

Safe Driving in LA: Report from the Greatest Intervehicular Accident Detection Test Ever

Gustavo Marfia, Marco Rocchetti, Alessandro Amoroso, and Giovanni Pau

Abstract—The UN Economic Commission’s Statistics of Road Traffic Accidents report of 2011 shows that every year, about 150 000 human beings lose their lives on the roads of the western world. Although it is a common belief that this figure could shrink with the use of new sensor and communication technologies, unfortunately, none such systems have hit the road to date. Ideally, if such technologies were put into place, vehicles could be part of a vehicular ad hoc network (VANET) capable of spreading relevant information about dangerous events (e.g., car accidents) to all approaching drivers. However, all this is mainly supported by simulation studies, as no practical results have been published to date, revealing the effective performances of such systems at work. In this paper, we fill this gap, presenting a detailed description of the greatest experiments (a few thousand throughout the streets of Los Angeles), to date, ever performed with an accident warning system specifically devised for highway scenarios. In particular, among all the possible candidate schemes, we ran a few thousand experiments with the accident warning system algorithm that was proven to be optimal in terms of bandwidth usage and covered distance in realistic scenarios. Our experiments confirm what has been observed before in theory and simulation, i.e., the use of such a system can reduce, by as much as 40%, the amount of vehicles involved in highway pileups.

Index Terms—Design for experiments, prototypes, vehicle safety, vehicular ad hoc network (VANET) testbed, vehicular and wireless technologies.

I. INTRODUCTION

THE SOCIETAL costs of car accidents in Europe, in year 2009, amounted to \$160 billion according to the Harmonised European Approaches for Transport Costing project commissioned by the European Union (EU) [1]. Such value only partially captures the plethora of problems that have their origin in such type of events, as it is rather difficult to quantify, for example, the costs that are imputable to the emotional distress symptoms or to the increased pollution levels that usually follow any motor vehicle accident.

Continuing with a policy that recognizes the negative impact that car accidents have on society, in July 2011, the Transport

Manuscript received January 30, 2012 revised May 7, 2012 and August 31, 2012; accepted October 17, 2012. Date of publication October 24, 2012; date of current version February 12, 2013. This work was supported by FIRB DAMASCO, PRIN ALTERNET, NSF/BBN GENI projects, and TOYOTA InfoTechnology Center, USA. The review of this paper was coordinated by Dr. P. Lin.

G. Marfia is with the Computer Science Department, University of Bologna, 40126 Bologna, Italy, and also with the University of California, Los Angeles, CA 90095 USA (e-mail: marfia@cs.unibo.it).

M. Rocchetti and A. Amoroso are with the Computer Science Department, University of Bologna, 40126 Bologna, Italy (e-mail: amoroso@cs.unibo.it; roccetti@cs.unibo.it).

G. Pau is with the Computer Science Department, University of California, Los Angeles, CA 90095 USA (e-mail: gpau@cs.ucla.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TVT.2012.2226484

Committee of the EU approved an ambitious initiative that aims at halving the number of road accident victims by 2020 [2]. The path that leads to such aim certainly goes through the use of advanced technologies and, in particular, through the adoption of those that can prevent vehicles from being involved in accidents rather than those that can simply improve their general outcome. In fact, until now, most in-car solutions have focused on reducing the damage suffered by passengers and hence ameliorate the general consequences of an accident instead of seeking for new ways apt to totally avoid such type of events. Now, as cars have evidently become more resistant to impacts, drivers more careful, and roads safer, it is natural to expect that less space for improvement is left to any traditional countermeasure.

In this application space, one of the most promising solutions lies in the realm of communication technologies and is based on the use of vehicular ad hoc networks (VANETs). When an accident occurs, involved vehicles can alert all approaching vehicles within a distance of a few hundreds of meters by sending a short-range alert message. This message can then recursively percolate throughout the platoon of vehicles that are traveling in the direction of the accident. Cars, on receiving such message, can smoothly stop and avoid being involved in the pileup. This scenario, as it is, appears simple and trivial to implement, but has been under investigation for the past decade, as it presents at least two important challenges. The first, with no doubt, regards the installation of a set of devices that compose the safety equipment on all vehicles: a Global Positioning System (GPS) receiver, an accelerometer sensor, and an 802.11 wireless interface. The second regards the adoption of the best alert distribution algorithm, whose requirement is to be reliable and fast, as it should be able to deliver alert information in real time. As such, the desired algorithm is capable of successfully dealing with 1) the shared nature of the 802.11 channels, avoiding falling into the well-known broadcast storm problem, and 2) the challenges of mobile wireless channels, which vary in time and space because of the fast and slow fading phenomena [3].

Now, although highway safety systems based on VANETs have been thoroughly studied in the past years, the almost totality of results has been obtained by means of computer simulation, with little or no experimentation performed on real-road scenarios. This biased research approach is certainly dictated by the fact that the resources required to perform realistic accident warning system tests are beyond the reach of most, if not all, research teams. In November 2011, the main car manufacturers (e.g., Daimler, GM, Ford, Toyota, etc.) performed a public demonstration in Orlando, FL, without, however, publishing any of the findings obtained from such a testing campaign [4].

The aim of this paper is to fill such gap, reporting on the results drawn from the most extensive accident warning system test performed to date, to the best of our knowledge, on the streets and highways of Los Angeles, in August 2011. The tests that have been conducted in Los Angeles involved two research teams, i.e., one from the University of Bologna, and one from the University of California, Los Angeles (UCLA), and entailed performing one entire month of laboratory and on-the-road experiments, for a total financial commitment that exceeded \$60 000, invested in equipment (including car leases) and software development costs. To the best of our knowledge, this experiment campaign is the most extensive, among those whose results have been made available to date to the scientific community, in terms of both 1) the number of kilometers driven by a VANET testbed on a highway scenario and 2) the number of accident warning system tests performed in reality (over 10 000 tests during that month). Evidence of the importance of this campaign is witnessed also by the widespread international media attention that it received (e.g., Wall Street Journal, BBC, Discovery News, etc.), where also a team composed of journalists, and not only scientists or researchers, followed one of the experiments held on the road in Los Angeles [5]–[8].

Compared with the work published in [9], where a small subset of our assessments have been first analyzed, here, we provide a thorough description of the tests that have been conducted on the road and of their results, this paper resulting novel in all of its parts. Differently from the tests performed more recently in Orlando by the world's premium automakers, instead, we produced publicly available data that can generally benefit the scientific community.

To do all this, we implemented the accident warning system broadcast algorithm that has been proven to be the optimal one [15], [16]. Differently from all previous proposals, this algorithm takes into account the fact that the vehicles' transmission ranges vary in time and that wireless links can be anything but symmetric, presenting situations where between two vehicles, only one can receive the messages transmitted by the other. With this in mind, in [15], it is individuated that each vehicle should choose as the best vehicle to retransmit a given alert message at each hop the farthest spanning relay, i.e., the relay that can retransmit the message the farthest away, providing analytic and simulative proof that this is the best possible strategy in a realistic situation.

Summarizing, the main contributions of this paper are the following: 1) the dissemination of experimental results drawn from the largest intervehicular accident warning system experiments ever performed in a live setting; 2) a confirmation, drawn from experimental data, of the theoretical and simulative results accomplished by the accident warning system at test; and 3) further evidence that the deployment of such system can reduce as much as 40% the number of vehicles involved in a pileup.

The remainder of this paper is organized as follows. In Section II, we discuss related work, whereas in Section III, we provide a brief description of the tested accident warning system. In Section IV, we sketch the experimental scenario, whereas in Section V, we comment on experiment results. We finally conclude with Section VI.

II. RELATED WORK

Since the most important contribution of this paper lies in the construction of a realistic testbed for the assessment of a VANET-based accident warning system, we begin this section by illustrating the most relevant experiments performed, to date, with real VANETs on the road. We then move on to describe the stream of work in the domain of VANET-based accident warning systems that fall closest to ours. In fact, a wealth of research has been carried out in the past decade in the space of VANET-based applications and safety systems for urban and highway scenarios [17]–[35], and here, we will be limited to describing those approaches that are most relevant to the one that is at test [15].

While a large share of VANET research has addressed, with no doubt, the simulation of such networks for safety purposes, the same cannot be said regarding their experimentation on the road. Realistic experiments with accident warning systems on highways have been for long impeded by the magnitude of the testbeds required to perform relevant experiments. Hence, we will here concentrate on those research efforts that have obtained most from the limited available vehicular resources. The authors of [26] carried out a prominent VANET multi-hop experimentation utilizing 802.11b interfaces installed on four different vehicles and running the optimized link state routing protocol to route communication messages between those vehicles. In their work, the authors present the assessment of different performance metrics (i.e., packet loss and delay) obtained while conducting VANET under evaluation through different scenarios (e.g., urban, highway). A similar approach can be found in [27], where the authors assessed their vehicular testbed in static and mobile one- and three-hop scenarios, and in [37], where authors add to these basic performance metrics an evaluation of the packet drop rate as a function of the distance between two moving vehicles. Clearly, the results of all these works are jeopardized by the fact that only a very limited amount of hops (i.e., at most three) can be traveled by any message throughout their testbeds. In our work, we broke such limit, exploiting a novel methodology that allowed us to reproduce the conditions that such systems can face in reality [17], [18]. In particular, we drove a platoon of four vehicles through city roads and highways and required the system that was running on the leading vehicle to send a message periodically. When the message reached the last vehicle of the platoon directly or through the intervention of one or both of the available relays, the message bounced back, i.e., was retransmitted in the opposite direction. The same mechanism was adopted every time the leading (i.e., the front) vehicle received a copy of the message it originally sent. In brief, bouncing back the message at every edge of the network, we were able to infinitely extend the platoon of vehicles for experimentation purposes (i.e., a message traveled back and forth indefinitely, as long as it was not lost). Utilizing the described mechanism, we were able to reach the most relevant result of this paper, which is the observation of alert messages that travelled very far away (e.g., kilometers) from the original place where a hypothetical accident occurred.

The seminal idea that concerns our work in the area of VANET-based accident warning systems has been presented in

[38], wherein an algorithm is proposed that limits the number of forwarders by operating the selection of a unique relay at each hop. In particular, the rationale is that choosing the farthest relay among all the potential ones also minimizes the number of hops required to reach all approaching vehicles. In fact, in case an accident warning system is implemented following such philosophy, two main variables influence the delay experienced by an alert message: 1) the number of hops (i.e., relays) required by a message to reach all vehicles and 2) the number of forwarders that intervene every time an alert message is received by a group of vehicles. A similar idea is implemented in other protocols that have been proposed successively [39]–[41]. One limitation of all these approaches derives from the unrealistic assumption that transmission ranges are constant and unique in time. A recent variant of the mentioned schemes overcomes such limitation as it accounts for the variable nature of the wireless channel (e.g., devising an automatic transmission range estimator), without dealing with it in an optimal way [42]. Not even this last scheme, however, provides a methodology that efficiently answers to the asymmetries that practically appear with wireless links (e.g., vehicle *A* can hear vehicle *B* but not *vice versa*).

In summary, to the best of our knowledge, 1) no extensive evaluation work has been performed on a real VANET-based safety system for the simple reason that putting together the vehicular resources that would be needed for such experimentation is costly, and 2) no accident warning system, with the sole exception of [15], can optimally deal with the challenges that wireless links pose to accident warning system schemes. In the following, we describe more in detail the motivations that led us to select [15] to manage the propagation of alert messages throughout our testbed.

III. ON AN OPTIMAL ACCIDENT ALERT SYSTEM: A SUMMARY

Rather than here providing a detailed description, which the interested reader can find published in [15], we supply, in the following, motivated evidence that explains why, among all the possible candidates anticipated in Section II, we selected just this mechanism to implement the accident warning system subjected to our experiments.

In fact, the algorithm, described in [15], widely demonstrated, both in simulation and in theory, that it can optimally deal with time-variant wireless channels and asymmetric transmission as it can disseminate an alert message throughout a platoon of vehicles requiring the minimum possible number of transmissions also when such situations arise. For this fundamental reason, we implemented [15] for testing as it exhibited the best characteristics in terms of expected performance in practical scenarios.

Now, the primary difference between the scheme published in [15] and all the other algorithms presented before resides in the fact that this algorithm can guarantee, whenever a route between a source vehicle and a destination vehicle is available, the delivery of an alert message along the minimum hop path, also in realistic scenarios. Such characteristic is not guaranteed by all the schemes that are described in Section II, as they

assume transmission ranges to be all equal and constant at all times, and they disregard the fact that communication links can be anything but symmetric. In practice, as transmission ranges can strongly vary with surrounding conditions (e.g., obstacles), the transmission range of the farthest relay may become very short. In such situation, there might exist another vehicle that is closer, for example, but that can relay an alert message much farther away. Hence, in general, at each hop, a vehicle should choose as its next relay not the farthest one but the farthest spanning relay, i.e., the one that can relay a message the farthest away in a given direction.

Differently from all those approaches where a vehicle chooses as the next relay the one that happens to be the farthest, the scheme proposed in [15] 1) constantly estimates the transmission ranges and the positions of all neighboring vehicles and 2) identifies as the next relay the farthest spanning relay, i.e., the relay that can retransmit that given alert message farthest away. In [15], proof is provided that demonstrates that choosing the farthest spanning relay at each hop provides a global optimal solution to the problem of finding the minimum set of relays required to broadcast a message throughout a platoon of vehicles.

The problem, at this point, becomes that of finding how each vehicle can identify its farthest spanning relay among the vehicles that follow it. This is performed by implementing a distributed software component termed Oracle service, which runs within the accident warning system of each vehicle. The role of the Oracle service is twofold: 1) It periodically exchanges control messages with the Oracle services that run on neighboring vehicles to discover their positions and transmission ranges and to resolve any asymmetric communications (i.e., discover any potential relay that cannot be heard), and 2) it provides to the accident warning system a list of relays ordered by how far they can retransmit a message. In such list, termed Out list, the topmost entry is the farthest spanning relay of that vehicle.

Finally, in [15], a mechanism is described that can be put to good use by a relay to designate the next relay. This is simply done by inserting the Out list provided by the Oracle service in the transmitted alert message. On receiving an alert message, a given vehicle *V* 1) checks whether its ID is in the list, and 2) if it is in the list creates a new alert message containing its own Out list and waits for a time proportional to its position in the list before transmitting. If *V*, before sending the relay message, receives another copy of the same message, it aborts the transmission, as evidently, another vehicle with a lower position on the list already relayed that message.

IV. ARCHITECTING THE GREATEST EXPERIMENT EVER

The greatest VANET-based accident warning system experiments that have ever been conducted, to this date, took place in August 2011 in Los Angeles, as a collaboration between two research groups: one based at the University of Bologna, Italy, and one at the UCLA.

In the remainder of this section, we provide information regarding the magnitude of the experiments that have been performed, details on the testbed methodology that we exploited,

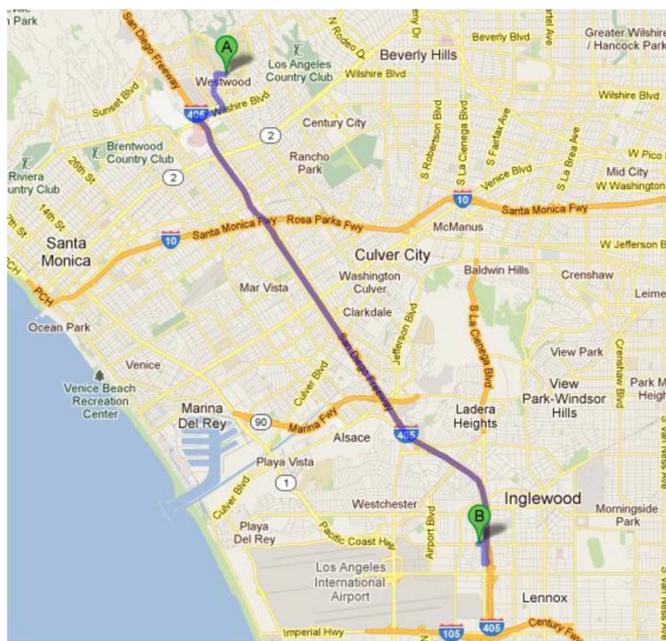


Fig. 1. Exemplar urban scenario.

and, finally, considerations regarding the hardware/software architecture of the equipment installed on each vehicle.

We begin by mentioning that the tests were performed on the UCLA campus and throughout the Los Angeles city area and included the traversal of all types of streets (i.e., almost a hundred, including illustrious streets such as Wilshire Blvd., Sunset Blvd., National Blvd., S. Monica Blvd., Sepulveda Blvd., Melrose Ave., Mulholland Dr. and I-405), experiencing heterogeneous traffic conditions, while recording logs and assessing the system. These road experiments lasted, in total, a week and comprise over 1000 km of traversed roads and over 10 000 fictitious dangerous situations advertised by the system at work. One single experiment was run every 10 s while the platoon of vehicles moved, as this value was regarded as sufficient to exceed the maximum lifetime of an alert message in an extreme communication scenario such as a vehicular one. In the following, we will concentrate on a large subset of such data, slightly more than 2000 experiments, which comprise all of the most relevant situations encountered while the system was at work, while the platoon was moving on the road.

In fact, we ran our experiments in a number of different road scenarios, which include urban and highway scenarios, with different types of vehicles, which consisted of both sedans and vans (a few pictures of our vehicles can be found in [8]). Although the accident warning system at test has been conceived for highway situations, it demonstrated on the road that it could easily adapt to urban ones, without requiring any modification. In particular, the experiments presented in this paper embraced three different scenarios: 1) densely urban; 2) suburban; and 3) rural. As an example, a densely urban scenario started at the UCLA campus and ended near the LAX airport, as depicted in Fig. 1. An exemplar suburban scenario was, instead, run around the UCLA campus, as it included Sunset Blvd., Hilgard Ave., and a portion of the I-405 highway that divides the UCLA campus from the Getty Museum area (see

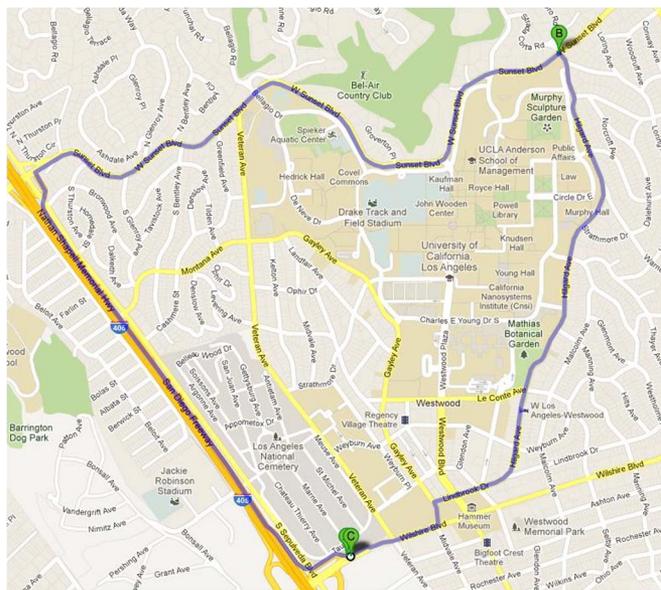


Fig. 2. Exemplar suburban scenario.

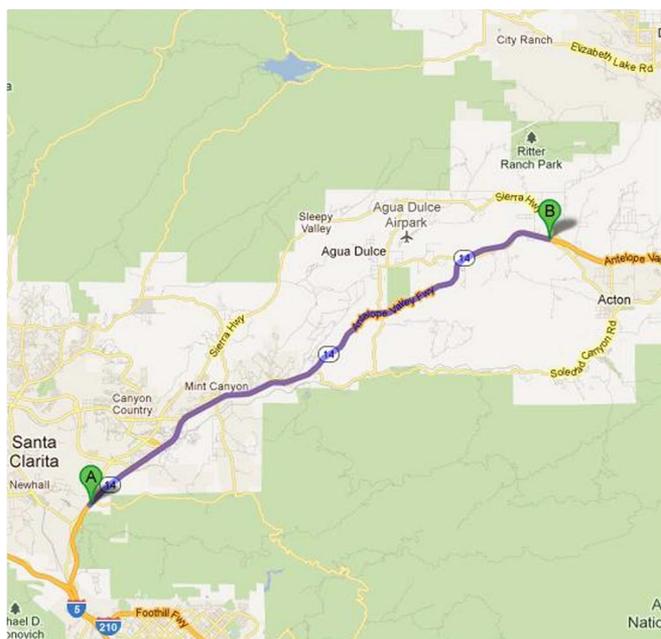


Fig. 3. Exemplar rural scenario.

Fig. 2). Finally, a sample rural experiment was located along the Antelope Valley Freeway, as shown in Fig. 3. In general, all test sites were chosen to assess the performance of our accident warning system as surroundings varied in terms of 1) traffic; 2) speeds; and 3) wireless channel conditions.

As to the methodology that was used to overcome the limitations imposed by the scarce availability of vehicular resources, it is worthwhile noticing that this amounts to a major problem that typically appears when experimenting with VANET-based accident alert systems. We already documented how this can be overcome using a so-called back and forth strategy described in [17] and extended in [18]. Although we do not want to repeat all the details provided in [17] and [18], this crucial issue needs to be addressed as follows. For example, if we send a message from the front vehicle to the back vehicle along a four-car-long

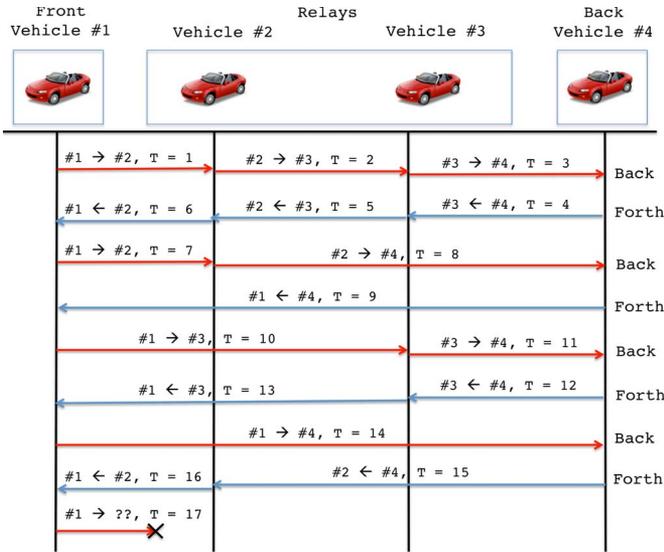


Fig. 4. Back and forth mechanism.

platoon, we could obtain at most three hops, yielding a distance of a few hundred meters. This clearly is unfit to test the performance of an alert system that has the potential to spread a warning kilometers away. Indeed, to overcome such a problem, we implemented a new experimentation procedure, as sketched in Fig. 4. In practice, periodically, the front vehicle sends an alert that keeps bouncing back every time any of the two edges (i.e., the back and front vehicles) of the platoon is reached. In other words, the front and back vehicles keep relaying the alert message they receive from a given direction along the opposite one. This is the reason why, when utilizing such methodology, the front and back vehicles are the only ones that cannot change relative positions within the platoon (i.e., the middle ones, instead, can overtake each other as often and as many times as they like). Although when utilizing such experimentation strategy only sparse network conditions are met, as, in our case, only four cars were part of the platoon at any time covering a stretch that ranged between approximately 30 (urban) and 2500 (highway) m, with this technique, an alert message can virtually travel for an indefinite number of hops. For example, Fig. 4 shows an alert message that traveled for 16 hops and was not received after the 17th transmission. Specifically, in Fig. 4, the value of T represents the number of hops traveled by the alert message so far, with the exception of the last hop (number seventeen), where the transmission finally failed. As the vehicles proceed along the road at speed, and an alert message keeps traveling between the edges of the platoon, potentially several times, this guarantees that different channel and environmental conditions are met, just as in a real deployment scenario. Clearly, in a real-world scenario, in most cases, a relay would have a wide selection of next-hop vehicles, whereas with our experimental system only two choices of next-hop vehicles were available at each retransmission step, yielding, hence, suboptimal message relay decisions and consequently suboptimal routes. Under this perspective, the results reported in this paper represent a conservative snapshot of the actual system working in reality. For future experimentation campaigns, such suboptimality, however, can be addressed by

Accident Warning System Architecture

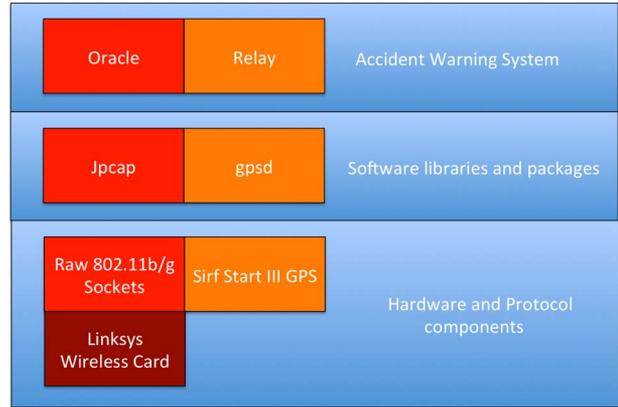


Fig. 5. Accident warning system hardware and software components.

increasing the number of vehicles involved in an experiment while still utilizing the back and forth strategy. Finally, for what concerns, instead, the channel and environmental variability that were experienced during our tests, the results presented in Section V-F will further witness that the adopted procedure is sound and that it can be consistently used to test the feasibility and performance of a highway accident warning system under realistic conditions.

We now sketch the hardware and software architecture of our implemented system, as shown in Fig. 5. Specifically, every car was equipped with one Dell laptop running Ubuntu Linux distribution version 11.04. Each laptop was connected to a Sirf Star III GPS receiver and to two high-gain antennas (i.e., one 5-dBi omnidirectional antenna by Hyperlink technologies and another 8-dBi omnidirectional antenna by l-com mounted at the center of the roof of each car [43]) via a Ubiquity Networks SRC PCMCIA/Cardbus 802.11g wireless card adapter exhibiting a 300-mW transmission power and a -93-dBm receiver sensitivity [44]. Our software system relied on two fundamental components: 1) *Jpcap* [45], a Java library that supports low-level networking operations, and 2) *gpsd* [46], a GPS monitor daemon. *Jpcap*, in particular, supports the use of raw sockets, which permitted us to bypass the overhead introduced by the IP layer. Hence, all messages, both Oracle messages and alert messages, were sent as broadcast layer-2 frames and were received and processed by all the other vehicles in the platoon. This means that vehicles sent messages at a fixed base rate as Auto Rate Fallback is disabled for broadcast frame transmissions, which was equal to 6 Mb/s under our particular configuration. The *gpsd* utility, instead, was utilized to read and translate the serial GPS stream in the JSON lightweight data interchange format, making it available through a local socket. For a more clear understanding, we summarized all of the described information in Table I.

Now, to conclude this section, it is worthwhile to spend a few words on how the results that have been observed utilizing the described setting can be put in relation with those that would have been achieved when adopting an accident warning system based on VANET 802.11p technologies. Any performance gap, in favor of any of the two telecommunication systems, heavily depends on the transmission ranges reached when opting for

TABLE I
EXPERIMENT INFORMATION

Vehicle types	One van and three sedans
Driven km	1,000
Number of experimented accidents	10,000
Accident frequency	1 every 10 seconds
Front-back distance	min = 30 m, max = 2500 m
Wireless Interfaces	802.11g Ubiquity Networks SRC PCMCIA/Cardbus
Wireless transmission power	300 mW
Wireless receiver sensitivity	-93 dBm
GPS receiver	Sirf Star III
Antennas	Omnidirectional 5 dBi gain by Hyperlink technologies and 8 dBi gain by I-com
Antennas position	Vehicle roof center
Transmission rate	6 Mbps
Auto Rate Fallback	Disabled
Operating System	Ubuntu Linux distribution v. 11.04
Employed software libraries	<i>gpsd</i> and <i>Jpcap</i>

any of the two solutions. To provide such value for 802.11p, we resorted to scientific literature, as well as to the standard. Interestingly, this is an ambiguous question whose answer is not clear, yet, to this date. In fact, when resorting to literature, we find that different studies experienced different transmission ranges, from as low as 70 m to almost 500 m [10]–[12]. Such uncertainty is confirmed by standardization entities, which have not yet finally determined how dedicated short range communication channels will be utilized [13]. Now, an indication of how far safety messages could be heard is given by governmental institutions, which report that they could reach as far as 1 km [14]. This said, how our results can be interpreted clearly depends on how 802.11p will be employed in its final version for vehicular safety applications. How far our messages traveled is clear in Fig. 11 (the ratio of lost messages exceeded the 20% percentage above 300 m). For what concerns a comparison with 802.11p, instead, if safety messages will be able to travel as far as 1 km, as the standardization and governmental documents that we cite report, we can conclude that with this study we took a conservative path, putting ourselves in a worst-case scenario.

V. EXPERIMENTAL OUTCOMES

This section is organized as follows. We will first provide the assessment of the most important performance figures that characterize an accident warning system and then discuss in detail and confute, supplying further experimental evidence, the possible objections that could be raised regarding those results. Hence, step by step, we will provide data supporting the feasibility and practicality of the accident warning system at test. As a final result, we will show that the accident warning system presented in [15] meets the expectations and can reduce by approximately 40% the number of vehicles involved in a pileup, compared with situations where it is not adopted.

A. How Far Could an Alert Be Spread?

In Fig. 6, we summarize the distances that were reached by alert messages while the vehicular testbed was operating. In a relevant number of cases, as shown in the figure, alert messages reached distances that well exceeded those necessary to correctly operate (e.g., a few kilometers).

Not only the sole distance but also the number of times alert messages were successfully relayed represent a relevant

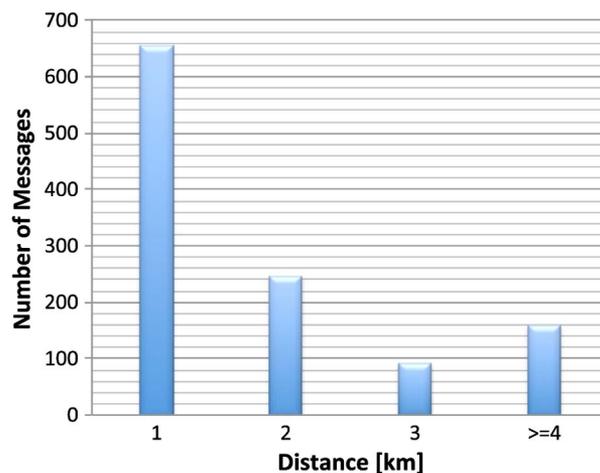


Fig. 6. Overall distribution of the distances reached by alert messages (one of them travelled as far as 36 km).

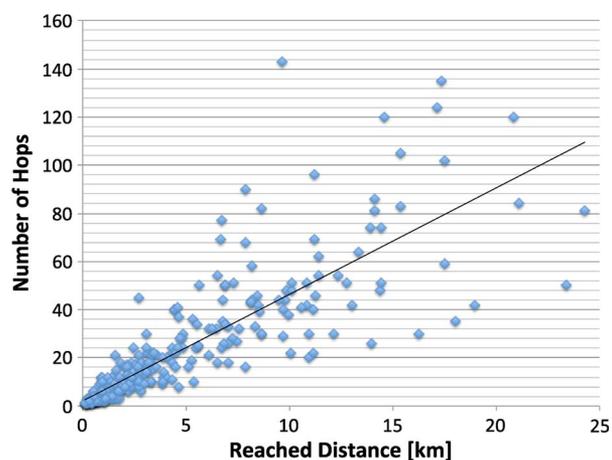


Fig. 7. Number of hops versus travelled distance when traveling at a speed above 54 km/h.

performance metric of a highway accident warning system. For this reason, we plotted both figures of merit in Fig. 7, whereas the platoon was moving at at least 54 km/h (at speed, as in a real case). The reached distance is computed as the total distance covered while a message travels back and forth, between the front and back vehicles. Hence, for example, if an alert message traveled 200 m, experiencing first a single transmission from the front to the back vehicle, and then another one from the

TABLE II
AVERAGE NUMBER OF HOPS PER REACHED DISTANCE WHILE VEHICLES MOVED AT MORE THAN 54 km/h

Distance [km]	1	2	3	4	5	6	7	8	9	10
Avg. #Hops	6.6	11	15.4	19.8	24.2	28.6	33	37.4	41.8	46.2

TABLE III
REACHED DISTANCE

Minimum Speed [km/h]	Average [m]	5-95 Percentile [m]
54	3,897	302-13,267
90	3,625	342-11,204

back vehicle to the front one, before being lost, this experiment would be represented in Fig. 7 with a point plotted in (2, 400). One of the situations where we observe the greatest reached distance (24.3 km) was obtained when forwarded 81 times, traveling an average of 300 m with each relay procedure. However, even farther distances than those shown in Fig. 7 were reached, as one message traveled as far as 36 km (for the sake of clarity, we set an upper limit on the x -axis to 25 km).

The points plotted in Fig. 7 exhibit an average traveled distance per hop of about 215 m, matching in general the approximate expected transmission range for 802.11g links. Also of interest is the regression line plotted in Fig. 7, which gives an approximate relation between the number of times an alert message was relayed and the distance, expressed in kilometers, that it reached. In Table II, we provide the average number of hops that applies to a given distance of interest, whereas our experiments were obtained from the regression line of Fig. 7.

B. After How Many Relays Did Alerts Vanish?

A second important question entails deriving for how long, on average, an alert message traveled before vanishing at some point. Clearly, an alert message can disappear, not being received by any of the designated relays, for a number of different reasons, which are generally related with the traversal of degraded wireless channels.

In Fig. 7, we provided the number of times a message was relayed when at least one hop was traveled (i.e., the message was successfully received by at least one vehicle). When at least one vehicle was reached, we found that the distance that was traveled by an alert message, on average, was high. In Table III, we report that for average platoon speeds that exceeded 54 km/h, the average reached distance was 3897 m (approximately 19 hops according to Table II), and this value did not significantly decrease when moving at higher speeds, being equal to 3625 m for speeds that exceeded 90 km/h.

When, instead, alert messages that immediately vanished (i.e., were not received by any relay) are also accounted for, we observe that the number of hops traveled by an alert message drops from an average slightly above 16, when the platoon moved at speeds confined between 54 and 90 km/h, to an average of about six hops, when the platoon exceeded a speed of 90 km/h (see Fig. 8). The reason for such drop, as will also be confirmed by the results discussed in the following sections, is due to the high amount of alert message that incurred in highly degraded wireless channels, especially when the sources and destination distances fell beyond a given value.

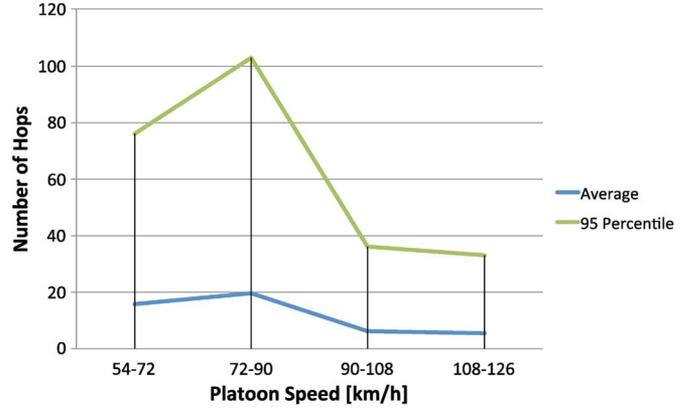


Fig. 8. Average number of times an alert message is relayed as a function of the average platoon speed.

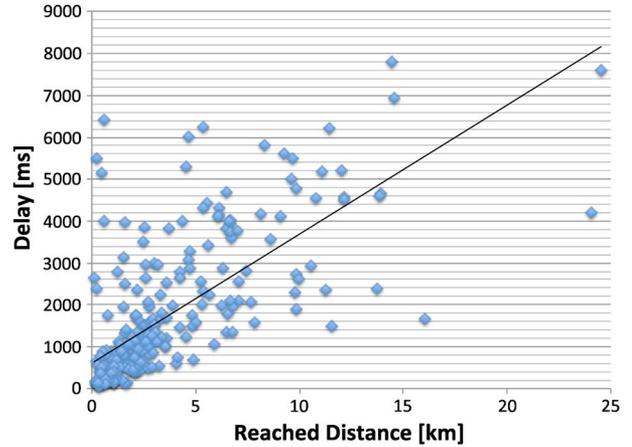


Fig. 9. Delay versus traveled distance when traveling at a speed above 54 km/h.

C. How Long Did it Take to Spread an Alert at a Given Distance?

To correctly analyze the delay incurred by an alert message when reaching a given distance, we first measured the overhead introduced by our software implementation. To do this, we measured during all of our on-the-road experiments at the front vehicle the time that elapsed between the reception of an alert message and the time at which that same message was successively relayed, finding that, during our on-the-road experiments, while traveling at a speed that exceeded 54 km/h, such a value amounted to 72 ms on average. This delay value, measured utilizing the laptop OS clock, was clearly dominated by our particular software implementation.

Now, keeping such information in mind, we are interested in knowing how long it took to reach a given distance, when traveling at a speed that exceeded the value of 54 km/h. Such information is shown in Fig. 9, where we provide a scatterplot where each single point denotes the reached distance and the time taken to reach such a distance.

TABLE IV
AVERAGE DELAY PER REACHED DISTANCE WHILE VEHICLES MOVED AT MORE THAN 54 km/h

Distance [km]	1	2	3	4	5	6	7	8	9	10
Avg. Delay [ms]	913.5	1221.6	1529.7	1837.8	2145.9	2454	2762.1	3070.2	3378.3	3686.4

TABLE V
AVERAGE NUMBER OF HOPS PER REACHED DISTANCE OVERALL EXPERIMENTS

Distance [km]	1	2	3	4	5	6	7	8	9	10
Avg. #Hops	13.6	18.4	23.2	28	32.8	37.6	42.4	47.2	52	56.8

TABLE VI
AVERAGE DELAY PER REACHED DISTANCE OVERALL EXPERIMENTS

Distance [km]	1	2	3	4	5	6	7	8	9	10
Avg. Delay [ms]	1199.4	1510.3	1821.2	2132.1	2443	2753.9	3064.8	3375.7	3686.6	3997.5

We again obtained and plotted in Fig. 9 a linear regression formula, this time computing the relation that best fits the samples that put in correspondence the traveled distance and the delay experienced by the alert messages. We then utilized such formula to compute the average delays per reached distance reported in Table IV.

The average delay experienced, per hop, can now be obtained, reminding that, according to Table III, on average, alert messages reached a distance of 3897 m (i.e., 3.9 km) when the platoon moved at speeds that exceeded 54 km/h. The average delay, hence, experienced by an alert message amounted to 1806 ms, also according to Table IV. The corresponding average number of times an alert message was relayed equals 19.3, in agreement with the values reported in Table II. This means that the average delay experienced per each relay procedure is equal to 92 ms. Recalling that the measured processing time (which included only the time spent within our software implementation at the front vehicle) amounted to 72 ms, we find a difference of 20 ms between these two values that cannot be justified by a cumulative contribution of Java Native Interface calls of *Jpcap*, OS scheduling procedures, and message transmission times. Leaving the solution of this puzzle to the next two sections, we here note that 72 ms would not be wasted in a final hardware implementation. Assuming, hence, an efficient hardware implementation processing time negligible, in such a situation, the average time required to receive an alert message with a single transmission would drop to 20 ms.

D. Is There a Correlation Between Vehicle Speed and Alert Relay Time?

Since the average time required for an alert message to be relayed exceeded our expectations (i.e., in simulation this value lied below 10 ms), we performed an additional analysis over all the data set that was collected without limiting to the samples that were recorded while the average speed of the platoon lied above 54 km/h.

The first difference that we encountered was that the average processing time measured at the front vehicle lowered to 68 ms; hence, we found a value that is 4 ms lower than the value computed before. We then reorganized the data and computed the regression formulas that best fit it according to the methodology used for the samples plotted in Figs. 7 and 9. In fact, the linear regression formulas also changed slightly and, consequently,

the average number of hops and the average delay per reached distance, as reported in Tables V and VI.

Visual inspection of the average values reported in Tables V and VI with those reported in Tables II and IV reveals that the former numbers increased. Nonetheless, when computing now the average time spent during each hop, we find that approximately 74 ms was required during each relay procedure. Finally, recalling that the measured processing time amounted to 68 ms, we now found that the time wasted in processing procedures at the front vehicle now amounted to the 92% of the total time and that 6 ms, on average, was realistically required for each relay procedure, hence, meeting our expectations.

E. Why Does Average Relay Time Increase While Traveling at Speed?

An interesting point now lies in understanding why, when traveling at an average speed that exceeded 54 km/h, the platoon experienced a higher alert message average relay delay. We found that the reason for such delay increase is due to the high amount of alert messages that vanished when the sources and destination distances fell beyond a given value. Here, we provide clues that corroborate such a fact.

The first is that when traveling at high speeds (i.e., on highways), the distance between vehicles was generally high. This fact clearly emerges from Fig. 10, which represents the speed of the platoon (gray line expressed in meters per second), and the relative distances between the front vehicle and the three other vehicles, as these evolved in time in the situation depicted in Fig. 3. During the time that the vehicles moved on the I-405 Freeway, the platoon stretch grew from 100 to 1200 m, with a consequent increase of the distance travelled by single transmissions.

Fig. 11, instead, provides the second fact, exhibiting the impact of source–destination distance on alert message losses: beyond 300 m, the loss probability reached relevant values, with a peak of almost 80% in the 500- to 600-m range. Such situation clearly jeopardizes the performance of any accident warning system operating in very sparse networks. In a case where four vehicles are lined, each at a distance of 300–400 m from the vehicle that is immediately in front, an alert message would be relayed six times with a probability of $(0.6)^6 = 0.046$, and, thus, almost never.

Now, the increased delay value per each relay procedure is easily explained considering that losses were frequent and that

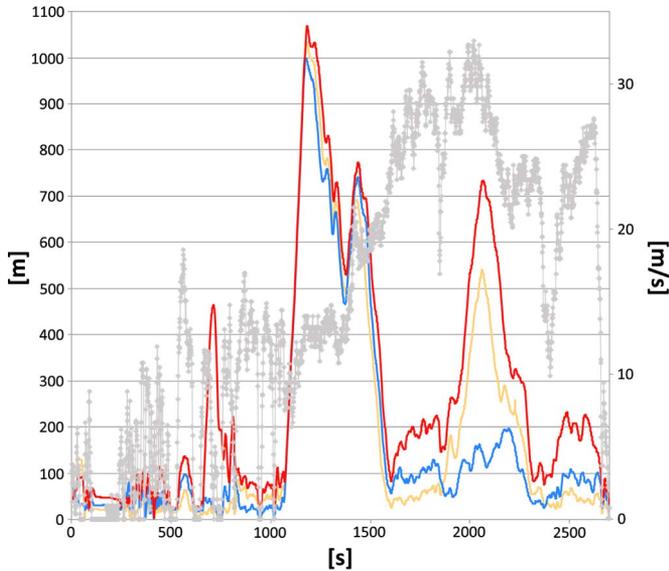


Fig. 10. Distances from the front vehicle (red, yellow, and blue lines) and average platoon speed (gray line) during the rural scenario experiment.

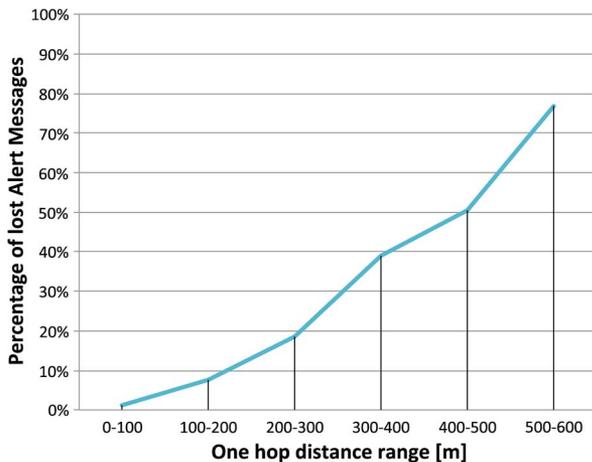


Fig. 11. Alert message loss rate as a function of the distance between the source and destination vehicles.

the system at test implements a redundancy mechanism where, if the transmission of the best relay is not heard by the second best relay, the second best relay transmits the message, which is assumed to be vanished, after waiting for 50 ms. The original mechanism includes up to 10 different relays to improve the robustness of the accident warning system. In our case, at most two other relays could intervene, with a total increase of at most 100 ms due to retransmissions.

This said, it is now easy to solve the puzzle left unanswered in Section V-C. The value of 72 ms, in fact, represents the time elapsed between the arrival and the departure of an alert message, as seen inside our accident warning system running on the front vehicle. Reminding here how the back-and-forth mechanism works, the front vehicle is the best relay at all times; thus, every time a message is received by the front vehicle, no retransmission timer is started. This means that Java Native Interface calls of *Jpcap*, OS scheduling procedures, and message transmission time contribute to those 20 ms as expected, but alert message retransmissions do as well. Taking now into account the delays reported in Sections V-C and D,

we observe that the 20 ms is heavily dominated by alert message retransmissions, which we here estimate amounting to at least $20 - 6 = 14$ ms, where 6 ms was the time computed for each relay procedure in the previous section.

Finally, the ground truth cannot be found without performing direct measurements of the delays that composed the total one that was experienced by an alert message during relaying operations. This operation has been performed recording the timestamps during the experimental campaign at all vehicles. In particular, we measured the contributions of the Java Native Interface calls of *Jpcap*, OS scheduling procedures, and message transmission times recording the timestamp right before a JNI call was operated and recording additional timestamps at all of the intermediate steps required before an alert message transited through the wireless interface. In addition, as a crosscheck, we captured the time instants at which that given alert message was first heard by other vehicles. Table VII reports the values that were measured from the experimental testbed for those messages that reached distances between 3 and 5 km (to be comparable with the values that we have so far indicated) while traveling at speed. These numbers roughly confirm those that have been obtained so far and reveal one additional interesting fact (considering that retransmission timeout times were set to 50 ms during the experimental campaign): When traveling at speed, one retransmission was required every three relay operations, on average.

Clearly, lower retransmission times potentially yield lower average delays, but on the other side, lower retransmission time values would have triggered unwanted retransmissions (and collisions), as the best relay could have been anticipated while still in the process of sending the received alert message. The value of 50 ms, hence, was set as the right compromise between these two needs. However, lower retransmission values can be experimented with more efficient implementations of our accident warning system.

F. Have Our Experiments Been Realistic, At Last?

The effective realism of the entire experimentation campaign lies upon the assumption that the hops that are traveled by an alert message between the vehicles of our testbed are representative of those that would be encountered in a practical situation. While it directly follows from the construction of the experiment that the wireless channel conditions that are experienced between distinct experiments radically differ (i.e., each experiment is characterized by a brand new accident warning message, which is sent every 10 s), as the platoon moved at speed, the same is not as obvious within a single experiment. We will address such issue here.

In a practical situation, we generally expect that the wireless channel encountered by an alert message can radically differ from transmission to transmission. This is intuitively true, as between different hops the sender and the receiver may be 1) traveling at radically different speeds and distances, 2) equipped with devices that exhibit diverse power and sensitivity characteristics, 3) experiencing increasing or no interference at all, 4) surrounded by reflecting objects (i.e., other cars, buildings, etc.), and 5) traveling in different or in the

TABLE VII
ALERT MESSAGE AVERAGE DELAY COMPONENTS PER VEHICLE WHEN TRAVELING AT SPEED

Our implementation [ms]	JNI call [ms]	OS scheduling [ms]	Transmission [ms]	Retransmission [ms]
73.4	< 1	4.7	< 1	15.2

same direction, also with respect to the propagation of an alert message. All these can be summarized by the fact that, in practice, we cannot expect the wireless channel encountered at one hop to be correlated with the wireless channel traversed at another, because vehicles are moving at speed [47]. Our scope, hence, is that of understanding to what extent the alert messages that traveled back and forth, from one edge to the other of the testbed platoon, found similar, or different, channel conditions when traversing a hop that was previously encountered. A reliable indicator of such similarity is given by the coherence time of the wireless channel between the vehicles of the platoon, i.e., the time interval over which the impulse response of those channels is considered correlated to their previous values [48].

Before, however, proceeding and resorting to the theory of wireless propagation, we would like to appeal to the principle of reality that lies behind any experimental work: our tests have been carried out on real roads with vehicles that moved at speed and with highly varying boundary conditions. This means that no two transmissions of alert messages encountered the same line-of-sight (LOS) and Non-LOS (NLOS) conditions. Hence, even when a source and a destination traveled exactly at the same speed (event that was rather unlikely), the relative positions of the surrounding objects rapidly changed, also rapidly changing propagation conditions.

Now, the average value of the channel coherence time can be found as the inverse of the Doppler shift, i.e., the change in frequency that an electromagnetic wave experiences as a result of a difference in speed between a given source and a given receiver. This value is simply found as $v_{s,r} \times f_0 / c$, where $v_{s,r}$ is the relative speed between the source and the receiver, f_0 is the carrier frequency of the transmitted signal, and c is the speed of light. At a differential speed of 1 m/s (i.e., 3.6 km/h) between a source and a receiver, the channel coherence time is approximately 125 ms for a carrier frequency of 2.4 GHz (i.e., the carrier frequency of employed 802.11g interfaces). At differential speeds of 2, 4, and 8 m/s, such value decreases to 62, 31, and 15 ms, respectively [49]. All of such values are confirmed in both theory and by previous experimental works, which also observe how channel coherence times are two to three times shorter in highway situations than in urban ones due to rapidly changing boundary conditions [50]–[53].

To further support these arguments, an additional set of experiments was performed with the scope of empirically estimating the longest channel coherence time that was experienced, thus in a worst case scenario for our experimental methodology. In particular, two vehicles in LOS were parked, in a static position, along the curb of Wilshire Blvd. One of the vehicles, the transmitting vehicle, steadily sent User Datagram Protocol packets matching the link capacity to the second vehicle, the receiver. The receiver could slowly perform small forward or backward movements (well below the speed of 1 m/s), just as when proceeding in line, never exceeding, however, a maximum distance of 5 m from its original position. In the meantime, intense traffic flowed through Wilshire, thus

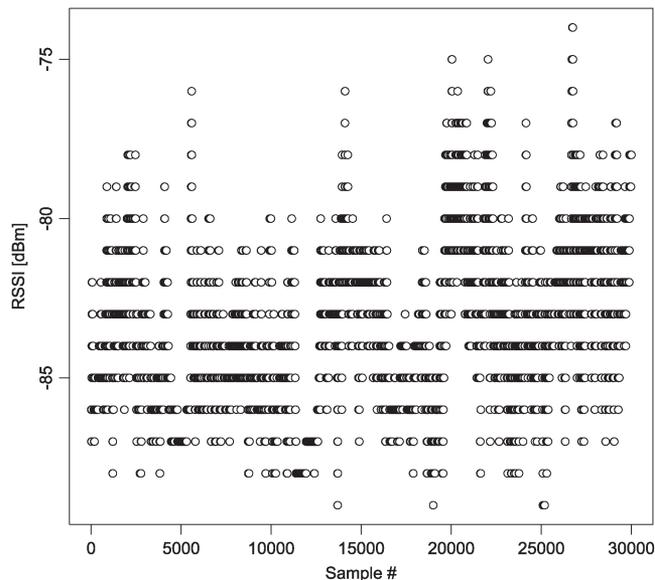


Fig. 12. RSSI values over a 7.5-min interval.

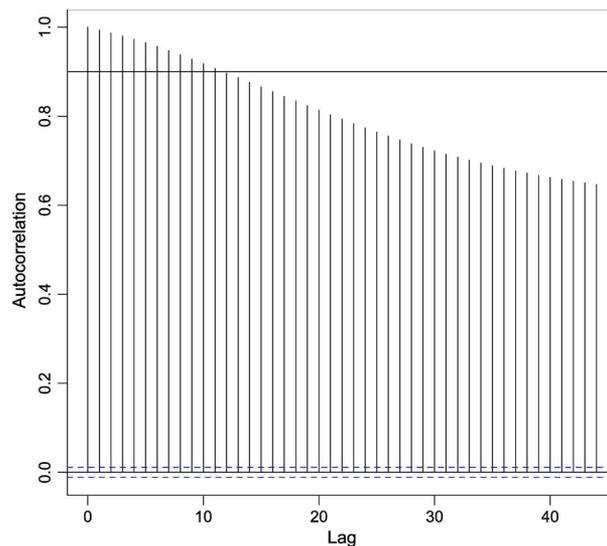


Fig. 13. Autocorrelation function of the RSSI plot shown in Fig. 12.

creating varying surrounding conditions. Ten of such experiments were performed, each lasting 15 minutes. During each experiment, the link received signal strength indication (RSSI) value was recorded every 15 ms at the receiver. In fact, the RSSI time series can be put to good use to empirically estimate the channel coherence time, as follows [54], [55]: the time instant where the autocorrelation function of the RSSI time series falls to 90% of its initial value. Fig. 12 shows the RSSI readings in the 7.5-min interval that exhibited the highest average autocorrelation. Fig. 13, instead, shows the autocorrelation function that refers to the same time interval. The aforementioned definition applied to the RSSI readings plotted in Fig. 12 returns a channel coherence time of about 180 ms (the horizontal line in correspondence of an autocorrelation value of

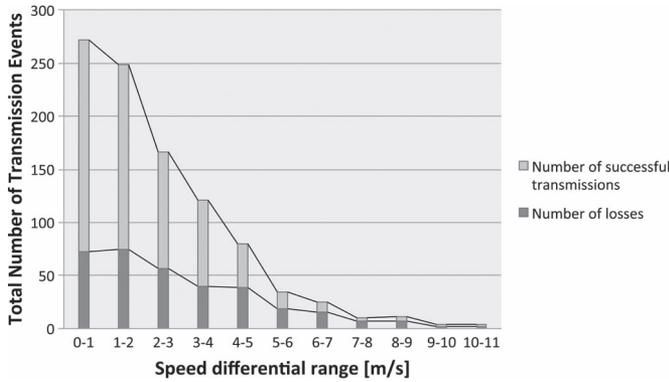


Fig. 14. Alert message loss rate as a function of the differential speed between the source and destination vehicles.

0.9 crosses the edge of 12th time lag in Fig. 13), agreeing with other works that performed similar analysis before in similar conditions [50]–[53].

Now, putting such information into context, we observe that the same source–destination pairs are typically repeated every time an alert message reaches the front vehicle after bouncing off the back vehicle. In fact, when, for example, an alert message sent by the front vehicle flows through a third vehicle, namely, V , to reach the back vehicle, a source–destination repetition will appear along the path travelled by that alert message when the front vehicle relays that same alert message again to V . Hence, the distribution of back and forth times (i.e., the round trip experienced by an alert message at the front vehicle) can indicate whether an alert message traversed the same source–destination pair experiencing the same channel conditions. Our results show that the average time elapsed between two subsequent transmissions of an alert message from the front vehicle amounts to 186 ms (the fifth percentile is equal to 104 ms and the overall minimum value to 52 ms). These values indicate that the chance that the same wireless channel conditions are met by a message during two subsequent loops is very low. To further support such argument, we analyzed the number of one-hop transmission events that occurred during all of our experiments as a function of the speed differential between a source and a destination pair, whereas the platoon moved at an average speed that exceeded 15 m/s (see Fig. 14). Observing the figure, the relative majority of transmissions occurred between a sender and a receiver that moved at a relative speed below 1 m/s. However, more interesting is the fact that the fraction of messages that got lost under given speed drift conditions kept quite constant up to a source–destination speed difference of 4 m/s (i.e., 26% between 0 and 1 m/s, 30% between 1 and 2 m/s, 34% between 2 and 3 m/s, and 33% between 3 and 4 m/s). Hence, for all those cases where a sufficient amount of experimental data is available, we found that the chance for an alert message of getting lost when the relative speed between the source and the destination was below 1 m/s (i.e., in a situation where the channel coherence is expected to be longer) was very similar to the chance of getting lost when the relative speed was instead higher. All this indicates, once more, that the experiment that has been carried out clearly resembles a practical case, as the channel conditions changed with the same modalities, just as in reality.

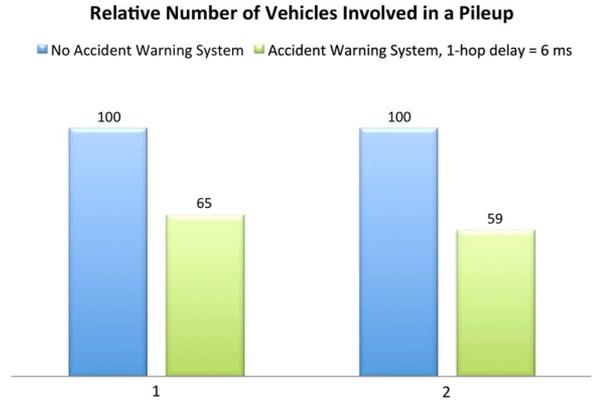


Fig. 15. Traffic scenarios: 1 = free flow traffic conditions, 2 = congested traffic conditions.

For the sake of completeness, we observe that, with our methodology, no transmission of alert messages between two vehicles that were traveling in opposite directions occurred. The worst-case scenario, from a wireless propagation point of view, is here represented by all those situations where two vehicles travel in the opposite directions at relative speeds that can be as high as 220 km/h (e.g., 110 km/h in the two respective directions). Although we believe that such type of event would not be as frequent as the event of a message being sent between two vehicles moving in the same direction, as only a few seconds, in fact, would be available for the discovery of a candidate relay and the successful transmission of an alert message to a vehicle that is traveling at a relative speed of 220 km/h, this case might as well occur. Interestingly, relevant previous work, which utilizes hardware settings that are very similar to ours, reveals the amount of data that can be exchanged between two vehicles that are moving in opposite directions on a highway [56]. The work that has been presented in the cited paper only focuses on the exchange of bulk data between two moving vehicles but can still be of interest to understand what contact time and what transfers of information can be achieved in such situation. In particular, in [56], average contact times that were slightly less than 15 s and average amounts of transferred data that were approximately equal to 8 MB were observed when utilizing packet sizes that lied between 1000 and 2304 B. A rough analysis of such numbers reveals that the transmission range conditions experienced in [56] are compatible with those observed during our experiments; in fact, vehicles traveling in opposite directions at a speed of 110 km/h would begin their transfer of information at a distance of $(110/3.6) * 7.5 \approx 230$ m from each other. Concluding, although we were not able to test our system in this last particular situation, the numbers that are available from relevant literature reveal that our accident warning system could be capable of also facing such adverse conditions, as alert messages weigh at most a few hundred of bytes, and the use of 802.11p, which is capable of achieving higher transmission ranges, and hence, contact times could not but improve its performance.

G. In Summary: How Many Vehicles Can Be Saved?

We are now able to conjecture how many vehicles could be saved by utilizing a system like the one that is described here. In

fact, we have found that, on average, the delay incurred by our system to deliver an alert message one hop away was confined between the values of 6 and 20 ms, depending on the particular wireless propagation scenario.

We simulated a vehicle crash on a three-lane highway and estimated how many vehicles, on average, could be involved with and without an intervehicular and also considering an ideal accident warning system, where all vehicles learn about the occurrence of an accident with no delay at all.

We considered two realistic traffic flow scenarios, a realistic pavement condition that corresponded to the dry conditions that we experienced in Los Angeles, realistic driver response times, and vehicle lengths as well. Vehicle speeds and related time-headway distributions are also realistic as drawn from a real measurement campaign [57]. Specifically, in the congestion-free scenario, vehicles move at 110 km/h and the time-headway distribution induces a linear density of 20 vehicles/km on a per lane basis. In the congested scenario, the speed is 40 km/h, and the linear density is 40 vehicles/km. The kinetic friction constant between the tires and the asphalt is set to 0.8 as the pavement is dry, as suggested in [58]. The drivers' response times are in the [0.75, 1.4] s range, and the vehicles' lengths are randomly drawn from the [3.5, 5] m interval according to [59] and [60]. The vehicle following model, used when no accident warning system is operating, entails that a driver brakes after a random response time since seeing the preceding vehicle braking.

We ran 100 simulations, whose results are as follows. Under free-flow conditions and without the use of any accident warning system, we find that the average number of vehicles involved in an accident equals 8.85. When, instead, our accident warning system is adopted, we find an average value equal to 5.76 when retransmitting an alert message with an average delay of 20 ms at a distance of 215 m. In the congested flow case, both values are lower and equal to 5.40 (no accident warning system) and 3.13 (20-ms delay). We plot such information in Fig. 15, normalizing to 100 the highest values pertaining to each respective experiment (i.e., 8.85 in the free-flow case and 5.40 in the congested case) and scaling the other accordingly within each experiment scenario. The leftmost group of bars, which correspond to the experiment labeled as 1, results from the free-flow scenario, whereas the rightmost group of bars, labeled as 2, results from a congested scenario. Within each of the two scenarios, the leftmost bar represents the base case of 100 vehicles involved in an accident when no accident warning system is deployed, whereas the fourth bar instead refers to those situations where alert messages experience a delay of 20 ms at each relay procedure. In summary, our simulation results, based on real delay values drawn from our experiments, definitely confirm that the use of an accident warning system like ours can potentially reduce the number of crashed vehicles by nearly 40%.

VI. CONCLUSION

In this paper, we have filled an important gap of VANET systems research, presenting a detailed description of the greatest experiments, to date, ever performed with an accident warning

system specifically devised for highway scenarios. In particular, we implemented and tested an accident warning system based on the multihop broadcast algorithm that has been proven to be optimal in terms of bandwidth usage and covered distance in realistic scenarios. We demonstrated, obtaining a confirmation drawn from experimental data of simulative and theoretical results, that such systems can really improve the conditions experienced on the road by drivers, as we collected further evidence that close to 40% of vehicles could be saved from being involved in a pileup when an accident occurs.

ACKNOWLEDGMENT

The authors wish to thank F. Bronzino for his help in carrying out the experiments and F. Callegati, M. D. De Amorim, M. Gerla, I. Seskar, and the anonymous reviewers that helped improve the quality of this paper.

REFERENCES

- [1] Accessed on Jan. 20, 2012. [Online]. Available: <http://heatco.ier.uni-stuttgart.de/>
- [2] Accessed on Jan. 20, 2012. [Online]. Available: <http://www.europarl.europa.eu>
- [3] M. Torrent-Moreno, D. Jiang, and H. Hartenstein, "Broadcast reception rates and effects of priority access in 802.11-based vehicular ad-hoc networks," in *Proc. 1st ACM Int. Workshop VANET*, 2004, pp. 10–18.
- [4] Accessed on Jan. 20, 2012. [Online]. Available: <http://www.itsworldcongress.org/>
- [5] Accessed on Jan. 20, 2012. [Online]. Available: <http://blogs.wsj.com/tech-europe/2011/07/13/car-to-car-communication-could-reduce-jams/>
- [6] Accessed on Jan. 20, 2012. [Online]. Available: <http://www.bbc.co.uk/news/technology-14125245>
- [7] Accessed on Jan. 20, 2012. [Online]. Available: <http://news.discovery.com/autos/accident-alert-app-taken-for-test-drive-110914.html>
- [8] Accessed on Jan. 20, 2012. [Online]. Available: <http://gadget.wired.it/news/applicazioni/2011/10/25/app-anti-incidente-ucla.html>
- [9] A. Amoroso, G. Marfia, M. Roccetti, and G. Pau, "To live and drive in L.A.: Measurements from a real intervehicular accident alert test," in *Proc. IEEE WCNCW*, 2012, pp. 328–332.
- [10] F. Martelli, M. E. Renda, G. Resta, and P. Santi, "A measurement-based study of beaconing performance in IEEE 802.11p vehicular networks," in *Proc. 31st IEEE INFOCOM*, 2012, pp. 1503–1511.
- [11] F. Bai, D. D. Stancil, and H. Krishnan, "Toward understanding characteristics of Dedicated Short Range Communications (DSRC) from a perspective of vehicular network engineers," in *Proc. 16th Annu. Int. Con. MobiCom*, 2010, pp. 329–340.
- [12] K. Hong, D. Xing, V. Rai, and J. Kenney, "Characterization of DSRC performance as a function of transmit power," in *Proc. 16th ACM Int. Workshop VANET*, New York, NY, 2009, pp. 63–68.
- [13] Soc. Automotive Eng. Int., Warrendale Marshall, PA, DSRC Implementation Guide, 2012.
- [14] M. McGurrin, "Vehicle information exchange needs for mobility applications," U.S. Dept. Transp., Res. Innovative Technol. Admin., Washington, DC, Final Rep. FHWA-JPO-12-021, Feb. 2012.
- [15] A. Amoroso, G. Marfia, and M. Roccetti, "Going realistic and optimal: A distributed multi-hop broadcast algorithm for vehicular safety," *Comput. Netw.*, vol. 55, no. 10, pp. 2504–2519, Jul. 2011.
- [16] A. Amoroso, G. Marfia, and M. Roccetti, "Optimal assessments in VANET: The Oracle," in *Proc. IFIP WD*, 2010, pp. 1–5.
- [17] A. Amoroso, G. Marfia, M. Roccetti, and G. Pau, "Creative testbeds for VANET research: A new methodology," in *Proc. IEEE CCNC*, 2012, pp. 477–481.
- [18] G. Marfia, M. Roccetti, A. Amoroso, M. Gerla, G. Pau, and J. H. Lim, "Cognitive cars: constructing a cognitive playground for VANET research testbeds," in *Proc. 4th Int. Conf. Cogn. Radio Adv. Spectr. Manag. (CogART)*, 2011, pp. 1–29.
- [19] I. Leontiadis, G. Marfia, D. Mack, G. Pau, C. Mascolo, and M. Gerla, "On the effectiveness of an opportunistic traffic management system for vehicular networks," *IEEE Trans. Intell. Transp. Syst.*, vol. 12, no. 4, pp. 1537–1548, Dec. 2011.

- [20] S. R. Cavalcanti, M. E. M. Campista, F. B. Abdesslem, L. H. M. K. Costa, M. D. de Amorim, and O. C. M. B. Duarte, "Veer: A trajectory-based peer selection algorithm for networks of vehicles," in *Proc. 9th IFIP Annu. Med-Hoc-Net Workshop*, 2010, pp. 1–8.
- [21] M. Gerla, B. Zhou, Y.-Z. Lee, F. Soldo, U. Lee, and G. Marfia, "Vehicular grid communications: The role of the Internet infrastructure," in *Proc. 2nd Annu. Int. Workshop Wireless Internet*, 2006, pp. 112–120.
- [22] S. Savasta, M. Pini, and G. Marfia, "Performance assessment of a commercial GPS receiver for networking applications," in *Proc. IEEE 5th CCNC*, 2008, pp. 613–617.
- [23] L. Bononi and M. Di Felice, "A cross-layered MAC and clustering scheme for efficient broadcast in VANETs," in *Proc. IEEE 2nd Int. Workshop MOVENET*, 2006, pp. 190–197.
- [24] C. E. Palazzi, F. Pezzoni, and P. M. Ruiz, "Delay-bounded data gathering in urban vehicular sensor networks," *Pervasive Mobile Comput.*, vol. 8, no. 2, pp. 180–193, Apr. 2012.
- [25] J. J. Blum and A. Eskandarian, "A reliable link-layer protocol for robust and scalable intervehicle communications," *IEEE Trans. Intell. Transp. Syst.*, vol. 8, no. 1, pp. 4–13, Mar. 2007.
- [26] M. Moske, H. Fussler, H. Hartenstein, and W. Franz, "Performance measurements of a vehicular ad hoc network," in *Proc. IEEE 59th VTC*, 2004, pp. 2116–2120.
- [27] J. Santa, M. Tsukada, T. Ernst, and A. F. Gomez-Skarmeta, "Experimental analysis of multi-hop routing in vehicular ad-hoc networks," in *Proc. 5th Int. Conf. TridentCom Workshops*, 2009, pp. 1–8.
- [28] N. Wisitpongphan, O. K. Tonguz, J. S. Parikh, P. Mudalige, F. Bai, and V. Sadekar, "Broadcast storm mitigation techniques in vehicular ad hoc networks," *IEEE Wireless Commun.*, vol. 14, no. 6, pp. 84–94, Dec. 2007.
- [29] F. Ye, M. Adams, and S. Roy, "V2V wireless communication protocol for rear-end collision avoidance on highways," in *Proc. IEEE Int. Conf. Commun. Workshops*, 2008, pp. 375–379.
- [30] R. A. Santos, R. M. Edwards, and A. Edwards, "Cluster-based location routing algorithm for vehicle to vehicle communication," in *Proc. IEEE Radio Wireless Conf.*, 2004, pp. 39–42.
- [31] 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 *Proc. 1st Int. MobiSys Workshop MobiOpp Netw.*, 2007, pp. 47–51.
- [32] S. Ferretti and M. Roccetti, "Fast delivery of game events with an optimistic synchronization mechanism in massive multiplayer online games," in *Proc. ACM SIGCHI Int. Conf. Adv. Comput. Entertainment Technol.*, 2005, pp. 405–412.
- [33] C. E. Palazzi, S. Ferretti, M. Roccetti, G. Pau, and M. Gerla, "How do you quickly choreograph inter-vehicular communications? A fast vehicle-to-vehicle multi-hop broadcast algorithm, explained," in *Proc. IEEE 4th Annu. CCNC*, 2007, pp. 960–964.
- [34] F. Malandrino, C. Casetti, C.-F. Chiasserini, and M. Fiore, "Content downloading in vehicular networks: What really matters," in *Proc. IEEE INFOCOM*, 2011, pp. 426–430.
- [35] M. Cesana, L. Fratta, M. Gerla, E. Giordano, and G. Pau, "C-VeT the UCLA campus vehicular testbed: Integration of VANET and Mesh networks," in *Proc. EW Conf.*, 2010, pp. 689–695.
- [36] G. Marfia and M. Roccetti, "Vehicular congestion detection and short-term forecasting: A new model with results," *IEEE Trans. Veh. Technol.*, vol. 60, no. 7, pp. 2936–2948, Sep. 2011.
- [37] H. Venkataraman, A. d'Ussel, T. Corre, C. H. Muntean, and G.-M. Muntean, "Performance analysis of real-time multimedia transmission in 802.11p based multihop hybrid vehicular networks," in *Proc. 6th Int. Wireless Commun. Mobile Comput. Conf.*, 2010, pp. 1151–1155.
- [38] G. Korkmaz, E. Ekici, F. Ozguner, and U. Ozguner, "Urban multi-hop broadcast protocol for inter-vehicle communication systems," in *Proc. 1st ACM Int. Workshop Veh. Ad Hoc Netw.*, 2004, pp. 76–85.
- [39] A. Zanella, G. Pierobon, and S. Merlin, "On the limiting performance of broadcast algorithms over unidimensional ad-hoc radio networks," in *Proc. Wireless Personal Multimedia Commun.*, 2004, pp. 165–169.
- [40] E. Fasolo, R. Furiato, and A. Zanella, "Smart broadcast algorithm for inter-vehicular communication," in *Proc. Wireless Personal Multimedia Commun.*, 2005, pp. 1–5.
- [41] S. Biswas, R. Tatchikou, and F. Dion, "Vehicle-to-vehicle wireless communication protocols for enhancing highway traffic safety," *IEEE Commun. Mag.*, vol. 44, no. 1, pp. 74–82, Jan. 2006.
- [42] C. E. Palazzi, S. Ferretti, and M. Roccetti, "An inter-vehicular communication architecture for safety and entertainment," *IEEE Trans. Intell. Transp. Syst.*, vol. 11, no. 1, pp. 90–99, Mar. 2009.
- [43] S. Kaul, K. Ramachandran, P. Shankar, S. Oh, M. Gruteser, I. Seskar, and T. Nadeem, "Effect of antenna placement and diversity on vehicular network communications," in *Proc. IEEE Conf. Sensor, Mesh Ad Hoc Commun. Netw.*, 2007, pp. 112–121.
- [44] *Super Range Cardbus*, Ubiquiti Netw., Inc., San Jose, CA, Apr. 25, 2012. [Online]. Available: http://www.metrix.net/datasheets/src_datasheet.pdf
- [45] Accessed on Jan. 20, 2012. [Online]. Available: <http://netresearch.ics.uci.edu/kfujiji/Jpcap/doc/>
- [46] Accessed on Jan. 20, 2012. [Online]. Available: <http://gpsd.berlios.de/>
- [47] M. Jerbi, P. Marlier, and S. M. Senouci, "Experimental assessment of V2V and I2V communications," in *Proc. IEEE Int. Conf. Mobile Adhoc Sens. Syst.*, Pisa, Italy, 2007, pp. 1–6.
- [48] L. Cheng, B. E. Henty, R. Cooper, D. Stancil, and F. Bai, "A measurement study of time-scaled 802.11a waveforms over the mobile-to-mobile vehicular channel at 5.9 GHz," *IEEE Commun. Mag.*, vol. 46, no. 5, pp. 84–91, May 2008.
- [49] T. S. Rappaport, *Wireless Communications: Principles and Practice*. Englewood Cliffs, NJ: Prentice-Hall, 1999.
- [50] G. Acosta-Marum and M. A. Ingram, "A BER-based partitioned model for a 2.4 GHz vehicle-to-vehicle expressway channel," *Int. J. Wireless Personal Commun.*, vol. 37, no. 3/4, pp. 421–443, May 2006.
- [51] A. Goldsmith, *Wireless Communications*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [52] H. Kreml, I. Seskar, and P. Spasojevic, "On spatial and temporal selectivity of a class of vehicular short range channels," in *Proc. 20th Virginia Tech Symp. Wireless Personal Commun.*, 2010, pp. 1–7.
- [53] I. N. L. Tan, W. Tang, K. Laberteaux, and A. Bahai, "Measurement and Analysis of Wireless Channel Impairments in DSRC Vehicular Communications," *Elect. Eng. Comput. Sci. Univ. Calif., Berkeley, CA, Tech. Rep. UCB/EECS-2008-33*, 2008.
- [54] D. D. Stancil, F. Bai, and L. Cheng, "Communication systems for car-2-x networks," in *Vehicular Networking*. Hoboken, NJ: Wiley, 2010.
- [55] J. Gross, O. Punal, and M. Emmelmann, "Multi-user OFDMA frame aggregation for future wireless local area networking," in *Proc. IFIP Netw.*, Aachen, Germany, 2009, pp. 220–223.
- [56] T. Zahn, G. O'Shea, and A. Rowstron, "Feasibility of content dissemination between devices in moving vehicles," in *Proc. 5th Int. CONEXT Technol.*, 2009, pp. 97–108.
- [57] L. Neubert, L. Santen, A. Schadschneider, and M. Schreckenberg, "Single-vehicle data of highway traffic: A statistical analysis," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 60, no. 6, pp. 6480–6490, Dec. 1999.
- [58] C.-G. Wallman and H. Astrom, "Friction Measurement Methods and the Correlation Between Road Friction and Traffic Safety," *Swedish Nat. Road Transp. Res. Inst., Linköping, Sweden, Tech. Rep.*, 2001.
- [59] M. Bilban, A. Vojvoda, and J. Jerman, "Age affects drivers response times," *J. Coll. Antropol.*, vol. 33, no. 2, pp. 467–471, Jun. 2009.
- [60] R. Akcelik, "On the validity of some traffic engineering folklore," presented at the 22nd Conf. Australian Inst. Transp. Res., Canberra, Australia, Dec. 2000.



Gustavo Marfia received the Laurea degree in telecommunications engineering from the University of Pisa, Pisa, Italy, in 2003 and the Ph.D. degree in computer science from the University of California, Los Angeles, in 2009.

He is currently a Postdoctoral Researcher with the Department of Computer Science, University of Bologna, Bologna, Italy. He is also with the University of California, Los Angeles. His research interests include biomedical, transportation, and multimedia networking systems, fields in which he has published

over 60 technical papers.

Dr. Marfia serves as an Associate Editor for the Elsevier *Journal of Visual Communication and Image Representation* and for the Association for Computing Machinery *Computers in Entertainment*.



Marco Rocchetti received the Italian Laurea degree in electronic engineering from the University of Bologna, Bologna, Italy.

Since 2000, he has been a Full Professor with the Computer Science Department, University of Bologna. His research interests include computer entertainment, intelligent transportation systems, and web-based applications, fields in which he has authored more than 250 technical papers.

Prof. Rocchetti serves as an Associate Editor for many international journals and is active in several

Italian and international projects.



Alessandro Amoroso received the Laurea degree in physics from the University of Bologna, Bologna, Italy, in 1987.

Since 1994, he has been with the Department of Computer Science, University of Bologna, where he is currently an Associate Professor in computer science. He has participated in several scientific projects, including the National Research Council, National Energy Board, and University of California at San Diego. His main research areas are mobile devices, multimedia systems, and distributed systems.

His recent research activity has included work on vehicular ad hoc networks.



Giovanni Pau received the Italian Laurea doctorate degree in computer science and the Ph.D. degree in computer engineering from the University of Bologna, Bologna, Italy, in 1998 and 2002, respectively.

He is currently an Adjunct Associate Professor with the Computer Science Department, University of California, Los Angeles (UCLA). He is currently leading the UCLA vehicular laboratory aimed at building a publicly available vehicular testbed for the research community. His main research interests lie

in the area of vehicular ad hoc networks and urban sensing.

Dr. Pau served as Vice Chair for North America for the IEEE Communications Society Multimedia Technical Committee and serves on the technical program committee of several IEEE and Association for Computing Machinery conferences. Finally, he is an Associate Editor for the Elsevier *International Journal of Ad Hoc Networks* and the Springer *International Journal on Peer to Peer Systems*.