

# Measuring IEEE 802.11p Performance for Active Safety Applications in Cooperative Vehicular Systems

Francesca Martelli, M. Elena Renda, and Paolo Santi  
IIT - CNR, Pisa – Italy

**Abstract**—In this paper, we present a measurement study of application layer performance in IEEE 802.11p vehicular networks. More specifically, our focus is on active safety applications, which are based on the exchange of beacon messages containing status information between close-by vehicles. We consider two performance metrics relevant to active safety applications: the first is application-layer goodput, which can be used to optimize congestion control techniques aimed at limiting the beaconing load on the wireless channel; the second is the beacon reception rate, which is useful to estimate the level of situation awareness achievable onboard vehicles. Our measurements were conducted using a prototypal, 802.11p compliant communication device developed by NEC, in both stationary and mobile V2V scenarios, and disclosed several useful insights on 802.11p application-level performance. To the best of our knowledge, the ones presented in this paper are the first application-level measurements of IEEE 802.11p based vehicular networks reported in the literature.

## I. INTRODUCTION

Given their potential of considerably reducing car accidents and improving traffic conditions, active safety applications are of primary importance in vehicular networking. This explains the extensive interest this class of applications has gathered in the research and industry community. Despite these intensive efforts, active safety application performance has been mostly evaluated through analysis and simulation. This is mainly due to the fact that the IEEE 802.11p standard for short range vehicular communications, upon which active safety application will be built, is still in draft form. Thus, 802.11p compliant hardware is still difficult to find, and often turns out to be rather expensive.

To the best of our knowledge, no application-level IEEE 802.11p measurement study has been reported in the literature so far. Existing 802.11p performance evaluation has been done only by means of analysis and/or simulation [1], [2], [3], or even testbeds for 802.11p physical layer measurements [4], [5]. Some work has been done on measuring performance of legacy IEEE 802.11a/b/g in vehicular environments [6]. The only IEEE 802.11p measurement-based papers we are aware of are concerned with characterization of PHY layer parameters of the 5.9 GHz radio channel used by IEEE 802.11p, such as multi-path propagation [7], received signal strength [8], delay spread [5], and so on. In a recent paper [9], Paier et al. investigate PHY layer performance of 802.11p infrastructure-to-vehicle links, mostly evaluating the communication range (defined as the range up to which a percentage  $p$  of MAC frames can be correctly received).

In this paper, we present for the first time the results of a measurement study of *application layer* performance in IEEE 802.11p vehicular networks. More specifically, given their importance as explained above, our focus is on *active safety applications*, which are based on the periodic exchange of beacon messages containing status information between close-by vehicles.

Our study considers two performance metrics relevant to active safety applications, namely *application-layer goodput* and *beacon delivery rate*. The former metric is important to, e.g., allow a careful optimization of congestion control techniques aimed at limiting the beaconing load on the wireless channel, such as the DFPAV approach proposed in [10]. In DFPAV, transmission power used to send beacons is tuned with the goal of keeping the resulting beaconing load below a certain threshold, called the maximum beaconing load (MBL) threshold. In turn, the MBL threshold is set in such a way that sufficient capacity is left on the channel to transmit event-driven emergency messages. Thus, knowing the application-layer goodput is fundamental to adequately tune the MBL threshold in the DFPAV approach.

The second considered metric is the beacon reception rate, i.e., the fraction of beacons that are correctly received by surrounding vehicles. This metric is very important to estimate the level of situation awareness achievable onboard vehicles. Furthermore, other congestion control mechanisms proposed in the literature builds upon communication models where a certain beacon reception probability is assumed [11]. Our results can be used to adequately tune such parameters in these congestion control techniques.

The measurements reported in this paper were conducted using a prototypal, 802.11p compliant communication device developed by NEC, in both stationary and mobile V2V scenarios, and disclosed several useful insights on 802.11p application-level performance, which are summarized in Section IV.

## II. HARDWARE AND SETUP

### A. Hardware

To perform our measurements, we used two IEEE 802.11p compliant devices, namely the LinkBird-MX v3 units produced by NEC. The LinkBird-MX units are embedded Linux machines (kernel 2.6.19) based on a 64 bits MIPS processor working at 266Mhz that can be configured to work both as RSU (Road Side Unit) and OBU (On Board Unit). LinkBird-MX units are endowed with several network interfaces: an



Fig. 1. The on-board unit setup.



Fig. 2. The NEC LinkBird unit with WiMo antenna.

Ethernet interface for communication with other onboard or roadside devices, a GPS connector, a wireless LAN interface based on IEEE 802.11p Draft 3.0, and an optional wireless LAN legacy (802.11a/b/g) interface. The characteristics of the 802.11p network interface are reported in Table I.

As it is known, both in Europe and US several channels have been reserved in the 5.9 GHz band for vehicular systems, and one of them is reserved for safety applications only (the control channel). Given our focus on measuring active safety application performance, we used the control channel for our measurement, which in Europe is channel 180 at 5.9 GHz.

LinkBird-MX units were connected to a WiMo antenna, whose characteristics and form factor fit well with vehicular applications. Technical details of the antenna used in our measurements are reported in Table II. The antennas were mounted on the roof of the vehicles for performing vehicle-to-vehicle radio measurements. The onboard equipment on each vehicle, composed of a LinkBird-MX unit, a GPS receiver, a laptop, and a rooftop antenna (not shown) is reported in Figure 1. A detailed view of the LinkBird-MX unit with the rooftop antenna is reported in Figure 2.

Parameter	Details
Frequency/Channel	5725 – 5925 MHz
Bandwidth	10 – 20 MHz
Version	IEEE802.11p Draft 3.0, July 2007
Transmit Power	Max 21 dBm (in Europe)
Bitrates (10MHz)	3, 4.5, 6, 9, 12, 18, 24, 27 Mbps
Bitrates (20MHz)	6, 9, 12, 18, 24, 36, 48, 54 Mbps

TABLE I  
CHARACTERISTICS OF LINKBIRD-MX 802.11P NETWORK INTERFACE.

Antenna type	rubber antenna
Model no.	18720.3h
Frequency range	5.5-5.8 GHz
Impedance	50 Ohm
Gain	5 dBi
Radiation	omni
Polarization	vertical
Length	108 mm

TABLE II  
ANTENNA DETAILS.

### B. Experiments setup

Together with the LinkBird-MX devices, NEC provides the Car-2-X SDK software toolkit, which abstracts the 802.11 protocol stack and offers Java API to interact with it. The abstracted protocol stack, called C2X stack, provides application developers with several communication primitives, including




Fig. 3. Transmission packet flow.

unicast, topologically- and geographically-scoped broadcast. Furthermore, the standard single-hop broadcast primitive, according to which a packet is received by all vehicles within the transmitter’s communication range, is also provided. Given our focus on active safety applications, and more specifically on the beacon exchange phase upon which such applications are built, all packets were transmitted in single-hop broadcast in our measurements.

We developed a very simple Java application, which generates UDP packets at a regular time interval. The time interval can be set by the user with a  $1ms$  granularity, allowing the application to generate as many as 1000 packets/sec. Other tunable parameters are the packet payload size, and the duration of the transmission phase. The Java application also counts the number of received packets while running.

Each LinkBird unit has been connected to a laptop by means of the Ethernet interface (used for transmitting/receiving packets) and a serial port (used for control purposes). Given the limited computational power of LinkBird units, we run the CBR application on the laptop computer. Packets generated by CBR application are sent to the LinkBird unit through the Ethernet interface and received by the C2X daemon running on the LinkBird, which forwards received packets to the 802.11p wireless card (see Figure 3). A similar, but reversed, packet flow is implemented at the receiver end of a communication.

We have performed two different measurement campaigns: a first campaign aimed at measuring the application-layer goodput in different conditions, and a second campaign aimed at evaluating V2V beaconing performance.

Application-layer goodput is measured as the total amount of data received by the Java application in a time interval, typically set to  $1sec$ . This value is computed by multiplying the number of packets received by the application in the time interval by the packet *payload* size. Note that application-layer goodput is necessarily lower than the raw PHY layer data rate,



Fig. 4. The aerial map of the area for mobile measurements.

due to the presence of PHY and MAC packet headers, the backoff time, and the various interframe spaces as defined in 802.11.

Beaconing performance is evaluated by computing the percentage of correctly received beacon packets over the total number of transmitted beacons, i.e., the well-known packet delivery rate (PDR) metric.

In both measurements campaigns, packets were transmitted using the lowest coding rate (1/2) with BPSK modulation, corresponding to a PHY layer raw data rate of  $3Mbps$  with  $10MHz$  channel bandwidth, and of  $6Mbps$  with  $20MHz$  bandwidth. Usage of the lowest available data rates is mandatory for active safety applications, due to the typical usage of single-hop broadcast packets (which can be sent only at the lowest data rate), and to the better communication reliability provided by the lowest data rate.

In goodput measurements, we considered packets with payload size set to 100, 500, 1000, and 1500 bytes, respectively. The first two payload values are representative of a typical beacon size with (500) and without (100) security-related overhead [12], [10], while the larger values were chosen to investigate the saturation limit of the wireless channel. In beaconing measurements, only the relevant payload sizes of 100 and 500 bytes were considered.

### III. PERFORMANCE MEASUREMENTS

#### A. Goodput measurements

We performed three sets of measurements, aimed at measuring throughput in (1) ideal conditions, and in a (2) relatively high or (3) relatively low SNR regime. For each set of measurements, we performed 10 experiments, whose average goodput values are reported in Figure 5.

Measurements for scenario (1) were performed indoor in our lab, with the two LinkBird units separated only by a few meters. Since 802.11p operates on a reserved frequency band sufficiently separated from 802.11a frequencies, we observed a negligible interference despite the several WLANs operating in the building. This is witnessed by the very high goodput values measured, as reported in Figure 5-a) and b). Goodput values of as high as  $2.5Mbps$  for the  $3Mbps$  PHY data rate, and  $4.9Mbps$  for the  $6Mbps$  PHY data rate have been measured

for the largest payload size. Application-layer goodput values correspond to about 83% of the PHY raw data rate, meaning that 802.11p design proved very effective in reducing MAC layer overhead with respect to the legacy 802.11 protocol. Clearly, the achievable goodput decreases with payload size, with a measured goodput value of about  $2.1Mbps$  (respectively,  $3Mbps$ ) with 500 bytes packets and  $3Mbps$  (respectively,  $6Mbps$ ) PHY data rate. For the smallest packet size value, the measured goodput drops to about  $0.75Mbps$  with both PHY data rates.

Measurements for scenarios (2) and (3) were performed outdoor, using two stationary vehicles for measurements. Vehicles were positioned along the side of a road surrounding the CNR Research Area in Pisa, and positioned at a distance of about  $30m$  for scenario (2) (relatively high SNR regime), and at about  $150m$  for scenario (3) (relatively low SNR regime). Note that  $30m$  approximately corresponds to the safety distance between vehicles traveling at  $100km/h$ , and it is thus representative of the situation in which two consecutive vehicles exchange packets in a relatively dense traffic scenario. On the other hand,  $150m$  is a typical distance up to which beacon reception is more relevant in active safety applications.

Goodput values obtained in scenario (2) and (3) are reported in Figure 5 c)–f). As seen from the plot, in the high SNR regime goodput values only slightly lower than in ideal conditions are observed at both data rates, with the exception of the largest packet size at  $3Mbps$ . With this relatively large packet size and relatively low data rate, the packet on-air time is relatively long (about  $4msec$ ), resulting in a very high likelihood of incurring in a bit error burst causing incorrect packet reception. This explains why the goodput value with 1500 bytes packets at  $3Mbps$  is close to 0. On the other hand, at  $6Mbps$  the packet on-air time is halved, and probability of successful packet reception considerably increases, resulting in the highest goodput value among the considered packet sizes.

In the low SNR regime, the goodput drop for relatively large packet sizes becomes evident: for both  $3Mbps$  and  $6Mbps$  data rates, the packet size resulting in the highest goodput values is 500 bytes; with larger packets, the probability of packet reception drops considerably, leading to lower goodput values, which are close to 0 in case of 1500 bytes packets. It is also worth observing that, as expected, the goodput values obtained are much lower than those obtained in ideal conditions, with a maximum goodput of about  $1.9Mbps$  at  $3Mbps$ , and of about  $0.7Mbps$  at  $6Mbps$ . The lower goodput value observed at  $6Mbps$  can be explained with the fact that our experiments were conducted in an environment open to vehicular traffic. Although vehicles were parked aside the road during experiments, movement of surrounding vehicles could not be controlled. When the separation distance between vehicles was  $150m$ , in some cases partial obstructions of the LOS between the vehicles were observed due to such uncontrolled vehicular movements, leading to a degraded communication performance.




Fig. 5. Measured goodput in the different experiments, with PHY data rate set to  $3Mbps$  and  $6Mbps$ .

PHY data rate	payload size	distance	Test 1		Test 2		Test 3		Avg
			A	B	A	B	A	B	
3 Mbs	100 byte	short	0.89	-	0.92	0.92	0.85	0.86	<b>0.89</b>
3 Mbs	500 byte	short	0.84	0.84	0.90	0.91	0.98	0.99	<b>0.91</b>
6 Mbs	100 byte	short	0.88	0.83	0.92	0.86	-	0.91	<b>0.88</b>
6 Mbs	500 byte	short	0.90	0.88	0.74	0.76	0.94	-	<b>0.84</b>
3 Mbs	100 byte	long	0.48	0.45	0.71	0.68	0.57	0.53	<b>0.57</b>
3 Mbs	500 byte	long	0.49	0.45	0.53	0.49	0.57	0.53	<b>0.51</b>
6 Mbs	100 byte	long	0.53	-	0.49	-	0.45	-	<b>0.49</b>
6 Mbs	500 byte	long	0.61	-	0.54	-	0.68	-	<b>0.61</b>

TABLE III  
MEASURED BEACON DELIVERY RATES.

## B. Beacon measurements

In order to measure the beaconing performance, we conducted an extensive number of tests with moving vehicles. The tests were performed on a stretch of near straight-line, 4 lanes road close to the CNR Research Area in Pisa, as depicted in Figure 4. The stretch of road is about  $1.5Km$  long, and it is delimited by two roundabouts. In each test, vehicles started from roundabout *A*, traveled up to roundabout *B*, and traveled back to the starting point. The cruising speed of vehicles was about  $70Km/h$ , corresponding to the road speed limit. Given the above configuration, the duration of each test is about  $3min$ . We then configured the CBR applications on both vehicles to send a packet (beacon) every  $100msec$  – corresponding to the recommended beaconing frequency for active safety applications [13]– over a  $3min$  time interval, corresponding to a total of 1800 beacons sent per vehicle for each test. At the end of each test, the number of received beacons on the two vehicles (named vehicle *A* and *B* in the following) were recorded, and used to compute the PDR values.

Two sets of tests have been performed: in the former set (named “short” distance in the following), vehicles *B* followed vehicle *A* at safety distance (approx.  $25m$ ), as visually evaluated by the driver. This scenario models a typical car following situation, with LOS between vehicles for most of the time – although occasional LOS obstructions due to

overtaking maneuvers have been observed. In the second set of measurements (named “long” distance in the following), vehicle *B* followed vehicle *A* at a distance of about  $150m$ , visually evaluated by the driver. In this scenario, relatively frequent LOS obstructions are expected, due to both the larger distance between vehicles, and presence of regular vehicular traffic.

For both sets of measurements, we considered beacons of 100 and 500 bytes, with PHY data rate set to  $3Mbps$  and  $6Mbps$ . For each parameter set, we performed 3 different tests, for a total of 24 tests overall. This corresponds to an overall traveled distance of  $72Km$  for the whole set of measurements.

The results of our measurements are reported in Table III. Missing entries refer to failed measurements in one of the vehicles. These failed measurements were due to a certain instability of the Ethernet connection between the laptop and the LinkBird unit, most probably caused by insufficient power supply when laptop batteries were running low.

The results of our measurement campaign disclose several interesting insights on beaconing performance. First, we observe the very good performance in the short distance scenario, with both beacon sizes and PHY data rates. It is also interesting to observe that, in this scenario, the wireless channel turned out to be almost perfectly symmetric. This is due to the prominent LOS conditions with relatively high SNR values observed in the short distance scenario. As expected,

the situation is quite different in the long distance scenario, where considerably degraded PDR values (around 0.6-0.5) are observed. Based on our direct observations during the tests, the degraded PDR performance is mostly due to the relatively frequent LOS obstructions, while the impact of distance (and, thus, of the lower SNR regime as compared to the short distance scenario) is less important: we have observed that, in clear LOS conditions, close to optimal PDR values are recorded also at a longer distance. It is also interesting to observe that the channel remains nearly symmetric also in the long distance scenario, indicating that LOS obstructions are very likely to corrupt communications in both directions. Finally, we observe that packet size has a negligible effect on PDR values in both the short and long distance scenario<sup>1</sup>.

#### IV. DISCUSSION AND FINAL CONSIDERATIONS

The results of our measurement campaign disclosed several useful insights which might help vehicular network and active safety application designers.

**Goodput:** our goodput measurements show that very high goodput values can be obtained not only in ideal conditions, but also at a relatively short distance approximately corresponding to safety distance in a highway. Furthermore, for packet size of 500 bytes (typical size of a beacon including security overhead) the measured goodput value showed only a minor drop – from 2.1Mbps to 1.9Mbps – also at a relatively long distance of 150m. Thus, relatively stable goodput values can be expected in the range of distances and packet sizes considered in active safety applications. Note that our measured goodput values can be directly used to feed and optimally tune congestion control mechanisms such as the DFPAV approach of [10], which are otherwise based on nominal 802.11p data rates.

**Symmetry of wireless channel:** our PDR measurements suggest that, based on our experience and restricted to the considered scenarios, the assumption of symmetric wireless channel is quite accurate for beaconing applications.

**LOS:** our PDR measurements have confirmed that 802.11p communication performance is severely impaired in NLOS conditions. We have experienced that even a slight curve in the road with a road side hedge is sufficient to cause a major drop in beacon reception rates. In the long distance scenario, we have observed that a small van in between the communicating vehicles is sufficient to significantly impair communication performance. In one case, we had a big truck positioned between the two vehicles, and the PDR dropped to 0.

**Design guidelines:** the previous observation about poor 802.11p performance in NLOS conditions brings several issues to the attention of the vehicular network designer and active safety application developer. Given vehicular network's goal of extending driver's awareness beyond human eye, and the poor 802.11p performance in NLOS conditions, actions need to be undertaken in order for vehicular networks to fulfill their

promises. Actions can be undertaken at a governmental level (e.g., enforcing adoption of 802.11p communication technology for trucks and big vans), at a wireless communication level (e.g., exploiting antenna diversity to improve SNR values), at the network planning stage (e.g., suitably deploying RSU with the purpose of extending vehicles' awareness of surrounding conditions in critical locations – e.g., along a curve), at a networking level (e.g., through aggregating and relaying in a multi-hop fashion vehicle status information contained in the beacons), or, preferably, at all these levels.

Before concluding this paper, we want to stress that the set of measurements presented in this paper is admittedly limited, and can be considered as a starting point towards a better understanding of active safety applications' on-the-field performance. More efforts are needed, e.g., to better characterize the beacon PDR performance at various distances, to evaluate the inter-packet reception time [14], etc. We are currently actively working on this.

Finally, we want to stress that our measurements were obtained in a non-congested wireless channel, with only two vehicles exchanging packets. As such, our results can be considered as quite optimistic. Investigating the effect of channel congestion on active safety application performance is another avenue for future measurement studies.

#### REFERENCES

- [1] S. Eichler, "Performance evaluation of the IEEE 802.11p wave communication standard," in *Proc. IEEE VTC Fall*, 2007, pp. 2199–2203.
- [2] K. Bilstrup, E. Uhlemann, E. Strom, and U. Bilstrup, "Evaluation of the IEEE 802.11p MAC method for vehicle-to-vehicle communication," in *Proc. IEEE VTC Fall*, 2008, pp. 1–5.
- [3] S. Wang and et al., "Improving the channel utilization of IEEE 802.11p/1609 networks," in *Proc. IEEE WCNC*, 2009, pp. 1–6.
- [4] L. Cheng, B. Henty, D. D. Stancil, F. Bai, and P. Mudalige, "A fully mobile, GPS enabled, vehicle-to-vehicle measurement platform for characterization of the 5.9 GHz DSRC channel," in *Proc. IEEE Symp. Antennas Propagation*, 2007, pp. 2005–2008.
- [5] I. Tan, W. Tang, K. Laberteaux, and A. Bahai, "Measurement and analysis of wireless channel impairments in DSRC vehicular communications," in *Proc. IEEE ICC*, 2008, pp. 4882–4888.
- [6] M. Wellens, B. Westphal, and P. Mahonen, "Performance evaluation of IEEE 802.11-based WLANs in vehicular scenarios," in *Proc. IEEE VTC-Spring*, 2007, pp. 1167–1171.
- [7] L. Cheng, B. Henty, R. Cooper, D. Stancil, and F. Bai, "Multi-path propagation measurements for vehicular networks at 5.9 GHz," in *Proc. IEEE WCNC 2008*, 2008, pp. 1239–1244.
- [8] L. Cheng, B. Henty, F. Bai, and D. Stancil, "Highway and rural propagation channel modeling for vehicle-to-vehicle communications at 5.9 GHz," in *Proc. IEEE Symp. Antennas Propagation*, 2008, pp. 1–4.
- [9] A. Paier and et al., "Average downstream performance of measured IEEE 802.11p infrastructure-to-vehicle links," in *Proc. Workshop on Veh. Connectivity*, 2010.
- [10] M. Torrent-Moreno, J. Mittag, P. Santi, and H. Hartenstein, "Vehicle-to-Vehicle Communication: Fair Transmit Power Control for Safety-Critical Information," *IEEE Trans. on Vehicular Technology*, vol. 58, pp. 3684–3703, 2009.
- [11] X. Yang, J. Liu, F. Zhao, and N. Vaidya, "A Vehicle-to-Vehicle Communication Protocol for Cooperative Collision Warning," in *Proc. MobiQuitous*, 2004, pp. 1–10.
- [12] M. Raya and J. Hubaux, "The security of vehicular networks," in *Proc. ACM SASN*, 2005, pp. 11–21.
- [13] E. T. ITS, "Intelligent Transportation Systems (ITS): Vehicular Communications, Basic Set of Applications, Definitions," *ETSI TR 102 638 V1.1.1*, June 2009.
- [14] T. ElBatt, S. Goel, G. Holland, H. Krishnan, and J. Parikh, "Cooperative Collision Warning Using Dedicated Short Range Wireless Communications," in *Proc. ACM VANET*, 2006, pp. 1–9.

<sup>1</sup>We stress that the relatively higher PDR values for relatively larger packet size in some tests have likely been caused by the different, not repeatable traffic conditions during the tests.