

ERA: Effective Rate Adaptation for WLANs

Saâd Biaz, Shaoen Wu

Computer Science and Software Engineering Department
Auburn University
Shelby Center for Engineering Technology, Suite 3101,
Auburn University, AL, 36849-5347, USA
{biazsaa, wushaoe}@auburn.edu

Abstract. Rate adaptation consists of using the optimal rate for a given channel quality: the poorer the channel quality, the lower the rate should be. Multiple rate adaptation schemes were proposed and studied so far. The first generation rate adaptation schemes perform well in collision free environments and manage quite well strict channel degradation. Under a congestion dominated environment, these schemes poorly perform because they do not differentiate losses due to channel degradation from those due to collisions and unnecessarily decrease the rate in response to collisions. A second generation rate adaptation schemes overcome this limitation. However, these most recent schemes usually require RTS/CTS control frames that are not in general used because they constitute an overhead that may heavily lower network performance especially when frame size is small. This work proposes an effective and practical rate adaptation scheme (*ERA*) that does not require RTS/CTS control frames. *ERA* judiciously uses fragmentation in compliance with IEEE 802.11 standard to diagnose the cause of frame loss and to promptly recover from frame losses. Through extensive simulations on *ns-2*, *ERA* exhibits a significant throughput improvement over other rate adaptation schemes.

1 Introduction

A wireless medium suffers frame losses from unstable channel conditions resulting from signal fading and/or interference. When channel conditions degrade, a transmitter should use more robust encoding schemes that enable more reliable communications, but yield lower data rates: *rate adaptation* is the determination of the optimal rate most appropriate for current estimated wireless channel conditions.

Rate adaptation in IEEE 802.11 networks has been studied at length in recent years and many rate adaptation schemes were proposed. Most first generation rate adaptation schemes implicitly assume that all packet losses are due to channel degradation with limited congestion losses [1–4]. Such an assumption is valid as long as the MAC layer deploys some collision avoidance mechanisms such as the use of RTS/CTS control frames to eliminate or minimize congestion losses. Some rate adaptation schemes [5, 6] rely on such RTS/CTS control frames to adapt the rate and minimize congestion losses. However, based on IEEE 802.11 standard, RTS/CTS control frames are optional and are recommended only when data frames are large (by default 2347 bytes in Annex D). Therefore RTS/CTS control frames are usually not enabled, thus limiting

RTS/CTS based rate adaptation schemes. Even with RTS/CTS control frames, congestion losses may still occur and mislead these rate adaptation schemes. Most rate adaptation schemes systematically respond to frame loss by decreasing their data rate. While such a response is appropriate for channel degradation, it is not for collisions for two reasons: 1) a lower rate may exacerbate medium congestion because of longer frame transmission duration and wider transmission range (more interference); 2) a lower rate is wasteful and unnecessary as channel conditions may well support a higher rate.

Thus, in recent years, researchers designed second generation rate adaptation schemes that diagnose the cause of a loss and appropriately react [7, 8]. But these schemes require the use of RTS/CTS control frames that are pure overhead. The design of a rate adaptation scheme differentiating losses without RTS/CTS frames is desirable, but challenging. The authors of this work proposed a previous scheme [9] that does not require RTS/CTS control frames, but that scheme may incorrectly diagnose the cause of a loss.

This paper proposes and evaluates a new effective rate adaptation scheme, dubbed **ERA**, that does not require RTS/CTS frames, is compatible with IEEE 802.11 standard, and efficiently works in *both* channel degradation and collisions dominated environments. The strategy in *ERA* yields two benefits: 1) an accurate diagnosis of the cause of a loss (channel degradation or congestion), and 2) a faster recovery from losses. *ERA* achieves these benefits by fragmenting the retransmitted lost frame in two fragments: a very short frame and the remainder of the retransmitted frame. For an accurate loss differentiation, *ERA* exploits the loss rate and the signal-to-noise ratio after a frame loss. *ERA* maintains the rate unchanged for collisions and judiciously adapts the rate for channel degradation based on a scheme described later. Through extensive *ns-2* simulations, *ERA* is quite effective under (congestion) collisions and channel degradation and outperforms previous rate adaptation schemes with loss differentiation. The primary contributions of this paper are:

- A simple, robust, and effective rate scheme with loss differentiation that does *not* require RTS/CTS frames.
- An efficient recovery strategy from losses.
- An extensive evaluation and analysis of ERA and other previous data rate adaptation schemes

The rest of this paper is organized as follows: The related work is presented in Section 2. Section 3 reviews the background information of IEEE 802.11 relevant to this work: distributed coordination function basic operation and fragmentation. Then we present the rationale of *ERA*. Section 4. Afterwards, the effective rate adaptation *ERA* scheme is described in Section 5. We evaluate the performance of *ERA* and other schemes in Section 6 through extensive *ns-2* simulations. Section 7 concludes this paper.

2 Related Work

This section does not cover all previous rate adaptation schemes, but it surveys the most typical schemes relevant to this work. First, we present first generation schemes: these schemes do not differentiate losses due to channel degradation from those due to collisions. Second, we present schemes with loss differentiation.

2.1 First Generation: Rate Adaptation without Loss Differentiation

Scheme *Auto Rate Fallback (ARF)* by Kamerman and Monteban [1] is the earliest rate adaptation proposed for IEEE 802.11 based wireless networks. It is simple and intuitive. If two consecutive transmissions fail, *ARF* decreases the rate and resets a timer. If either the timer expires or the sender achieves ten successful transmissions consecutively at rate r_{old} , *ARF* increases the rate to r_{new} and the timer is reset. Note that if a transmission fails immediately after a rate increase, *ARF* falls back to the prior rate r_{old} . Under a stable environment, *ARF* keeps increasing its rate leading to periodical failures and rate oscillations.

Scheme *Adaptive Auto Rate Fallback (AARF)* by Lacage, Manshaei, and Turletti [2] mitigates *ARF* oscillations by varying the threshold number N of successful transmissions required to increase the rate. *AARF* increases the data rate after N consecutive successful transmissions and decreases the data rate if the transmission fails twice consecutively. If the first transmission fails at the newly increased rate, *AARF* falls back to the prior rate just like *ARF*. In addition, *AARF* doubles the threshold required for the next rate increase to $2N$.

In general, most data rate adaptation schemes such as **SampleRate** [10] or **Onoe** (MadWifi Linux driver implementation) collect statistics about losses in the most recent past and adjust the rate according to some rules.

2.2 Second Generation: Rate Adaptation with Loss Differentiation

Recent rate adaptation schemes focus on the differentiation of the cause of frame losses between channel degradation and collision.

The scheme dubbed *Robust Rate Adaptation Algorithm (RRAA)* by Wong, *et al.* [8] is a frame loss rate driven scheme with loss differentiation. It keeps track of short term frame loss rates, and accordingly adapts the rate. In order to address the hidden terminal problem, Wong proposed the “Adaptive RTS Filter” that uses RTS frames only to assess new channel conditions when needed and to minimize collisions. “Adaptive RTS Filter” *turns on* RTS frames to prevent collision losses such that the frame loss rate reflects channel degradation.

Kim *et al.* designed a *Collision-Aware Rate Adaptation (CARA)* scheme that uses RTS/CTS control frames whenever a frame loss occurs. *CARA* assumes that since all RTS frames are transmitted at the lowest rate, then they are resilient to channel fading and would fail only due to collisions. On the other hand, if RTS is transmitted successfully, but the transmission of the following data frame at higher rate fails then this failure is attributed to channel degradation, and *CARA* decreases the rate as in *ARF*. *CARA* does not consider the scenario where an RTS frame could be lost due to a receiver node moving out of radio range.

In parallel and independently from *RRAA* and *CARA*, we proposed a *Loss Differentiation Rate Adaptation (LDRA)* [9]. *LDRA* retransmits a lost frame at the lowest (basic) rate to diagnose the cause of a frame loss, and appropriately adjusts the rate. *LDRA* strategy is analytically justified and particularly efficient in low SNR environments. But *LDRA* does not perform well under a congested environment and relies upon the signal-to-noise ratio (SNR). Previous experiments [10, 11] show a weak correlation between delivery probability at a given rate and SNR.

This work addresses the design of a rate adaptation scheme with 1) loss differentiation, 2) a judicious retransmission strategy, and 3) some collision avoidance. The key challenge is that the new scheme does *not* use RTS/CTS control frames, but yet delivers better performance than other schemes proposed so far.

3 IEEE 802.11 Standard

This section presents key aspects of IEEE 802.11 standard relevant to this paper.

3.1 Distributed Coordination Function

This work is limited to the widely deployed IEEE 802.11 distributed coordination function (DCF) [12] and considers the DCF basic protocol without RTS/CTS frames preceding data frames. IEEE 802.11 standard discourages the use of RTS/CTS frames when the data frame length is less than the *RTS threshold* (2347 recommended). Garg and Kappes [13] showed that, for VoIP traffic with 160 byte frames, the data frame efficiency drops to about 12% in 802.11b [14] networks at 11 Mbps when RTS/CTS control frames are used. Thus, RTS/CTS frames are usually not used.

3.2 Fragmentation in IEEE 802.11

Fragmentation is *usually* performed for packets too large or unfit for current channel conditions. Under a lossy environment, a 1500 byte frame may yield poor performance and its fragmentation into optimally sized frames may dramatically improve performance. As illustrated in Figure 1, a sender divides a frame F into two fragments $Frag_1$ and $Frag_2$. Fragment $Frag_1$ bears a network allocation vector (NAV) NAV_1 : NAV_1 reflects the time duration that the sender requires to complete the transmission of the entire frame F (till the reception of ACK_2). All stations that hear fragment $Frag_1$ will remain silent for duration NAV_1 . The receiver of $Frag_1$ will acknowledge it with an ack ACK_1 that bears a new NAV that also covers the complete transmission of the entire frame F . Therefore, in a collision dominated environment, if the first fragment is successfully acknowledged, the channel is cleared for all remaining fragments of the same frame. *The probability of collision for the whole frame is reduced to the probability of collision of its first fragment.*




Fig. 1. Transmission of fragments in IEEE 802.11




Fig. 2. Throughput under strict channel degradation




Fig. 3. Throughput under pure collision

4 Rationale

Despite the fact that SNR is not a good indicator [10, 11], *LDRA* yields a significant throughput improvement when there are no collision losses and no hidden terminal problem. *The strength of LDRA is due to its retransmission strategy that provides prompt recovery.* Figure 2 illustrates this strength by plotting the throughput between two static stations communicating at different distances. The Ricean model was used to create a variable channel quality. In presence of channel degradation without collisions, Figure 2 also illustrates as expected that 1) rate adaptation is beneficial (the scheme that maintains a constant rate (**CONSTANT**) yields a poor throughput, close to null, because its transmissions at the highest rate almost fail all the time.), 2) a scheme without loss differentiation (*ARF*) performs as good as *CARA* (loss differentiation based).

On the contrary, the **CONSTANT** rate scheme yields the best throughput in a congested network as shown on Figure 3. These results suggest that *decreasing the rate for collisions hurts performance.* The good performance of *CARA* in this figure illustrates another fact: *a rate adaptation scheme not only should not decrease its rate in response to collisions, but it should proactively diffuse congestion: CARA uses RTS frames to avoid collisions.*

Based on the above observations, a good rate adaptation scheme should 1) accurately distinguish losses due to channel degradation from those due to collisions, 2) appropriately react (decrease rate only in response to channel degradation), 3) develop an efficient strategy for recovery from losses, and 4) diffuse congestion when collisions occur. So far, RTS/CTS frames were used [7, 8] to probe the channel and/or alleviate congestion to yield efficient rate adaptation schemes. However, since RTS/CTS control frames may represent a significant overhead and are usually not used, it is particularly desirable, though challenging to develop, such a rate adaptation scheme *without* using RTS/CTS frames.

In order to address the challenge of not using RTS/CTS frames, we plan to use fragmentation to probe the channel, to diagnose the cause of a loss, and to develop an efficient retransmission strategy for loss recovery.

5 ERA: An Effective Rate Adaptation

When a frame is lost, *ERA* deploys effective strategies to *diagnose* the channel, to *adjust* the rate, and eventually to *recover* from the loss. We outline the key elements of *ERA* while the precise scheme is described by the flowchart in Figure 4.

5.1 ERA Frame Loss Diagnosis

Let r_{fail} be the rate used for the current lost packet and $r_{previous}$ be the rate of the latest successful transmission before the current failed transmission. We have to consider two cases:

- $r_{fail} > r_{previous}$: This means that we just increased the rate to r_{fail} after a successful transmission at rate $r_{previous}$. It seems that the channel does not support the new higher rate r_{fail} . We conclude that the loss is likely due to channel degradation at rate r_{fail} : the wireless channel condition can only support $r_{previous}$.

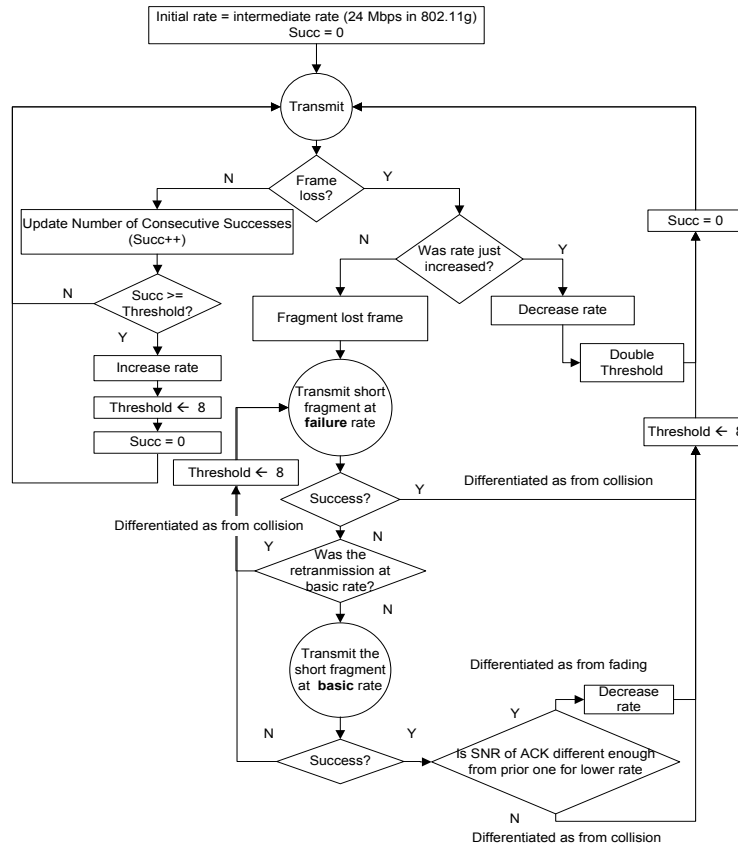


Fig. 4. ERA flowchart

- $r_{fail} = r_{previous}$: Since previous transmissions succeeded at the same rate r_{fail} , the loss may be due to a channel degradation as well as a collision. We assume the channel dynamic is fairly stable in a short recovery process, which is reasonable because the time on recovering a frame loss is not long compared to the channel fading. In order to differentiate, we fragment the lost frame into two fragments: a very short frame and the remainder fragment. The first (short) frame is retransmitted at the same rate r_{fail} . **If the transmission is successful** for this short frame, then the loss was most likely due to a collision. This inference is reasonable because of two reasons: 1) the probability of delivery is higher for a smaller frame in a collision prone environment, 2) the probability of delivery is not significantly higher for a smaller frame in a degraded channel. **Otherwise**, if the transmission of the short frame is not successful, then the loss is likely due to channel degradation, and further procedure shown in flowchart on Figure 4 refines this diagnosis.

A **key feature** of ERA consists of using the rate r_{fail} to retransmit the short frame rather than using the basic (lowest) rate immediately after the *first* failure. This feature allows to refine the diagnosis loss in contrast with other schemes: the other schemes use short frames sent at the basic rate immediately after the first failure; if the transmission is successful, it is unclear if the transmission is successful due to the use of a *lower* (basic) rate (channel degradation) or the use of a short frame (collisions).

5.2 ERA Rate Adjustment

As shown on the flowchart in Figure 4, ERA decreases the rate when a transmission fails immediately after a rate increase. Another instance of rate decrease occurs when

ERA successfully retransmits the short frame at the basic rate following two consecutive transmission failures at rate r_{fail} . *ERA* uses the same strategy as *AARF* to increase the rate: the rate is increased after *Threshold* consecutive successful transmissions. The value *Threshold* is the number of consecutive successful transmissions required to increase the rate. *Threshold* should not be too high to allow a sender to quickly reach the highest available rate and should not be too small to avoid rate oscillations and frequent losses due to overly optimistic increases. The value *Threshold* is dynamically adjusted similarly to *AARF*: for each frame loss due to channel degradation, *ERA* doubles the value *Threshold* up to 64.

5.3 ERA Frame Loss Recovery

The strategy that *ERA* adopts for frame loss recovery is efficient because it addresses both channel degradation and collisions. To tackle collisions, the fragmentation into a very short frame followed by the remaining fragment increases the probability of delivery. To tackle **channel degradation**, *ERA* uses our *LDRA* strategy: use the basic rate in case of the failure of the short fragment (at rate r_{fail}) likely due to channel degradation.

6 Performance Evaluation

The performances of *ARF*, *CARA*, and *ERA* are respectively evaluated and compared under different scenarios: strictly congested network (no losses due to channel degradation), strict channel degradation (collision free), and the more usual case where mixed losses may occur due to channel degradation as well as collisions. *CARA* was chosen for comparison as a representative of the schemes, such as *RRAA*, that use RTS/CTS control frames to address collisions in rate adaptation.

Performance evaluation is conducted through extensive simulations using *ns-2* network simulator. IEEE 802.11g [15] is used as the MAC layer because it has more predefined rates. For the physical layer, we use the parameters of Cisco Aironet 802.11a/b/g cardBus wireless LAN adapter [16] and assume symmetric links. Table 1 summarizes the network configurations and parameters used in these simulations; and Table 2 describes the main four network topologies and the number of stations (including the access point) in each one. Figure 5 shows the *static* topology with 16 stations and an access point. The evaluation is limited to one hop communication: all stations are within the radio coverage of the same access point. Note that these client stations are not necessary within the radio coverage of each other. The traffic is always from the stations to the access point.

Simulator	<i>ns-2</i> v.30
MAC	CSMA/CA DCF basic access
MAC data rates	6, 9, 12, 18, 24, 36, 48, 54 Mbps
Channel fading	Two-ray ground or Ricean
Ricean fading factor (K)	4
Traffic	UDP Constant Bit Rates (CBR)
Frame length	100, 500, 1000, 1500 Bytes
The short fragment length	Same as RTS

Table 1. Simulation parameters

Topologie	Number of Stations	Motion Model
#1	2	static
#2	2	random movement
#3	17	static
#4	17	random movement

Table 2. Simulation topologies

Each result figure plots the data averaged from multiple runs with random starting time for different evaluation environments. Each run of simulation is of 2-minute saturated UDP traffic. We evaluate *ERA* in three environments: channel degradation dominated, collision dominated, and mixed (frame losses due to collisions as well as channel degradation). First, we present the experiments made in a mixed environment.

6.1 Mixed Environments: Collisions and Channel Degradation

First, we evaluate the presumed strength of *ERA*, i.e., its effectiveness to promptly recover from losses. Then we measure the throughput (by Kframes/s) for the static network on Figure 5 and for a network of mobile stations.

ERA Retransmission/Recovery Effectiveness We measure the effectiveness by collecting the fraction of frames losses recovered after only ONE retransmission. For example, if there are 100 losses and forty of them were recovered after only one retransmission, then the fraction is 0.40. The experiments are conducted over a static network with 16 stations and one access point located at the center as shown on Figure 5. We use topology #3 (See Table 2). We adopt the Ricean fading model to induce dynamic channel conditions. By varying the size of the network layout area, we get a variable of loss ratios due to collisions: the collision ratio is the number of frame losses due to collisions over the total number of frame losses. If the collision ratio is null, this means that the environment is collision free (channel degradation dominated), which is emulated with only one client station. A strictly collision dominated environment has a collision ratio of one, which happens when all the stations are so close to the access point that there is no transmission failure due to channel fading even with Ricean model. Figure 6 plots the fraction of recovered losses after only one retransmission on the y -axis with the collision ratio on the x -axis. Under most conditions, *ERA* succeeds in recovering 80% of the losses with only one retransmission of the whole lost frame and outperforms *CARA* under all circumstances. In a channel degradation dominated environment (close to null collision ratio), *CARA* does not recover quickly because it often addresses the loss as a collision.




Fig. 5. Topology of 17 static stations




Fig. 6. Fraction of recovered losses with ONE retransmission




Fig. 7. The mean of retransmission

For the same experiment, we also compute the mean of the number of retransmissions plotted on the y -axis of Figure 7. The strength of *ERA* stems from its refined

mechanism to diagnose the cause of a frame loss and adopt the best strategy to recover. But for *CARA*, the success of RTS/CTS at the *basic* rate can not indicate channel fading or collision.

Throughput for Static Network To evaluate *ARF*, *CARA*, and *ERA* under realistic channel conditions with both channel degradation and collision, we use the Ricean fading model on a static network of 16 stations (see Figure 5) and vary the area size to adjust the congestion level. We use topology #3 (See Table 2). CBR traffic with 1000 bytes packets is sent from the 16 stations to the access point. Figure 8 plots on the y -axis the throughput achieved by *ARF*, *CARA*, and *ERA*, respectively. As we observe, *ERA* performs better in heavily congested environments (in smaller areas). This is because *ERA* implicitly assumes the loss is due to collision at first (by transmitting the short fragment at the current rate). Consequently, it favors the loss differentiation in collision dominated environments.




Fig. 8. throughput under both channel degradation and collision

Fig. 9. throughput of mobile stations

Fig. 10. Network parameters for mobile network simulations

Throughput for Mobile Nodes The throughputs of *ARF*, *CARA*, and *ERA* are respectively evaluated on a mobile network of 17 stations: all 16 client stations are *mobile* with random velocities to random destinations with short pause between two movements in a 350mX350m area. The access point, located at the center of the network, is static. We use topology #4 (See Table 2). The Ricean fading model is used to simulate channel degradation. Table in Figure 10 shows the parameters used in these simulations. All stations keep transmitting to the access point UDP CBR traffic.

The first experiment evaluates the throughputs for all schemes with a CBR traffic with different packet lengths: 500 bytes, 1000 bytes and 1500 bytes. Figure 9 plots the throughputs for *ARF*, *CARA*, and *ERA* for different packet lengths. The horizontal axis represents the packet length used for each run. With longer frames, there are more collisions, therefore *ERA* has more opportunities to deploy its strength.

We also collect the adaptation transient dynamics for *ARF*, *CARA*, and *ERA* for the mobile stations. Figure 11 displays three plots for *ARF*, *CARA*, and *ERA*. Each plot corresponds to a two-second period of dynamics for a randomly chosen mobile client station. The x -axis is the time. The y -axis shows the data rate in Mbps. Each plot presents two curves: the first curve labeled “Channel Rate” (symbol \times) is the rate that the channel can support at that instant (assuming perfect knowledge) and the second curve (symbol $+$) is the rate selected by *ARF*, *CARA*, or *ERA*. We observe that in such

conditions (mixed environment with both collision and channel degradation), 1) the volatile nature of the channel makes it difficult for any rate adaptation scheme to select the optimal supported rate 2) the plot for *ERA* has more points (which means more frame transmissions) than the two others (*ARF* and *CARA*) because *ERA* results in less congestion losses, and 3) *ERA* selects a rate that is closer to the optimal achievable rate. *CARA* outperforms *ARF* mainly due to its ability to reduce collisions with *on* RTS frames.




Fig. 11. Adaptation dynamics

With above extensive simulations on different network topologies and configurations, *ERA* robustly exhibits a dramatic performance improvement. These simulation results illustrate *ERA*'s effectiveness mixed environments (channel degradation and collisions). In order to understand better *ERA*, we evaluated it separately in a strictly channel degradation dominated environment and then in a strictly collisions dominated environment.

6.2 Channel Degradation Dominated Environment

For these experiments, we use topology #1 (See Table 2) with one **static** station at a constant distance from an access point. First, we evaluate the throughput performance throughout the experiments of *ARF*, *CARA*, and *ERA* in stable channel environments. Second, we measure the throughput for these three schemes in an environment with dynamic channel conditions simulated with the more general Ricean fading model. Finally, we measure the impact of packet length on the throughput of the three schemes for a station at 300m from an access point. Due to the space limitation, we only present the result gathered in the stable environment and leave the result of the other two scenarios in our technical report [17].

Stable Channel Conditions In general, a user sits at a desk and has a long work session where a given constant rate will be most often supported. We somewhat create such environment with only one fixed station and one access point. The two-ray ground fading model is used to ensure that a given constant rate is supported: the supported rate under a two-ray ground fading model depends only on distance. Figure 12 plots on the y -axis the throughput achieved by *ARF*, *CARA*, and *ERA* respectively when the client is at a constant distance (x -axis) from the access point. *ERA* outperforms the other schemes due to 1) the effectiveness of its recovery strategy and 2) the mechanism inherited from *AARF*: adjusting the threshold number (*Threshold*) of consecutive successful transmissions required to increase the rate.




Fig. 12. throughput under stable channel conditions




Fig. 13. throughput under different levels of collision




Fig. 14. Impact of Packet length on throughput improvement

6.3 Collisions Dominated Environment

For this environment, all simulations are carried out in a network with 17 *static* stations corresponding to topology #3 (See Table 2): one access point and 16 client stations. The objective is to evaluate the performance of *ERA* and compare it to *ARF* and *CARA* performances, respectively, in a collision dominated environment. We use the two-ray ground fading model such that the best supported rate is constant (stable channel conditions): with the two-ray ground fading model, the quality of the channel is determined only by the distance. All 16 client stations are uniformly placed in a rectangular area as shown in Figure 5. The congestion level in the network is varied by adjusting the size of the rectangle. All stations keep sending 10 Mbps CBR traffic to the access point at the center of the rectangular area. We measure the throughput and evaluate the impact of the packet length on the performance.

Impact of the Congestion Level Figure 13 plots on the y -axis the throughputs of *ARF*, *CARA*, and *ERA*, respectively. The x -axis the length of the rectangle. The packet size is set to 1000 bytes.

ERA performs best again, followed by *CARA*, and then *ARF*. Especially, as the collision increases (smaller rectangular area), more improvement is achieved by *ERA* due to a more accurate loss differentiation and to the judicious recovery strategy.

Impact of Packet Length on ERA For each simulation run, each station sends CBR traffic with packets of length 100 bytes, 500 bytes, 1000 bytes, or 1500 bytes. All stations send packets of the same size. Figure 14 plots the percentage throughput improvement ($100 * \frac{ERA-CARA}{CARA}$). When the frame size increases, the improvement increases because collisions are more likely with longer frames. The improvement is more significant for smaller rectangles because there are more collisions and *ERA* has more opportunity to take advantage of its superiority.

7 Conclusions and Future Work

We propose a new effective rate adaptation scheme that diagnoses quite accurately the cause of a loss, appropriately recover it, and judiciously adjust the rate without requiring RTS/CTS control frames. Extensive simulations with *ns-2* illustrate that *ERA* performs well under channel degradation and/or collisions dominated environments. *ERA* outperforms *ARF* (a first generation scheme) and *CARA*, a recent second generation scheme

which uses RTS control frames. The strength of *ERA* is based on 1) the fragmentation of a lost frame into a very short frame followed by the remainder, 2) a refined technique to diagnose the cause of a frame loss, and 3) an adaptive threshold to increase rate inherited from *AARF*.

The fragmentation of the lost frame yields a faster recovery in channel degradation and/or collisions dominated environments. The overhead of fragmentation is well lower than using RTS/CTS control frames. We are currently implementing and evaluating *ERA* on Linux operated laptops. The performance of *ERA* implementation will be compared to the performance of most recent rate adaptation schemes.

8 Acknowledgement

This work is funded by the National Science Foundation through grants NeTS CNS#0435320 and CRCD/EI CNS#0417565, and by the Vodafone Foundation through a fellowship awarded to Shaoen Wu.

References

1. Kamerman, A., Monteban, L.: WaveLAN II: A high-performance wireless LAN for the unlicensed band. Bell Labs Technical Journal (1997) 118–133
2. M. Lacage, M. Manshahi, and T. Turletti: IEEE 802.11 Rate Adaptation: A Practical Approach. In: MSWiM04. (2004) 126 – 134
3. D. Qiao, S. Choi, and K. Shin: Goodput Analysis and Link Adaptation for IEEE 802.11a Wireless LANs. IEEE TRANSACTIONS ON MOBILE COMPUTING **1** (2002) 278–292
4. Pavon, J., Choi, S.: Link Adaptation Strategy for IEEE 802.11 WLAN via Received Signal Strength Measurement. In: ICC. (2003) 1108–1123
5. Holland, G., Vaidya, N., Bahl, P.: A rate-adaptive MAC protocol for multi-hop wireless networks. In: ACM MOBICOM’01. (2001)
6. B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly: Opportunistic Media Access for Multirate Ad Hoc Networks. In: MOBICOM02. (2002) 24 – 35
7. J. Kim, S. Kim, S. Choi and D. Qiao: CARA: Collision-Aware Rate Adaptation for IEEE 802.11 WLANs. In: IEEE INFOCOM’06, Barcelona, Spain (2006)
8. S. Wong, H. Yang, S. Lu and V. Bharghavan: Robust Rate Adaptation for 802.11 Wireless Networks. In: MobiCom’06, Angeles, California, USA. (2006) 146 – 157
9. Shaoen Wu and Saad Biaz: Loss Differentiated Rate Adaptation in Wireless Networks. Technical Report CSSE07-02, Auburn University (2007)
10. Bicket, J.C.: Bit-rate Selection in Wireless Networks. Master’s thesis, Massachusetts Institute of Technology (2005)
11. D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris: Link-Level Measurements from an 802.11b Mesh Networks. In: ACM SIGCOMM. (2004)
12. IEEE802.11: <http://standards.ieee.org/getieee802/download/802.11-1999.pdf> (1999)
13. Garg, S., Kappes, M.: An Experimental Study of Throughput for UDP and VoIP Traffic in IEEE 802.11b Networks. In: WCNC2003. Volume 3. (2003) 1748– 1753
14. IEEE802.11b: <http://standards.ieee.org/getieee802/download/802.11b-1999.pdf> (1999)
15. IEEE802.11g: (<http://standards.ieee.org/getieee802/download/802.11g-2003.pdf>)
16. Cisco Systems: (http://www.cisco.com/en/US/products/hw/wireless/ps4555/products_data_sheet09186a00801ebc29.html)
17. Shaoen Wu and Saad Biaz: ERA: Efficient Rate Adaption Algorithm with Fragmentation. Technical Report CSSE07-04, Auburn University (2007)