Performance Evaluation of IEEE 802.11-based WLANs in Vehicular Scenarios

Matthias Wellens, Burkhard Westphal and Petri Mähönen
Department of Wireless Networks, RWTH Aachen University
Kackertstrasse 9, D-52072 Aachen, Germany
Email: {mwe, bwe, pma}@mobnets.rwth-aachen.de

Abstract—**Communication between cars can be exploited for various applications. Security enhancements as well as interactive games played by occupants of different cars are only two examples. The IEEE 802.11p standard will specify physical and MAC-layer for such scenarios. As one of the design goals is to partially reuse existing WLAN technology we investigate in this paper the performance of IEEE 802.11a, 802.11b and 802.11g devices in car-to-car (C2C) as well as car-to-roadside (C2R) scenarios. The determined performance is mostly independent of the car velocity up to the maximum measured speed of 180 km/h. The major impact factors found are the C2C distance, the availability of line-of-sight between sender and receiver and the rate adaptation algorithm.**

I. Introduction

Recently car safety has been continuously improved. In order to achieve further substantial advance it would be extremely useful if car accidents could be forecasted. Car to Car (C2C) and Car to Roadside (C2R) communication is a promising technology in this respect because it enables the car to retrieve information about its surroundings and to take appropriate actions to secure the occupants if an accident is forecasted. Examples of typical use cases are traffic jams starting after a curve or slippery areas on the road about which cars inform each other automatically. Additionally, several other applications such as adaptive navigation systems are possible because other cars can be informed about situations, such as temporarily blocked roads, and the navigation system can redirect the driver so that the jammed street can be avoided.

Several research projects have considered car communication aspects. The FleetNet project [1] evaluated position-based addressing and routing schemes. Furthermore, also lower layer aspects were considered and in [2] the authors found UTRA TDD to be a more appropriate radio technology for inter-vehicular communications than IEEE 802.11b. Singh *et. al.* reported in [3] about WLAN measurements in different vehicular scenarios and showed in [4] how link layer information can be used to enhance inter-vehicular ad-hoc routing. Bergamo *et. al.* investigated the WLAN 802.11b performance in C2R as well as in C2C scenarios [5]. They fixed the data rate to 2 Mbps during all tests and focused on jitter and handover aspects in more detail. Ott *et. al.* performed C2R measurements with WLAN equipment [6] and used their results as a basis for a delay tolerant network architecture, the Drive-Thru Internet [7].

On the standardization side the IEEE task group 802.11p [8] specifies MAC enhancements and a new physical layer for car communications. One of the design goals during the 802.11p development is to base the new standard on existing technologies available for WLANs used in static environments and only improve certain features where necessary.

This paper presents results from an extensive measurement campaign evaluating the performance of present off-the-shelf WLAN hardware for C2C as well as for C2R scenarios. We used the three major flavors of IEEE 802.11 compliant physical layers 802.11a, b, and g. We also included 802.11a although it is rarely used today because it works in the 5 GHz ISM-band and 802.11p is planned to work in the licensed 5.9 GHz band. However, also the deployment of standard WLAN hardware could be commercially interesting for low-cost, not security-related applications such as games played by occupants of different cars. We evaluated the performance under various different parameter settings and investigated the impact of the velocity and type of environment.

The remainder of this paper is structured as follows. In section II we will present the measurement setup in terms of deployed hardware, software and the measurement environments. We continue in section III with presenting the main findings of the measurement campaign and put those in a larger context in section IV. We conclude the paper in section V.

II. Measurement Setup

We measured one hop performance using two notebooks and up to two additional notebooks as passive sniffers to track all traffic sent over the WLAN channel. We switched off any internal WLAN device and used a PC-card with an external antenna plugged to it. We chose the NEC WARP Star WL 54 AG utilizing an Atheros chipset implementing 802.11a, b, and g. Additionally, we used the Z-Com XI-325HP+ based on the Prism 2.5 chipset implementing only 802.11b. The measurement laptops run Red Hat Fedora Core 4 and the MadWiFi Rev 1208 device driver when using the NEC cards. When running the Z-Com cards we used the wlan-ng device driver. The sniffers run Windows XP and the Airopeek WLAN sniffer software that tracks all MAC-layer frames including management and control frames. All laptops were equipped with GPS receivers allowing us to analyze the velocity and position of each device at any time during the measurement campaign.

1550-2252/$25.00

WLAN card	WLAN channel	Linux tx power	Measured tx power (EIRP)
NEC	11	18 dBm	4 dBm
NEC	36	18 dBm	5 dBm
NEC	161	18 dBm	-28 dBm
Z-Com	11	N/A	4 dBm

TABLE I
MEASURED TRANSMIT POWER OF THE DEPLOYED DEVICES.

WLAN channel	Frequency [GHz]	S11 [dB]	S12 [dB]
11	2.462	-19	-28
161	5.805	-6	-39

TABLE II
S11 AND S12 PARAMETERS OF USED ANTENNAS AND CABLES.

We always run one notebook in managed mode, meaning that it acted as a WLAN access point (AP), in the following called server. The other notebook was set to default configuration, in the following called client. We used iperf to generate UDP as well as TCP traffic. Self-developed scripts controlled the measurements and adapted the parameter settings based on predefined configuration files. In addition to throughput measurements we tracked the signal to noise ratio (SNR) as provided by the device drivers and the frame error rate (FER). As the FER is based on the sniffer trace it does not perfectly resemble the conditions at the client. However, the sniffer also lists erroneously received frames so that only frames for which the detection of the physical layer preamble failed at the sniffer interface were not considered. The analysis of the FER traces showed that these cases happen rarely.

Reference measurements with an Agilent Technologies E4440A spectrum analyzer showed that the transmit power numbers reported by the MadWifi device driver are not correct. Taking into account losses of the cables and the adapters between PC-card and spectrum analyzer of about 5 dB still leads to differences of up to 41 dB as presented in table I. We performed all measurements in the 2.4 GHz band in WLAN channel 11 at 2.462 GHz. We chose the highest channel 161 at 5.805 GHz for the 802.11a measurements because 802.11p will work in the 5.9 GHz band and thus at a very similar frequency.

We used external antennas with 5 dBi gain for all measurements and investigated their parameters including the finally used cables with an Agilent Technologies E5071B network analyzer. For the 2.4 GHz antennas we used magnetic foots to mount the antennas to the roof of the cars. As the respective 5 GHz magnetic foot had a very high attenuation we chose better cables and mounted the antennas directly to roof racks. However, table II shows that the performance of the 2.4 GHz antenna is still much better although better cables and no magnetic foots were used in the 5 GHz case.

As cars we used a BMW 525d for the C2R measurements and two Ford Mondeo estate cars for the C2C measurements.

Standard	TCP 1460 B	UDP 125 B	UDP 750 B	UDP 1250 B
802.11a	29.4 Mbps	8.1 Mbps	27.5 Mbps	34.1 Mbps
802.11b	6.7 Mbps	2.3 Mbps	6.3 Mbps	7.3 Mbps
802.11g	29.7 Mbps	8.1 Mbps	27.3 Mbps	33.8 Mbps

TABLE III
REFERENCE MEASUREMENTS IN STATIC INDOOR ENVIRONMENT.

A. *Geographical locations*

We performed the measurements in four different environments. For the C2R scenario we had the opportunity to measure on the German highway A44 which was closed a couple of weeks before for tearing it away as preparation of coal mining. We got access to about 2 km of still fully functional road and could perform measurements without any other car traffic and thus in a rather controllable environment. The server was placed in the middle. Nearly at the whole road of 2 km length line of sight (LOS) between client and server was available. However, three bridges influenced the wave propagation and a slightly curved track blocked the LOS at one end.

The C2C measurements were performed in three different environments in the surroundings of the city of Aachen. We chose one urban scenario with higher buildings, lots of car traffic and less vegetation. The suburban setup was in a residential quarter with smaller houses, gardens in front and less car traffic. Finally, we also measured C2C communication in highway scenario.

III. RESULTS

As a starting point we report the results of reference measurements in our laboratory to present the maximum achievable throughput with our setup. Table III lists these results for different payload sizes and a distance of 1 m between server and client.

A. *Car to Roadside measurements*

We begin with reporting about our results from the C2R measurements run at the A44 area. We performed measurements driving in both directions and also transmitting in both directions. As expected all four cases resulted in similar results so that we will not specify the driving or transmitting direction for each of the following graphs. Additionally, we performed all tests several times also leading to similar results.

Figure 1 shows the UDP goodput measured with 802.11g and 802.11b also comparing different payload sizes. Besides a short period at the beginning when still speeding up the car we drove constantly 120 km/h. The maximum achievable goodput is only slightly smaller than the result in a static environment showing that WLAN usage at such velocities is feasible. However, the reception performance especially for large packets is worse compared to smaller packets so that the goodput improvement because of smaller overhead is partially spend for a higher number of retransmissions.

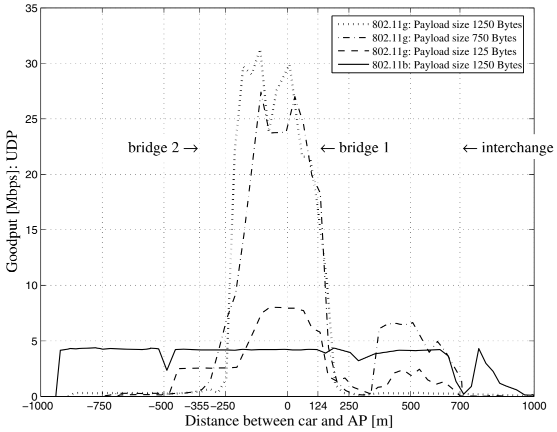

Fig. 1. Goodput measured using UDP datagrams with variable payload size and both 802.11b and 802.11g.

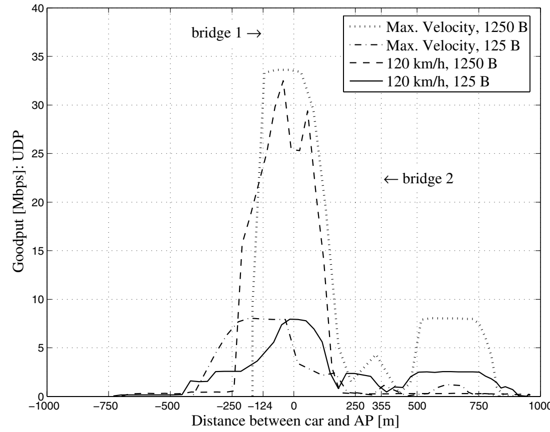

Fig. 3. Goodput measurements using UDP datagrams with various payload sizes and different velocities.

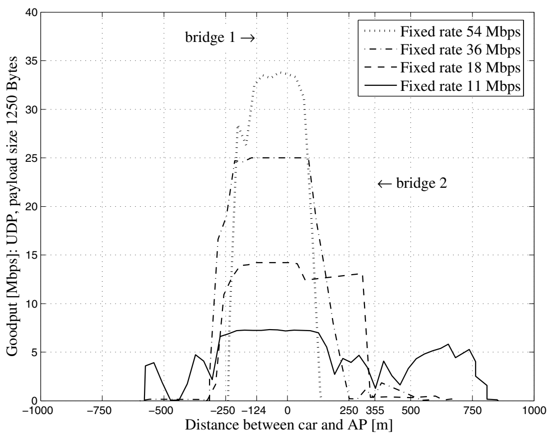

Fig. 2. Goodput measured using UDP datagrams with 1250 Bytes payload size and fixed physical layer rate during each measurement.

The SNR measurements over distance led to a symmetrical graph as expected. As the impact of the bridges is clearly visible on the goodput graph multipath effects can be determined as reasons for the lowered performance.

The performance of the 802.11b card is more stable over the whole distance due to mainly two reasons. Firstly, the card uses two antennas leading to a diversity gain. Secondly, the area reachable with the maximum data rate is larger so that no rate adaptation algorithm has any impact on the performance.

Figure 2 shows the same measurement in the other direction but setting the 802.11g card to fixed physical layer data rates. It can be seen that the communication range is larger for less aggressive coding and modulation schemes (in the following called PHY mode). When comparing both presented graphs we can determine that the automatic rate adaptation algorithm applied is not able to adapt fast enough so that the maximum possible goodput is achieved. The MadWiFi driver uses the SampleRate algorithm as described in [9], which blocks a rate after four consecutive transmission failures for a period of 10 sec. The SampleRate algorithm is optimized for static mesh networks and our results show that the period of 10 sec is too long for vehicular environments. When the distance to the server is still large the algorithm tries faster data rates but blocks those because of transmission failures. When the blocking period elapses the client is already very near to the server so that fast PHY modes can be used without using intermediate levels before.

In the case of TCP traffic the effect is less severe because less packets are sent before a reliable connection is available. Thus, less PHY modes are blocked when reaching the communication range so that the rate adaptation algorithm is able to choose a faster rate.

In order to evaluate the impact of the vehicle speed we did the same measurements accelerating to the maximum possible speed. As the available track-length was limited we could not drive at a constant speed during the whole measurement but passed the AP at a speed of about 180 km/h. The results are similar to the ones presented as can be seen at figure 3. Another measurement series led to slightly better results for slower velocity but no general impact can be determined.

B. *Car to Car measurements*

Based on the results from the A44 measurement series we evaluated the impact of the SampleRate configuration during the C2C measurement campaign. We changed the blocking period from 10 sec to 1 sec during the suburban measurements and figure 4 shows the goodput results.

When considering also the FER-results as shown in figure 5 the more aggressive behaviour of the rate selection is clearly visible. The system accepts the higher FER and thus the higher number of required retransmissions because the achieved goodput is also clearly higher.

Figure 6 shows the TCP goodput over time measured in the suburban environment with the 802.11g setup. The graph also shows the C2C distance and the amount of traffic generated

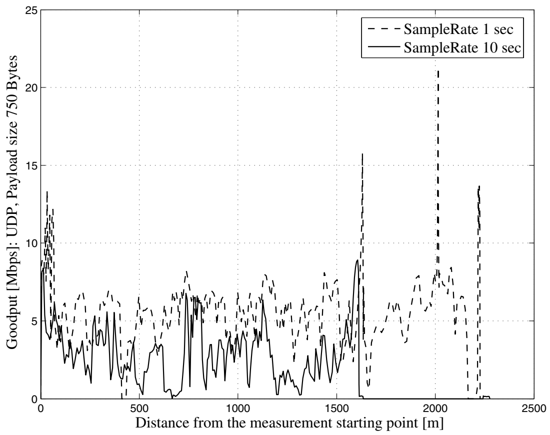

Fig. 4. Goodput measurements in suburban environment comparing different settings for the SampleRate algorithm.

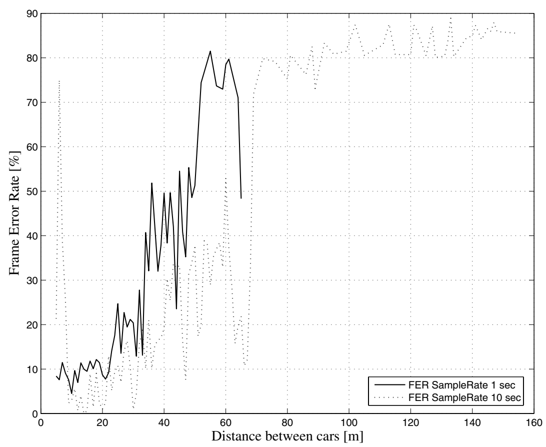

Fig. 5. Frame Error Rate measurements in suburban environment comparing different settings for the SampleRate algorithm.

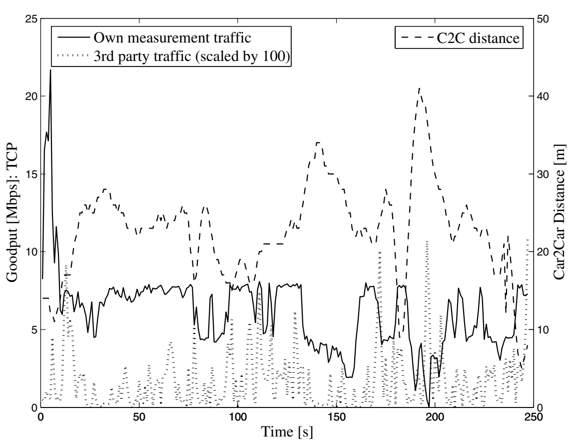

Fig. 6. TCP goodput measurements using 802.11g equipments in the suburban environment taking into account 3rd party traffic and the car2car distance.

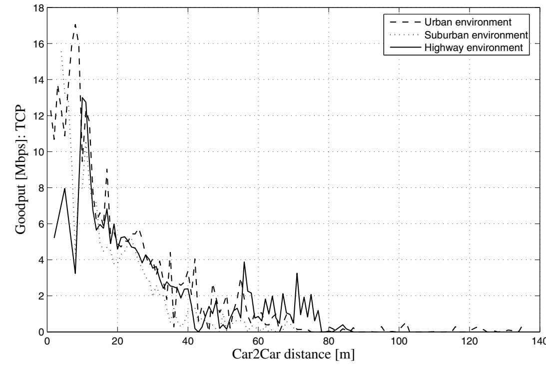

Fig. 7. Comparison of C2C TCP goodput measured in different environments.

by other parties using the same channel. As the amount of traffic received from other parties was rather small it is scaled up by a factor of 100. The figure shows the impact of the C2C distance and a direct relation between C2C distance and TCP goodput can be determined. The impact of the third party traffic is nearly negligible because of its small amount. Similar results were measured in the urban environment.

The same measurements were also performed with 802.11a in the 5 GHz band. In contrast to theory the SNR measured is more or less the same as for 802.11g in the 2.4 GHz band so that the higher attenuation expected for the 5 GHz band is not noticeable for the measured C2C distances of up to about 70 m. Also the FER results look similar. The achievable goodput was better for the 802.11a case due to the fact that no other traffic was received in the same channel and probably no inter-channel interference was present. In the 802.11g case additional APs might work on nearby channels but frames from these channels are not received by the sniffers. However, such frames would severely worsen the receiver performance.

The second larger difference is system stability. The 5 GHz system is much worse so that several tries were needed before reliable communication between both cars was possible. In several cases the client could not authenticate with the server so that the link layer connection establishment failed.

Summarizing the C2C measurements figure 7 shows the TCP goodput measured in the three environments. No clear difference can be determined but in the highway environment a longer communication range is achievable. In the highway scenario usually less obstacles block the LOS leading to a larger communication range.

IV. Discussion

Overall, WLAN technology proved to work also at vehicular speed. No clear performance difference could be determined and the major impact factors are availability of LOS and the C2C distance. In contrast to former measurements [3] we did not find big differences between scenarios. Only the highway environment leads to larger communication ranges.

The conditions during the first measurements on the A44 highway did not vary drastically and neither additional car traffic nor further data traffic was present. Thus, the scenario was stable and our results were reproducible between different runs. During our C2C measurements the conditions continuously changed leading to worse and less reproducible results. Also the communication range, which we reached in the A44 area was much larger than in the other environments. The maximum of about 1000 m fits to other reported measurements although it is difficult to compare these because the transmission power is rarely reported in detail and, e.g., the regulations allow higher transmission powers in North America.

Our A44 measurements show which performance will be achievable if the system adapts to the changing environment. As already found in [3] the usual statement that larger packets lead to a more efficient system does not hold for vehicular environments. The receiver performance decreases with larger frame sizes so that optimal settings could be chosen based on the current situation.

We did not switch off the multi-rate features of the 802.11 devices as done in former studies [5] and send standard unicast traffic. The results clearly show that the rate adaptation mechanism that is finally deployed has to be context-aware in several means in order to provide good performance. We showed that faster adaptations are more suitable for vehicular environments and taking into account the existing results for static scenarios [9] the parameters should be adapted based on the vehicle speed accessible via GPS-receivers. Additionally, the networking environment could be taken into account enabling the system to react, e.g., to higher interference levels or the fact that more important information related to security-aspects is to be transferred.

Such classification of traffic is already foreseen for 802.11p where different transmit power levels are used for different channels. Additionally, no authentication procedure is used. The latter is not only important in the security context. Our experiments showed that authentication takes a considerable amount of time and will prevent communication if failed.

Although the initial goal of the 802.11p development was to reuse existing technology several aspects have to be revisited. The MAC layer was already adapted and several aspects of the physical layer were changed, e.g. transmission power and channel bandwidth. When analysing the hardware only smaller parts should be reused. In usual local scenarios sophisticated antennas or longer cables are not needed. However, our practical experience shows that especially these building blocks have considerable impact on the overall system performance. Similar findings were published earlier [10] were the impact of better antennas was shown for 802.11g in vehicular environments. Additionally, we showed that especially the 5 GHz band is problematic when cheap and off-the-shelf hardware is used. Firstly, most available hardware works better in the lower 5 GHz band but 802.11p will work at 5.9 GHz and secondly, the robustness of the technology has to be improved. Although 802.11a performed better in one test the 802.11g setup was much more stable and thus would be still preferable.

V. Conclusion

In this paper we presented the results of an extensive measurement campaign evaluating the performance of IEEE 802.11a, b, and g in car communication scenarios. We showed that the velocity has a negligible impact up to the maximum tested speed of 180 km/h. Also the environment does only have a small impact on the communication range. The deployed hardware is of considerable importance leading to the requirement of more careful system design when comparing car communication to standard WLAN networking.

Additionally, we showed that the rate adaptation can be improved when reacting faster to changing channel conditions and that car-to-car distance and the availability of line-of-sight communication are the most important performance factors. If the proposed improvements for the rate adaptation are implemented and antennas are carefully chosen also WLANs based on legacy devices working at 2.4 GHz could be used for not security-related applications. Although our measurements showed that the danger of interference is nearly negligible it has to be taken into account for security-related applications. Therefore, those should be implemented using 802.11p-compliant devices as they work in separate licensed bands.

Acknowledgment

The authors would like to thank RWTH Aachen University for providing financial support partially through the UMIC excellence cluster and the RWE AG and the Bilfinger & Berger GmbH for offering us access to the closed highway and the support during the preparation phase. Additionally, the authors would like to thank Gero Schmidt-Kärst and Ossi Raivio for their support during the measurements.

References

[1] H. Hartenstein, B. Bochow, A. Ebner, M. Lott, M. Radimirsch, and D.Vollmer, "Position-Aware Ad Hoc Wireless Networks for Inter-Vehicle Communications: the Fleetnet Project, Poster," in *Proc. of ACM MobiHoc*, Long Beach, USA, October 2001, pp. 259–262.

[2] A. Ebner, H. Rohling, L. Wischhof, R. Halfmann, and M. Lott, "Performance of UTRA TDD Ad-Hoc and IEEE 802.11b in Vehicular Environments," in *Proc. of IEEE VTC '03*, vol. 2, Jeju, South Korea, Spring 2003, pp. 960–964.

[3] J. Singh, N. Bambos, B. Srinivasan, and D. Clawin, "Wireless LAN Performance under Varied Stress Conditions in Vehicular Traffic Scenarios," in *Proc. of IEEE VTC '02*, vol. 2, Vancouver, Canada, Fall 2002, pp. 743–747.

[4] J. P. Singh, N. Bambos, B. Srinivasan, D. Clawin, and Y. Yan, "Proposal and Demonstration of Link Connectivity Assessment based Enhancements to Routing in Mobile Ad-hoc Networks," in *Proc. of IEEE VTC '03*, vol. 5, Orlando, USA, Fall 2003, pp. 2834–2838.

[5] P. Bergamo, D. Maniezzo, K. Yao, M. Cesana, G. Pau, M. Gerla, and D. Whiteman, "IEEE 802.11 Wireless Network under Aggressive Mobility Scenarios," in *Proc. of ITC 2003*, Las Vegas, USA, 2003.

[6] J. Ott and D. Kutscher, "Drive-thru Internet: IEEE 802.11b for "Automobile" Users," in *Proc. of IEEE INFOCOM*, vol. 1, March 2004, pp. 362–373.

[7] ——, "The "Drive-Thru" Architecture: WLAN-based Internet Access on the Road," in *In Proc. of IEEE VTC '04*, Spring 2004.

[8] *Status of Project IEEE 802.11p*, http://grouper.ieee.org/groups/802/11/Reports/tgp_update.htm [Cited on: 2nd of January 2006], 2006.

[9] J. C. Bicket, "Bit-rate selection in wireless networks," Master's thesis, Massachusetts Institute of Technology, February 2005.

[10] J. Ott and D. Kutscher, "Drive-thru Internet: Bringing Local-Area Wireless Connectivity to Wide-Area Mobile Users," in *Proc. of Upperside WiMax Summit 2004*, May 2004.