Behind the Scenes: Lessons Learned from the Greatest Intervehicular Accident Detection Test Ever

G. Marfia*, M. Roccetti**, A. Amoroso**
*Department for Life Quality Studies
** Department of Computer Science and Engineering
University of Bologna
Bologna, Italy
{gustavo.marfia, marco.roccetti, alessandro.amoroso}@unibo.it

M. Gerla
Department of Computer Science
University of California, Los Angeles
Los Angeles, USA
gerla@cs.ucla.edu

Abstract—**This paper describes a few aspects of the work that led to the first experiments ever carried out on the road with an accident warning system based on vehicle-to-vehicle (V2V) communications. The mentioned experiments were performed in Los Angeles as part of a collaboration project between two teams, one from the University of Bologna and one from UCLA. In essence, driving along a few of the most crowded streets of LA and mounting on each vehicle an 802.11 based communication system, an accident warning dissemination algorithm was took for a test ride. The tested algorithm demonstrated being both in theory and simulation the optimal one, as it could successfully deal with realistic asymmetric links always choosing the optimal relay. However, a practical assessment was lacking. This required an important research and financial investment, as an accident warning system based on V2V technologies was never tested before. On the basis of that pioneering experience, we here revisit its "behind the scenes" results and the lessons we learned. In fact, we believe that those experiments provide new insights and point new possible directions of work. The contribution of this paper, hence, goes beyond the sharing of information, as its scope is also that of providing a vision of how vehicular research can be applied in practical terms when taken on the road.**

Index Terms—**Vehicular networks, accident warning system, experimentation campaign.**

I. Introduction

When thinking of the future of the automobile industry we all imagine cars that can fly and move over multiple imaginary lanes in the sky, as well as systems that can safely and rapidly take passengers from one point to the other of the globe. Many small steps have been taken, so far, that will eventually lead to that result: modern vehicles now carry connection devices (e.g., Bluetooth, WiFi, etc.), screens for passengers and radars for front-rear collision prevention. In the near future we may expect that vehicles will also be able to drive themselves and remove human drivers from the burden of keeping control and taking responsibilities.

The idea of having cars *talk* to each other very well fits such vision. We can foresee multiple reasons why it may be interesting having cars *talk* (e.g., passengers exchanging online gaming information, chatting, downloading files, etc.). Safety is certainly one of its most interesting applications as, when an accident occurs, cars can rapidly relay such information to approaching ones. Although the number of accidents on the road is constantly decreasing, certainly thanks to the many safety countermeasures that have been taken throughout the past years, this remains a fascinating opportunity. In fact, the need for safety on roads is a recurrent argument that periodically fills the pages of newspapers, regularly appearing when major accidents occur.

What does in practice mean implement systems where cars *talk*? According to the path that industry and academia have taken in the past years, it means implementing protocols that rely on the 802.11 standard. In practice, when exchanging information, vehicles can talk only one at a time (i.e., in a half duplex fashion), as two vehicles that transmit information at the same time would interfere with each other. This particular property of the 802.11p amendment (i.e., the one that regulates vehicular communications) has laid down the path to numerous research efforts during the past fifteen years. In fact, the problem of letting cars *talk* and report the occurrence of an accident has posed the problem of how this could be done making cars *talk* one at a time, avoiding interference and hence message losses.

In the past fifteen years many have also thought of ways of disseminating the notice of danger events at a fast pace, as soon as they occur. This is very critical for accidents on highways, which are often due to harsh weather conditions (e.g., reduced visibility) and expose approaching drivers to a high probability of being involved in pileups of vehicles straight ahead. Clearly, these are the typical situations where knowing what occurs far ahead can save human lives. Now, although such application has fostered a long going debate on how communications should be devised between cars, most of the speculations that have so far been made built up over theoretical and simulative arguments.

When considering a highway scenario, the general approach has been that of simplifying the problem and assuming messages propagating in a linear fashion, in a single direction (i.e., one-dimensional network). Under such assumption, an optimal way of transferring a warning message from a vehicle involved in an accident to an entire mesh of vehicles has been found and works as follows [1], [2]. Such scheme simply takes into account that wireless links, especially in highly dynamic situations as vehicular ones, can be anything but symmetric. Hence, when a vehicle, say Kitt, talks, it is possible that another one, say Karr, hears what it says, but not vice versa. This also means that in case Kitt is involved in an accident, there might exist an unknown relay, e.g., Karr, which is the optimal one. The scheme presented in [1], [2] provides a mechanism capable of overcoming such shortcoming and hence of reaching the optimal relay at each hop.

Now, how could we ideally test a system like the one that we just described in reality? If we were an important automaker, (e.g., Toyota, Volkswagen, Mercedes or Audi) we would ask 1,000 commuting employees to get all in their vehicle-to-vehicle (V2V) enabled cars at the same time and create a realistic 2 km, say, solid *experimental cloud*. Unfortunately, we are not automakers, so we can at most think of deploying enough cars to stretch for 2 Km, two by two or three by three side by side, at a distance of say, 50 to 100 m. This, however, would require from 100 to 200 cars, practically doable, but a great cost! In addition, this experiment setting assumes that only one alarm is triggered at a time, and propagates through the vehicles in the minimum connected mesh.

To overcome such problem the following solution was thought [3], [4], [5]: fold the mesh onto itself and reduce it to 2 sets of vehicles ping-ponging alarms to each other. What does this mean? It means that no matter how many vehicles are put in place (e.g., two in the worst case), messages propagate back and forth (i.e., towards the front when reaching the back and vice versa), hence never stopping when reaching the front or the back of the mesh, thus emulating a real situation where an indefinite number of vehicles are placed along the path between the vehicle that launched the alarm and the last vehicle that can be physically reached. Clearly, an alternative approach is that given by emulating thousands of cars (e.g., with virtualization), but such approach misses the opportunity to capture realistic motion in actual traffic and realistic radio propagation.

This said, utilizing the abovementioned methodology, in August 2011 the largest ever documented accident warning system experiments have been carried out on the street of Los Angeles. These experiments have been designed and carried out thanks to the collaboration of two teams: one residing at the Computer Science Dept. of the University of Bologna and one at the Computer Science Dept. of the UCLA. In particular, the experiments took the mechanism presented in [1], [2] for a test ride, verifying its feasibility, scalability, speed and performance in general. The whole set of results has been published in [3], [4], [5]. Here, rather than focusing on the specific scientific results that have been already published elsewhere, we will summarize the lessons learned and briefly present our vision of future vehicular networking systems.

The remainder of this article is organized as follows. In Section II we will briefly report the main result that has been obtained during the experiments. We will then move on to a critical discussion in Section III, unveiling some of the lessons learned while planning and executing our experiments. We finally conclude with Section IV.

II. Main Result

No need to say, the main result that has been obtained with our experimentation campaign on the road has been that of being capable, utilizing real data, of estimating how many vehicles could be saved from being involved in an accident when using our accident warning system. Interestingly, the result that was initially conjectured has been confirmed: it could be possible to reduce the number of vehicles involved in a pileup by 40%, on average.

In fact, the time taken by a message to propagate from an accident was first estimated in simulation and then confirmed by our real experiments. In [1], we found that the end-to-end time required by an accident warning message to travel between two vehicles moving at a distance of 8 km one from the other was approximately equal to 125 ms. Our experimentation campaign in Los Angeles gave us the opportunity of estimating the amount of time taken to traverse a single vehicular hop. This value, as explained in [5], is confined between 6 and 20 ms, depending on the wireless propagation scenario. Again, considering a distance of 8 km between the origin and the destination we found values confined between 220 and 750 ms, two to six times more than the value found in simulation. This discrepancy is well justified by the fact that in reality transmission ranges fell well below than the value of 350 m that was employed in simulations (i.e., 40% of messages were lost when transmitter and receiver travelled at a distance between 300 and 400 m). When taking into account this discrepancy (i.e., a distance of 8 km was covered with approximately 23 hops in simulation and 37 hops in reality), we instead find that the end-to-end times were confined between 134 ms and 457 ms, thus much closer to what was estimated with our simulations.

Now, when we first conjectured the amount of vehicles that could be saved in [1], we considered a situation where vehicles flowed on a highway, with dry pavement conditions. However, we did not have any real estimate regarding the average delay that could be experienced by a safety message while percolating a generic relay of a vehicular network. Hence, taking the value of 6 ms, we computed the amount of vehicles that would be warned, in time to stop. We hence confirmed that, on average, 40% of the cars that would be normally involved in an accident could be saved. This is clearly a relevant value and well justifies the amount of effort and research that has been so far devoted to such research field.

III. Discussion

We now move on to discuss two different aspects that emerged from our experiments. Both of these aspects were not included in [TVT] and are related to two very different fields.

The first pertains a few observations that we were able to make without, however, collecting the minimum amount of data required to make any rigorous scientific speculations. We will here report this information, considering it may be useful for future research efforts.

The second field is related to all the events that occurred after we performed our experiments and after the notice of these experiments spread through mass media. We find that also this aspect, very closely related to technology transfer processes, has a very significant role in such type of applied research.

A. Empirical Observations

During our experiments in Los Angeles we had the opportunity of seeing a number of different phenomena occur in an empirical way, considering that we did not plan their study and did not hence design our experimentation platform for their specific collection/observation.

Beware of access points. This may be the title of our first observation. In fact, the development of our testbed has proceeded as follows. Before getting to test our accident warning system on the road, we performed a thorough set of preliminary tests in the lab. The moment that these tests confirmed how logically our accident warning system should behave, under many different scenarios, we moved outside, but still remaining in a static scenario. While we performed an additional amount of tests in front of the UCLA Computer Science Department building, we still found everything working as expected. Interestingly, when we were ready to start with on the road experiments, the system collapsed. The moment we mounted all the equipment on our vehicles and started the system, nothing worked (i.e., no messages flowed between cars). After a few hours where all the software and all the hardware was checked many and many times, we fortuitously discovered why this was happening. This was really a fortunate event, as it happened when we were moving back to the lab, accepting the defeat. In fact, after a few hours where nobody could figure out what was happening, we all agreed that we should move all the equipment back to the lab, for further investigation. Hence, one of the vehicles moved, reaching a distance somewhere halfway between the first and the last vehicle of our mesh. Approximately at that moment, everything started working.

Hence, what was the problem? We later discovered that the location where we set our starting point was not the best possible one. In fact, at the top of parking structure number 9 we could later count over a hundred access points interfering for wireless channel access. We parked our cars in a location that was right in the middle, being disturbed by all the sources that were transmitting in the three different buildings that were surrounding that location. Only when our vehicles got closer our wireless interfaces were then capable of distinguishing accident warning system messages from interfering ones. The interested reader will now say that once 802.11p will be mounted on cars, such interference problems will not occur, as such standard occupies a different range of frequencies (i.e., 5.9 GHz). However, when vehicular networks will become reality, also roadside units will be deployed and we can easily expect interference from both vehicles and access points. An efficient utilization of spectrum resources is hence key in such environment and for such application.

Another interesting point regards the relationship that the published results have with the results that the same system would have utilizing 802.11p interfaces. Clearly, an in depth investigation would require a repetition of the same experiments with 802.11p interfaces. However, we are still in the condition of making some conjectures regarding the performance gap that can be expected between an 802.11p implementation of our system and the one that is described in [5]. Any performance gap, in favor of one standard or the other, heavily depends on the transmission ranges reached when opting for any of the two solutions. In order to provide such value for 802.11p we resort to the 802.11p standard, as well as to the most relevant scientific literature, searching for the transmission ranges involved when specifically sending safety messages. Interestingly, we found that this is a rather ambiguous matter and that there is no unique answer, to this date, to this search. Beyond any other consideration, the first fact is that not even in the available literature the problem whether, and to which extent, 802.11p performs better than 802.11g (or vice versa) has been determined. For example, three of the most recent and relevant scientific works that have recently appeared do not provide an answer to this question [6], [7], [8], as they show that different transmission ranges can be experienced with 802.11p, that go from as low as 70 m in [6] to as high as 500 m in [7]. To be noticed the fact that we shot at about 300 m with 802.11g. Putting aside the literature (that did not give us a definite response), we started, in our experiments, from another viewpoint: that provided by the standard. It must be admitted that even the standard is not yet completely clear. In fact, the Society of Automotive Engineers (SAE) J2735 standard, which defines the Dedicated Short Range Communications (DSRC) Message Set Dictionary [9], reports the following: *"The DSRC structure has one control channel and multiple service channels. Every 100 mSec, all OBUs listen to the service channel for broadcasts. ... . Each of the service channels is intended for use for differing power levels. For example, one channel may be allocated a higher power level and be used by fire trucks to request signal priority, whereas another channel may be used for payment and need low power to prevent adjacent lane reads. <u>However, the precise use of each of the channels has yet to be determined</u>."* So, although the channel (and hence the power level, as DSRC allocates different power levels to different channels) has yet to be determined, we find that a direction has been taken by a relevant governmental agency, the US DOT Research and Innovative Technology Administration in its recent technical document [10] states: *"Connected Vehicle to Vehicle (V2V) safety applications heavily rely on the Basic Safety Message (BSM), which is one of the messages defined in the Society of Automotive standard J2735, Dedicated Short Range Communications (DSRC) Message Set Dictionary, November 2009. The BSM is broadcast from vehicles over the 5.9 GHz DSRC band. <u>Transmission range is on the order of 1,000 meters</u>."* This said, we would like that at least one point were clear. From the perspective of our experiments, we did not want to enter the debate between 802.11g and 802.11p, as no one can see a clear end, but our final intent was to use a solid and practical protocol able to guarantee to reach an inter-vehicular distance (300 m or similar) that was lower than that indicated by the standard. This said, how our results can be interpreted clearly depends on how 802.11p will be employed in its final version for vehicular applications. How far our messages traveled is clear (the ratio of lost messages exceeded the 20% percentage above 300 m). For what concerns the 802.11p standard, instead, if safety messages will be able to travel as far as 1 km, as the cited standard and institutional papers indicate, the results reported in our study can clearly represent a worst case scenario for the performance that could be achieved employing a fully functional DSRC based accident warning system.

Our third point is related to the propagation of messages in tunnels. During our experiments, in fact, it happened more than once that we traversed tunnels (e.g., when moving from one side to the other of I-405). Clearly, this topic deserves a thorough analysis utilizing the instrumentation that is typically employed in wireless propagation research for channel characterization [11]. We were, however, capable to make the following observation: it was very interesting to notice that in tunnels safety messages never got lost. In fact, when in a tunnel our system kept on running, never stopping because of a message loss. We did not insert these results in [5], as we believed that they deserved a specific analysis on their own. However, for what concerns the research of safety measures in tunnels (i.e., a very hot topic after the Frejus accident at the border between France and Italy), this is an opportunity that could be worth exploiting.

Finally, the run of these experiments have also been the occasion where a novel testing methodology has been introduced. In fact, in order to overcome the problem that only a few vehicular resources were available, we utilized the *back and forth* methodology, i.e., a message kept being sent back and forth through the mesh of vehicles, as if there were a mirror at the first and at the last vehicle. We have long argued why this procedure provides a sound methodology to perform realistic experiments in a real setting in [5], [12]. We would here like to briefly describe how this idea was initially conceived and developed. Starting from the aim of being able to send a message for virtually an infinite number of hops, we begun thinking that a possibility could have been that of sending messages from the first vehicle (i.e., front vehicle) to the vehicles that followed. Once a message got to the last vehicle (i.e., the back vehicle), we imagined that we could relay it back to the front utilizing other communication technologies (e.g., cellular networks) and then loop such communication pattern as long as messages reached the back vehicle. Although this option provided an experimentation architecture where messages always flowed in the same direction (i.e., counter-traffic), we found it not realistic, as in a realistic scenario vehicles moving in the other direction could also relay messages. We hence opted for a solution where we could see a combination of the two propagation directions.

B. Technology Transfer

A fundamental step in all applied science and engineering projects is technology transfer. Hence, the natural question that one could pose is: did any major player in vehicular safety acquire your algorithm/system? The answer is no. In fact, during the months that followed our experimentation campaign, we received many interested and curious calls. We had many meetings and often had the impression that a safety system like the one that we experimented could see life. This did not happen, we will try to explain why we believe it did not.

The vehicular safety realm is composed of many important players, including automakers, insurances, highway management companies, state institutions, communication equipment firms, and fleet managements ones, to list only a few. All of these players can clearly see the importance of safety on the roads. Can also see the importance of providing cars with communication equipment and software systems that may reduce the number of vehicles involved in pileups. However, traditionally, safety measures are demanded (or also imposed through state laws) to automakers, and are not seen as part of their core business by other industry players. Considering this is the background, and also considering the hard hit that most automakers have received in the past few years from market drops, the deployment of such type of systems is probably not at the top of the list of priorities. We interpret this might be the situation, also considering that some of the most important automakers declare that the introduction of V2V communications will happen no earlier than the second half of the decade [13]. Other players, that could receive significant earnings and benefits from such type of technology (e.g., insurance and highway management companies) evidently do not foresee the advantages that they could have and feel these type of activities way too off from their traditional businesses.

This said, our experiments have been conducted with the true spirit of a research: move a bit further away the boundaries of human knowledge. Following this spirit, we are open to share our data, software and results with other research groups that aim at designing systems that may improve the degree of safety of people on the road.

IV. Conclusions

The contribution of this paper has been that of sharing both information and a vision of how vehicular research can be applied in practical terms when taken on the road. We hence have not aimed at providing new results and new theories, but rather we tried to provide additional insights regarding the "behind the scenes" aspects of the first ever documented large scale accident warning system experiments that were performed on the road in 2011 in Los Angeles, CA.

Acknowledgment

We acknowledge the FIRB DAMASCO, the PRIN ALTER-NET, the NSF/BBN GENI projects and the TOYOTA InfoTechnology Center, U.S.A., for their support.

.

References

[1] M. Roccetti, G. Marfia, A. Amoroso, "An Optimal 1D Vehicular Accident Warning Algorithm for Realistic Scenarios," *Proc. IEEE Symposium on Computers and Communications (ISCC'10)*, IEEE, Riccione, June 2009.

[2] 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, July 2011, pp. 2504-2519.

[3] G. Marfia, M. Roccetti, A. Amoroso, "From Theory to Reality: Testing an Accident Warning System on a Real Highway Scenario," *Proc. 7th ACM Asian Internet Engineering Conference (AINTEC'11)*, Bangkok, November 2011.

[4] A. Amoroso, G. Marfia, M. Roccetti, G. Pau, "To Live and Drive in L.A.: Measurements from A Real Intervehicular Accident Alert Test," *Proc. IEEE Wireless Communications and Networking Conference - Wireless Vehicular Communications and Networks (WCNC'12)*, Paris, April 2012.

[5] G. Marfia, M. Roccetti, A. Amoroso, G. Pau, ``Safe Driving in LA: Report from the Greatest Intervehicular Accident Detection Test Ever'', *IEEE Transactions on Vehicular Technology*, IEEE Vehicular Technology Society, Vol. 62, n. 2, February 2013.

[6] F. Martelli, M.E. Renda, G. Resta, P. Santi, "A measurement-based study of beaconing performance in IEEE 802.11p vehicular networks," *Proc. IEEE INFOCOM 2012,* pp. 1503-1511, 25-30 March 2012.

[7] F. Bai, D. D. Stancil, H. Krishnan, "Toward understanding characteristics of dedicated short range communications (DSRC) from a perspective of vehicular network engineers," *Proc. of the sixteenth annual international conference on Mobile computing and networking* (MobiCom'10), pp. 329-340.

[8] K. Hong, D. Xing, V. Rai, J. Kenney, "Characterization of DSRC performance as a function of transmit power," *Proc. of the sixth ACM international workshop on VehiculAr InterNETworking* (VANET '09), pp. 63-68.

[9] Available, online, accessed on January 18th 2013: http://www.sae.org/standardsdev/dsrc/DSRCImplementationGuide.pdf.

[10] M. McGurring, "Vehicle Information Exchange Needs for Mobility Applications," Final Report FHWA-JPO-12-021, February 2012.

[11] J.-M. Molina-Garcia-Pardo, M. Lienard, and P. Degauque. 2009. Propagation in tunnels: experimental investigations and channel modeling in a wide frequency band for MIMO applications. *EURASIP J. Wirel. Commun. Netw.* 2009, Article 7 (January 2009), 9 pages.

[12] 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, October 2011.

[13] Available, online: http://rumors.automobilemag.com/nhtsa-launches-largest-vehicle-to-vehicle-demo-yet-in-michigan-164743.html#axzz2IhW7O0p1