<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications 2012 (PIMRC 2012)*.

Citation for the original published paper:

Schmidt, F., Hithnawi, A., Punyal, O., Gross, J., Wehrle, K. (2012)

A Receiver based 802.11 Rate Adaptation Scheme with On-Demand Feedback.

In:

<http://dx.doi.org/10.1109/PIMRC.2012.6362818>

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-136847>

A Receiver-Based 802.11 Rate Adaptation Scheme with On-Demand Feedback

Florian Schmidt*, Anwar Hithnawi*, Oscar Puñal†, James Gross†, Klaus Wehrle*

*Communication and Distributed Systems Group

†Mobile Network Performance Group

RWTH Aachen University, Germany

Email: {schmidt|hithnawi|wehrle}@comsys.rwth-aachen.de, {punal|gross}@umic.rwth-aachen.de

Abstract—Classical 802.11 rate adaptation algorithms rely on feedback from the receiver to correctly choose a sending rate, typically in the form of acknowledgments (ACKs). In the absence of such frames, novel techniques are required for rate selection.

We present a novel On-Demand Feedback Rate Adaptation algorithm (OFRA) that works with ACK-less traffic. Feedback information is sent on-demand using a control frame to explicitly inform the transmitter about which bit rate to use on subsequent data frames. This approach guarantees standard conformity and exhibits fast and accurate bit rate adaptation at the cost of a modest overhead increase. We evaluate the performance of OFRA against various state-of-the-art rate adaptation schemes by means of simulations. If ACK frames are to be transmitted, OFRA performs better than related work in most considered scenarios, and on par in the others. In the absence of ACKs, OFRA provides large goodput gains under good channel conditions and comparable goodput in other situations.

I. INTRODUCTION

The IEEE 802.11 wireless local area network (WLAN) standard remains the most popular way to exchange data over wireless links. One of the fundamental problems of any wireless technology is the volatile nature of the channel, which requires adaptation to its time-changing properties. To this end, the 802.11 standard defines a set of modulation and coding schemes (MCS), each of which has a nominal transmission rate that allows a trade-off between robustness and speed of the transmission; hence the term *rate adaptation* (RA). However, there is no single standardized system in place to adapt to the most efficient rate at any given point in time. Instead, numerous RA algorithms have been proposed [1]–[12] over the years.

One fundamental problem of rate adaptation is scarcity of information. The sender needs to adapt the transmission rate; however, the information about reception quality is only available at the receiver, and needs to be fed back to the sender by some means. Typically, RA algorithms employ acknowledgment frames (ACKs) as feedback mechanism. ACKs are sent by 802.11 for each successfully received data frame under normal circumstances, so their use for rate adaptation purposes does not introduce any additional overhead. However, sending of acknowledgments has been made optional with the 802.11e QoS extensions. Not sending acknowledgments allows using that time to send additional data packets, increasing spectral efficiency. This can be beneficial for traffic that does not benefit much from retransmissions, such as in

vehicular networks or for multimedia streaming that might also employ error-tolerance at the receiver [13]. In the absence of acknowledgments, most well-known RA schemes cease to function, because they lack information transfer from receiver to sender.

We therefore propose a novel **On-Demand Feedback Rate Adaptation** algorithm (OFRA) with the following properties:

- To the best of our knowledge, it is the first rate adaptation scheme that works with ACK-less traffic without limitations such as requiring significant bidirectional traffic.
- It sends feedback from the receiver to the sender on-demand, that is, only if the channel conditions change significantly, instead of for every data frame.
- It is backward compatible as the sending of feedback does not interfere with the correct functioning of legacy nodes.

The rest of the paper is structured as follows: In Section II, we present related work and compare it to OFRA. The overall design of OFRA is discussed in Section III. We describe the simulator setup used to evaluate OFRA and present results in Section IV. Finally, we discuss possible future work and conclude the paper in Section V.

II. RELATED WORK

Rate adaptation algorithms for 802.11 networks have been intensively studied in literature [1]–[12]. These are typically categorized based on the metrics that trigger adaptation and on the communication peer that performs the selection of the rate. For instance, *frame-loss* based schemes select the rate depending on the delivery success of previously transmitted frames [1]–[6] and are widely used by commercial 802.11 devices mostly due to their ease of implementation.

Despite the popularity of frame-loss based schemes, it has been shown [11] that they do not provide a timely and efficient reaction to fast changes in the wireless channel conditions. To overcome this limitation, multiple works propose to adapt the rate according to the observed signal quality, which is commonly characterized by the signal-to-noise ratio (SNR). SNR can be measured at the transmitter [7], [8] or at the receiver [9]–[11]. Transmitter-based schemes extract the channel quality information from packets generated at the receiver exploiting and assuming channel reciprocity. However, the latter has been determined [11] as too optimistic for certain types of propagation environments (e.g., vehicular channel).

Fig. 1. Design of the structural components of OFRA.

Furthermore, the transmitter relies on incoming control frames (CTS or ACK) or DATA frames. This either introduces large overhead (RTS/CTS handshake) or, in case of unidirectional ACK-less traffic, such approaches simply do not work. On the other hand, transmitter-based approaches avoid the overhead associated with explicit feedback. Receiver-based algorithms such as [9], [12] use the RTS/CTS handshake to gather channel information. While they perform the rate adaptation based on *fresh* channel estimations (prior to the data packet), they require modification of the CTS frame to include explicit feedback about the rate selected by the receiver and, hence, lose standard conformity. Alternatively, the authors in [11] propose a receiver-based scheme that conveys feedback information implicitly in ACKs by exploiting the different bit rates available for these frames, which is easily implementable on commodity 802.11 devices. On the other hand, the scheme relies on the existence of ACKs and may induce ACK errors as these frames are potentially sent with aggressive modulations. In addition, only one bit rate level increase is allowed at a time, which depending on the communication conditions may represent a significant shortcoming. For instance, in the simulation scenarios we evaluate in this paper, we observe that in up to 25% of the situations in which a bit rate change is advisable, a jump by more than one bit rate is optimal.

III. SYSTEM DESIGN

One of the main features of OFRA is that, as opposed to most other rate adaptation algorithms, it is a *receiver-based* system. This means the receiver uses the rich information available from frame receptions, and feeds back instructional information to the transmitter about which rate to use. The overall structure of OFRA is depicted in Figure 1. At the receiver, the channel is *estimated* by examining the SNR of each incoming correct data frame.¹ Based on this, the future state of the channel is *predicted*. For the work presented in this paper, we opted for a simple prediction that assumes that the SNR of the next data frame will be equal to the SNR of the frame that was just received. Nevertheless, more sophisticated sliding-window or derivative prediction schemes are possible.

The receiver then *selects* the optimal rate based on lookup tables that, for different frame sizes, contain information about which bit rate is optimal at which SNR. These tables have to be created only once and off-line, that is, before the rate adaptation algorithm is used for the first time. Due to the

¹While it would be beneficial to also react to erroneous frames for more timely feedback, we decided against this. In such a frame, the sender MAC address could be corrupted, which would lead to incorrect feedback.

Fig. 2. Frame format of the OFRA feedback control frame used by the receiver to instruct the sender which MCS to use. The feedback information is encoded in 3 bits for 802.11a/g.

Fig. 3. OFRA's on-demand feedback is sent more often in fast channels, facilitating fast adaption, and less often in slow channels, reducing overhead.

monotonous nature of SNR-to-rate mapping (a higher SNR results in a same or higher optimum bit rate), these tables contain, for a set of frame sizes, the thresholds at which to switch between rates. These tables are therefore small and lookups fast. For intermediate frame sizes, lookup table results are interpolated.

This information then needs to be *fed back* to the sender. Because transferring channel quality information feedback is expensive, we decided to send this feedback on-demand, that is, only if the channel conditions change sufficiently to warrant a new bit rate selection. Furthermore, because one of OFRA's design goals and contributions is support for ACK-less traffic, we cannot rely on encoding our feedback into the acknowledgment frames. We therefore created a special *feedback frame*. 802.11 defines several frame types (management, control, data), each of which has several sub-types (see [14], Sec. 7.1.3.1.2). Our feedback frame is a control frame with a subtype identifier that is reserved in the standard for future use. This way, we maintain backward compatibility with nodes that do not support OFRA. Figure 2 shows the format of our feedback control frame. Most of the fields in the frame are mandated by the standard; the actual feedback information is encoded into one octet, which consists of a version field, several so-far-unused bits, and the actual information in the 3-bit "MCS Select" subfield.

Feedback frames are sent at the lowest base rate (6 Mbps for 802.11a/g), and as other frames of the control frame type they are not acknowledged. The receiver implicitly deduces the correct reception of the feedback frame by inspecting the bit rate of subsequent data frames. If the receiver still perceives the need to change the bit rate, it will send another feedback frame. Figure 3 illustrates via a quarter-second excerpt from

TABLE I
SIMULATION PARAMETERS

Packet length ξ	1500 Byte
Carrier frequency f_c	5.2 GHz
Available power $\mathcal{P}$	40 mW
Noise power N_0	-101 dBm
Environmental speed v	[0.72, 4.32, 14.12] kph
Path loss exponent γ	3
Reference loss $ H_{pl(d_0=1m)} ^2$	46.67 dBm

our simulations how the frequency of feedback changes with the channel speed.

At the sender's side, data frames for communication with a previously unknown receiver are conservatively sent at the lowest base rate initially until the first feedback frame is received. Whenever the sender receives a feedback frame, it adapts the bit rate as indicated by the receiver and employs it on subsequent data frame transmissions to that receiver.

IV. EVALUATION

A. System Model

We focus on an 802.11a/g network in infrastructure mode, where an access point (AP) serves a different number of clients. Data packets of fixed size ξ [bits] are transmitted at carrier frequency f_c [Hz] making use of a total power budget $\mathcal{P}$. The chosen bit rate is selected out of the set $\{6, 9, 12, 18, 24, 36, 48, 54\}$ Mbps. Transmitted signals are attenuated due to path loss $|H_{pl}|^2$ and multi-path fading $|H_{fad}(t)|^2$. Assuming a (constant) background noise power of N_0 , the signal-to-noise ratio can be obtained as $\text{SNR}(t) = \mathcal{P} \cdot |H_{pl}|^2 \cdot |H_{fad}(t)|^2 / N_0$. Both transmitter and receiver are static, which translates in a time-invariant path loss attenuation. On the other hand, we assume a scatterers' speed of v [kph] which induces a Doppler frequency shift that results in time variant multi-path fading attenuation $|H_{fad}(t)|^2$.

B. Simulation Model

In our simulations we account for path loss attenuation by means of the *log-distance* model as follows:

$$|H_{pl}|^2 = |H_{pl(d_0)}|^2 + 10\gamma \cdot \log \frac{d}{d_0}, \quad (1)$$

where $|H_{pl(d_0)}|^2$ denotes the path loss at reference distance d_0 and γ represents the path loss exponent. Furthermore, we model Rayleigh distributed multi-path fading by means of the stochastic sum-of-sinusoids channel model [15]. The model provides time and frequency correlated fading instances based on the discrete tapped-delay-line propagation model. The model has been further parameterized with the cluster values provided in [16] (Table 1) yielding an RMS delay of 25 ns. Further information about the exact parameterization of our simulations is provided in Table I.

C. Simulation Methodology

For our evaluation we use the ns-3 simulator [17], which contains a complete network stack and an accurate 802.11a MAC implementation. As channel model we implement the sum-of-sinusoids model described in Section IV-B.

We select goodput as our primary metric, which is defined as the rate (in Mbps) of correctly decoded payload packets. Packet error rate and overhead are chosen as secondary metrics. The latter is defined as the percentage of time needed for the transmission of *overhead* (i.e., ACKs, OFRA's feedback frames, RTS/CTS, and their corresponding inter-frame spaces) compared to the overall simulation time. In case of an erroneous data reception the ACK frame is not transmitted, however, the channel has been previously reserved for a time span covering the hypothetical ACK transmission. Therefore, the overhead associated with ACKs is considered also in the cases where these frames are not transmitted:

$$\text{Overhead} = \frac{t_{ACK} + t_{feedback} + t_{RTS/CTS}}{t_{sim}} \quad (2)$$

We investigated the setups depicted in Figure 4. In the single-link topology, one station (STA) sends *saturated* UDP traffic to the access point (AP) over 60 seconds. This traffic profile is characterized by a constant full buffer state, where there is always a packet ready to be transmitted. We evaluated this setup for distances of 10, 20, 30, 45, 65, and 75 m between STA and AP. We also investigated a multi-link topology with 4 and with 8 nodes, where each STA communicates with the AP using the same traffic profile as in the single-link case. The nodes are evenly distributed along the circumference of a disc with the AP as the center. We used the same distances as in the single-link topology as radius for the disc in our evaluation. We repeated each of these simulation setups 20 times with different random seeds to provide confidence intervals to our results. For the evaluation we compare the performance of OFRA with and without ACK frames with four different state-of-the-art algorithms. The implementations of these four algorithms are part of the default ns-3 distribution. The considered schemes are briefly described below.

- **ARF**: Transmitter increases the rate after 10 consecutive successful data transmissions and lowers the rate after 2 consecutive failed data transmissions [1].
- **Minstrel**: 10% of the traffic consists of probe packets sent at different rates than the current one to *search* for more appropriate rates [6].
- **CARA**: It extends ARF by transmitting RTS/CTS frames whenever the data transmission has failed, hence, protecting against collisions [4].
- **RRAA**: Transmitter adapts the rate based on the success rate observed over a *short* time window. RTS/CTS are dynamically transmitted to protect against collisions [3].
- **OFRA-Ack & OFRA-NoAck**: OFRA scheme with and without ACK frames, respectively.

Fig. 4. Topologies used for the ns-3 evaluation. For the 8-node case, we introduced an additional 4 nodes at distance d from the AP and equidistantly between the STAs into case 4b.

D. Simulation Results

In the following, we present the results from our ns-3 simulations of OFRA. While we evaluated OFRA in several scenarios against various other RA algorithms, due to space constraints, we present only a subset of those results. Specifically, we only present distances of 10, 30, and 65 m, and only slow (0.72 kph) and fast (14.15 kph) channel speeds. We excluded the 75 m case because at this distance, intermittent connection losses occurred, which skewed the results. The other non-included cases do not show any unexpected results that cannot be inferred from those presented here. Specifically, the 8-node topology results were very similar to the 4-node case. Note that in the 4-node case hidden nodes can be observed at 30 m depending on the instantaneous multipath attenuation and at 65 m there is always at least one hidden node. For each of the presented results, we show the mean value over 20 simulation runs, and the 95% confidence intervals. Furthermore, for comparability reasons, we present two versions of OFRA: one with acknowledged traffic, and one without. While OFRA does not require ACKs, other RA algorithms do; to fairly compare metrics such as goodput and overhead, it is therefore necessary to compare them to OFRA working on acknowledged traffic. However, because support for ACK-less traffic is one of the key contributions of OFRA, we also present a setup with OFRA and NoAck traffic to highlight the behavior of such a setup in comparison.

1) *SNR Safety Margin*: Before we present the comparison of OFRA to other RA algorithms, we discuss a design decision that stems from the fact that SNR-based RA suffers from over-optimistic rate selection [11], especially in fast-changing channels. If the channel changes fast, the channel conditions experienced during the reception of a data packet may differ significantly from the conditions observed during the subsequent transmission. Hence, *aggressively* selecting the bit rate based on the theoretic SNR thresholds, without taking the (fast) channel fluctuations into account, leads to an increase in packet errors. We could clearly see the effect of this in our results and therefore investigated a static “safety margin”. This means that for rate adaptation decisions, the bit rate switching thresholds are up-shifted by a certain amount of SNR dBs, which leads to a more conservative rate selection. We present results for OFRA’s NoAck variant and for safety margins of

Fig. 5. Goodput and error rate performance for OFRA-NoAck with different safety margins values in the fast-channel scenario. Introducing a 2 dB margin reduces errors significantly.

0, 1, and 2 dB SNR in Figure 5. In both the single-node and the 4-node topology, the safety margin significantly reduces the frame error rate, while increasing goodput: although it reduces the average speed at which data is sent, it increases the robustness, so fewer frames are discarded. In the single-node topology, a 1 dB margin outperforms 0 dB (no margin) by a large amount. A 2 dB margin further reduces the error rate, at the expense of a slightly reduced goodput. In the 4-node topology, the 2 dB safety margin outperforms both 1 dB and 0 dB, in both error rate and goodput. The higher performance of the 2 dB margin in this case stems from the fact that, while the channel speed itself has not changed, four station contend for medium access. This increases the time between two consecutive frames sent by a single station, which leads to a larger variance in channel conditions. Based on these results, specifically considering the large reduction in error rate, we decided to use the 2 dB setting to compare OFRA to other RA algorithms in the subsequent sections.

2) *Single-Node Topology*: The least challenging setup, a single node with a slow channel (Figure 6a), is managed well by all RA algorithms. It also clearly shows the potential benefit of NoAck transmissions: Using the time otherwise spent for acknowledgments on further data transfers increases the goodput by up to 30%. As the distance increases, this advantage is reduced because lower rates are selected that cannot send as much additional traffic in the time freed by disabling ACKs. Noticeable is the extremely low error rate for both OFRA variants compared to the other algorithms. This underlines the effectiveness of our feedback scheme. The same holds for the fast channel scenario (Figure 6b). As expected, the packet error rate increases due to the higher channel speed, which does not prevent OFRA-Ack and -NoAck from outperforming the comparison RAs.

Our overhead metric illustrates the percentage of time that a certain scheme is busy transmitting control information instead of payload information. This metric allows us to quantify the overhead burden introduced by OFRA’s feedback frames and compare it with the overhead of related work schemes. In

Fig. 6. Aggregated goodput, error rate, and overhead results for the single-node topology. Even with acknowledgments, OFRA always provides similar or better goodput and lower error rates than the comparison algorithms. OFRA-NoAck shows a significantly higher throughput at low distances.

Figure 6a we observe that OFRA-NoAck exhibits an almost negligible overhead that corresponds to the low number of required feedback frames as the channel changes slowly. OFRA-Ack introduces around 10% overhead, mainly due to the ACK transmitted for every correctly received data frame. Note that all schemes (except OFRA-NoAck) introduce a similar amount of overhead in the single node case. The impact of overhead decreases under challenging conditions (i.e., large distances) as, in general, the relative channel occupancy time of control frames compared to the one of data frames is inversely proportional to the (data) bit rate. This pattern can be observed in all presented results. In a faster channel (Figure 6b), OFRA's overhead increases due to the larger amount of feedback frames required and the more accurate bit rate selection.

3) *4-Node Topology*: In the 4-node topology, the performance of the RA algorithms is more varied even in a slow channel (Figure 8a). Concerning goodput, ARF shows the overall worst performance, and Minstrel strong degradation as distance increases. OFRA outperforms all comparison RAs at least slightly. Note that it even outperforms CARA, which protects against collisions due to the hidden-node problem, an issue that OFRA does not target at specifically.

The results in the fast-channel setup (Figure 8b) are again comparable to those in the slow-channel. One interesting result is the performance of RRAA, which seems to be least influenced by channel speed, though it also shows the lowest goodput at short distances already. Again, both OFRA variants outperform all comparison algorithms in terms of goodput.

OFRA-Ack still shows an error rate comparable to the other schemes, even though it does not clearly outperform them by the large amount visible in the single-node topology. The benefits of OFRA-NoAck observed in the single-node

Fig. 7. CDF of burst error lengths for different RA algorithms (fast channel, 4-node scenario, 65m distance).

topology are, however, severely reduced in the multi-node scenario. Especially in the fast channel scenario at larger distances, the high frame error rate limits its performance. The main contributing factor to this are error bursts, which are the result of three factors working together: (1) With fast changes in the channel and several stations contending for the medium, the channel coherence between consecutive frames from a station is significantly reduced. (2) Currently, OFRA only sends feedback when receiving correct frames. Therefore, in the presence of prolonged *high-attenuation* phases, no feedback is sent that would lead to a lower rate choice and the probable recovery from error bursts. (3) Under normal circumstances, frame losses increase the contention window, which reduces the number of frames sent during such a period of low channel quality. In the NoAck case, the contention window is always kept at its minimum size, which results into a larger amount of data frames that are subsequently lost

Fig. 8. Aggregated goodput, error rate, and overhead results for the 4-node topology. As in Figure 6, OFRA provides at least the same goodput as the comparison algorithms, at typically similar error rates. OFRA-NoAck outperforms all other algorithms at short distances, but suffer from a stronger increase in error rate, which degrades its performance at larger distances.

Fig. 9. Accuracy of rate selection for 4-node topology in the slow channel. OFRA minimizes overselection and keeps correct selection high.

Fig. 10. Accuracy of rate selection for 4-node topology in the fast channel. OFRA-Ack has the lowest overselection and highest correct selection, while OFRA-NoAck's high accuracy degrades with the transmission distance.

due to bit-errors. This behavior, and the resulting larger burst error length, is highlighted in Figure 7. We will discuss and suggest solutions to overcome this problem in Section V.

With respect to overhead, we do not observe significant changes for OFRA compared to the single-node case. The overhead introduced by CARA (RTS/CTS after every corrupted frame) is magnified in this scenario, as the packet error rate is larger due to collisions. We observe that the overhead of ARF is reduced, which is a side-effect of the, in general, lower bit rate selected by this scheme; as fewer data frames are transmitted, fewer ACK frames are correspondingly sent. The situation changes slightly in the fast channel case (Figures 8b). Here we observe the larger impact of OFRA's feedback frames, which are transmitted more often due to the faster changing channel conditions. This is highlighted by the, now noticeable, overhead of OFRA-NoAck.

Finally, OFRA's high performance in most cases can be explained by the high accuracy in selecting the correct bit rate for data transmissions. Conversely, the lower performance of OFRA-NoAck in fast-channel large-distance scenarios is reflected in the lower accuracy. To measure the accuracy we check, once the simulation is finished, the SNR for each frame after it was received, and which bit rate would have been the optimal for this SNR. With this information it can be then decided if the RA performed accurately or if it under/over-selected the rate. In a slow channel, OFRA accurately selects the bit rate and is only outperformed by CARA at the cost of a larger error rate and overhead (see Figure 9). The conservatively selected safety margin of 2 dB keeps the rate over-selection at negligibly low values. In a fast channel on the other hand, as shown in Figure 10, the percentage of over-selection increases for all schemes. Under these challenging

conditions, OFRA's effective feedback mechanism and accurate rate selection yield the overall best performance.

V. CONCLUSION AND FUTURE WORK

OFRA is characterized by a continuously good performance under different communication conditions: different channel speeds, different number of transmitting nodes, and different average SNR values. Furthermore, it supports rate adaptation for unacknowledged traffic, which so far has not been provided by state-of-the-art algorithms. In most scenarios, OFRA shows a very high accuracy in selecting the most efficient bit rate, which leads to high goodput and low error rate.

We identify the following aspects as future work:

(1) *Feedback to erroneous frames*: OFRA currently only reacts with feedback to correctly received frames. This can lead to larger error bursts than in other RA schemes. We are currently investigating the feasibility of also reacting to erroneous frames. The main problem here is to make sure that, due to a corruption in the MAC address field, we do not send feedback to the wrong node. However, it has been shown before [13] that the heuristic repair of addresses is feasible.

(2) *Support for 802.11 networks*: We expect that due to the higher number of available rates, OFRA's bit rate selection scheme will work very well in 802.11n networks. OFRA's feedback frame layout (see Figure 2) easily allows extension to accommodate the additional rates available in 802.11n by specifying a new version of the feedback frame.

(3) *Dynamic lookup table thresholds*: In the current implementation of OFRA, the lookup tables that are used to decide on the optimum bit rate based on SNR are calculate once. Algorithms such as CHARM [7] show that dynamic adaptation of thresholds during runtime can further increase performance. Such an adaptation could be integrated into OFRA without difficulty, and further improve performance.

(4) *Dynamic safety margins*: For this paper, we decided to employ a fixed safety margin of 2 dB between optimal bit rate and feedback instruction, to reduce overselection (see Section IV-D1). From those results, it is clear that the optimal margin depends on conditions such as channel speed and number of stations in the network. Therefore, we expect further improvement if the margin is adapted to environmental conditions. The number of stations and the channel utilization can be signaled by the AP, as defined in the 802.11e extensions (see [14], Sec. 7.3.2.18). This information could be used to adapt the safety margin dynamically.

Overall, this paper shows that OFRA, despite the overhead introduced by its feedback, is a performant bit rate adaptation scheme with high goodput and low error rates in most scenarios. So far, ACK-less traffic was limited in its use due to the reliance on fixed-rate transmission. OFRA opens up possibilities to further investigate the use of ACK-less traffic in scenarios such as error-tolerant transmissions and vehicular networks.

ACKNOWLEDGMENTS

This research was funded in part by the DFG Cluster of Excellence on Ultra High-Speed Mobile Information and Communication (UMIC).

REFERENCES

- [1] A. Kamerman and L. Monteban., "WaveLAN-II: A High-Performance Wireless LAN for the Unlicensed Band," *AT&T Bell Laboratories Technical Journal*, 1997.
- [2] M. Lacage, M. H. Manshaei, and T. Turletti, "IEEE 802.11 Rate Adaptation: A Practical Approach," in *Proc. MSWIM*, 2004.
- [3] S. H. Y. Wong, H. Yang, S. Lu, and V. Bharghavan, "Robust Rate Adaptation for 802.11 Wireless Networks," in *Proc. MobiCom*, 2006.
- [4] S. Kim, L. Verma, S. Choi, and D. Qiao, "Collision-Aware Rate Adaptation in Multi-Rate WLANs: Design and Implementation," *Computer Networks*, 2010.
- [5] Onoe Bit-rate Selection Algorithms. [Online]. Available: <http://madwifi-project.org/wiki/UserDocs/RateControl>
- [6] Minstrel Specifications. [Online]. Available: <http://linuxwireless.org/en/developers/Documentation/mac80211/RateControl/minstrel>
- [7] G. Judd, X. Wang, and P. Steenkiste, "Efficient Channel-Aware Rate Adaptation in Dynamic Environments," in *Proc. MobiSys*, 2008.
- [8] J. Zhang, K. Tan, J. Zhao, H. Wu, and Y. Zhang, "A Practical SNR-Guided Rate Adaptation," in *Proc. INFOCOM*, 2008.
- [9] G. Holland, N. Vaidya, and P. Bahl, "A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks," in *Proc. MobiCom*, 2001.
- [10] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic Media Access for Multirate Ad Hoc Networks," in *Proc. MobiCom*, 2002.
- [11] X. Chen, P. Gangwal, and D. Qiao, "Practical Rate Adaptation in Mobile Environments," in *Proc. PerCom*, 2009.
- [12] J. Camp and E. Knightly, "Modulation Rate Adaptation in Urban and Vehicular Environments: Cross-Layer Implementation and Experimental Evaluation," *IEEE/ACM Transactions on Networking*, 2010.
- [13] F. Schmidt, M. H. Alizai, I. Aktaş, and K. Wehrle, "Reflector: Heuristic Header Error Recovery for Error-Tolerant Transmissions," in *Proc. CoNEXT*, 2011.
- [14] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification*, IEEE Std. 802.11, 2007.
- [15] C.-X. Wang, M. Patzold, and Q. Yao, "Stochastic Modeling and Simulation of Frequency-Correlated Wideband Fading Channels," *IEEE Trans. Vehicular Technology*, 2007.
- [16] X. Zhao, J. Kivinen, and P. Vainikainen, "Tapped Delay Line Channel Models at 5.3 GHz in Indoor Environments," in *Proc. VTC-Fall*, 2000.
- [17] The ns-3 network simulator. [Online]. Available: <http://www.nsnam.org/>