An Experimental Analysis of IEEE 802.11a/b/g Cards in Ad-hoc Mode[†]

K. Kosek, S. Szott, M. Natkaniec, K. Łoziak
AGH - University of Science and Technology
Department of Telecommunications
al. Mickiewicza 30, 30-059 Kraków, Poland
{kosek, szott, natkaniec, loziak}@kt.agh.edu.pl

***Abstract*- This article presents experimental results of cooperation between wireless IEEE 802.11a/b/g cards coming from different vendors, working in ad-hoc mode. The throughput efficiency of each of the cards was analysed in two different configuration scenarios consisting of two or three cards working together. Emphasis was put on the multirate adaptation abilities of the cards, which utilize such modulation techniques as DBPSK, DQPSK, CCK, and OFDM. The achieved results, presented in the form of figures, show that cards manufactured by independent vendors perform differently. Therefore, choosing the optimum configuration, according to the user's requirements, is possible.**

I. Introduction

Wireless networks are one of the youngest and most dynamically evolving technologies in telecommunications. Some Wireless Local Area Network (WLAN) solutions enable infrastructure-less communication in the form of spontaneous, ad-hoc net-works.

IEEE standards 802.11 [5], 802.11a [6], 802.11b [7], and 802.11g [9] enable operating at multiple speeds in the uplink and downlink direction, heavily dependant on the physical properties of the radio channel. The distance between the sending and receiving WLAN card influences these physical properties and it is correlated with the rate of operation.

IEEE 802.11 commercial cards produced by different vendors were analyzed in [14], where MAC implementation and hardware delays were tested. It is clear from the conducted research that a remarkable amount of unfairness among different cards is a result of specific hardware/firmware implementations. Therefore, it can be assumed that a similar considerable difference, in terms of perceived throughput, will appear in multirate IEEE 802.11a/b/g environments.

The aim of this work is to show the behaviour and performance of multirate IEEE 802.11a/b/g cards in ad-hoc mode made by the following vendors: Cisco, Linksys, Lucent, D-Link Proxim, and 3COM. In the first experiment, the TCP throughput of the cards was tested with the TTCP tool. In the second scenario FTP traffic carried by the ad-hoc network was analysed.

The rest of the paper is organised as follows. Section 2 describes the PHY layer of the IEEE 802.11 family. The ad-hoc mode of the MAC layer is presented in Section 3. Section 4 gives a brief overview of rate selection methods. After this basic theoretical introduction, a description of the testbed is presented in Section 5 and the results of the measurements are given in Section 6. The paper concludes with Section 7.

II. The PHY Layer in IEEE 802.11a/b/g

The IEEE 802.11 family of standards has become the most popular wireless transmission method. A brief description of the characteristics of the legacy 802.11 and the 802.11a/b/g standards is given below.

The 802.11 standard utilizes several different modulation techniques and transmission strategies. First of all, the ISM (Industrial, Scientific and Medical Applications) band is used, with the frequency range being 2400 – 2483.5 MHz. Furthermore, two spread spectrum techniques are utilized: DSSS (Direct Sequence Spread Spectrum) and FHSS (Frequency Hopping Spread Spectrum). The DSSS technology divides the total bandwidth into 5 channels (each 26 MHz wide), whereas, the FHSS technology divides the bandwidth into 79 channels (each 1 MHz wide). DSSS relies on BPSK (Binary Phase Shift Keying) modulation which allows for transmission rates of 1 or 2 Mbit/s. FHSS relies on a two- or four-level GFSK (Gaussian Frequency-Shift Keying) modulation, with transmission rates of 1 or 2 Mbit/s accordingly. Additionally, it takes advantage of three hopping sequences – each consisting of 22 hops. The range of the system is 20-50 m indoors, and up to hundreds of meters outdoors.

The 802.11a standard specifies the physical layer and allows for transmission rates of 54 Mbit/s in the 5 GHz range. One of the innovative features of this standard is the use of OFDM (Orthogonal Frequency Division Multiplexing), which enables high transfer rates. Each channel is divided into 52 sub-carriers (each 300 kHz wide). In total, 802.11a uses 300 MHz in the 5 GHz UNII (Unlicensed National Information Infrastructure) band. Equipment conformant to the 802.11a standard must work at the following transmission rates: 6, 12, and 25 Mbit/s. Optional rates include: 9, 18, 36, 48 and 54 Mbit/s. These differences are caused by using different modulation techniques and different FEC (Forward Error Correction) levels. In order to achieve the highest possible rate of 54 Mbit/s, each of the sub-carriers must use 64-level QAM (Quadrature Amplitude Modulation).

The 802.11b standard specifies a physical layer, which enables a maximum trans-mission rate of 11 Mbit/s in the

[†] The work described in this paper is based on results of IST FP6 Integrated Project DAIDALOS. DAIDALOS receives research funding from the European Community's Sixth Framework Programme. Apart from this, the European Commission has no responsibility for the content of this paper. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

2.4 GHz band. The other allowed rates are 1, 2, and 5.5 Mbit/s. Their use depends on channel conditions between the communicating stations (distance, bit error rate, etc.). The Barker code is used as the spreading sequence both in the basic IEEE 802.11 standard, as well as in the 1 and 2 Mbit/s rates of IEEE 802.11b. For higher speeds (5.5 and 11 Mbit/s), CCK (Complementary Code Keying) is used. The chip speed of 802.11 is always 11 Mc/s, which results in a speed of 1 Mbit/s for the DPSK modulation and 2 Mbit/s for QPSK. The 802.11b standard allows for dynamic changes in the transmission rate. When channel conditions are optimal, the maximum rate of 11 Mbit/s is used. However, if interference or distance in-creases, the rate automatically decreases to 5.5, 2 or 1 Mbit/s. If conditions become better, the rate will increase.

The 802.11g standard's specification of the physical layer allows 802.11a trans-mission rates to be achieved in the 2.4 GHz band and ensures compatibility with 802.11b devices. To ensure this interoperability, 802.11g is more complex – it needs to use older modulation techniques to communicate with legacy equipment but on the other hand uses OFDM to achieve higher transmission rates. The necessary elements include: (a) CCK which assures backward compatibility with 802.11b devices which can detect transmission and avoid collisions. The payload may be of varying sizes (from 63 B to 1500 B), depending on the transmission rate and the amount of transmitted data, and (b) an OFDM implementation for the 2.4 GHz band with a transmission rate of 54 Mbit/s (a solution provided by Intersil). If 802.11b devices are found, a CCK-modulated CTS (Clear To Send) packet is sent before the OFDM transmission. This "warning" informs slower stations to ignore the subsequent transmission.

III. The IEEE 802.11 MAC Layer in Ad-hoc Mode

The basic access method in ad-hoc mode is CSMA/CA (Carrier Sense Multiple Ac-cess / Collision Avoidance) used by DCF [5]. This access method is enhanced with Virtual CS (Virtual Carrier Sense) and NAV (Net Allocation Vector).

The MAC protocol distinguishes three important time periods: SIFS, PIFS, DIFS (Short-, PCF-, DCF- Inter Frame Space) which have lengths corresponding to the following rule: DIFS>PIFS>SIFS. When stations sense that the medium is free, they begin to measure these periods in order to estimate when they can begin their own transmission. The protocol also identifies three priorities of transmission, according to these periods.

Each station, ready to transmit, senses the medium to check whether any other station is currently transmitting. If the medium is free, it begins to transmit. Otherwise, it waits for the current transmission to finish and then waits until the medium has been free for one DIFS period. Afterwards, it calculates a backoff value – i.e., the time slot in which it will begin its transmission [13]. The backoff is randomly chosen (within certain limits) and is used to decrease the probability of two or more stations transmit-ting simultaneously (which would lead to a collision). The backoff value is decreased with time. If another stations begins transmission before the backoff value reaches zero, its countdown is suspended until the medium is once again free for a DIFS period. When the backoff value of a station reaches zero, the station may begin trans-mission in one of three modes of operation: (a) frame transmission, (b) frame trans-mission with RTS/CTS (Request/Clear To Send), or (c) frame transmission with broadcast/multicast. A more detailed description of DCF in 802.11 can be found in [5], [13].

IV. Rate Selection in IEEE 802.11

As mentioned before, the 802.11 standard provides medium access through CSMA/CA. However, the standard does not specify the method of automatic rate se-lection in the presence of multirate capable devices. As a consequence, there are several existing methods of choosing the appropriate rate and vendors of 802.11 equipment are free to choose or design their own.

The basic requirement of any WLAN card to utilize high transmission rates is for the received signal to be greater than a given threshold. These thresholds depend on the WLAN card receiver sensitivity, which indicates the amount of signal a card needs to receive in order to work correctly at a given speed level. These values are usually between -80dBm and -95dBm, however, some cheap cards can reach -70 dBm. The difference between a very good card and a bad one at a given rate can be as much as 32 times. This means that the very good card needs 32 times less signal strength as the bad card to work at the same rate. Receiver sensitivity is measured in dBm @BER (Bit Error Rate) 10E-5 or 8% FER (Frame Error Rate), according to the IEEE 802.11 standard (8% FER for DSSS devices stands for an MPDU (MAC Protocol Data Unit) having 1024 octets. 8/100 (errors/frames) / (1024 (octets) * 8 (bits) = approx. 1/100,000 errors/bits = 10E-5 BER). Table 1 shows sample values for seven different cards (although for many devices the specifications are confidential). It can be clearly seen that the quality of the card can determine its performance in a multi-rate environment (c.f., Section 6).

Furthermore, available transmission rates are not linear. Therefore, e.g., an 11 Mbit/s link with a delivery ratio of just above 50% is always better than a 5 Mbit/s link.

TABLE I
Receiver Thresholds (Sensitivity) of IEEE 802.11 Cards from Different Vendors

Rate (Mbit/s)	Receiver Thresholds (dBm)					
	Proxim 8480-WD	Linksys WPC-11	Cisco CB21 AG	Lucent Silver/ Gold	3Com 3CRPA G 175	Cisco Aironet 350
11	-91	-80	-90	-82	-86	-85
5.5	-94	N/A	-92	-87	-88	-89
2	-95	N/A	-93	-91	-91	-91
1	-96	N/A	-94	-94	-93	-94

Statistics-based algorithms are used to determine the best rate and, according to [1], one of the most commonly used (and possibly the first) is the ARF (Auto Rate Fallback) protocol. It was developed for Lucent's WavelanII devices

[10] and it utilizes the link-layer ACK frames (i.e., FER) to determine the quality of the channel. After successful reception of a given number of consecutive acks from a neighbouring node, the transmission rate is increased. Similarly, after a consecutive number of ACKs has been lost, the rate is decreased. This protocol requires no changes in the 802.11 standard because the sender imposes the transmission rate. However, ARF is not the optimal strategy because it is very slow to adapt to the channel conditions and, even if the channel conditions are stable, it will unnecessarily try to change the rate. Furthermore, it can mistake collisions for channel losses.

A slight improvement over ARF is a retry-based approach [10], [15]. Based on ARF, it differs in that down-scaling is performed after a number of unsuccessful retransmissions. This results in a very short response time to deteriorating links. However, the protocol behaviour is pessimistic. The rate will increase only after a FER threshold has been reached which takes longer than the down-scaling procedure. The Atheros AR5000 chipset (used in 802.11a cards) implements a throughput-based approach [3]: a constant small fraction (10%) of the data is sent using data rates one level higher and one level lower than the currently used. After a certain time, the decision which rate to use is based on their performance. This makes the algorithm slow to adapt to changes in the channel condition.

The MadWiFi driver used for current Atheros chipsets includes three different rate adaptation algorithms: Onoe [12], AMRR (Adaptive Multi Rate Retry) [11], and SampleRate [2]. The default algorithm is Onoe, which is similar to ARF but not as sensitive to individual packet loss. It looks for the highest bitrate that has a loss rate less than 50%. AMRR (based on Adaptive ARF) uses Binary Exponential backoff and works well for high latency systems while SampleRate uses aggressive probe packets to estimate the optimum transmission rate.

Numerous SNR-based alternatives to the ARF statistics-based approach have been proposed. One of them is RBAR (Receiver Based Auto Rate) [4]. In this protocol, the receiver side can determine the rate. It measures the SNR of the received RTS and informs the sender about the desired rate in the CTS packet. This solution requires changes to the 802.11 standard and requires RTS/CTS even when there are no hidden nodes, but on the other hand, it allows faster adaptability then ARF.

OAR (Opportunistic Auto Rate) is a similar protocol (the receiver decides upon the rate) but utilizes the coherence times of good channel conditions to send high-rate multi-packet bursts, similar to the TXOP (transmission opportunity) feature of 802.11e [8]. Overhead in OAR is low because there is no contention period or sending of RTS/CTS frames in these bursts. Changing the burst size can also increase fairness (in terms of bandwidth allocation time) within the network. However, the downside to these advantages is that OAR also requires modifications to the 802.11 standard.

Both RBAR and OAR suffer from using preselected SNR thresholds – they may not perform well under different channel conditions.

V. TESTBED DESCRIPTION

The efficiency measurements of 802.11a/b/g wireless cards of different vendors were performed under typical office conditions. The following cards were tested: Cisco AIR-CB21AG, Cisco AIR-PCM350, Linksys WPC-11, Lucent Silver PC24E, D-Link DWL-650+, Proxim 8480-WD, and 3COM 3CRPAG175. All of them worked in ad-hoc mode and their output power set to 30mW. Two different scenarios were considered.

In the first case, the testbed consisted of one analysing station and two stations equipped with the tested WLAN cards. These two stations were close to each other at first and, during the experiment, the distance between them was increased. The experiment was performed 2048 times with the use of the TTCP application which generated packets of size 8192 B for every measurement point. OmniPeek Personal was used to capture and analyze these packets.

In the second scenario the testbed consisted of three stations (c.f. Fig. 1): one FTP server (Station C) and two clients (Stations A and B). The clients, after connecting to the server, began downloading a file with an average size of 1 GB. Station A was stationary whereas B was mobile and increased its distance from the server station C. Over 50 thousand FTP frames sent from the server to the clients were captured.

For all experiments, emphasis was put on eliminating all near sources of interference in the 2.4 GHz and 5 GHz bands (e.g., access points, microwave ovens or Bluetooth devices). However, in EU-countries around 90% of electric devices operate in the 2.4GHz band, therefore, interferences from remote sources are unavoidable.

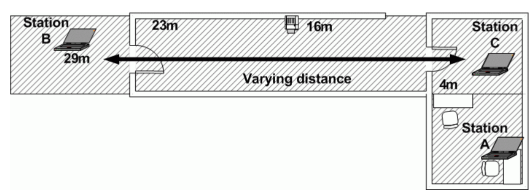

Fig. 1. Testbed

VI. SCENARIOS

In the first scenario the effective throughput was considered as a function of distance. The experiment was focused on the analysis of the point-to-point connection efficiency of two wireless cards operating in ad-hoc mode. Most of all, the effective speed of the transmission between stations A and B (c.f. Fig. 1) was measured. Furthermore, not only was the effective throughput measured but also all of the exchanged frames were filtered by the analysing station. These analyses were performed in order to explore and evaluate the 802.11 multirate transmission effects.

Fig. 3 and 4 represent the transmission speed as a function of distance from the ad-hoc station B to the server station C for 802.11b and 802.11a/g accordingly.

The results from Fig. 3 show that among many cards supporting 802.11b, the D-Link DWL-650+ card appears to be the best at lower distances. Furthermore, with the increase in distance, the Cisco AIR-PCM350 card

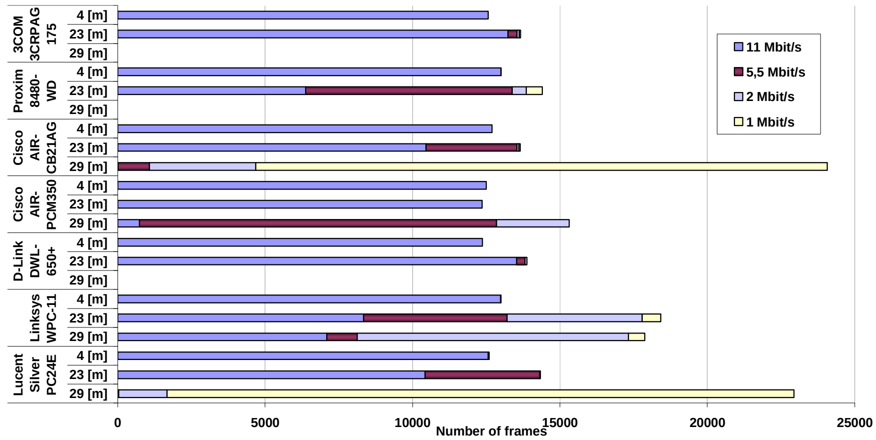

Fig. 2. TCP Scenario: Multirate for IEEE 802.11b cards

performance overcame all others (e.g., for distances exceeding 30 meters, its gain was 1.5 or higher).

Only two of the tested cards supported the 802.11g standard, i.e., Proxim 8480-WD and 3Com 3CRPAG175. As can be seen in Fig. 4, the cards operating in mode g achieved much lower transmission speeds than those working in mode a. This is due to more interferences occurring in the 2.4GHz band as not all of them could be eliminated (as mentioned previously). To summarize, Proxim 8480-WD (mode g) was the worst card, though, on average, almost the same transmission speeds were achieved by all three cards.

To summarize, the analysis of Fig. 3 and 4 allows choosing the appropriate device depending on the user's needs. The cards with the best short-distance efficiency are: D-Link DWL-650+, Cisco AIR-PCM350 and Cisco AIRCB21AG (mode a). On the other hand, the cards which assure the best efficiency with the increase in distance are: Proxim 8480-WD (mode a) and Cisco AIR-PCM350. The obvious conclusion is that the Cisco AIR-PCM350 works best between the 802.11b cards. However, among cards supporting all modes of operation, the Cisco AIR-CB21AG is the best choice.

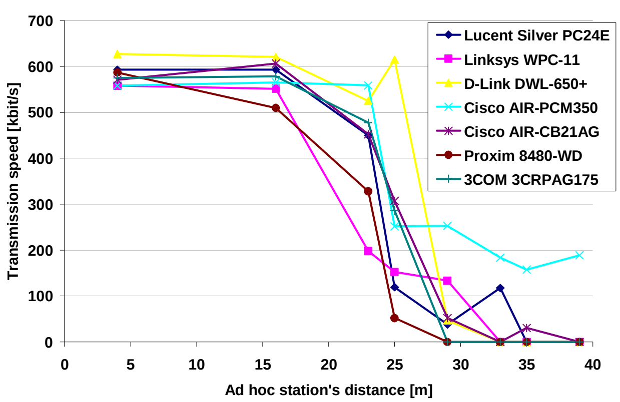

Fig. 3. TCP Scenario: Transmission speed as a function of distance for 802.11b cards

The explanation for the previous results comes from the number of frames transmitted with different speeds for the IEEE 802.11a/b/g standards as presented in Fig. 2, 5, and 6. For all test series, the number of frames transmitted at the highest speed decreased dramatically with distance. This effect was caused by the multirate mechanism and in all cases was dependant on vendor implementations.

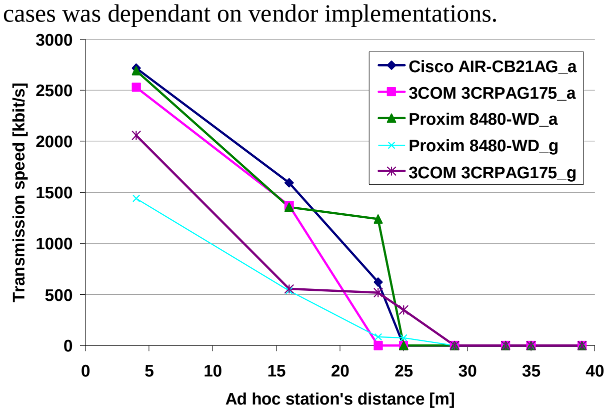

Fig. 4. TCP Scenario: Transmission speed as a function of distance for 802.11a /g cards

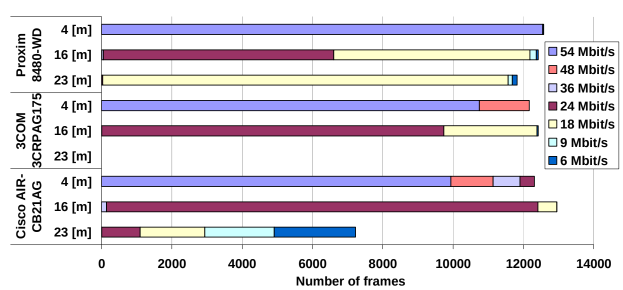

Fig. 5. TCP Scenario: Multirate for IEEE 802.11a cards

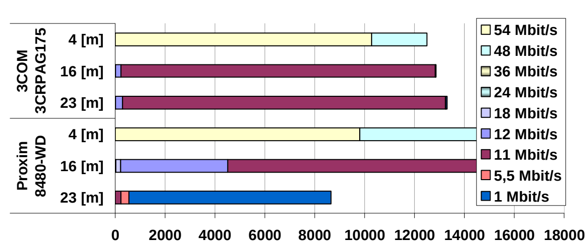

Fig. 6. TCP Scenario: Multirate for 802.11g cards

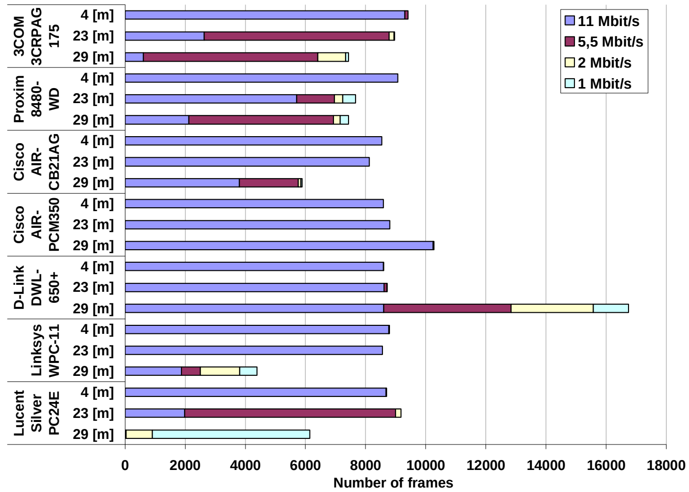

Fig. 7. FTP Scenario: Multirate for mobile 802.11b cards

In the second scenario, different speeds of sending frames over the physical layer were analysed. Two main characteristics were observed. First of all, equal medium sharing occurred when both clients (Stations A and B, c.f. Fig. 1) were close to the server and the transmission rates chosen by both stations were optimal. Secondly, the more Station B moved further away the more the transmission speed decreased (c.f. Fig. 7). Overall, Cisco AIR-PCM350 showed the best performance (almost regardless of distance), whereas Lucent Silver PC24E turned out to be the worst.

VII. Conclusions and Future Research

This paper presents the behaviour and performance analysis of IEEE 802.11a/b/g cards made by different vendors, working in ad-hoc mode. In particular, the multirate adaptability of the cards was analysed, when two or three stations were communicating.

The cards were tested in typical office conditions which resulted in an unpredictable radio channel (multipath propagation phenomenon). These conditions affected card performance as much as, or perhaps even more, than the transmission distance.

The obtained results show, that the performance of a WLAN card highly depends on its manufacturer. Some cards turned out to be significantly worse than others, because they implement the multirate functionality differently. Furthermore, the authors are convinced that the sensitivity of the cards also had a significant impact on the correct reception of packets.

Therefore, to achieve high performance, it is crucial to implement an appropriate algorithm which can choose the best transmission rate. If the rate is chosen too high, the frame error rate increases which leads to more retransmissions and, as a consequence, network performance decreases. If the card is not able to quickly adapt to varying radio channel conditions or if it chooses a rate which is too low, the degradation of network performance will also occur. Thus, high adaptability with the utilization of short periods of good conditions (in the form of TXOP or similar proposals) seems to be a good solution.

To summarize, when buying a WLAN card it is important to take into account not only the transmission rates but also other parameters (e.g., sensitivity, output power) as well as laboratory tests. Therefore, to facilitate the user's final choice, card behaviour in different scenarios has been presented in this paper.

Future research should provide further insight into the issues of multirate adaptation and card behaviour. A mathematical model describing the problem needs to be formulated and compared with experimental results. The problem of achieving multirate compatibility between cards belonging to different vendors can be studied as well.

References

[1] Awerbuch, B., Holmer, D., and Rubens, H.: High throughput route selection in multi-rate ad hoc wireless networks. Technical report, Johns Hopkins University, Computer Science Department (2003)

[2] Bicket, J.: Bit-rate Selection in Wireless Networks. Master's thesis, MIT (2005)

[3] Haratcherev, I., Taal, J., Langendoen, K., Lagendijk, R., and Sips, H.: Automatic IEEE 802.11 rate control for streaming applications: Research Articles. Wirel. Commun. Mob. Comput. 5, 4 (2005)

[4] Holland, G., Vaidya, N.H., and Bahl, P.: A rate-adaptive MAC protocol for multi-hop wireless networks. Mobile Computing and Networking (2001)

[5] IEEE 802.11 Standard for Wireless LAN: Medium Access Control (MAC) and Physical Layer (PHY) Specification, New York, IEEE Inc. (1999)

[6] IEEE 802.11a: High-Speed Physical Layer in the 5 GHz Band, New York, IEEE Inc. (1999)

[7] IEEE 802.11b: High-Speed Physical Layer Ext. in the 2.4 GHz Band, New York, IEEE Inc. (1999)

[8] IEEE 802.11e: Medium Access Method (MAC) Quality of Service Enhancements, New York, IEEE Inc. (2005)

[9] IEEE 802.11g: Further Higher-Speed Physical Layer Extension in the 2.4 GHz Band, New York, IEEE Inc., 2003

[10] Kamerman, A., and Monteban, L.: WaveLAN-II: A high-performance wireless lan for the unlicensed band. Bell Labs Technical Journal (1997)

[11] Lacage, M. , Manshaei, M.H., Turletti, T.: IEEE 802.11 Rate Adaptation: A Practical Approach. INRIA Research Report, No 5208 (2004)

[12] Madwifi: Multiband Atheros Driver for WiFi. http://madwifi.sourceforge.net/

[13] Natkaniec M., Pach A.R., Analysis of Backoff mechanism in IEEE 802.11 standard – ISCC'2000 – IEEE Symposium on Computers and Comm., An-tibes – Juans les Pins, 3-6 July 2000 France

[14] Di Stefano, A., Terrazzino, G., Scalia, L., Tinnirello, I., Bianchi, G., and Giaconia, C.: An Experimental Testbed and Methodology for Characterizing IEEE 802.11 Network Cards. Proceedings of the 2006 international Symposium on World of Wireless, Mobile and Multimedia Networks (2006)

[15] Vegt, A.J. van der: Auto Rate Fallback Algorithm for the IEEE 802.11a Standard. Technical report, Utrecht University (2002)