

Performance of IEEE 802.11a in Vehicular Contexts

David N. Cottingham, Ian J. Wassell, Robert K. Harle
Computer Laboratory, University of Cambridge
Cambridge CB3 0FD, United Kingdom
E-mail: {david.cottingham, ian.wassell, robert.harle}@cl.cam.ac.uk

Abstract—A key component of intelligent transportation is the provision of adequate network infrastructure to support vehicle-to-vehicle and vehicle-to-roadside communication. In this paper we report on performance evaluations carried out using the IEEE 802.11a protocol at 5.2 GHz between a moving vehicle and a fixed base station. We concentrate our evaluation on realistic urban speeds and environments, observing that performance at very low speeds is degraded due to the presence of null zones. We vary the modulation scheme and analyse the spread of resulting throughputs. Our results have implications for multimedia and other real-time applications that will utilise vehicle-to-roadside connectivity.

Keywords: IEEE 802.11a, WAVE, Wireless Access for the Vehicular Environment, wireless LAN, mobile/wireless systems for transportation, experimental performance ¹

I. INTRODUCTION

The application space for vehicle-to-vehicle and vehicle-to-roadside communication is vast. Uses range from toll collection, through intersection collision avoidance, to obtaining up-to-date traffic information [1]. As a consequence, various different communications technologies have been proposed for enabling this new generation of sentient vehicles. These include the various versions of the IEEE 802.11 standard [7], as well as the 802.20 Mobile Broadband for Wireless Access (MBWA) standard, and the 802.16e Mobile WiMAX draft.

In the United States, the Dedicated Short Range Communication (DSRC) standard for vehicular communication will utilise 802.11p Wireless Access for Vehicular Environments (WAVE) technology, which is very similar to the 802.11a standard. The former will utilise spectrum at 5.9 GHz, whilst the latter exists in the 5.2 GHz range. Hence, evaluation of 802.11a hardware, which is widely available, will yield a good indication of the performance of forthcoming DSRC devices.

The IEEE 802.11a protocol [8] differs from the 802.11b “WiFi” protocol in both the frequency range used and the multiplexing scheme. Whilst the former uses Orthogonal Frequency Division Multiplexing, the latter uses Direct Sequence Spread Spectrum techniques in the 2.4 GHz band. Maximum data rates are quoted as 54 Mbps and 11 Mbps respectively, 802.11a achieving substantially higher throughput in part due to offering 64-QAM modulation.

¹©2007 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

In practice, data rates achieved at higher protocol layers are lower than those given above, as the overhead of encapsulation is significant. For example, a TCP or UDP packet must be placed in an IP datagram, and subsequently in a MAC layer frame for output by the radio transmitter. Hence, maximum UDP data rates for 802.11a are in the region of 30 Mbps. A comprehensive overview of the protocol’s MAC layer can be found in [10].

The remainder of this paper analyses the performance of the 802.11a protocol in urban environments. First, we outline related work in Section II, then describe the experimental set-up in Section III. The experimentally determined performance of the technology at low speeds is then described in Section IV. We conclude with a summary and suggestions for further work in Section V.

II. RELATED WORK

The performance of the IEEE 802.11b standard has been evaluated extensively in different contexts and with varying goals in mind. In contrast, evaluation of the IEEE 802.11a standard has been less comprehensively addressed.

There have been several investigations into the performance of IEEE 802.11b at different speeds, such as [5], where tests were performed in an interference-free environment at speeds ranging from 5 to 75 mph. Another study by Bergamo et al. [2] examined performance under “aggressive mobility scenarios”, which included free space propagation at high speeds. However, with semi-free space propagation (i.e., with some objects occasionally obstructing the line of sight between antennas), the performance dropped significantly, due to brief occlusion periods where packet loss occurred, coupled with an increase in jitter. The effect of range and packet size are further examined in [11].

Work has also been carried out analysing the theoretical [14][12] and practical [4] performance for intervehicular networks. This again has centred around IEEE 802.11b, mostly in motorway environments. The Drive-Thru Internet project [9] has analysed these environments in detail, and proposed a disconnection-tolerant architecture that is optimised for such short-lived connections.

The performance of the DSRC standard has also been simulated [15][13]. The results suggest that delays are highly variable, contingent on whether there is line of sight communication, and the transmission range (assuming that a lower range implies multipath components cannot be resolved, resulting in frequency flat-fading). The Bit Error Rate (BER)

is found, by simulation, to vary significantly with the relative speed of the transmitter and receiver. High vehicle speeds mean that the channel characteristics may change in the middle of a packet transmission, which increases the BER as training sequences only exist at the start of each packet.

Previous work has mostly taken place under desert or test track conditions, at high speeds. Our work differs from these in that it analyses the “a” version of 802.11 in a realistic urban environment at low (<50 km/h) speeds. These velocities are particularly important as they are the norm in a congested urban environment, where there is likely to be the greatest density of access points. It is also the case that in cities users are more likely to require extra information to be downloaded to their vehicles, such as more detailed map or traffic information.

III. EXPERIMENTAL SET-UP

In order to ascertain the performance of 802.11a in various different scenarios, we used laptops as the transmitter and receiver for all tests. One machine was connected using 100 Mbps wired Ethernet to a Cisco 1200 wireless access point (AP) which had an 802.11a module installed, transmitting at a frequency of 5.32 GHz. The other laptop used an NEC Warpstar 802.11a/b/g PCMCIA interface card, connected using 2 metres of low RF-loss cable to an omnidirectional aerial, that provided a rated gain of 7 dBi for frequencies in the 5 GHz band.

All tests were performed using the iperf² network measurement tool, with the server being the laptop with the 802.11a PCMCIA card, and the client connected to the access point. In each test UDP datagrams of 1470 bytes in size were transmitted from the client to the server, and the average throughput (of uncorrupted packets received by the server) for each 0.5 second interval calculated. The two attempted throughputs (i.e., the rate at which the client sent packets to the access point) were 10 and 30 Mbps.

The reasons for using UDP, and not TCP were:

- TCP infers congestion from packet loss and scales back its send window accordingly. For these experiments, the aim was to ascertain how throughput is affected by the wireless propagation characteristics of the channel, rather than due to protocol-induced throughput reductions.
- TCP uses an ARQ mechanism to retransmit lost packets, causing the sending of normal data to cease temporarily. Again, the desire was to analyse the performance of the underlying data-link layer rather than ARQ-based protocols being carried over it.
- Although TCP is used for many network applications, the majority of transmissions to vehicles are likely to be broadcast, best-effort streams (e.g., traffic updates).

For tests of performance at varying speeds, the access point was placed on a metal stepladder, at a height of approximately 1.8 metres above the road, with the laptop connected to it by 3 metres of Ethernet cable. The access point’s antenna was

placed with its largest face orthogonal to the direction of the road, to achieve the best possible propagation characteristics. The laptop acting as the server was placed inside a small van, and the antenna externally attached to the roof rack. The test environment was a straight, two-lane road, with low levels of traffic. There are few buildings close to the road, and these are only on one side, with the other being bounded by fields. The distance between the start and end points of each drive past the access point was approximately 460 metres, with the access point being placed in a lay-by 270 metres from one end. The end points were chosen such that the vehicle was out of wireless coverage when it reached the extremes.

IV. PERFORMANCE AT LOW SPEEDS

We performed a large number of experiments assessing the performance of a connection from a moving vehicle to an access point beside the road, concentrating on two speeds, 7 km/h and 45 km/h. The former is characteristic of a heavily congested inner city road and the latter of a similar road in a free flow situation. The speeds were kept as constant as possible using the vehicle’s speedometer, and were verified using an onboard GPS unit. The results from two sample runs at an attempted UDP throughput of 10 Mbps are shown³ in Figures 1 and 2, whilst those for 30 Mbps are shown in Figures 3 and 4.

At 10 Mbps there is no difference in the throughput at the two speeds, with both traces showing that the throughput is either zero or approximately 100% of that attempted⁴. In contrast, at 30 Mbps the connection is less “binary”, and exhibits a much greater variability in throughput at 7 km/h.

A. Throughput Variability

For many multimedia applications, such as Voice over IP, not only is total data transmitted or received important, but also the variability of the throughput over the lifetime of the connection. Figures 5 to 8 show histograms of the throughputs achieved at the two test speeds, with attempted throughputs of 10 and 30 Mbps.

The variation in throughput when using 30 Mbps can be attributed to the modulation scheme used being 64-QAM, rather than QPSK, for 10 Mbps. The former is more sensitive to interference, and hence more likely to incur packet loss for a given amount of noise.

The histograms of throughput shown in Figures 7 and 8 are evidence that travelling at lower speeds, whilst providing a greater connected time period, also gives a less certain throughput. This is a significant observation, given that other related work has found that at higher speeds there is no effect on the throughput achieved. Our assertions concerning the causes of this effect are outlined in the following Section.

³The traces occasionally show that the throughput exceeds 10 Mbps. This always occurs directly after a loss, and is due to some packets that were from one interval arriving in the following one. Were the intervals to be 1, rather than 0.5, seconds in length, this would not occur.

⁴Figure 1 shows only a few isolated points that are at low throughputs, but that each appears as a pair of near-vertical lines on the graph.

²<http://dast.nlanr.net/Projects/Iperf/>




Fig. 1. Throughput for two oppositely-directed passes of an 802.11a AP at 7 km/h at 10 Mbps

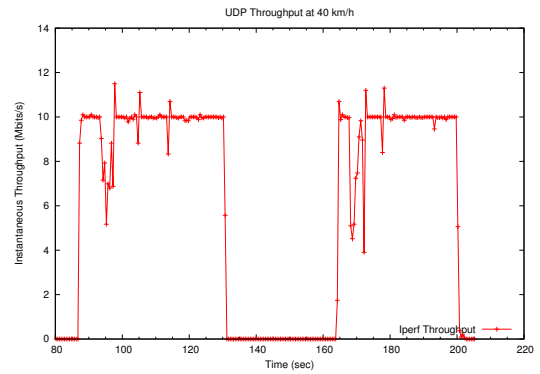


Fig. 2. Throughput for two oppositely-directed passes of an 802.11a AP at 45 km/h at 10 Mbps

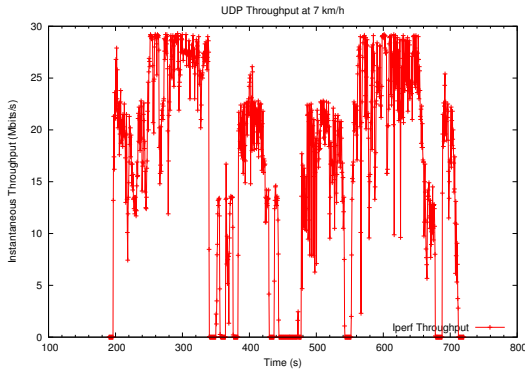


Fig. 3. Throughput for two oppositely-directed passes of an 802.11a AP at 7 km/h at 30 Mbps




Fig. 4. Throughput for two oppositely-directed passes of an 802.11a AP at 45 km/h at 30 Mbps

B. Null Zones

In any channel that can be modelled as having Rayleigh fading, the received signal strength (RSS) will vary over distance, with regions of poor reception being present where transmitted waves interfere destructively. In environments where reflections (multipath effects) are present, such as those due to urban street furniture, such interference effects are pronounced. In contrast, in environments such as the desert or a road of very low utilisation, their effects are rarely observed.

Whilst intuitively it is clear that a lower received signal strength will result in a lower throughput, (as expected from the Shannon-Hartley theorem), this is complicated by the use of modulation schemes of varying robustness. If the transmission includes sufficient error correction, a constant data rate can be maintained for large variations in RSS. We therefore carried out experiments in an office corridor to ascertain how IEEE 802.11a performed when the receiver antenna was placed in a known null zone.

Using a spectrum analyser connected to an omnidirectional antenna, we were able to identify a location where the RSS was significantly lower than its surroundings. These zones have dimensions of the order of tens of millimetres in size, and therefore require careful antenna placement. We then moved the connection from the signal analyser to an 802.11a receiver card, maintaining the antenna in the same position. Figure 9 shows the variation in throughput as the antenna entered the null zone and was then removed. The first section

of the graph represents the time when the antenna was moved slightly to find the null zone once more. The second portion is performance with the antenna stationary inside the null zone ($\leq 33\%$ degradation), whilst the remainder of the trace is the performance with the antenna outside the null zone. We note that the throughput does not return to its normal rate immediately on the antenna exiting the null zone; this is due to the receiver taking a small amount of time to adjust to the new channel characteristics. Subsequent temporary variations in the throughput are attributed to people walking along the corridor where the experiment was carried out (confirmed by other experiments).

The above experiment was performed with an attempted send rate of 30 Mbps. A similar test was performed using a 10 Mbps send rate where the effect of the null zone on throughput was not evident, due to the more robust modulation scheme (QPSK).

When the receiver antenna is in a null zone a greater proportion of packets are lost than when it is not. This will therefore affect the performance of a wireless connection at low speeds for a vehicle passing an access point. The propagation environment, and hence the number and position of null zones, will only be affected by how the objects in the area are positioned and the frequency of transmission, (provided that there is no time-varying RF interference). Hence, for any vehicle speed, we can say that the range of the AP is constant, and that the number and total length of the null zones within




Fig. 5. Spread of throughputs when passing an 802.11a AP at 7 km/h at 10 Mbps




Fig. 6. Spread of throughputs when passing an 802.11a AP at 45 km/h at 10 Mbps




Fig. 7. Spread of throughputs when passing an 802.11a AP at 7 km/h at 30 Mbps




Fig. 8. Spread of throughputs when passing an 802.11a AP at 45 km/h at 30 Mbps

that coverage area is also fixed. Therefore as speed increases the total time spent in null zones decreases. This therefore means that the absolute packet loss decreases at the expense of less time being spent in the fixed-length coverage region. Reduced packet loss implies a greater average throughput, and hence a greater transfer size. Hence, it is important to ascertain how the spread of throughputs changes with different vehicle speeds.

Another point of interest is that as a vehicle moves faster, there is a greater likelihood that the proportion of a frame affected by a null zone is smaller. There will come a point where the error correction in the frame will be able to correct for the loss, and hence higher layers do not see a packet drop. However, this phenomenon is unlikely to have a significant effect.

C. Connection Times

Tables I and II summarise the data from more extensive experiments, only subsets of which are presented in graph format in the figures in this paper. All data given in the tables are derived using the following criteria:

- Zone 1 of a run (the production zone) is defined as the longest time period within which the connection's throughput exceeds 10% of the attempted send rate, with no interruptions greater than 5 seconds, excepting a single interruption of up to 10 seconds. The reasoning behind this choice of criterion is that the majority of runs showed

such an interruption in their production zone, but at apparently random times. To not allow this interruption would therefore have made the results meaningless. The cause of such interruptions may be due to occasional shadowing of the access point by other vehicles.

- Zone 2 is defined as the time period, regardless of interruption, that the connection's achieved throughput was greater than 10% of the attempted send rate.

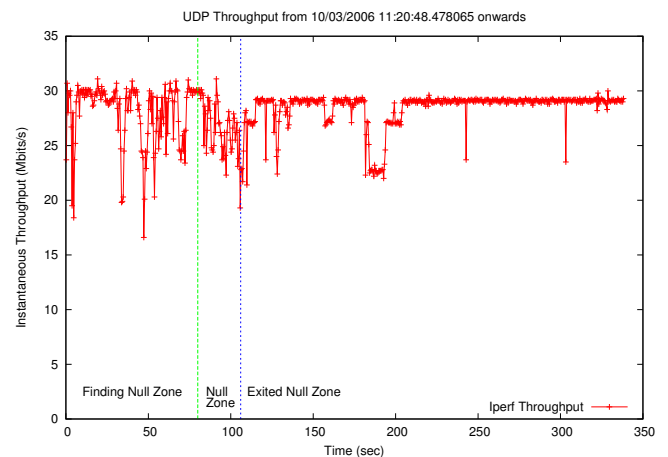


Fig. 9. Throughput variation with receiving antenna in a null zone, at an attempted send rate of 30 Mbps

- Zone 3, sporadic connectivity, is similar to zone 2, save that any achieved throughput above zero may be counted.

The data given are all the mean connection times, with their corresponding standard deviations, in seconds.

Speed (km/h)	Zone 1 (s)	Zone 2 (s)	Zone 3 (s)	Runs
7 km/h	190.2 ±9.9	227.5 ±31.3	228.7 ±31.9	3
45 km/h	36.3 ±9.9	36.3 ±9.9	37.6 ±11.7	7

TABLE I

PERIODS OF CONNECTIVITY AT A THROUGHPUT OF 10 MBPS.

Speed (km/h)	Zone 1 (s)	Zone 2 (s)	Zone 3 (s)	Runs
7 km/h	180.6 ±52.2	243.75 ±8.6	246.7 ±8.8	6
45 km/h	46.0 ±8.2	46.0 ±8.2	46.7 ±8.5	5

TABLE II

PERIODS OF CONNECTIVITY AT A THROUGHPUT OF 30 MBPS.

The data show that, as expected, the length of each connection zone for a given speed is approximately constant, (the means are well within two standard deviations of each other) irrespective of the attempted send rate. Because of the way in which zone 1 is defined, the standard deviation can be relatively high, whereas the period for which throughput is approximately 10% of the attempted send rate is a more reliable indicator of performance.

Another observation is that the connection times achieved at 45 km/h were less than would be expected were they to be directly proportional to the speed of the vehicle (i.e., the time achieved at 7 km/h decreased in proportion to the change in speed). The explanation for this is that whilst the connected time period is proportional to the speed, this is only true once the AP has been detected and the station registered. This is dependent on the length of time taken for the MAC association frames and acknowledgments to be sent to/from the AP. In the course of our experiments, we found the re-association time (in this case taken to be the time from which the connection is lost to when it is usable for sending UDP packets once more) to take 8 seconds. This then means that at higher speeds the vehicle covers a greater distance within the coverage region whilst association is taking place, leaving less time for useful communication.

D. Antenna Positioning

There appeared to be no significant effect due to the position of the antenna on the vehicle, confirming the results reported in [6]. Two locations were used, one at the front of the vehicle nearest the passenger side, attached to the roof rack, and the other diametrically opposite, on the back roof rack. Table III shows the respective results.

V. CONCLUSION

This paper has presented the results of an evaluation of the performance of the IEEE 802.11a wireless networking standard. In particular, we have shown that at low speeds, a vehicle

Speed (km/h)	Zone 1 (s)	Zone 2 (s)	Zone 3 (s)	Runs
Front	194.7 ±48.7	240.2 ±7	242.5 ±4	3
Back	180.5 ±65.6	247.3 ±9.9	250.8 ±11.1	3

TABLE III

PERIODS OF CONNECTIVITY AT 7 KM/H, THROUGHPUT OF 30 MBPS.

passing an access point will experience a greater variation in throughput than that at higher speeds due to the effect of null zones in the propagation environment. This is important when considering performance in urban environments, where there is likely to be a greater density of wireless deployment than motorway scenarios. We have also outlined how connection times are not directly proportional to the speed of travel, but instead are also dependant on the association time.

Our current work is focusing on improving the performance of 802.11a by utilising more than one receiving antenna, to negate the effect of null zones. We also intend to perform further experiments examining how TCP is affected by the variation in throughput, given its assumption that packet loss implies congestion, as part of the sentient vehicles project [3] at the University of Cambridge Computer Laboratory.

ACKNOWLEDGMENTS

We gratefully acknowledge the help of Jonathan Davies, and the support of Andy Hopper with various aspects of this work.

REFERENCES

- [1] AMI-C. AMI-C use cases. Technical Report 1001, AMI-C, January 2003.
- [2] 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. International Telemetry Conference*, October 2003.
- [3] Jonathan J. Davies, David N. Cottingham, and Brian D. Jones. A sensor platform for sentient transportation research. In *Proc. 1st European Conference on Smart Sensing and Context*, volume LNCS 4272, pages 226–229, October 2006.
- [4] 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. IEEE VTC*, April 2003.
- [5] R. Gass, J. Scott, and C. Diot. Measurements of in-motion 802.11 networking. In *Proc. IEEE WMCSA*, 2006.
- [6] Martin Grade, Klaus Meier, Bernd Rech, and Andreas Lübke. Physical IEEE 802.11 – measurements in automotive environment. In *Proc. ITS*, June 2005.
- [7] IEEE Working Group 11. *IEEE Std 802.11-1999 Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. Number ISO/IEC 8802-11:1999(E). IEEE, 1999.
- [8] IEEE Working Group 11. *IEEE Std 802.11-1999 Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High Speed Physical Layer in the 5 GHz Band*. IEEE, 1999.
- [9] Jörg Ott and Dirk Kutscher. Drive-thru Internet: IEEE 802.11b for “automobile” users. In *Proceedings of IEEE INFOCOM*, March 2004.
- [10] Daji Qiao and Sunghyun Choi. Goodput enhancement of IEEE 802.11a wireless LAN via link adaptation. In *Proc. IEEE ICC*, volume 7, pages 1995–2000, June 2001.
- [11] Jatinder Pal Singh, Nicholas Bambos, Bhaskar Srinivasan, and Detlef Clawin. Wireless LAN performance under varied stress conditions in vehicular traffic scenarios. In *Proc. IEEE VTC*, volume 2, Fall 2002.
- [12] C. Steger, P. Radosavljevic, and P. Frantz. Performance of IEEE 802.11b Wireless LAN in an Emulated Mobile Channel. In *Proc. IEEE VTC*, April 2003.
- [13] Marc Torrent-Moreno, Daniel Jiang, and Hannes Hartenstein. Broadcast reception rates and effects of priority access in 802.11-based vehicular ad-hoc networks. In *Proc. ACM VANET*, pages 10–18, 2004.

- [14] H. Wu, M. Palekar, R. Fujimoto, R. Guensler, M. Hunter, J. Lee, and J. Ko. An empirical study of short range communications for vehicles. In *Proc. ACM VANET*, pages 83–84, 2005.
- [15] J. Yin, T. ElBatt, G. Yeung, B. Ryu, S. Habermas, H. Krishnan, and T. Talty. Performance evaluation of safety applications over DSRC vehicular ad hoc networks. In *Proc. ACM VANET*, pages 1–9, 2004.