

On The Capacity Of A Linear Vehicular Network

F. El Ali⁽¹⁾⁽²⁾, B. Ducourthial⁽¹⁾⁽²⁾, S.-M. Senouci⁽³⁾⁽⁴⁾

(1) Université de Technologie de Compiègne

(2) CNRS Heudiasyc UMR6599, BP 20529, 60205 Compiègne Cedex, France

(3) Orange Labs, Lannion, France (4) Now at University of Bourgogne, ISAT Nevers, France

(corresponding author: Bertrand.Ducourthial@utc.fr)

Abstract—Intelligent Transport Systems envision many applications relying only on vehicle-to-vehicle communications. Depending on the application (road safety, driver information, infotainment...), the requirements are different in terms of throughput, delay and loss rate. This paper explores the performances issues of a convoy of vehicles on the road, in order to estimate the capacity of such linear vehicular network. Among other results, we show that, while the loss rate is important, it is possible to rely on the vehicular network to relay informations issued from on-board sensors.

I. INTRODUCTION

A. Motivation

Intelligent Transport Systems aim to improve infrastructure management and road experience for users. Vehicular networks are supposed to help in designing efficient applications for the vehicles, the drivers or the passengers (safety, info traffic, ...). As these applications are different in their needs and constraints, they require different communications performances. For instance, emergency applications like alerts are mainly concerned by the delay for the first reception. The loss rate is important for applications that regularly relay data issued from on-board sensors. Passengers applications would require large throughput.

While network capacity can be evaluated by simulation (eg. [11]), such studies suffer from the lack of realism. In fact simulations differ from the road tests since low layers, traffic road and environmental conditions are hard to simulate. In this paper, we present results of extensive measurement campaign evaluating loss rate, loss events, delays and throughput in a convoy of up to seven vehicles.

B. Related works

Road tests have been related in some previous papers [9], [8], [4], [13], [10], [15], [12].

In [9], communication between two vehicles have been studied. The authors show that IEEE 802.11 offers enough capacity for applications like P2P or FTP. In [4], performances were measured using few cars connecting to WiFi access points under typical driving conditions, while caching was optimized to speed up IP addresses acquisition. The authors show that this type of networks is adequate for a variety of applications, particularly those that can tolerate intermittent connectivity.

In [13], UDP and TCP performances were measured in vehicles moving at different speeds that pass one or more

IEEE 802.11 access points at the roadside. The authors discuss the interest of scattered WLAN cells by devices in fast moving vehicles, and provide an analysis of the expected performances.

In [10], experiments with vehicles on a rural highway communicating with an access point have been studied. The authors show that current protocols achieve 50% of the overall maximal throughput. For the authors, conditions in the vicinity of the access point are predictable, which could be exploited to improve vehicular opportunistic access.

In [15], the performance of IEEE 802.11a, 802.11b and 802.11g devices for vehicles-to-vehicles and vehicles-to-infrastructure are investigated. They show that the main impact factors is the inter-vehicle distance, the line-of-sight between sender and receiver and the adaptation rate algorithm.

Other FOT (Field Operational Test) projects like Pre-Drive C2X and SIM-TD prepare a large scale field trial for vehicular communication technology [1], [2].

C. Contribution

Our study deals with the vehicular network capacity in a convoy of vehicles. We observe on road performances, but we also analyze them, and try to find improvements to the given results. It relies on road tests in real conditions using up to 7 cars. For the wireless communications, the IEEE 802.11 standard has been used. Every communication has been done in broadcast, at 2Mb/s. Local broadcast are well adapted to dynamic networks because they avoid to discover the neighbors and to exchange addresses while the neighborhoods are unstable. Our results can be extrapolate to other protocols and further standards (eg. 802.11p).

All our measures have been done at the application level. For this purpose a light protocol stack has been used, and a very simple application has been designed. This application generates packets over a simple multi-hop protocol over UDP over IP. Measuring at the application level gives realistic measures for end-to-end communications. Results can be compared to the requirements of applications to be deployed on vehicular networks.

We show that losses are important, but they can be reduced by repeated transmissions. While the delay increases with the number of hops and the number of retransmissions, it is still possible to rely on the vehicle-to-vehicle communications for

real time applications relying on data produced by embedded sensors.

This paper is organized as follows. In Section II, the experimental platform is presented. In Section III, road tests scenarios are detailed, as well as the measured metrics. In Sections IV, V, VI, results and analyzes for losses, delays and throughputs are presented.

II. EXPERIMENTAL PLATFORM

A. Hardware

All our tests have been done using a platform embedded in 5 or 7 vehicles depending on the tests. This platform is composed by cheap equipments: PC Dell mini-9 DP118 under Ubuntu v8.04, GPS receiver BU-353-W, external USB WiFi card Alfa AWUS036EH and D-Link roof antenna (Figure 1).




Fig. 1. (a): Platform used for the tests, (b): Road experiments

B. Software

This study relies on the Airplug Software Suite, that allows to perform road-tests [7], emulation [3] and simulation [14]. For the road tests, we used the Airplug middleware, the conditional transmissions (HOP application) [6], and a test program generating packets and performing measures at the reception (TST application), that we detail in the following.

Airplug is a light robust middleware for ad hoc networks, dedicated to study dynamic networks [5], [7]. All communications (either intra- or inter- vehicles) are managed by Airplug through message passing. A message coming from a given application can be sent to many other applications, remotely or locally. Airplug implements an addressing scheme adapted to dynamic neighborhood: the destination of a message is composed by two fields: an area (local to the vehicle or the surrounding of the vehicle) and the name of the destination application.

With Airplug, the inter-applications communications are done in the simplest and more robust way possible, by using the standard inputs and outputs. Applications are then simply developed as independent processes, which enforced the robustness. As Airplug also manages the network interfaces, applications access the network in the same way they do to communicate with other local applications, simply by writing on their standard output.

Conditional transmissions is an efficient routing-like protocol, where logical conditions replace addresses [6]. A message is sent with two conditions namely, CUP and CFW. When a vehicle receives a message, it evaluates the conditions. If CUP is true, it forwards the message to the upper layer (for the receiving application). If CFW is true, the message is forwarded to nearby cars. All sorts of logical conditions can be used: duration, date, direction of vehicles, geographical positions or even IP addresses. One of the most interesting ones appears to be the trajectory correlation, allowing to determine whether the receiving car follows the sender or not. By dynamically evaluating conditions at receptions, the protocol accommodates better to the dynamic than other protocols relying on addresses (including geographical ones) because many control messages are avoided. Conditional transmissions are implemented as an Airplug compatible application called HOP [6], which has been studied in Network Simulator and tested on the road. For this paper, we use HOP in order to ensure that messages are relayed from car to car in the convoy.

In order to measure performances in the convoy of vehicles, we designed a specific Airplug compatible application named TST. On the sending vehicle (first vehicle of the convoy), TST generates packets of a given size, with a given inter-packet gap (IPG). On the other vehicles, TST receives the packets and computes some metrics (see below). For this purpose, the TST packets include all necessary informations. The complete architecture used for the test is drawn in Figure 2.




Fig. 2. Architecture used for the tests

III. SCENARIOS AND METRICS

A. Scenarios

The previously described architecture has been used with up to 7 vehicles. The measures have been taken for communications between the first car in the convoy (sender) and the last one (receiver). The inter-vehicle distance varies between 325 m and 400 m. The speed is about 76km/h in average.

During the tests, there is no control nor estimation about the range of the wireless antennas. However, the HOP software allows to force the hops in order to obtain a linear convoy and results expressed in terms of number of hops. Without such feature, some messages may reach vehicle $i+2$ from vehicle i without being relayed by vehicle $i+1$ and the results depends on the distance between the cars instead of the number of hops. Both tests have been done; in this paper we report only results depending on the number of hops.

Test	1	2	3	4	5	6
IPG (ms)	100	100	100	100	250	500
# packets	994	611	255	294	170	137
Duration (s)	113	63	25	37	43	68

Fig. 3. Table of tests, with corresponding IPG and number of packets and durations

Among the different tests performed on the road, we retain 6 of them, described in Table 3, where IPG denotes the Inter Packet Gap. We focus in this paper on tests where the IPG is 100 ms because in lab experiments showed that such a value gives the best results with our platform regarding the loss rate.

B. Metrics

This paper deals with the capacity of the vehicular networks, in order to show their ability to support different kind of applications. We then focus on losses and loss events, delay and throughput. The delay is measured from the packet departure from the first vehicle until its arrival to the last vehicle. The loss rate at vehicle i is the ratio of unreceived packets by vehicle i over the total number of packets vehicle i ought to be received.

IV. LOSSES PERFORMANCES IN THE CONVOY

In this section, the losses in the wireless network are discussed and studied. We also study the effect of retransmissions.

A. Results

We study the losses in a network by means of the percentage of reception at each site. Figure 4 shows the variation of the percentage of reception with the number of hops, as well as the average (straight line). We observe that the experiments were very different. This percentage decreases with the number of hops, with a variation of 10% for the first hop, and between 4% and 5% for the rest of the hops. The standard deviation remains acceptable (between 8 and 11%), except for the first hop (21%).

By comparison with the impact of the environmental conditions, the IPG (inter packet gap) influence appears to be very limited. This observation is different from the results in [11] obtained by simulation, where the IPG shows to have influence. The simulator used for the simulations (ns-2) did not follow accurately the environmental conditions.

B. Analysis

The first thing that we can notice is that after 5 hops, an average good reception rate of 10% is guaranteed. It is important to note that there is no retransmission in our tests for the lost packets (there is no acknowledgement). This allows to show the real numbers of losses and the importance of retransmissions either done by unicast WiFi (layer 2) or TCP (layer 4). The retransmissions can be either handled by the protocol or by the application itself. Note that retransmissions are not always pertinent, especially in case




Fig. 4. Percentage of reception at each car of the convoy for 4 experiments with IPG=100 ms

of data regularly produced by sensors. We study in Figure 5 the influence of multiple transmissions.

To increase the success of reception, a message can be sent several times. We then draw a graph (cf. Figure 5) showing the percentage of success of reception based on the number of times a message has been sent and of the number of hops it travels (experiment with IPG=100 ms). For example, if we want to send a packet and to guaranty its reception with 80% of success, then to reach one hop we should send it 2 to 3 times, to reach 2 hops we should send it 4 to 5 times, and to reach 3 hops we should send it 6 times at least. Further than 3 hops, we should send it 5 to 6 times but with only 70% of success at the reception.

These results may appear to be very unacceptable. However, they may be convenient for some applications, generating data periodically. We study now the frequency of such data production allowing to reach vehicles several hops far with a given probability of success. For instance, suppose that an application generates a new message every 1 s, then with an inter packet gap of 100 ms, the same message can be repeated 10 times in order to reach cars as far as possible. The figure displays results for such applications (a single sensor). However the packet size is equal to 1000 Bytes while sensors messages are short. This means that several messages produced by sensors can be aggregated in a single larger message. By sending such a message every 100 ms with redundant data, it is then possible to reach far vehicles. For instance, Figure 6 illustrates the content of the messages that can be sent. These messages are the concatenation of many short messages generated by different sensors. The last received message for each sensor will be placed in the message to send.

C. Loss events

The previous study shows that losses are very important but it is possible to send data successfully by repeating the packets. However such a technique may fail if several successive packets were lost. We then study the engraving of these loss events.

We use these informations to get an idea on how packet losses are distributed, and on how many packets are lost at a time in a loss event (packets lost successively). A loss event is a loss that occurs after a sequence of good receptions. These events happen to be uniformly distributed, with a




Fig. 5. Maximum number of hops a message can travel with a certain number of retransmissions for the packet




Fig. 6. Example of the concatenation of messages of different applications to be deployed in the network

tendency to lose one to two packets every time. A loss of a burst of packets at a time is very rare (cf. Figure 7).

The most important loss events induce the loss of 1 to 5 packets successively. But the percentage of occurrence of these events is important, and the cumulative average of these percentages show that in about 80% of the loss events, 1 to 5 successive packets are lost. As we can see, most of the loss events concern 1 or 2 packets, 80% of these events concern less than 5 packets. However, one may ask the question about the spreading out of the packet losses over the time. These distribution is shown in Figure 8 for Test 1 (cf. Figure 3). The values on the Y axis indicate the number of the last car that received the packet in question. For example, the packet with the sequence number 250 has never been received by car 2, meaning that the last vehicle having this packet is car 1. We




Fig. 7. Event loss occurrence in function of the number of packets lost during this event




Fig. 8. Graph showing the last car to receive a certain packet (indicated to by the number of sequence). The car 1 is the sending car.




Fig. 9. Delays for three experiments with IPG=100 ms

can observe that the important losses follow a certain pattern grouped into periods of 5 s containing 50 packets. These losses result then mainly from the environmental conditions in which the tests were made.

Hence, even if loss events concern mainly one or two packets, they are grouped by periods of 5 s on some given vehicles. This means that external factors affect locally the communications. Among such environmental factors we may quote an increasing of the inter-vehicle distance, the presence of a convoy of trucks, a roundabout, etc.

V. DELAY PERFORMANCES

In this section, we study the delay for original and repeated packets.

A. Results

In the conditions in which our tests were made, environmental and traffic conditions had a big influence on the results, and hence the delay.

As we can see in the Figure 9, when the number of hops grows, the delay grows too. Even if the values vary from an experiment to another, there is a tendency shown by the mean in Figure 9. The standard deviation goes from 59 s on the first hop to 1439 s on the last hop, representing 30% on the first hop and growing to 66% on the last one.

We must note that, like for the losses, the delay is not affected by the IPG.

B. Analysis

For these results and calculations, only the received packets were treated which means that they are highly affected by losses.

As said before, losses highly affect the communication in the convoy, but sending several times the packets allow to




Fig. 10. Delays needed to repeat a packet if repetition is needed in function of the number of hops to travel for Test 1

ensure good reception rates. However, these repetitions will naturally increase the delays between the first packet sending and the arrival of the information several hops later, that may require several packets sending.

We can see for instance that during Test 1 (Figure 10), the delay can exceed security distance when a packet has to be repeated more than five times to reach the third vehicle. By security distance, we mean the time needed for the driver to react to dangerous situations that may appear on the road. It is fixed to 2 s in France. This means that, if the required percentage of success is 90%, then the packet should be sent five time, leading to a total delay larger than the time needed by the driver to react.

The results depend of the experiments, but we can notice that 3 repetitions are allowed in both experiments to reach up to the sixth vehicle.

VI. THROUGHPUT PERFORMANCES

A. Results

Figure 11 shows the throughput for different experiments. The standard deviation is important for the first hop (17 Kb/s), and varies between 5 and 7 Kb/s for the other hops. This represents 38% of the mean value for the first hop, and reaches 83% for the last hop.

B. Analysis

It is important to note that when sending a message of 1000 bytes every 100 ms, the theoretical throughput is 80 Kb/s. Hence in Figure 11, the mean throughput is close to 50% of the maximal theoretical throughput after the first hop, and close to 25% after the second hop. These results are acceptable for road safety applications that have few amount of data to transmit but are very sensible to the delay and percentage of success.

VII. CONCLUSION

This paper shows some on-road experimental results of a linear IEEE 802.11 vehicular network. We analyse losses, loss events, delays and throughputs.

We show that the 802.11b communications offer in real conditions, performances that suit safety applications such as road safety and driving assistance applications. Although losses are important, they can be reduced by repeated




Fig. 11. Throughputs for three experiments with IPG=100 ms

transmissions. Delays increase with the number of hops and the number of retransmissions but a limited number of retransmissions is allowed to ensure the reception of data few hops farther, while still respecting the safety delay. For this type of communications, our study indicates that good reception rates and acceptable delays can be obtained by repeating packets and aggregating data produced by on-board sensors.

As a perspective, our work may lead to some optimized data transfert protocols for linear vehicular networks.

REFERENCES

- [1] Pre-Drive C2X.
- [2] SIM-TD Field test of car2X communication.
- [3] A. Buisset, B. Ducourthial, F. El Ali, and S. Khalfallah. Vehicular networks emulation. *ICCCN 2010*, 2-5 août 2010.
- [4] V. Bychkovsky, B. Hull, Balakrishnan Miu, A., H., and S. Madden. A measurement study of vehicular internet access using in situ wi-fi networks. *International Conference on Mobile Computing and Networking (MobiCom)*, Los Angeles, CA, USA, page 50–61, 2006.
- [5] B. Ducourthial. About efficiency in wireless communication frameworks on vehicular networks. *Proceeding of the ACM WIN-ITS workshop colocated with IEEE ACM QShine'07*, 2007.
- [6] B. Ducourthial, Y. Khaled, and M. Shawky. Conditional transmissions: performances study of a new communication strategy in VANET. *IEEE TVT*, 56(6):3348–3357, November 2007.
- [7] B. Ducourthial and S. Khalfallah. A platform for road experiments. *Proc. of the 69th IEEE VTC 2009-Spring, Barcelona*, 2009.
- [8] R. Gass, J. Scott, and C. Diot. Measurements of in-motion 802.11 networking. *IEEE Workshop on Mobile Computing Systems & Applications (WMCSA)*, page 69–74, April 2006.
- [9] M. Goncalves Rubinstein, F. Ben Abdesslem, S. Rodrigues Cavalcanti, M. Elias Mitre Campista, R. dos Santos Alves, L. Henrique Maciel Kosmalski Costa, M. Dias de Amorim, and O. Carlos Muniz Bandeira Duarte. Measuring the capacity of in-car to in-car vehicular networks. *Proc. of the 69th IEEE VTC 2009-Spring, Barcelona*, 2009.
- [10] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal. Vehicular opportunistic communication under the microscope. *ACM International Conference on Mobile Systems, Applications, and Services (MobiSys)*, San Juan, Puerto Rico, June 2007.
- [11] Y. Khaled, B. Ducourthial, and M. Shawky. IEEE 802.11 performances for inter-vehicle communication networks. *Proceedings of the 61st Semiannual Vehicular Technology Conference (VTC 2005-Spring)*, Stockholm, Sweden, May-June 2005.
- [12] J. Luo and J.-P. Hubaux. A survey of inter-vehicle communication. *School of Computer and Communication Sciences, EPFL, Switzerland*, technical report IC/2004/24.
- [13] J. Ott and D. Kutscher. Drive-thru internet: IEEE 802.11b for “automobile” users. *IEEE Conference on Computer Communications (INFOCOM)*, Hong Kong, March 2004.
- [14] Khalfallah S. and B. Ducourthial. Bridging the Gap between Simulation and Experimentation in Vehicular Networks. *2010 IEEE 72nd Vehicular Technology Conference (VTC2010-Fall)*, 6–9 September 2010.
- [15] M. Wellens, B. Westphal, and P. Mahonen. Performance evaluation of IEEE 802.11-based w lans in vehicular scenarios. *IEEE Vehicular Technology Conference, Dublin, Ireland*, Apr. 2007.