To Live and Drive in L.A.: Measurements from A Real Intervehicular Accident Alert Test

Alessandro Amoroso, Gustavo Marfia, Marco Roccetti
Computer Science Department
University of Bologna
Mura A. Zamboni 7, 40127 Bologna, Italy

Giovanni Pau
Computer Science Department
University of California, Los Angeles
420 Westwood Plaza, 90095-1596 Los Angeles, CA

Abstract—**Although the number of accidents that occur on roads is slowly decreasing in time, as both cars and streets become progressively safer, the Statistics of Road Traffic Accidents study from the United Nations reveals that during the last decade every year an average of 150,000 people have lost their lives and 5.5 million have suffered injuries on the roads of western countries. Among the many proposals that have been made, during the years, to combat such phenomenon, an important place is taken by highway accident warning technologies. Such type of system can play an important role, especially in those cities, like Los Angeles, Seoul and Shanghai, specifically designed for cars, more than for human beings. In fact, when an accident occurs on a highway of a city like Los Angeles, it is vital to warn as rapidly as possible all the approaching vehicles, in order to give them the time to stop before crashing on any unexpected obstacle. Recently, an inter-vehicular accident warning system has been devised, theoretically proven to be optimal in terms of bandwidth usage and covered distance. In this paper, we present preliminary results that assess the feasibility of such system. The presented results and measurements were taken from real experiments, performed on Los Angeles freeways and roads in August 2011. To the best of our knowledge, these are the first real experiments of such magnitude performed in a live setting with an inter-vehicular accident warning system.**

Keywords-VANETs; Vehicular overlay networks; Experimental testbeds; Accident Warning Systems.

I. INTRODUCTION

Car accidents have an enormous impact in terms of societal costs. Recent studies have estimated that, only in the US, in economic terms car accidents cost over $200B per year. Moreover, this monetary value only captures part of the problems that are caused by such events, not being able to quantify the social distress that follow them and the damages produced by the pollution that they cause where hundreds, or even thousands of cars remain stuck in traffic, in car-dependent cities like Los Angeles, Seoul and Shanghai. Now, although the number of road accidents that occur is slowly decreasing year after year, due to a number of factors that are easy to enumerate (i.e., better roads, safer cars, stricter rules, etc.), no new solution has succeeded, to this date, in achieving an abrupt reduction in terms of the number of crashes, injured and casualties. Airbags, for example, which are often considered one of the most important safety devices installed in vehicles, are accounted for an accident fatality reduction of 12% and, in general, can only smoothen the harsh effects of a crash, rather than prevent the occurrence of an accident.

Time is ripe to perform the next step in vehicular safety, which is to prevent accidents, rather than adopting countermeasures that can ameliorate their outcome. Many systems that aim at performing such type of operation have been proposed, including radars that measure the safety distance between vehicles and alert the driver when this distance falls below a certain value. However, with such solutions it is not possible to foresee what is going on far ahead, but only control the state (e.g., distance and speed) of neighboring vehicles. In the era of wireless communication technologies, clearly, one of the most exciting ideas that have been proposed has been that of adopting a vehicle-to-vehicle (V2V) strategy to advertise accident occurrences to approaching vehicles. In brief, assuming all vehicles were equipped with wireless interfaces that could put them into direct communication between each other, when an accident occurs, the involved vehicle could send an alert message to all the vehicles that are traveling in its vicinity. That alert message could then be recursively relayed to all approaching vehicles through those vehicles that heard it and, hop by hop, also reach vehicles that are miles away and completely unaware of what has just occurred far ahead of them, along their chosen path.

Such class of systems, often referred to as *accident warning systems,* have been thought on top of the use of vehicular ad hoc networks (VANETs), communication networks where vehicles talk between each other, without needing the intervention of any external communication infrastructure. In fact, the fastest way to convey an alert message between a vehicle that has been involved in an accident and all the vehicles that follow is through neighboring vehicles, as transmitting that alert message through an infrastructure would add delays that could result to be fatal. However, such type of systems cannot be implemented just telling all vehicles to rebroadcast an alert message that has been heard, as this would lead to the well known broadcast storm problem, situation where transmissions interfere with each other and hence no message finally gets delivered. Recently, the authors of [1]-[3] were able to demonstrate that an accident warning system can be devised, that is optimal in terms of the bandwidth resources required to distribute an accident warning and the distance reached by accident alerts. The peculiarity of the accident warning system they proposed is that, based on a distributed oracle mechanism that sends and receives information regarding a vehicle and its neighbors, it is possible to propagate an alert message utilizing the minimum number of hops, even in situations where obstacles, and hence realistic asymmetric communication conditions, are present.

Now, although many research works have confidently proposed V2V accident warning systems as a new safety standard for vehicular environments, all results that have been published to this date are simulation results. The reason for this is clear: it is unfeasible to build a vehicular network composed of tens or hundreds of vehicles that can act as relays of a message that could travel many kilometers away. Hence, to date, no working system has been presented that can actually transport a safety message from a source to a destination.

In order to address the problem of assessing in reality the optimal accident warning system described in [3], we implemented and took it for a test run around the streets of Los Angeles in August 2011, one of the cities that most would benefit from the implementation of such type of technology. The most impressive result of this experimentation campaign has been the observation of alert messages that travelled for many kilometers from the spot where an assumed accident occurred to the final vehicle that received it. Hence, the particular contribution of this paper is that of describing the first real experiments of such magnitude performed in a live setting with an inter-vehicular accident warning system. This has confirmed our finding that, adopting such mechanism, the number of vehicles involved in a pile-up can greatly be reduced as anticipated in [3].

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

II. Related Work: Vehicular Testbeds

We will here avoid any discussion regarding previously proposed accident warning system focusing, instead, on the main approaches that have, in the past, tackled the problem of testing VANET-based algorithms within real scenarios. Specifically, while a wealth of research has so far been produced in the domain of VANET research, and on the application of VANET technologies to accident warning systems, unfortunately little advances have been so far performed on the experimental side that concerns the testing of accident warning systems in realistic scenarios; the main problem being that of transmitting alert messages several hops away with testbeds composed of only a few vehicles. For example, in [4] the authors perform one of the few VANET multi-hop experimentations, mounting 802.11b interfaces on top of a platoon composed of four vehicles and utilizing the Optimized Link State Routing protocol (OLSR) to deliver packets between any two vehicles within the platoon. The authors drove the platoon of vehicles over urban and highway environments to assess the performance of the vehicular networks in terms of delay, bandwidth, packet loss and distance between nodes. Compared to our experience, this vehicular testbed falls short in testing multi-hop scenarios that go beyond those imposed by the scarcity of available vehicles. Instead, utilizing a peculiar methodology named *back and forth* and explained in Section IV, in fact, we have been able to conduct experiments where several kilometers were travelled by alert messages, even using only a few vehicular resources [1], [2].

Similarly, the authors of [5] built a vehicular testbed where cars were equipped with computing and communication devices, and experimented a position-based routing scheme implemented on Linux-based laptops. Within this work, they were able to test and measure the performances of static and mobile l-hop and 3-hop scenarios. As before, this work is limited by the fact that any communication packet could only experience a few hops when traveling between a sender towards a given receiver.

Concluding, it is difficult to find in literature any work that assesses the performance of a mobile system as a vehicular network with an increasing number of hops between the sender and any given receiver [6]. Although many different experiences exist that moved in such direction, ours is, to the best of our knowledge, the first work that reports on experiments performed with an accident warning system that sent its messages several kilometers away.

III. Background: An Optimal Accident Warning System

Since the first proposals that have envisioned a connected vehicular grid capable of providing a mobile communication infrastructure open to a variety of services, safety applications have been seen as key driver that could, alone, justify the amount of investments required to make VANETs come true [7]. As such, many have been the proposals of new broadcast protocols, tailored for specific highway vehicular scenarios, where it could be possible to timely propagate an accident alert, hence avoiding the involvement of approaching vehicles in pile-ups [8]-[15].

What differentiates the proposal advanced by the authors of [3] from any previous initiative, is that it guarantees the use of the minimum number of hops to distribute an alert message throughout a platoon of vehicles. Such property is not only guaranteed when the transmission ranges are constant and uniform for all vehicles, but also within all those situations (e.g., obstacles, different antennas, etc.) where communication links are asymmetric and transmission ranges rapidly vary in time. Now, while several proposals exist where a warning message is relayed by a so-called *farthest relay*, i.e., the vehicle that travels farthest in a given direction, this may be no the right choice in reality, because of the particular transmission conditions experienced by that vehicle. In fact, it may happen that the transmission range of the farthest relay is reduced due to the presence of an obstacle, or because the hardware mounted by that particular vehicle is defective or scarcer in transmission power, with respect, for example, to a vehicle that precedes it. The right choice, instead, is to choose the *farthest spanning relay*, i.e., which is not the farthest vehicle but the one that can relay a message the farthest away in a given direction. All this is made possible in [3] thanks to an Oracle service implementing a distributed algorithm that provides each vehicle with the information concerning which are the *farthest spanning relays*, namely the relays that can retransmit a message farthest away from its source, in the forward and backward directions. A vehicle, when indicating that the farthest spanning relay will retransmit the message it is sending, achieves two goals: (a) avoids any broadcast storm problem, and, (b) guarantees that that message will be sent the farthest away. In the following we will show the results that have been achieved with the particular accident warning system engineered in [3] on the streets of Los Angeles. The importance of this contribution is also given by the fact that, to this date, no system that is based on multi-hop communications within vehicular scenarios has ever been pushed to such level of performance assessments before, within real settings.

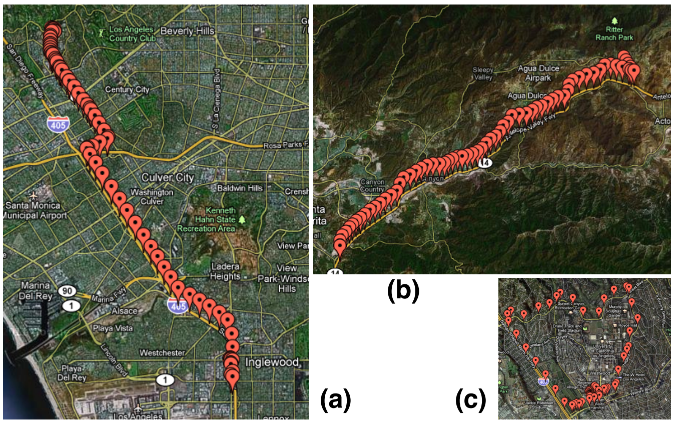

Figure 1. Example of routes: (a) urban, (b) suburban and countryside, and (c) urban and suburban.

IV. Experimental Setting

In the following we will briefly describe the setup of the vehicular testbed that was utilized for our experimentation campaign. In particular, we will discuss: (a) the technique exploited to implement a vehicular testbed where alert messages can travel for, virtually, an infinite distance even in the presence of scarce vehicular resources, (b) what type of equipment was installed on each vehicle and (c) which type of roads were traversed while driving through Los Angeles in August 2011. As already mentioned, the main problem here is that of allowing an alert message to travel, virtually, an infinite number of hops with a vehicular testbed composed of only a few vehicles. We made all this possible devising a new method apt to experiment with vehicular networks in real settings. Specifically, we exploited a testing methodology, termed *back and forth*, conceived to reproduce the real conditions that would be experienced when extending the number of hops, the number of interfering transmissions and the channel conditions that can be encountered by a communication packet containing an alert message, just as it happens in reality. In practice, a front and a back vehicle characterized the vehicular testbed, both being in charge of relaying back any received alert message. Alert messages, hence, travelled back and forth between the front and the back vehicles, as long as they could. Hence, even with a small platoon of vehicles, our mechanism allowed alert messages which were sent by the first vehicle of the platoon, to reach the last one through intermediate ones and immediately bounce back, as long as they did not get lost for any given reason [1], [2]. In such setting where the front vehicle was chosen as the first one in the platoon and the back one was the last one, the remaining vehicles were subject to the unique constraint of travelling between the front and the back vehicles, regardless of their relative positions. This testing mechanism was triggered, within each experiment, by the front vehicle, that sent an alert message, with a frequency of one every ten second, with the whole platoon managing the alert propagation activity through the use of farthest spanning relays.

As to the experimental equipment deployed on each vehicle, it was comprised of: a portable personal computer mounting an 802.11b/g card with enabled Auto Rate Fallback (ARF), two omnidirectional antennas and a Sirf Star III GPS receiver. Each personal computer ran our accident warning system that was implemented in Java. The broadcast of both control and alert messages was implemented utilizing raw sockets, hence inserting application layer data directly into layer-2 frames. We carried out hundreds of experiments for almost thirty days in August 2011, running our software while driving a platoon composed of four cars through different types of roads in Los Angeles and in its immediate surroundings. We chose three different paths to address two main different scenarios, one with prevailing urban and suburban surroundings and one with prevailing countryside surroundings. The first path ran through urban areas of Los Angeles, starting from the UCLA campus, moving along Westwood and National Blvd. and ending near the LAX airport once the I-405 South Fwy was left (Figure 1.a). The second path was performed moving from the UCLA campus and reaching the initial south-west part of the Mojave Desert, while driving on the I-405 North and on the I-110 Fwys (Figure 1.b). Finally, the third path is composed by roads that included both urban (I-405 North Fwy, Hilgard Ave. and Wilshire Blvd.) and suburban ones (Sunset Blvd.), and is shown in Figure 1.c. The three mentioned paths represented important exemplar cases to assess the performance of our accident warning system as surrounding conditions varied in terms of: (a) traffic, (b) speed, and, (c) varying wireless channels.

V. Experimental Results

In our experiments we measured several characteristics of the alert system. Due to space constraints imposed to this article, we will here discuss only the following.

A. Average covered distance

Differently from any other experience performed before in vehicular environments, ours were the first experiments where the number of hops performed by communication packets were so high that single alert messages reached vehicles that were travelling several kilometers away.

In particular, we here provide as a measure of relevant importance the distance covered by accident warnings distributed by our system. As to this, in Figure 2, we counted the number of alert messages, as a function of the distance that they covered to spread the warnings throughout the entire platoon of vehicles. As shown in the Figure, many are the alert messages that travelled over 1 km from the spot where an accident was simulated to occur. In other words, many are the vehicles that could potentially be warned of a dangerous situation that occurred far ahead of them. Most importantly, Figure 2 shows that over 400 alert messages travelled more than 2 km, thus giving an important advantage to all those vehicles that were approaching an area where an accident occurred.

To be noticed is also the fact that we assessed our system utilizing a vehicular testbed that moved on highways where simulated accidents were created at a high frequency. Considering that alert messages on average have spanned as far as several kilometers from the spot where the accident happened, a terrific practical evidence was obtained that using these alerts approaching vehicles had time to stop before being involved in a pile-up. We are hence not so far from the 40% reduction of involved vehicles we anticipated with the our simulated and theoretical results of [3].

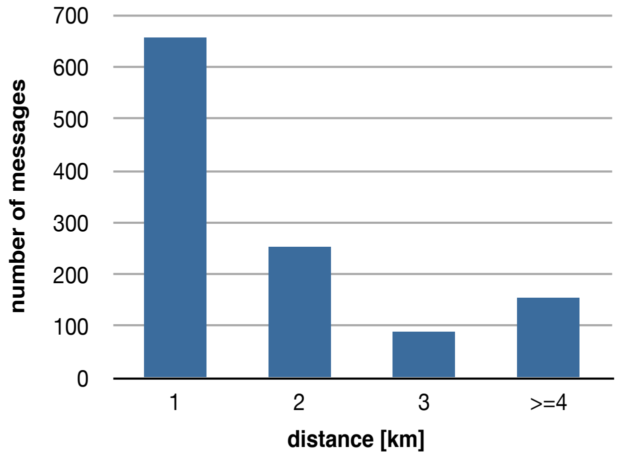

Figure 2. Distribution of the alert messages as a function of their reached distance (one of them travelled as far as 36 km).

B. *Robustness to faults*

In our experiments alert messages eventually stopped being relayed because of several reasons. This is strictly related to experimental conditions, such as vehicle positions and surrounding environment, and as such it is very difficult to analyze. Nonetheless, our alert system incorporates a mechanism capable of dealing with the loss of an alert message. In fact, each alert message contains also the list of its potential relays. Essentially, each vehicle in the relay list starts a timer after receiving that message.

When that vehicle detects that such given alert message was relayed by a vehicle preceding it on the relay list, it aborts any relaying activity. If, instead, that vehicle does not detect a relaying activity being initiated by any of its predecessors on the relay list after a certain period of time, it takes charge of relaying that message [3]. To assess how frequently the best relay actually sent a given alert message, we count in Figure 3 how many times the first, the second and the third relay, in the list, have performed the corresponding relay activity, respectively. The shown pattern reveals that the first relay was able to retransmit the messages almost in all cases and rarely failed in doing so.

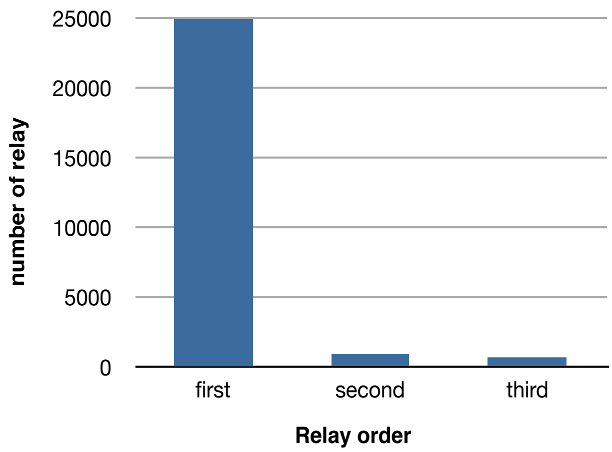

Figure 3. Actual relaying vehicles.

The alert messages got lost only when the whole platoon was not able to perform the relaying activity. This was a very rare case that occurred, typically, when either the distance between subsequent vehicles was too large with respect to the communication capability of the mounted antennas, or when all the vehicles of the platoon were to close to each other, suffering from cross interferences.

C. *Asymmetric communications*

The distributed Oracle is that entity that, running on each vehicle, communicates with the neighboring ones and discovers which are the farthest spanning relays in the backward and forward directions [3]. The Oracle is also that mechanism that lets each vehicle know which are those other vehicles that can receive its relayed messages, even in those cases where communication channels are asymmetric and, hence, in those cases where our vehicle cannot reach them directly (i.e., receiving their transmissions).

Now, we were interested to assess the ability of the Oracle of managing asymmetric communications. To this aim, Figure 4, in particular, shows an example of an asymmetric communication between the vehicles *back* and *front*. The message *m1* from *back* reaches both vehicles *v1* and *front*, but the message *m2* from *front* cannot reach *back* unless *v1* relayed it. We measured the time, in seconds, required by the Oracle to detect a situation where the communication channel was asymmetric. Referring to the example shown in Figure 4, we measured the time between the reception of message *m1* by the *front* vehicle and the reception of message *m2* by the *back* vehicle. Figure 5 shows the aggregate number of asymmetric communications that occurred during the experimentations as a function of the time, in seconds, that the *back* vehicle required to be notified of the existence of the *front* one. The yielded result is that, in most cases, the objective of detecting asymmetric conditions is achieved in no more than six seconds.

VI. Conclusion

In this paper we have addressed the problem of assessing how a real accident warning system performs, on the streets of a chaotic city as Los Angeles is. In particular, we developed a vehicular testbed that was driven in Los Angeles in August 2011 to assess the performance of the inter-vehicular accident warning system of [3].

The results of these experiments have shown that alert messages can travel for many kilometers from the spot where an accident occurred to approaching vehicles. To the best of our knowledge, these are the first experiments where a multi-hop inter-vehicular communication system of such magnitude has been tested. Considering the covered distances, we can confirm that adopting such mechanism, the number of vehicles involved in a pile-up could greatly be reduced.

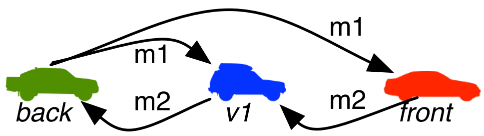

Figure 4. Example of asymmetric communication between the *back* and *front* vehicles.

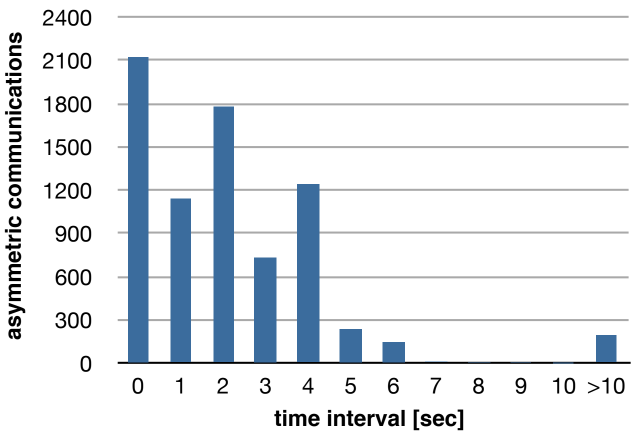

Figure 5. Time, in seconds, required to discover asymmetric communications.

Not only, the real additional benefit of our experimental campaign has been also to show that an efficient method/principle exists for the devise and implementation of a broad class of inter-vehicular services, ranging from on-board safety to vehicular mobility management and exchange of digital data (e.g., entertainment, news, commercials, pollution and roadwork info, parking availability, etc.). We hence conclude by confirming that the real good news is that this can be done without resorting to any external communication infrastructure, but only allowing cars to talk to each other.

Acknowledgment

We acknowledge the FIRB DAMASCO, the PRIN ALTER-NET and the NSF/BBN GENI projects for supporting this work.

References

[1] A. Amoroso, G. Marfia, M. Roccetti, G. Pau, "Creative Testbeds for VANET Research: A New Methodology," *Proc. 9th IEEE Communications and Networking Conference (CCNC 2012)*, Las Vegas (USA), January 2012.

[2] G. Marfia, M. Roccetti, A. Amoroso, M. Gerla, G. Pau, J. H. Lim, "Cognitive Cars: Constructing a Cognitive Playground for VANET Research Testbeds," *Proc. 4th International Conference on Cognitive Radio and Advanced Spectrum Management (CogART 2011)*, Barcelona, ACM, October 2011

[3] A. Amoroso, G. Marfia, M. Roccetti, "Going Realistic and Optimal: A Distributed Multi-Hop Broadcast Algorithm for Vehicular Safety," *Computer Networks*, Elsevier, vol. 55, n. 10, Jul 2011, pp. 2504-25.

[4] M. Moske, H. Fussler, H. Hartenstein, W. Franz, "Performance measurements of a vehicular ad hoc network," *Proc. 59th IEEE Vehicular Technology Conference (VTC 2004)*, Milan (Italy), May 2004.

[5] J. Santa, M. Tsukada, T. Ernst, A.F. Gomez-Skarmeta, "Experimental analysis of multi-hop routing in vehicular ad-hoc networks," *Proc. 5th IEEE International Conference on Testbeds and Research Infrastructures for Development of Networks and Communities and Workshops (TridentCom 2009)*, Washington DC (USA), April 2009.

[6] I. Leontiadis, G. Marfia, D. Mack, G. Pau, C. Mascolo, M. Gerla, "On the Effectiveness of an Opportunistic Traffic Management System for Vehicular Networks," *IEEE Transactions on Intelligent Transportation Systems*, to appear.

[7] M. Gerla, B. Zhou, Y.-Z. Lee, F. Soldo, U. Lee, G. Marfia, "Vehicular grid communications: the role of the internet infrastructure," *Proc. 2nd Annual International Workshop on Wireless Internet (WICON '06), Boston* (USA), August 2006.

[8] N. Wisitpongphan, O.K. Tonguz, J.S. Parikh, P. Mudalige, F. Bai, V. Sadekar, Broadcast storm mitigation techniques in vehicular ad hoc networks, IEEE Wireless Communications 14 (6) (2007) 84–94.

[9] A. Zanella, G. Pierobon, S. Merlin, On the limiting performance of broadcast algorithms over unidimensional ad-hoc radio networks, in: Proceedings of Wireless Personal Multimedia Communications, 2004, pp. 165–169.

[10] E. Fasolo, R. Furiato, A. Zanella, Smart broadcast algorithm for inter-vehicular communication, in:Proceedings of 2005 Wireless Personal Multimedia Communication, September 2005.

[11] J.J. Blum, A. Eskandarian, A reliable link-layer protocol for robust and scalable intervehicle communications, IEEE Transactions on Intelligent Transportation Systems 8 (1) (2007) 4–13.

[12] G. Korkmaz, E. Ekici, F. Ozguner, U. Ozguner. Urban multi-hop broadcast protocol for inter-vehicle communication systems, Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks, October 2004, pp. 76–85.

[13] C.E. Palazzi, S. Ferretti, M. Roccetti, An inter-vehicular communication architecture for safety and entertainment, IEEE Transactions on Intelligent Transportation Systems 11 (1) (2010) 90–99.

[14] F. Ye, M. Adams, S. Roy, V2V wireless communication protocol for rear-end collision avoidance on highways, in: Proceedings of IEEE 2008 Vehi-Mobi Workshop, May 2008, pp. 375–379.

[15] S. Biswas, R. Tatchikou, F. Dion, Vehicle-to-vehicle wireless communication protocols for enhancing highway traffic safety, IEEE Communication Magazine 44 (1) (2006) 74–82.

[15] R.A. Santos, R.M. Edwards, A. Edwards, Cluster-based location routing algorithm for vehicle to vehicle communication, in: Proceedings of 2004 IEEE Radio and Wireless Conference, September 2004, pp. 39–4