},$$

where \bar{v} denotes the average speed of the vehicles and the factor 2 is due to the bi-directionality of traffic.

Given the Poisson assumption, one can compute an upper bound on the average fraction $E[\phi(t)]$ of vehicles that are connected to no other vehicles in the stationary regime, that is, when the road network is at a state that the rate of vehicles entering the road network is the same as the rate of vehicle leaving it:

$$E[\phi(t)] = \exp(-2\lambda\rho r).$$

Fig. 5 compares the analytical curve with simulation results, showing the limitations of the Poisson approximation: When the density of vehicles is not very low, the intersections start shaping the vehicle spatial distribution. Vehicles that are stopped at an intersection are typically not isolated, whereas vehicles leaving an intersection are spaced apart and they are likely to be isolated.

To compute $\theta(t)$, the fraction of vehicles (nodes) belonging to the largest connected component, we start by looking at the probability p that a road segment (i.e., the portion of road between two consecutive intersections) is covered by a sequence of connected vehicles. We can then see the network as an edge percolation model, where each road segment is *covered* (or *open*) with probability p .

Percolation theory allows us to address the case of a network with infinite size, with $N \rightarrow \infty$. Then it is well known that (see, for example, [10]):

$$E[\theta(t)] = 0 \text{ if } p < \frac{1}{2}$$

and

$$E[\theta(t)] > 0 \text{ if } p > \frac{1}{2},$$




Fig. 4: Vehicle density vs. vehicle flow in the scenarios (a) NL-R, (b) SL-R.




Fig. 5: Theoretical and simulated fraction of isolated vehicles in a scenario with no traffic lights and Anti-diametric destinations. $N = 10$ and $L = 400\text{m}$ long.

which is termed the *percolation phenomenon*. In other words, there is a *critical density* above which a giant connected component appears (*super-critical phase*). Below this density, all connected components have a finite size.

If the network is large but finite, its connectivity behavior is not much different. Below the critical density, all connected components are relatively small, and the largest of them may only contain a small fraction of the vehicles. On the contrary, above the threshold, a large connected component forms, containing typically more than half of the vehicles. Therefore, a “soft” transition is expected, with the fraction of vehicles in the largest cluster suddenly shifts from low values to values close to one.

We emphasize that these results are valid only if each and every road segment is covered independently with the probability p . In practice, however, the location of vehicles near intersections may influence the connectivity of adjacent segments. To cope with this dependence, we use two conditions, one necessary and one sufficient for connectivity. This

provides us an upper and a lower bound on the connectivity respectively.

The necessary condition is obtained by making the optimistic assumption that there is at least one vehicle located at the center of each intersection. Thus the condition for the connectivity of a segment is to have a sequence of connected vehicles linking the two intersections. This depends only on the location of the vehicles on the segment itself and is thus independent of any other segment. We denote the probability this condition is fulfilled for a given segment by p_u , which we compute with the help of Theorem 1 in [8]:

$$p_u = \begin{cases} 1 & \text{if } 0 \leq L < r \\ \sum_{i=0}^{\lfloor x/r \rfloor} \frac{(-\lambda e^{-\lambda r} (x - ir))^i}{i!} & \text{if } L \geq r \\ -e^{-\lambda r} \sum_{i=0}^{\lfloor L/r \rfloor - 1} \frac{(-\lambda e^{-\lambda r} (L - (i+1)r))^i}{i!} & \text{if } L \geq r \end{cases}$$

The sufficient condition is obtained by assuming that each vehicle is the center of a disk of radius $r/2$ and requiring that the segment completely is covered by disks. Consecutive vehicles must be at least at distance at most r , to be connected. However, the ends of the segment might be covered by disks from adjacent segments, which implies the coverage also depends on the location of vehicles in neighboring segments. To avoid this dependence, we strengthen the condition, by requiring that the segment be covered by disks centered on the same segment. This leads to an independent condition that is sufficient for connectivity; it is fulfilled with probability p_l , which can also be computed from Theorem 1 in [8]:

$$p_l = \frac{\sum_{i=0}^{\lfloor L/r \rfloor + 1} \frac{(-\lambda e^{-\lambda r} (L - (i-1)r))^i}{i!}}{-e^{-\lambda r} \sum_{i=0}^{\lfloor L/r \rfloor} \frac{(-\lambda e^{-\lambda r} (L - ir))^i}{i!}}$$

Fig. 6 shows a comparison of the theoretical values of $E[\theta(t)]$, using the necessary and sufficient conditions, and




Fig. 6: Fraction of vehicles in the largest cluster; two theoretical curves and the simulation-based one, $N = 10$ for the NL-A scenario.




Fig. 7: Proportion of vehicles in the largest cluster vs. vehicle density.

simulation results. We see that despite the inaccuracy of the Poisson model in predicting detailed connectivity (isolated vehicles), global connectivity matches the model qualitatively. As expected, due to the finite size of the simulated network, the transition is much softer than the one in the theoretical model (for infinite networks). However, the threshold at which the fraction of vehicles in the largest cluster jumps is well predicted by our model.

2) *Simulation Analysis:* The theoretical analysis provided above hinges on the assumption that vehicles are uniformly distributed across the network area. But this is not the case for SL and GW, for the settings with traffic lights, and more so for the SL-R and GW-R scenarios. This is why we analyze the SL and GW cases only through simulations, both for anti-diametric and random destination scenarios. Fig. 7 shows the average proportion of vehicles in the largest cluster, as a function of vehicle density for SL and GW. We observe in Fig. 6 a major difference between SL-A and GW-A and NL-A. For NL-A, $E[\theta(t)]$ increases with density (λ), and reaches 1, with almost all vehicles connected. On the contrary, for traffic-light regulated transportation, that is, for GW-A and SL-A, $E[\theta(t)] < 0.35$.




Fig. 8: Example of a typical biggest cluster in the scenario SL-A, corresponding to the “release” of the vehicles in the vertical direction.

This is explained if we observe the vehicle distribution and density for each scenario. Recall that for $\lambda < 60$ vehicles/Km, vehicles are equally distributed across central and peripheral roads. The network percolates when reaching the critical density: the vehicles form one giant cluster. For the scenarios with traffic lights, the largest cluster size is bounded: $0.2E[n(t)]$ and $0.35E[n(t)]$ for the SL-A and GW-A settings respectively. Interestingly, even though higher density, λ , is achieved (with increasing f values) for SL-A than the density achieved for NL-A (as shown in Sec. IV-A), the SL-A largest cluster size is lower than for NL-A. This can be explained by considering the traffic dynamics: For SL, there exist basically two traffic phases: traffic lights green, first, in the horizontal and then in the vertical direction. This results in the formation of lengthy horizontal and vertical disconnected clusters, as those shown in a snapshot of the network connectivity in Fig. 8. These “parallel” clusters of vehicles in motion never merge into a large and eventually giant cluster in the range of scenarios we investigate. Note that the network percolates in the SL-R scenario because the network reaches congestion, with the mean travel time being 25% higher than that in SL-A). As the central road segments become really congested, vehicles traffic becomes almost static in the center of the domain. On the other hand, for the GW scenarios, the central road density is low and well under the critical density; in the sub-critical phase, connectivity remains poor as the network does *not* percolate.

Fig. 9 shows the logarithm of the average proportion of isolated vehicles, ϕ , as a function of λ . A straightforward observation is that increased density drives ϕ to zero. Furthermore, we observe that $E[\phi(t)]$ is lower for settings with traffic lights (SL-A and GW-A), compared to no traffic lights (NL-A) setting. Vehicles held at junctions by traffic lights act as “bridges” for incident road segments. This is true even for low densities, with no vehicles approaching from other directions, as vehicles arriving at traffic lights still stop. In contrast, this is not the case for simple intersections and free-flow traffic. Note that for high densities the average proportion of isolated vehicles is higher for the SL-R scenario: congestion occurs in




Fig. 9: Proportion of isolated vehicles vs. vehicle density.




Fig. 10: Critical range vs. vehicle density.

the central road segments and vehicles entering the network are more likely to be isolated.

C. Critical Range

The critical range is shown in Fig. 10 and it is, as one would expect, a decreasing function of vehicle density. We notice however that for the traffic lights scenario the critical range surprisingly flattens when we increase the vehicle density from 110 to 150 vehicles/Km. The critical range is the mean over the duration of each of the randomly seeded simulation runs. This implies that full connectivity is not guaranteed at all times. Recall that increasing the in-flow rate f does not lead to an ever-increasing density. For example, for NL-A, the maximum average density achievable is approximately 110 vehicles/Km, and the corresponding curve in Fig. 10 is truncated accordingly.

With traffic lights, the vehicle density increases but the critical range does continue to decrease. Rather, there is a minimum value of a critical range of approximately 225 meters in the scenarios investigated here for SL and GW. The reason for this and the faster convergence of the critical range for GW is two-fold: (i) Central road density remains low and constant while in the peripheral roads it continues to increase,

and (ii) gaps appear in connectivity (if the critical range were low) due to traffic lights. In other words, being interested in the critical range, we should consider the sparser part of the network, which is that of the central road segments. At the same time, an increase in the density in central roads does not necessarily result in lower critical range values: traffic lights essentially hold vehicles stopped, in longer queues, while the higher communication range values are needed to connect those stopped vehicles with the ones moving freely (not currently stopped by a traffic light). The designer of vehicular communication protocols should take into account this stop-and-go behavior, which is more pronounced in the SL scenario, as it causes inter-vehicle distance fluctuations and thus requires higher communication ranges to ensure connectivity.

D. VC Market Penetration

In this section, we study the impact of market penetration and the resultant background traffic, that is, the non-equipped vehicles, and their impact on connectivity. We model market penetration by declaring that each vehicle starting its trip in the simulation is equipped with vehicular communication hardware and software with probability ρ ; we measure the connectivity level via simulations. As expected, the connectivity improves for increased market penetration: Increasing the proportion of equipped vehicles increases the probability to have an equipped neighbor, and this way the overall connectivity level improves. However, increasing the background traffic does not only affect the proportion of equipped vehicles but impacts the overall traffic.

In Figs. 11, 12, for scenarios with anti-diametric destinations: Background traffic has no significant effect on the connectivity level. Even though the average proportion of isolated vehicles tends to increase or remain constant, the average proportion of vehicles in the biggest cluster increases somewhat, but the gain is not significant. In Figs. 13, 14, for scenarios with random destinations: we obtain unexpected and quite interesting behavior of the average proportion of vehicles in the largest cluster. It increases significantly with background traffic, especially in the case where we have traffic lights at intersections.

The question is how can the background traffic have an effect on the connectivity, and why we have this phenomenon only for the case when we have scenarios with distributed destinations. The key factor is the effect of vehicle trips and thus the distribution of destinations on the vehicle traffic and consequently their spatial distribution. Looking closely at the scenario with random destinations and no traffic lights (corresponding to the Fig. 13), the proportion of vehicles in the biggest cluster decreases from ($\lambda = 24$ vehicles/Km, $\rho = 1$, $E[\theta(t)] = 0.4$) to ($\lambda = 48$ vehicles/Km, $\rho = 0.51$, $E[\theta(t)] = 0.34$) and increases to ($\lambda = 163$ vehicles/Km, $\rho = 0.15$, $E[\theta(t)] = 0.6$). It appears that around the point corresponding to ($\lambda = 48$ vehicles/Km, $\rho = 0.51$, $E[\theta(t)] = 0.34$) a significant change takes place. We plot the spatial vehicle distribution over time around this point:




Fig. 11: Scenario with no traffic lights and anti-diametric destinations.




Fig. 13: Scenario with no traffic lights and random destinations.




Fig. 12: Scenario with traffic lights and anti-diametric destinations.




Fig. 14: Scenario with synchronized traffic lights and random destinations.

We divide the maps in small zones and count the number of vehicles that went through in this zone over time.

We see clearly that the spatial distribution of vehicles changes. In Fig. 15, there is significant concentration and increased vehicle density at intersections. While, in Fig. 16, the central roads are congested and with an increased vehicle density. This is due to the fact that, in the distributed destinations scenarios, vehicles make longer trips and are more likely to turn at intersections. This incurs a delay in the queues at intersections and leads to denser central roads (Fig. 4). We see in Fig. 16 that the vehicle traffic is essentially forming a square in the central zone. Therefore, equipped vehicles are more likely to be in range of each other and form large clusters, unlike the case of having vehicles mostly at intersections. Overall, the background traffic affects the mobility and distribution of equipped vehicles, it has an impact, sometimes positive, on the connectivity level and more precisely on the proportion of vehicles in the largest cluster.

V. CONNECTIVITY WITH FIXED ROAD SIDE UNITS

We investigate the effects of *road side units* (RSUs) on connectivity of VANETs. We quantify improvements in connectivity due to RSUs using the measures defined in the previous sections, considering NL-A and SL-A scenarios, in order to compare the connectivity levels between RSU-assisted configurations and without RSUs present in the network. We

do not dwell on the exact form of communication via RSUs, as this will depend on the applications that will be deployed in specific systems. For the purpose of this investigation, it is assumed that all RSUs are connected over wired or other fixed communication links, e.g., the Internet. Hence, any two vehicles out of range of each other but within range still communicate as long as they both have an RSU in range. Again, how this will be done depends on future applications; as one example, a vehicle can pass a hazardous traffic alert to an RSU (possibly over multiple hops, i.e., vehicles), and the road side infrastructure can relay this alert via other RSUs to vehicles that are otherwise disconnected. The RSUs are placed at intersections and have a distance of 400m between each other. We assume that they have the same communication range as the vehicles. The placement of RSUs at the intersections is motivated by the already existing infrastructure (traffic lights, electricity, etc.) at these locations and has not been optimized with respect to connectivity.

A. Impact of RSUs on the biggest cluster

We observed in Sec. IV that the network percolates for the scenarios with simple intersections while it does not do so when there are traffic lights. We naturally expect to see a more significant connectivity gain when introducing RSUs to the scenario with traffic lights. The results in Fig. 17 and 18




Fig. 15: Vehicle distribution for $\lambda = 37$ vehicles/Km and a NL-R scenario; the z-axis is the mean number of vehicles at each time step.

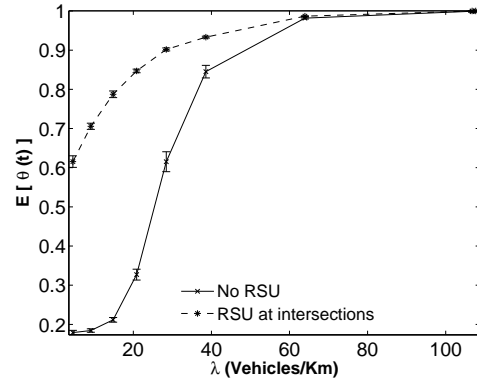


Fig. 17: Proportion of vehicles in the biggest cluster vs. vehicle density in scenario NL-A.




Fig. 16: Vehicle distribution for $\lambda = 65$ vehicles/Km and a NL-R scenario; the z-axis is the mean number of vehicles at each time step.




Fig. 18: Proportion of vehicles in the biggest cluster vs. vehicle density in scenario SL-A.

confirm our expectations. When we consider Scenario NL-A at the subcritical phase, we notice that there is a significant improvement in connectivity. At the supercritical phase, the connectivity improves slightly just above the critical density. However, the difference vanishes at higher densities.

The difference in connectivity gain between the two phases can be explained as follows. In the subcritical phase, the vehicles form small clusters disconnected from each other; i.e., vehicles belonging to the same cluster can communicate but they are isolated from other vehicles. Nonetheless, a vehicle connected to a RSU essentially acts a relay for all other vehicles in the cluster and links its peers to other clusters; this forms a big cluster consisting of small interconnected ones and the proportion of vehicles in the largest cluster increases. In the super-critical phase, a percolating cluster, containing the majority of vehicles communicate in a pure ad-hoc way, already exists. This percolating cluster is almost surely connected with an RSU. Thus, the slight connectivity improvement results from the isolated small clusters which are connected to an RSU and can then join the main percolating component. In the SL-A scenarios, the proportion of vehicles in the biggest cluster increases drastically with the assistance of RSUs: the network never percolates in a pure ad hoc (vehicle to vehicle) manner but the RSUs link the isolated




Fig. 19: Proportion of isolated vehicles vs. vehicle density in both NL-A and SL-A.

clusters to form a percolating one.

B. Impact of RSUs on the isolated vehicles

Fig. 19 shows the fraction of isolated vehicles as a function vehicle density in pure and hybrid (with RSUs) ad-hoc networks. We notice that the proportion of isolated vehicles does not change significantly when introducing RSUs at

intersections. Knowing that the intersections are the most congested spots in the domain, it follows that the isolated vehicles are more likely to be in the middle of the road or at the entering points of the domain. Therefore, the RSUs placed at intersections do not benefit these isolated vehicles. For improved connectivity in such cases a different RSU placement strategy can be used; for example, one placing RSUs at the middle of road segments.

VI. CONCLUSION

In this paper, we provide a framework for the study of vehicular connectivity in urban scenarios, and we show the use of percolation theory for that purpose. We provide an extensive set of simulations that reveal the impact of main vehicle and transportation parameters and factors, such as vehicle density, traffic light, “background” vehicle traffic (in the case of low market penetration rate for vehicular communications), and RSUs that can facilitate vehicle connectivity. These investigations result in a range of interesting findings, among which we summarize here the basic ones.

- 1) We can infer the global connectivity level based on the local vehicle density, under certain conditions (near-uniform distribution of vehicle traffic). We confirm through simulations the existence of a critical density (shown by percolation theory), above which the connectivity significantly improves.
- 2) A well-connected vehicle-to-vehicle network can be formed at all times, even with relatively sparsely placed vehicles. In fact, good (or even full) vehicular connectivity is possible without road congestion (due to high vehicle density) and relatively low communication range, for example, 25% of a road segment length.
- 3) Traffic lights have an important impact on connectivity. In particular, on the one hand, the commonly observed accumulation of vehicles at red traffic lights can be beneficial for connectivity as it can create vehicle-to-vehicle meeting points. On the other hand, this clustering has the drawback of increasing the distance between equipped vehicles and its fluctuations. Thus, the largest cluster size remains low.
- 4) For low market penetration, vehicular communication connectivity can be significantly different from that encountered when all vehicles are equipped with communicating radios. This is especially due to the role of unequipped vehicles acting as “background” traffic, that is physically occupying space and altering the spatial distribution, mobility, and finally connectivity of the equipped vehicles.
- 5) RSUs do *not* significantly improve connectivity in all scenarios. On the one hand, RSUs at intersections do not reduce the proportion of isolated vehicles. On the other hand, if the vehicular communication network is at its super-critical phase (but not fully connected yet), RSUs do not significantly increase the connectivity (i.e., the size of the largest cluster).

These results show that connectivity cannot be taken for granted. In fact, it should be carefully assessed before the deployment of any real-world application. The approach provided by this paper shows how to do so. This is a significant first step, bringing in microscopic simulation and theoretical results. In terms of future work, we intend to integrate a sophisticated radio model, the use of directional antennas, and the influence of the MAC layer. At the same time, we will produce additional theoretical approximations, on aspects such as the critical range, or modeling the traffic light regimes, for example, calculating the connectivity probability for those scenarios, as well as those with RSUs.

REFERENCES

- [1] M. Artimy, W. Phillips, and W. Robertson. Connectivity with static transmission range in vehicular ad hoc networks. In *CNSR*, 2005.
- [2] M. Artimy, W. Phillips, and W. Robertson. Connectivity in inter-vehicle ad hoc networks. *ECE Canadian Conference*, May 2004.
- [3] M. M. Artimy, W. J. Phillips, and W. Robertson. Assignment of dynamic transmission range based on estimation of vehicle density. In *ACM VANET*, 2005.
- [4] F. Bai, N. Sadagopan, and A. Helmy. IMPORTANT: A framework to systematically analyze the impact of mobility on performance of routing protocols for adhoc networks. In *INFOCOM*, 2003.
- [5] R. Cohen, K. Erez, D. Avraham, and S. Havlin. Resilience of the internet to random breakdowns. *Phys. Rev. Lett.*, 85(21):4626–4628, Nov 2000.
- [6] R. Cohen, K. Erez, D. Avraham, and S. Havlin. Breakdown of the internet under intentional attack. *Phys. Rev. Lett.*, 86(16):3682–3685, April 2001.
- [7] M. Bonert D. Krajzewicz and P. Wagner. The Open Source Traffic Simulation Package - SUMO. 2006.
- [8] O. Dousse, P. Thiran, and M. Hasler. Connectivity in ad-hoc and hybrid networks. In *IEEE Infocom*.
- [9] M. Fiore and J. Harri. The networking shape of vehicular mobility. In *ACM Mobihoc*, May 2008.
- [10] G. R. Grimmett. *Percolation*. Springer, 1999.
- [11] I. W. H. Ho, K. K. Leung, J. W. Polak, and R. Mangharam. Node connectivity in vehicular ad hoc networks with structured mobility. In *IEEE LCN*, 2007.
- [12] R. Mangharam, D. Weller, R. Rajkumar, P. Mudalige, and F. Bai. Groovenet: A hybrid simulator for vehicle-to-vehicle networks. *Mobiquitous 2006*, July 2006.
- [13] G. Marfia, G. Pau, E. De Sena, E. Giordano, and M. Gerla. Evaluating vehicle network strategies for downtown portland: opportunistic infrastructure and the importance of realistic mobility models. In *ACM MobiOpp*, 2007.
- [14] G. Marfia, G. Pau, E. De Sena, E. Giordano, and M. Gerla. Vanet: On mobility scenarios and urban infrastructure. a case study. *Mobile Networking for Vehicular Environments*, May 2007.
- [15] K. Nagel and M. Schreckenberg. A cellular automaton model for freeway traffic. *Journal de Physique I*, (115):2221, 1992.