Understanding, Modeling and Taming Mobile Malware Epidemics in a Large-scale Vehicular Network

Oscar Trullols-Cruces*, Marco Fiore†‡, Jose M. Barcelo-Ordinas*

* Universitat Politècnica de Catalunya
Barcelona, Spain
Email: {trullols,joseb}@ac.upc.edu

† Université de Lyon, INRIA,
INSA-Lyon, CITI-INRIA,
F-69621, Villeurbanne, France

‡ CNR – IEIIT
Torino, Italy
Email: marco.fiore@ieiit.cnr.it

***Abstract*—The large-scale adoption of vehicle-to-vehicle (V2V) communication technologies risks to significantly widen the attack surface available to mobile malware targeting critical automobile operations. Given that outbreaks of vehicular computer worms self-propagating through V2V links could pose a significant threat to road traffic safety, it is important to understand the dynamics of such epidemics and to prepare adequate countermeasures. In this paper we perform a comprehensive characterization of the infection process of variously behaving vehicular worms on a road traffic scenario of unprecedented scale and heterogeneity. We then propose a simple yet effective data-driven model of the worm epidemics, and we show how it can be leveraged for smart patching infected vehicles through the cellular network in presence of a vehicular worm outbreak.**

I. Introduction

Mobile malware spreading through wireless connectivity first appeared almost a decade ago [1], but major outbreaks have been prevented to date by the low penetration of smart devices and the heterogeneity of their operating systems [2]. As the number of communication-enabled devices grows and the OS market becomes more stable, with two or three major competitors remaining, the research community has already started assessing the risks yielded by large-scale diffusions of so-called mobile worms in the near future. Simulative and experimental studies have considered diverse scenarios, like campuses [3] and urban areas [4], as well as different infection vectors, ranging from metropolitan WiFi hot-spot associations [5] to text messaging in cellular networks [6].

One of the scenarios where mobile malware could cause the most damage is the automotive one. Indeed, vehicles feature today a wide range of Electronic Control Units (ECUs) interconnected by a bus, e.g., the Controller Area Network (CAN), that directly determine most of the cars' automatic behaviors. Experimental tests have proven that not only ECUs are extremely fragile to the injection of non-compliant random messages over the CAN, but that a knowledgeable adversary can exploit them to bypass the driver input and take control over key automotive functions, such as disabling brakes or stopping the car engine [7]. Lives could be thus put at stake, if a remote attack was run against a moving vehicle's ECUs. What is worse is that the above has been proved to be feasible even remotely, by exploiting the Tire Pressure Monitoring Systems (TPMS) [8] or the CD player, Bluetooth and cellular interfaces [9]. In that sense, the forthcoming IEEE 802.11p-based WAVE interfaces, allowing direct vehicle-to-vehicle (V2V) communication, risk to significantly widen the range of attack surfaces available to adversaries.

Previous works on malware diffusion in vehicular environments have considered highway-only scenarios covering a few tens of road kilometers at most. In this paper, we extend such works by proposing a more comprehensive study of the spread dynamics of a mobile worm that exploits WAVE V2V communication to self-propagate. Namely, we provide the following contributions:

- we consider a vehicular environment encompassing a geographical region of 10.000 km^2 and including more than 3.600 km of highway, regional and urban roads. Such a scenario is orders of magnitude larger than those considered in the current literature on vehicular worm diffusion, and more heterogeneous in terms of the road traffic description, comprising congested road arteries, moderately trafficked routes and underutilized rural streets at once. Deriving results in such a vast and variegated scenario allows for a better understanding of the actual level of danger of vehicular malware;
- we characterize the features of a generic vehicular worm, and assess their impact on the malware spreading and survival. Our analysis accounts for different WAVE technology penetration rates, time and location of the infection origin, as well as for the diverse factors that determine the worm efficiency in passing from one vehicle to another;
- we provide a model of the worm spreading speed, that builds on statistical road traffic data commonly available at transportation authorities. Such a data-driven approach (i) leads to a model that is significantly simpler than traditional mathematical descriptions of epidemics in vehicular environments, and (ii) unlike the latter, does not require unrealistic assumptions on the inter-vehicle arrival and spacing processes. More importantly, and despite its low complexity, our model proves very effective in representing the worm propagation process over the complex heterogeneous road network we consider.
- we leverage our model for the smart patching of a vehicular worm outbreak through the cellular network. Results show that our approach achieves a 100% healing of infected nodes, with unnecessary patching limited to less than 20% in the worst case.

After a discussion of the related literature, in Sec. II, we identify the features of a generic vehicular worm in Sec. III. Simulative results on the worm spreading are presented in Sec. IV. Our worm spread model, introduced and validated in Sec. V, is leveraged in Sec. VI for the smart containment of the worm epidemics. Finally, Sec. VII concludes the paper.

978-1-4673-5828-6/13/$31.00

II. Related Work

Worm epidemics. As first pointed out by Mickens and Noble [10], spreading models of human viruses from traditional epidemiology cannot be directly adapted to the diffusion of worms in mobile network environments. In fact, the worm spreading process differs even depending on the mobile network scenario considered, and, when focusing on vehicular environments, the literature is relatively thin.

The seminal work conducted by Khayam and Radha [11] introduces a first analytical study of vehicular worm spreading in a highway environment. In order to make the model tractable, their analysis relies on average values rather than on a complete description of the network connectivity. That way, however, the model fails to capture the complexity of the vehicular network topology, and overestimates the infection rapidity. This is also noted by Nekovee [12], who adopts instead a *frozen network* approach, studying individual snapshots of the road traffic. There, the worm epidemics is simulated as subsequent transfers among static vehicles within communication range. Although the vehicular density is derived through realistic microscopic mobility models, the author employs a uniform distribution of vehicles, an assumption which has been later shown not to hold in the real world [13]–[15].

Such a problem is overcome by Chen and Shakya [13], who also adopt frozen-network approach, but populate the snapshots according to realistic inter-vehicle spacing distributions fitted on real-world data collected by the Berkeley Highway Laboratory. The availability of such dual-loop detector data for different daily traffic conditions allows to explore the impact of daytime on the worm spreading. However, the lack of temporal correlation between the snapshots does not allow to study the propagation of worms over time; this, in turn, precludes the possibility of leveraging the model in systems where car positions change during the spreading process, e.g., in presence of roads longer than a few kilometers or of worms that take more than a few milliseconds to self-propagate. For the same reason, this technique cannot capture diffusion through carry-and-forward, where vehicles physically transport the malware until the latter can infect other cars during occasional contacts.

The diffusion-reaction and advection models employed by Hoh and Gruteser [16] describe the spatio-temporal spreading of a mobile malware. Such an approach naturally describes the system evolution over time, and thus avoids the limitations of frozen-network analyses. However, given the particular non-random nature of the road traffic mobility, calculating the diffusion coefficient in presence of realistic vehicular movements is not always possible, which limits the precision of the model and its applicability to complex heterogeneous scenarios such as the one we consider.

Our work overcomes the limitations outlined above. First, by describing the worm propagation speed rather than relying on vehicular network snapshots, the data-driven model we propose implicitly includes the time dimension. Therefore, our model can mimic the worm spreading through both multi-hop connected forwarding as well as carry-and-forwarding over temporarily network disconnections. Second, by leveraging statistical road traffic data commonly available at transportation authorities, our model does not require any complex calibration, but its single formulation is shown to closely reproduce the spreading processes for the whole space of system parameters.

Furthermore, our evaluation is conducted on significant larger scales than those considered in previous works on vehicular worm spreading. The simulation and modeling of a scenario with over 3.600 km of heterogeneous roads allows a more comprehensive analysis of the malware epidemics than those performed on a single 10-km highway corridor.

Epidemic dissemination. The spreading of generic vehicular worms can be also assimilated to the epidemic dissemination of information in vehicular networks, making the literature on the latter topic also relevant to our work. Within such a context, many analytical models have been proposed, but they all rely on the assumption that the inter-vehicle spacing is distributed deterministically [17] or exponentially [18], [19]. Although the use of these distributions simplifies the mathematical modeling, real-world experimental assessments indicate that vehicle inter-distances are typically not deterministic nor exponential [13]–[15]. In addition, some works also consider independent car speed [20], [21], breaking the well-established car-following paradigm observed in the real world [22]. Finally, all of the previous works are limited to single highway scenarios.

Our simple data-driven model does not rely on the assumption that the vehicle inter-distance and speed follow mathematically tractable distributions. Also, our study extends to a 10.000-km^2 region, much larger than those considered in the epidemic dissemination literature. Finally, as we deal with malware rather than normal content, we address worm containment, which is not a concern in epidemic dissemination.

Non-epidemic dissemination. A number of works address the problem of defining practical protocols for the efficient dissemination of information in vehicular networks. These are however of limited interest in the context of our work. Indeed, the goal of a malware designed for vehicular environments is to self-propagate as largely and rapidly as possible: this makes a simple and uncontrolled epidemic diffusion the obvious choice, as it guarantees minimum delays. The cost, paid in terms of overhead induced by the spreading process, is a primary concern for non-epidemic dissemination protocols, that thus add protocol complexity to reduce it. However, the overhead is very minor concern for a rapid malware whose objective is to harm the system.

III. Worm epidemics in vehicular networks

In this section, we characterize the features of a generic malware designed to self-propagate in vehicular environments through WAVE V2V communication.

Worms are programs that self-propagate across a communication network through security flaws common to large groups of network nodes; they are thus different from computer viruses in that the latter need the intervention of the user to propagate. Worms can be classified on the basis of several factors [23]: the *target discovery*, i.e., the way they discover targets to propagate to; the *carrier*, i.e., the infection mechanism used for the self-propagation; the *activation*, i.e., the technique by which the worm's code starts its activity on the infected host; the *payload*, i.e., the set of routines undertaken

by the worm, that clearly depend on the nature and objective of the attacker.

Our interest is on the worm epidemics within the vehicular network. Therefore, in this paper we focus on the first two aspects above, the type of target discovery and the kind of carrier employed by the worm, as they mainly drive the malware self-propagation process. Our study is instead activation- and payload-independent, since we do not delve into the kind of damage caused by the worm nor the motivation behind the attacks – although the discussion in Sec. I hints at how dangerous the outcome could be.

Target discovery. The target discovery in a vehicular network is bounded by the relatively short range of V2V communication, which limits the set of potential worm targets to cars in geographic proximity of the one the worm resides in. Thus, it is the physical mobility of cars that allows the worm to enlarge its target set, by exploiting links established between vehicles that come into contact during their trips.

Such a *mobility-driven geographic target discovery* occurs in a way that significantly differs from that of standard Internet worms, that have to perform global or local scans for IP addresses to infect. In vehicular networks, the target discovery is implicitly (and involuntarily) supported by the forthcoming V2V communication standards, that mandate the period broadcast of beacon messages by all vehicles with sub-second periodicity: this is the case for, e.g., SAE J2735 heartbeats – part of the WAVE stack – and ETSI ITS Cooperative Awareness Messages (CAMs). A worm could then simply leverage the information collected by the vehicle via these messages to determine its current target set.

Carrier. We envision two possible carrier mechanisms. In the first case, the worm is designed in a way that it can self-propagate through broadcast messages, thus infecting all of its neighbors at once. We refer to this mechanism as *broadcast carrier*. In the second scenario, the worm can only propagate itself to one neighboring vehicle at a time, and we tag such a paradigm as *unicast carrier*. We argue that, in the case of a unicast carrier, no real decision has to be taken on which communication neighbor to attack: unlike what happens in the Internet, where the choice is between hundreds of millions of machines, the number of cars concurrently in range of a worm is generally low. As an example, in the scenario we will consider in our analysis, that will be detailed in the next section and whose road topology is illustrated in Fig. 1(a), the distribution of the number of one-hop neighbors (i.e., the vehicular network node degree) follows the curves in Fig. 1(b). There, we can observe how the degree ranges from a few units to a few tens of vehicles at most, and that a car typically has less than 10 neighbors half of time. Such a small target set size does not allow for an actual selection of a target node subset, and a rapid malware would simply infect all of its neighbors. This leads in the end to an epidemic spreading of the worm even in the unicast carrier case, as we will see in our analysis.

The carrier is also characterized by a second aspect, i.e., the number of transmissions (either broadcast or unicast) required to complete the infection. This value depends on the length of the worm code and on the way it is hidden in the messages. We translate this aspect to a second parameter, referred to as *carrier latency* and indicated as τ in the following. The carrier latency is the amount of time a worm needs to self-propagate to all of its neighbors (in the broadcast case) or to one neighbor (in the unicast case). We remark that τ accounts for eventual protocol-related delays, due, e.g., to association or session establishment procedures, wireless channel contention or lost message retransmissions.

SIR model. Considering the worm epidemics from the viewpoint of the whole network, and borrowing the terminology from epidemiology, in this paper we will adopt a Susceptible, Infected, Recovered (SIR) model with Immunization. According to this model, a clean node is *susceptible* to become *infected* by the worm, but it is *healed* if it receives a dedicated cure, i.e., a patch, that prevents it from contracting the infection again. The same cure can be delivered even to a susceptible node, which is then *immunized*, i.e., it cannot be infected by the worm. We also denote the first infected vehicle as *patient zero* and its location at the time it was first infected as the *origin* of the worm infection.

The population affected by the SIR model with Immunization is formed by all the communication-enabled vehicles circulating in the geographical area of interest that suffer from the security flaw exploited by the worm to propagate. We thus characterize the population through a *penetration rate* parameter, indicated as ρ, that indicates the fraction of vehicles participating in the vehicular network and susceptible of being infected from the worm.

IV. Simulation results

We run a comprehensive simulation campaign in order to unveil the major features of worm epidemics in large-scale vehicular networks, as well as the impact that the different system parameters have on them.

Our reference scenario encompasses the whole Canton of Zurich, an area of 10.000 km^2 in Switzerland. The region, whose 3.683-km road layout is portrayed in Fig. 1(a), comprises the urban and suburban neighborhoods of Zurich, several smaller towns nearby, as well as the highways, freeways and minor regional roadways interconnecting them. The mobility of vehicles in the area has been synthetically generated by means of the multi-agent microscopic traffic simulator (MMTS) developed at ETH Zurich. The MMTS queuing-based mesoscopic modeling approach has been proven to reproduce real-world large-scale traffic flow dynamics and small-scale car-to-car interactions [24]. That of the Canton of Zurich is in fact an unescapable choice, as it is the only mobility dataset we can leverage for a large-scale study of vehicular worm epidemics. Indeed, no other synthetic or real-world mobility dataset that is publicly available covers today a similarly wide region in a comparably realistic manner.

From a network simulation viewpoint, the scale of the scenario, where up to 36.000 vehicles travel concurrently for a time span of several hours, prevents the use of a traditional network simulator, such as ns-3 or OMNeT++. Instead, we developed a dedicated simulator, that avoids the detailed processing of messages through the whole network stack and adopts a simple R-radius disc modeling of the radio-frequency signal propagation[1]. Such a design makes simulations of very

[1]Our simulator is available at `http://trullols.site.ac.upc.edu`.

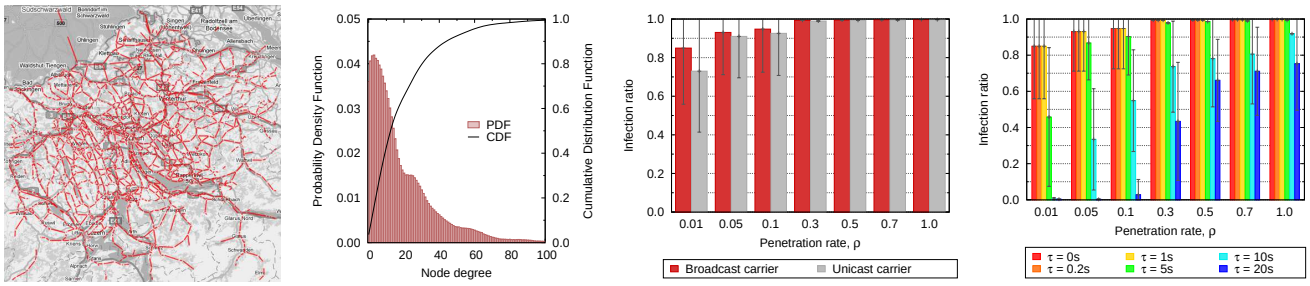

(a) Reference road topology (b) Node degree distributions (c) Carrier mechanism (d) Carrier latency, τ.

Fig. 1: (a) Road topology scenario. (b) Node degree distributions in the vehicular network. (c) Broadcast versus unicast carrier, in terms of infection ratio and under different WAVE penetration rates, ρ. (d) Impact of diverse carrier latencies versus ρ.

large-scale vehicular networks computationally feasible, providing significant qualitative insights into the system behavior in presence of different carrier types, penetration rates, V2V communication ranges and infection origins. All simulation results are averaged over 20 runs. Finally, we remark that for the moment our focus is on the understanding of the worm propagation in the vehicular environment. We will address malware patching later on, in Sec. VI.

A. Worm carrier

We first study the impact of the worm carrier. Let us first assume that patient zero originates in downtown Zurich, i.e., at the center of the map in Fig. 1(a) approximately, at 3 pm, when the road traffic intensity is at its peak. In Fig. 1(c), we focus on the carrier mechanism, setting the carrier latency τ at 1 s and comparing the results achieved by a broadcast carrier against those obtained by a unicast carrier. The performance is evaluated in terms of *infection ratio*, i.e., the fraction of vehicles the worm has infected after four hours from the injection time[2]. The x axis reports the combined WAVE technology and security flaw penetration rate ρ, that grows from 0.01 to 1. Error bars represent the standard deviation.

We observe that a broadcast carrier achieves a higher infection ratio than a unicast one. This is expected, since the latter mechanism requires the worm to self-propagate multiple times, each requiring a time τ, to reach all the nodes that a broadcast-carrier worm can reach with a single infection in a time τ. However, the difference is noticeable at very low penetration rates only, since the two carrier mechanisms perform basically the same once 5% or more of the automobiles are susceptible of contracting the worm. Furthermore, even for $\rho \leq 0.05$ the performance gap is marginal.

As far as the impact of ρ is concerned, higher penetration rates clearly lead to a more connected network of susceptible vehicles, which in turn facilitates the spreading of the worm. However, it is surprising to note how very high infection ratios are achieved even in very sparse networks comprising 1 to 5 percent of the cars. In fact, a $\rho = 0.3$ is largely sufficient to achieve a complete infection of the network. This phenomenon is imputable to the fact that the high velocity of cars can compensate for the reduced penetration rate, generating many V2V contacts and facilitating the worm self-propagation in a carry-and-forward fashion.

[2]As we will see, four hours are largely sufficient to our analysis.

In Fig. 1(d), we focus on the broadcast carrier case and study the impact of the carrier latency τ, in presence of different penetration rates. More precisely, we consider a τ ranging from 0 s (which represents an ideal upper bound to the worm spreading performance, since the worm infection is instantaneous) to 20 s. We remark that, even in presence of low penetration rates, a sufficiently fast worm (i.e., one capable of infecting its 1-hop neighborhood in one second or less) can still successfully infect the vast majority of the vehicles in a very large region such as the one we considered. In fact, even when $\tau \leq 1$ s, at least 85% of the vehicles are infected under any penetration rate. Longer carrier latencies appear instead to be more dependent on ρ: when $\tau \geq 10$ s the worm is unable of infecting the whole network even when all the vehicles are susceptible of contracting it.

The observations above let us conclude that: (i) unicast-carrier worm are as dangerous as broadcast-carrier ones; (ii) worms do not need a lot of vehicles to successfully spread through a large area, due to the fast dynamics of road traffic that tend to facilitate the malware propagation; (iii) worms do not need to be extremely fast in infecting neighboring vehicles, as a 1-second carrier delay (a perfectly realistic value, considering that worms occupy a few tens of KBytes of code and the V2V basic data transfer rate is in the order of a few Mbps) proves to be largely sufficient to vehiculate the worm to the whole network in all conditions.

B. Worm epidemics over time and survivability

The percentage of infected nodes is not the only metric of interest in the analysis of the worm propagation. The time needed for the worm to reach different regions of the vehicular network is also an important factor. Another relevant aspect is the *worm survivability*, defined as the period of time during which the infection can self-sustain in the vehicular network. These metrics can be studied by observing the dynamics of the epidemics over time.

In Fig. 2, each plot refers to a specific penetration rate ρ, and portrays the evolution of the infection for different values of the carrier latency τ. When $\rho = 0.01$, in Fig. 2(a), only rapid malware with carrier latencies τ of 1 s or less can propagate through most (although not all) of the network. The bell-shaped infection ratio for $\tau = 5$ s is explained by the aggregated road traffic volume, also depicted in the figure: the worm is not fast enough to infect the whole network before the traffic peak ends, at around 4.30pm, i.e., 1 hour 30 minutes after the infection started. As a result, the infection stays limited

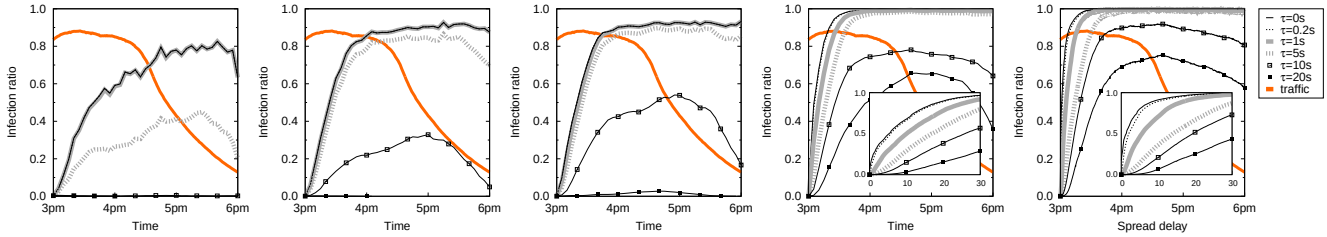

(a) ρ = 0.01 (b) ρ = 0.05 (c) ρ = 0.1 (d) ρ = 0.5 (e) ρ = 1.0

Fig. 2: Worm epidemics and survivability as a function of the penetration rate, ρ, using the broadcast carrier mechanism.

to the surroundings of the injection area, and then dies out when the traffic becomes sparser due to vehicles leaving the area or stopping. Slower worms do not even start to spread in the system. Increasing ρ to 0.05, in in Fig. 2(b), also allows slightly slower worms, characterized by a τ in the order of a few seconds, to infect an even larger majority of the vehicles. Namely, worms with $\tau \leq 5$ s perform similarly and achieve a 95% infection ratio with a linear growth during the first 45 minutes from the worm injection. This clearly makes such worms extremely dangerous, since, in order to be effective, a patch should be provided to network nodes within the very few minutes after the worm injection. The bell shape now characterizes the diffusion of malware with a carrier latency of 10 s, for the same reasons discussed above. Slower worms find it still difficult to spread at such a low ρ.

An larger participation of 10% of the cars in the network, in Fig. 2(c), does not affect the behavior of worms characterized by a $\tau \leq 5$ s, and only favors slower worms. On the other hand, as the population of susceptible vehicles grows to 50% of the overall road traffic, in Fig. 2(d), we remark two effects. First, the infection evolutions of the faster worms start to separate, as highlighted in the inset plot, which details the spreading process during the first 30 minutes from the worm injection time. Indeed, very fast worms (i.e., with $\tau < 1$ s) were previously limited by the lack of multi-hop connectivity, and had to rely on carry-and-forward to find new susceptible vehicles. As a result, their performance, hitting the bar imposed by the limited network connectivity, was similar to that of slower worms (e.g., with $\tau = 5$ s). Now, fast malware can take advantage of the presence of larger connected clusters of vehicles, and spread over 95% of the network in some 20 minutes. Moreover, the growth is now faster than linear, with 50% of the nodes being infected in less than 6 minutes. As a second remark, the higher penetration rate has a largely beneficial effect on the spread dynamics of slower worms, as now the curves for $\tau \geq 10$ s depict the infection of very wide portions of the network. However, we can still notice that the worm does not self-sustain, since its infection rate tends to drop once the traffic peak ends.

In Fig. 2(e) we consider the case where all the vehicles are communication-enabled, i.e., a maximally connected network. The effects already observed in the previous plot are here exacerbated, with the faster worms capable of reaching 50% of the network in less than 2 minutes and spreading over 95% of the network in 10 minutes. Slower worms also take advantage from the increased network connectivity, although not yet reaching a complete infection ($\tau = 5$ s) or even self-sustainability ($\tau \geq 10$ s).

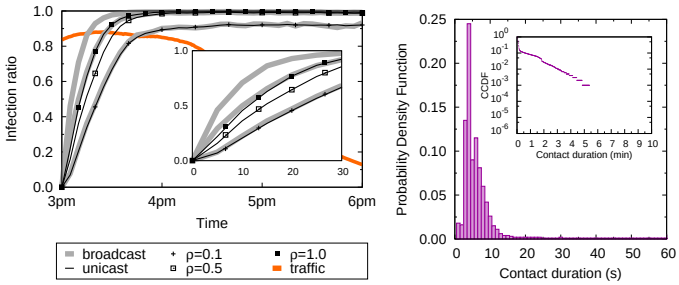

(a) Carrier mechanism (b) V2V contact duration

Fig. 3: Worm epidemics and survivability versus the carrier mechanism (a) and V2V contact duration distribution (b).

A similar temporal analysis can be done when comparing different carrier mechanisms. Fig. 3(a) depicts the infection evolution of broadcast and unicast worms, in presence of varying penetration rates, when $\tau = 1$ s. The inset plot shows again the detail of the first thirty minutes of the spreading process. We can note that, unlike what seen for the final infection ratio in Fig. 1(c), unicast and broadcast carriers differ in terms of delay. However, such a difference is mostly remarkable at high penetration rates. As the penetration rate decreases, the difference in the time needed to infect the network is reduced, and the two paradigms match when 10% or less of the cars are involved in the network.

C. Carrier latency and contact time

The previous results show a striking difference in the spread process of worms characterized by various carrier latencies. In particular, values of τ of 1 s or less seem to result in high infection rates no matter the number of vehicles involved in the network; moreover, such values of τ allow the infection to occur much faster as the penetration rate increases. On the other hand, worms with a $\tau \geq 10$ s need high penetration rates to diffuse and take a lot of time to do so. Values of τ in between those seem to result in intermediate behaviors.

The physical reason behind these performance lies in the vehicle-to-vehicle contact duration distribution, in Fig. 3(b). Most contacts among moving vehicles are very short: more than 70% of them last less than 5 seconds, and less than 10% of the contacts are longer than 10 seconds. However, the distribution is heavy-tailed, with a 5% of the contacts lasting one minute or more, as from the inset plot. Our conclusion is that fast worms, capable to spread from one vehicle to another in one second or less, can exploit any contact occurring in the network. Conversely, a worm characterized by a τ of 5 s will be only able to leverage 20% of the contacts, and one with τ =

10 s will propagate through a mere 8% of the actual V2V links. In other words, fast worms enjoy a more connected network to spread through.

D. Summary

Summarizing our findings, we can conclude that, no matter the penetration rate and carrier mechanism considered, a reasonably fast worm can be extremely dangerous. More precisely, we observed how a worm that fits a few IP packets, and that could thus be transmitted over the wireless medium in less than one second (accounting for channel contention and losses), can easily infect a vast majority of the tens of thousands of vehicles traveling in a very large urban, suburban and rural region. Even worse, such infection would occur in a time in the order of few tens of minutes at most, making it hard to counter the infection. The physical reason behind such an impressive performance of the worm diffusion lies in (i) the high number of short-lived connections generated by the movement of vehicles, and (ii) the elevate mobility of nodes in the vehicular network. Both factors contribute to create an ideal environment for a fast worm to self-propagate. We also conducted tests on the impact of the infection origin, in terms of geographical location and injection time. Although we omit[3] these results due to space limitations, we found that the location of patient zero has a dramatic effect on the epidemics, due to the spatial heterogeneity of road traffic. Conversely, the injection time only has a minor impact on the worm survivability.

V. Modeling the worm epidemics

The simulation results presented in the previous section illustrate the epidemic behavior of a generic vehicular worm in a large-scale scenario. In this section, we propose a model capable of faithfully mimicking such a behavior. Our broadcast-carrier worm propagation model is data-driven, in that it is based on commonly available road traffic statistics. Although simple in its expression, the model can capture the exact impact of the wireless communication radius R, the penetration rate ρ and the carrier latency τ on the large-scale worm propagation delay. In addition, and unlike many of the models discussed in Sec. II, our model does not require any assumption on the distribution of vehicle inter-spacing, speed or inter-arrival time (IAT). For the sake of completeness, we mention that we found the latter to be an exponential/normal mixture in the Canton of Zurich dataset, as portrayed in Fig. 4. This result matches the real-world observations in [14], and invalidates the typical deterministic and Poisson arrival assumptions employed in the literature.

The model we propose builds on (i) information about the road topology and (ii) statistical information about the road traffic. In the first category, we need, for each lane i, knowledge of its length l_i and a list of the other roads it intersects with: data that can be easily extracted from road map services such as, e.g., OpenStreetMap. As for the second category, the model requires information on the average travel speed $v_i(t)$ on a road lane i at time t and mean inter-arrival time $a_i(t)$ at road lane i and time t. Such road traffic metrics are commonly collected by transportation authorities and automobile service operators through induction loops, infra-red counters, traffic monitoring cameras, and, more recently, floating car data systems. Therefore, such historical or statistical data is currently available for large portions of the road topology and its public disclosure is growing, fostered by open data initiatives.

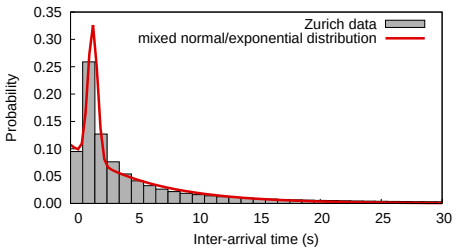

Fig. 4: IAT distribution in the ETH Canton of Zurich dataset.

Road traffic information is by its own nature time-varying, i.e., the average speed and IAT are not the same during the day or on different days of the week, which is why we consider $v_i(t)$ and $a_i(t)$ be dependent on time. The aforementioned statistical data necessarily reflects this aspect, with a finite yet representative time granularity[4]. Also, note that although higher order statistics may be available, our model only requires knowledge of the mean values of $v_i(t)$ and $a_i(t)$ at each time period. Leveraging the data above, we next discuss the modeling of the worm propagation speed along a single road, in Sec. V-A, and then extend the result to a network-wide worm propagation, in Sec. V-B.

A. Per-road worm propagation

Our goal is initially to model the *worm propagation speed*, $s_i(t)$, along a lane i characterized by average road traffic parameters $v_i(t)$, $a_i(t)$ at time t, accounting for the technology-related parameters R, ρ and τ. For the sake of clarity, in the following we refer to a generic time instant and drop the time notation. We start from the consideration that the worm propagation speed mainly depends on the network connectivity level. Namely, the malware can propagate wirelessly, and thus at a high speed, in a well-connected vehicular network where multi-hop communication can take place. Conversely, the worm propagation is slowed down when communication opportunities are scarce. Focusing on the two extreme cases, we can state that: (i) in complete absence of vehicle-to-vehicle connectivity, the worm propagates at the vehicular speed v_i, as it is physically carried by isolated cars; (ii) in presence of a complete road coverage by a very dense multi-hop vehicular network, the worm instantaneously[5] jumps of a distance equal to the communication range R at each carrier latency, the latter requiring a time τ during which the worm still moves at the vehicular speed v_i. Therefore, the worm propagation speed has an expression of the type distribution of the form

$$s_i = v_i + \frac{R}{\tau} f(a_i, v_i, \rho, R, \tau) \quad (1)$$

[3]For a detailed description, see Technical Report UPC-DAC-RR-2012-19 at https://www.ac.upc.edu/app/research-reports/html/2012/19/report.pdf

[4]In our evaluation, we assume that statistical data on the road traffic is aggregated and updated with a time granularity of 15 minutes, largely sufficient to capture the time variability of the Zurich road traffic metrics.

[5]We assume the RF signal propagation delay to be negligible, since it is in the order of nanoseconds, a value at least three orders of magnitude lower than the duration of the other events involved in the process.

where $f(\cdot)$ is a function of the different system parameters that represents the vehicular network connectivity level. Such a function assumes values between 0 (absence of connectivity) and 1 (fully connected network). In order to characterize the exact expression of $f(\cdot)$, we observe the impact of the different parameters on the network connectivity. Considering the simplified case of vehicles moving along one single road direction, the average distance between two subsequent vehicles is given by $a_i \cdot v_i$, i.e., the distance traveled by the first vehicle before the following one enters the same road. The technology penetration rate can be accounted for by assuming that the first vehicle is equipped with a communication interface, and ρ can be seen as the probability that the following vehicle is communication-enabled as well. Then, an average of $1/\rho$ vehicles must enter the road before a second car equipped with a radio interface actually appears on the road. The average distance between to vehicles that are participating in the network is then $\frac{a_i v_i}{\rho}$. The connectivity is determined by the relationship between the distance above and the transmission range R. In particular, it is the ratio between the two values, $\frac{a_i v_i}{\rho R}$, that matters: the lower the ratio, the higher the network connectivity, and vice-versa.

The discussion above refers however to the case of vehicles all moving in a same direction. The presence of an opposite vehicular flow can be accounted for through a factor K, that divides R. In other words, if vehicles in the other direction of movement can be leveraged for the worm propagation, a range R that is K times smaller provides the same connectivity achieved by a range R in a single-direction scenario. Alternatively, the introduction of K can be interpreted as the fact that one can allow a distance K times larger between two communication-enabled vehicles and still achieve the same level of connectivity. As discussed later, we found the value of K to be almost invariant to the whole range of road traffic and communication settings we evaluated, and we thus treat it as a constant in the following. Finally, τ has no major impact on the network connectivity expressed by $f(\cdot)$, since it is an application-level parameter. The only case where τ can indirectly affect the network connectivity is that of a carrier latency so large to be comparable to the time required to travel along a whole road segment between two intersections. However, the latter is at least several tens of seconds, while the values of τ of interest to our study are significantly shorter. Therefore, we neglect the impact of τ in the following.

Summarizing our discussion above, the worm propagation speed can be expressed as

$$s_i = v_i + \frac{R}{\tau} f\left(\frac{a_i\, v_i}{\rho\, R/K}\right). \quad (2)$$

We still have to identify a proper expression for the function $f(\cdot)$ and a value for the parameter K. To that end, we employ the data from our worm propagation simulations. Interestingly, by fitting the expression in (2) to the data, we observe that a single function $f(x) = exp(-x^2)$ and a single value $K = 3$ fit the data for any combination of the road traffic and communication parameters. Therefore, the worm propagation speed along lane i can be finally expressed as:

$$s_i = v_i + \frac{R}{\tau}\, exp\left[-\left(\frac{a_i\, v_i}{\rho\, R/3}\right)^2\right]. \quad (3)$$

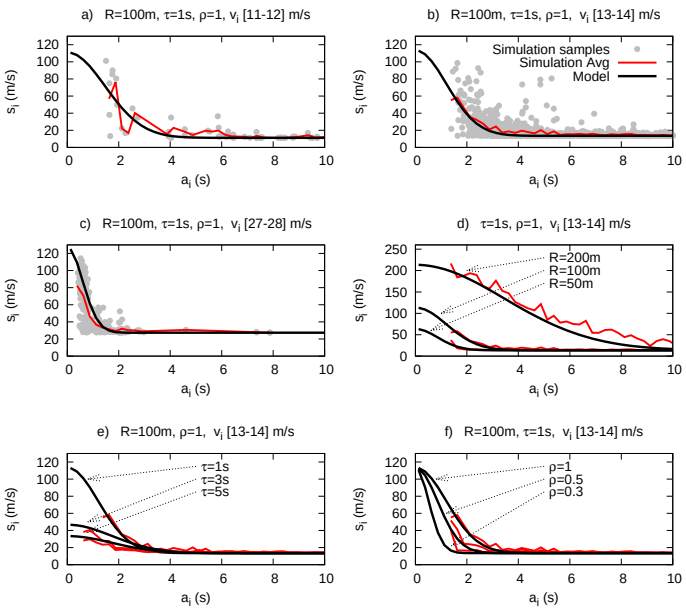

Fig. 5: Per-road worm propagation speed, s_i, for different combinations of the road traffic parameters v_i, a_i and technology parameters R, ρ, τ.

This single simple expression can thus be employed to describe the average worm propagation speed along any road in the Zurich scenario, given that its road traffic statistics, i.e., the average vehicular speed v_i and the average inter-arrival time a_i, are known. The equation in (3) allows then to evaluate the impact of the propagation settings R, ρ and τ, since it holds for any combination of the same.

Examples of the correctness of the model are provided in Fig. 5. Plots in Fig 5(a), (b) and (c) aggregate the results for roads with similar average vehicular speeds, ranging between 11 m/s (less than 40 km/h) and 28 m/s (over 100 km/h). Each plot displays a scatterplot of the worm propagation speed measured at each lane i, versus the road average inter-arrival time, a_i, with baseline parameters R = 100 m, τ = 1 s and ρ = 1. The red curve represents the average behavior observed over all roads, while the black curve is the result provided by our model. The plots in Fig 5(d), (e) and (f) show instead the worm propagation speed for different values of R, ρ and τ. There, for the sake of clarity, the scattered simulation samples are not drawn and only the average curves are reported. It is however clear that our data-driven model can faithfully mimic the average behavior of the worm propagation speed, in any road traffic condition. Of course, the model does not capture the random variability around the mean that is observed for specific roads. This is due to the fact that we only consider the average values of v_i and a_i in our study, and not higher-order moments of their distributions. This is, however, an intentional choice, that allows us to keep the model very simple, still obtaining excellent results when considering the network-wide worm propagation process, as discussed in the next section.

B. Road network-wide worm propagation

We now leverage the worm propagation speed expression in (3) to describe the propagation process over the whole road network. In particular, the propagation is characterized

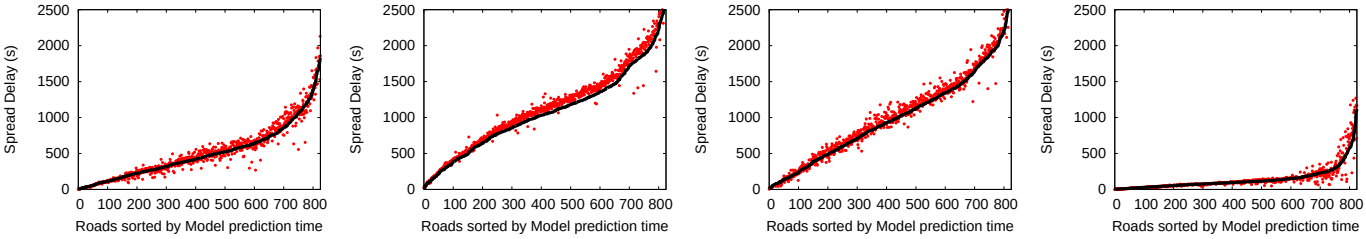

(a) Default (b) $\tau = 10$. (c) $\rho = 0.3$. (d) $R = 300$ m.

Fig. 6: Road network-wide spread delay with (a) R = 100 m, τ = 1 s, ρ = 1, and when (b) τ = 10 s, (c) ρ = 0.3, (d) R = 300 m.

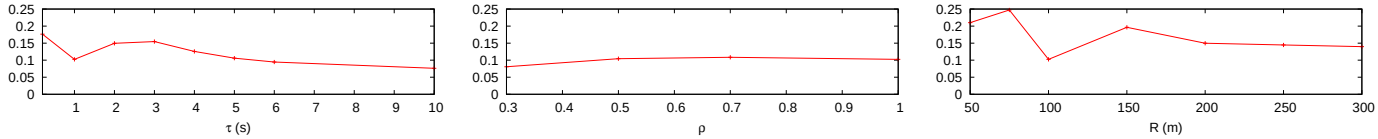

Fig. 7: Relative Error, η, versus the carrier latency τ (left), penetration rate ρ (middle) and communication range R (right).

in terms of *spread delay*, i.e., the time that a worm takes to reach a specific location after the infection of patient zero.

Let us first represent the road layout as a graph G=($\mathcal{V}, \mathcal{E}$), where the set of vertices $\mathcal{V}$ represents the intersections and the set of edges $\mathcal{E}$ represents the roads joining such intersections. Knowing the worm propagation speed s_i along a road segment i, the spread delay from one end of the road to the other can be derived as w_i=$\frac{l_i}{s_i}$. Each edge in $\mathcal{E}$ is then associated to a weight matching its spread delay w_i. Note that the resulting weighted graph is time-varying, since the worm propagation speeds along each road, and thus the weights derived from them, change over time.

Given the infection time t and the location on the road topology of patient zero, calculating the spread delay from the origin point to any other point of the region reduces to a single-source shortest path problem on the weighted graph associated to time t. A standard Dijkstra's algorithm can be used to rapidly solve the problem. Then, the spread delay to a given location on the road network is given by the cost of the shortest path to its corresponding vertex or edge on the graph. Indeed, such a cost maps to the sum of the spread delays along the fastest path from the infection origin up to the location of interest.

The qualitative evaluation of the model is presented in Fig. 6, where each plot portrays the road network-wide spread delay measured in simulation (dots) and computed by our data-driven model (solid line). Graph vertices (i.e., road intersections) are ranked along the x axis according to the worm spread delay determined by the model. Fig.6(a) refers to the case of τ = 1 s, ρ = 1 and R = 100 m. Although there are a few outliers, a vast majority of the delays needed to reach the different road intersections in simulation is correctly reproduced by the model. We can observe that the quality of the result is the same when the different systems parameters τ, ρ and R are varied, in Fig. 6(b), Fig. 6(c) and Fig. 6(d), respectively.

A more complete picture of the model reliability is provided in Fig. 7, that shows the average relative error η between the simulation results and the model outcome for the whole parameter space. Notably, the error remains below 0.18% for all values of $\tau \in [0.2, 10]$, stays below 0.1% for any value of the penetration rate ρ, and is at most 0.25% for short communication ranges below 50 m. Overall, the error is negligible in all cases, with a correct prediction for more than 99% of road intersections in all situations.

VI. Containing the Worm Epidemics

The results in the previous sections outline the dangerousness of vehicular worms, and motivate the development of solutions for the rapid containment of their outbreaks. The typical techniques proposed in the literature are preemptive immunization and interactive patching [11], [12]. In the first case, a subset of the vehicles is preemptively immunized so as to prevent the propagation of the worm. However, preemptive immunization makes sense only in the case of frozen networks, where immunized nodes can disrupt the network connectivity exploitable by the malware. In presence of a more complete time-evolving analysis, worms can easily overcome the obstacle of preemptively immunized vehicles thanks to the car mobility over time. In the case of interactive patching, a patch to the worm is released in the network and diffused through V2V communication in an epidemic fashion. In other words, the patching follows a spreading similar to that of the worm itself. However, resorting to V2V communication to contain the malware epidemics does not seem a sensible choice, as it implies high delays and a probability of success that cannot be certain.

We consider instead that cellular communication can be leveraged to distribute the patch in a rapid and reliable manner. Indeed, vehicles already start to be equipped with 3G/4G radios, whose diffusion will anticipate that of WAVE communication interfaces. Therefore, that of a complete cellular coverage of WAVE-enabled vehicles seems a reasonable assumption. The problem then becomes that of determining which vehicles to patch. Indeed, simply patching all the vehicles through cellular network downloads can be uselessly expensive. If a rough estimation of the area and time at which the infection started is available, a *smart cellular-based patching* can be adopted, limiting the immunization to vehicles actually interested by the infection.

We thus propose a smart cellular-based patching based on our data-driven model of the worm epidemics. Namely, the model is exploited to determine the region within which the

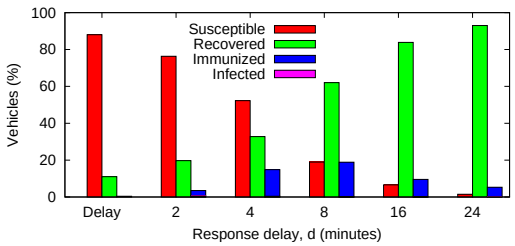

Fig. 8: Infected, susceptible, healed and immune node ratios.

worm may have spread within the time elapsed from the estimated infection instant. Then, only vehicles within such a region are immunized through the cellular network.

Fig. 8 shows the results of the smart patching in the Canton of Zurich reference scenario, considering that the epidemics starts in the center of Zurich during the peak traffic time. This is a worst-case scenario, since it provides the vehicular malware with maximum network connectivity and thus ideal self-propagation conditions. The response delay d is the estimated time between the patient zero appearance and the instant at which the smart patching is run: clearly, longer response delays imply the infection of larger portions of the road network. For each value of d, along the x axis, we report the number of vehicles belonging to different mutually exclusive categories: *susceptible* nodes were not infected and did not receive the patch, *recovered* nodes were infected and later recovered upon receiving the patch through the cellular network, *immunized* nodes were not infected yet received the patch, and *infected* nodes were infected but did not received the patch. Clearly, the goal of a smart cellular-based patching is to recover all infected nodes, leaving no infected vehicles and reducing the number of unnecessarily immunized nodes to a minimum.

The results show that for a response time of 2 minutes, the model correctly predicts the nodes to be patched, with a negligible number of unnecessarily immunized vehicles. As the response delay increases, the percentage of vehicles infected by the malware grows, leading to the necessity of patching a larger portion of the road traffic. Yet, our model allows to successfully patch all infected vehicles, with a percentage of unnecessarily immunized nodes that stays below 20% even in the worst case, when d = 16 minutes. More importantly, in all cases the percentage of nodes that remain infected after the smart patching is zero, proving once more the quality of our worm spreading model and its utility towards an efficient containment of malware outbreaks.

VII. Conclusions

We presented an extensive study of malware spreading in vehicular networks through WAVE V2V communication. Our simulative analysis outlined the high level of danger of vehicular worms, that are shown to be able to spread through very large areas, infecting tens of thousands of vehicles, in a few tens minutes at most. We found that the high mobility of vehicular nodes and the elevate number of short-lived V2V contacts they generate are the key reason behind such a result. We then presented a simple yet very effective data-driven model of the worm propagation process, and leveraged it for the smart patching of infected vehicles through cellular communication.

Acknowledgment

This work is partially supported by grants TIN2010-21378-C02-01 and SGR2009-1167.

References

[1] P. Ferrie, P. Szor, R. Stanev, R. Mouritzen, *"Security responses: Symbos.cabir"*, Symantec Corporation, 2004.
[2] J. Kleinberg, *"The wireless epidemic"*, Nature, 449, 2007
[3] E. Anderson, K. Eustice, S. Markstrum, M. Hanson, P. Reiher, *"Mobile Contagion: Simulation of Infection & Defense"*, Symp. on Measurement, Modelling, and Simulation of Malware, Monterey, USA, 2005.
[4] J. Su, K.W. Chan, A.G. Miklas, K. Po, A. Akhavan, S. Saroiu, E. de Lara, A. Goel, *"A Preliminary Investigation of Worm Infections in a Bluetooth Environment"*, ACM WORM, Fairfax, VA, USA, 2006.
[5] P. Akritidis, C.W. Yung, V.T. Lam, S. Sidiroglou, K.G. Anagnostakis, *"Proximity Breeds Danger: Emerging Threats in Metro-area Wireless Networks"*, USENIX Security, Boston, MA, USA, 2007.
[6] P. Wang, M.C. Gonzalez, C.A. Hidalgo, A.-L.Barabasi, *"Understanding the spreading patterns of mobile phones viruses"*, Science 324, 2009.
[7] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, *"Experimental Security Analysis of a Modern Automobile"*, IEEE S&P, Oakland, CA, 2010.
[8] I. Rouf, R. Miller, H. Mustafa, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe, I. Seskar, *"Security and Privacy Vulnerabilities of In-Car Wireless Networks: A Tire Pressure Monitoring System Case Study"*, USENIX Security, Washington, DC, USA, 2010.
[9] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, T. Kohno, *"Comprehensive Experimental Analyses of Automotive Attack Surfaces"*, USENIX Security, 2011.
[10] J.W. Mickens, B.D. Noble, *"Modeling Epidemic Spreading in Mobile Environments"*, ACM WiSe, Cologne, Germany, 2005.
[11] S.A. Khayam, H. Radha, *"Analyzing the Spread of Active Worms over VANET"*, ACM VANET, Philadelphia, PA, 2004.
[12] M. Nekovee, *"Modeling the Spread of Worm Epidemics in Vehicular Ad HocNetworks"*, IEEE VTC-Spring, Melbourne, 2006.
[13] L. Cheng, R. Shakya, *"VANET Worm Spreading from Traffic Modeling"*, IEEE RWS, New Orleans, LA, USA, 2010.
[14] M. Gramaglia, P. Serrano, J.A. Hernandez, M. Calderon, C.J. Bernardos, *"New Insights from the Analysis of Free Flow Vehicular Traffic in Highways"*, IEEE WoWMoM, Lucca, Italy, 2011.
[15] L. Cheng, S. Panichpapiboon, *"Effects of intervehicle spacing distributions on connectivity of VANET: a case study from measured highway traffic"*, IEEE Communications Magazine, 50(10), 2012.
[16] B. Hoh, M. Gruteser, *"Computer Ecology: Responding to Mobile Worms with Location-Based Quarantine Boundaries"*, WSPWN, Miami, FL, USA, 2006.
[17] Y.P. Fallah, C.-L. Huang, R. Sengupta, H. Krishnan, *"Analysis of Information Dissemination in Vehicular Ad-Hoc Networks with Application to Cooperative Vehicle Safety Systems"*, IEEE Trans. Vehicular Technology, 60(1), 2011.
[18] A. Agarwal, D. Starobinski, T.D.C. Little, *"Analytical Model for Message Propagation in Delay Tolerant Vehicular Ad Hoc Networks"*, IEEE VTC-Spring, Singapore, 2008.
[19] E. Baccelli, P. Jacquet, B. Mans, G. Rodolakis, *"Information propagation speed in bidirectional vehicular delay tolerant networks"*, IEEE Infocom, Shanghai, China, 2011.
[20] H. Wu, R.M. Fujimoto, G.F. Riley, M. Hunter, *"Spatial Propagation of Information in Vehicular Networks"*, IEEE Trans. Vehicular Technology, 58(1), 2009.
[21] Z. Zhang, G. Mao, B.D.O. Anderson, On the Information Propagation Process in Mobile Vehicular Ad Hoc Networks IEEE Trans. Vehicular Technology, 60(5), 2011.
[22] M. Fiore, *"Vehicular Mobility Models"*, in S. Olariu, M. Weigle (Editors), Vehicular Networks: from Theory to Practice, Chapman and Hall/CRC, 2009.
[23] N. Weaver, V. Paxson, S. Staniford, R. Cunningham, *"A Taxonomy of Computer Worms"*, ACM WORM, Washington, DC, USA, 2003.
[24] B. Raney, N. Cetin, A. Völlmy, M. Vrtic, K. Axhausen, K. Nagel, *"An agent-based microsimulation model of Swiss travel: First results"*, Networks and Spatial Economics, 3(1), 2003.