

Scalability Analysis of Rate Adaptation Techniques in Congested IEEE 802.11 Networks: An ORBIT Testbed Comparative Study

Kishore Ramachandran, Haris Kremono, Marco Gruteser, Predrag Spasojević and Ivan Šeškar
WINLAB, Rutgers University
Rt 1 Tech Center, North Brunswick, NJ 08902-3390, USA
Email: {kishore, harisk, gruteser, spasojev, seskar}@winlab.rutgers.edu

Abstract

Recent real-world measurements in dense congested radio environments have pointed out the inefficiency of frame error-based bit-rate adaptation mechanisms, which significantly reduce network capacity by misinterpreting frame errors due to collisions. These effects are likely to be amplified with the heavy use of media applications. Fortunately, traditional SNR-based rate adaptation, and the more recently proposed throughput-based, and collision-aware rate adaptation algorithms are expected to provide more robust performance in these scenarios. To our knowledge, however, their performance has never been experimentally validated in a congested environment.

In this paper, we report our implementation experiences with rate adaptation in a dense, congested IEEE 802.11 network. We find that throughput-based adaptation, contrary to expectations, also suffers from poor bitrate selection. Due to an increase in physical layer capture, while using lower bitrates, nodes can increase their individual throughput at the expense of cumulative network throughput. SNR-based rate adaptation performs well in static environments but the lack of sufficient precision in RSSI measurements makes accurate rate selection in dynamic radio environments difficult. The use of RTS/CTS, in the spirit of collision-aware rate adaptation, shows throughput improvements for frame error-based algorithms and, additionally, for throughput-based algorithms as well. However, results are below expectations likely due to RTS/CTS implementation issues on the Atheros 5212 platform.

1 Introduction

The rapid proliferation of IEEE 802.11 devices is leading to ever denser wireless LAN deployments. Empirical observations from such dense, congested deployments have shown that inefficiencies of traditional frame-error based

rate adaptation algorithms, which seek to vary bitrate to match link conditions, lead to throughput well below the theoretical MAC capacity [1, 2]. These algorithms misinterpret a short-term accumulation of frame errors due to collision as a reduction in the signal-to-noise ratio (SNR).

Emerging multimedia applications such as voice- and video-over-wifi are expected to increase the challenges faced by IEEE 802.11 rate adaptation. In current WLAN usage, these effects are mitigated by the heavy use of flow-control and the dominance of downlink traffic, both of which reduce overall collision rates. In contrast, most media applications use UDP-like transport protocols and the emergence of WiFi-enabled phones, cameras and portable video servers points towards an increase in upstream data sources. Moreover, approaches to reduce congestion such as capacity over-provisioning, admission control, or traffic prioritizing are difficult to deploy in (i) meeting places where flash crowds of WLAN users congregate with a diverse set of legacy devices [2]; or, (ii) unplanned urban deployments of access points and clients [3]. These challenges motivate our detailed study of MAC layer scalability.

A number of rate adaption algorithms are expected to be more robust to frame collisions in dense deployments. These include the traditional SNR-based adaptation [4, 5] and the more recent proposals of throughput-based rate adaptation [6] and collision-aware rate adaptation (CARA) [7] (which relies on the judicious use of RTS/CTS to prevent DATA frame collisions). To our knowledge, however, none of these algorithms have been experimentally studied in a dense congested environment.

This work seeks to fill this void. We present a comparative analysis of the behavior of SNR-based, throughput-based, and RTS/CTS-based algorithms in a congested 802.11 infrastructure network scenario. To allow in-depth analysis of bitrate choices and direct comparisons between the result from different algorithms, we conducted the experiments on the ORBIT testbed [8], which provides a controlled and repeatable radio environment. We do not claim that the results will be representative of dynamic channel

environments, but aim to contribute a case study to better understand algorithm resiliency in collision-dominated environments. Specifically, key contributions include:

- Validating simulation results of frame error-based rate adaptation algorithms using RTS/CTS, in the spirit of CARA, through experiments on real 802.11 systems. We also show that throughput-based algorithms benefit from the use of RTS/CTS.
- Analyzing the effect of rate-diversity and physical layer capture on throughput-based rate adaptation and the associated fairness implications. In practice, some packets may be decoded even with interference from simultaneous transmissions. We show that rate choices influence capture probabilities.
- Determining whether RSSI measurements are affected by frame collisions, subsequently influencing rate choices. We show that no such effect exists.

The remainder of this paper is organized as follows. The next section explains different approaches to rate adaptation in 802.11 systems and how they might be affected by collisions. Section 3 details the controlled experimental environment used for analyzing and comparing 802.11 system performance. Section 4 describes the experimental results and compares the performance of different rate adaptation algorithms. These results are analyzed in section 5, which discusses practical considerations for improving rate adaptation performance. We conclude in section 6.

2 Adapting Bitrate to Changing Channel Environments

In the following subsections, we first review the bit-rate adaptation problem followed by a discussion on collisions and their effect on rate adaptation.

2.1 Rate Adaptation

Rate adaptation enables IEEE 802.11 radios to cope with time-varying channel environments. The IEEE 802.11 standard mandates twelve bitrates between 1 and 54Mbps. Generally, higher bitrates correspond to higher nominal throughput but require higher signal-to-noise ratios (SNR) for correct demodulation. In an SNR-limited environment, higher bitrates will suffer from frame errors, limiting the effective goodput. In such an environment lower bitrates may provide higher effective goodput than high rates. Rate adaptation aims to dynamically adjust the transmission rate to maximize goodput depending on channel conditions. Since the standard does not specify any particular rate adaptation mechanism, manufacturers use different proprietary

implementations. Published rate adaptation mechanisms can be classified into frame error-based, throughput-based and SNR-based adaptation.

Frame-error based adaptation: *Auto Rate Fallback (ARF)* [9], developed for WaveLAN-II 802.11 cards, and *Adaptive Auto Rate Fallback (AARF)* [10] use fixed and dynamic frame error thresholds to increase/decrease the bit-rate. *ONOE* [11], a frame-error based algorithm used in the MADWIFI driver for Atheros-based wireless NICs aims at selecting the highest bit-rate with less than 50% frame loss rate. Periodically, for each destination station, the algorithm maintains a credit score that it increments if less than 10% of packets required a retransmission and no packets were dropped in the last time period. If the credit score surpasses a threshold (default 10), the bit-rate is raised. If each data packet required at least one retransmission, the bit-rate is lowered and the credit score is reset to zero. The current implementation also uses Atheros' multi-rate retry feature, which allows algorithms to select different rates for retransmissions of frames. *Adaptive Multi-Rate Retry (AMRR)* [10], a modification of ONOE, adaptively raises the threshold for rate increases to prevent frequent attempts at bit rates higher than the optimal one in an SNR-limited channel. The recently proposed *Robust Rate Adaptation Algorithm* [12] builds on top of ARF by using a combination of short-term loss estimation and selective use of RTS/CTS. However, the authors themselves acknowledge the potential degradation in performance when the number of stations in the network increases, due to a lack of samples used to infer the channel quality.

Throughput-based adaptation: The *SampleRate* [6] algorithm selects the rate that minimizes mean packet transmission time. Initially, the lossless packet transmission times are calculated for each bit rate and an initial rate is chosen (36Mbps). Hereafter, for each successfully sent packet, the transmission time is updated (using an exponentially weighted moving average (EWMA)) based on the number of retransmissions, packet length and protocol timing overheads. The algorithm also periodically attempts transmission at bitrates whose lossless transmission time is lower than the measured time on the current rate. If these sample transmissions indeed show lower mean transmission time, the algorithm switches the rate.

SNR-based rate adaptation: Since the frame-error rate on a collision-free channel is determined by the receiver's SNR, these algorithms measure channel SNR and select the appropriate rate based on a precomputed table. In general, 802.11 implementations usually only provide the received signal strength indicator (RSSI). This indicator reflects the amount of energy measured on the channel during the reception of the PLCP header.¹ Receiver Based Auto Rate

¹According to the standard, it is measured between the beginning of the start frame delimiter (SFD) and the end of the PLCP header error check

(RBAR) [5] defines a closed loop rate adaptation mechanism through which the receiver can inform the sender of the most suitable rate choice. Specifically, the receiver selects the bit rate based on the RSSI of RTS frames and piggybacks this information on the CTS frame. Pavon and Choi [4] propose a hybrid approach that utilizes the RSSI of acknowledgment frames to choose the bit rate. This algorithm attempts to address asymmetric channels through recalibration of the SNR thresholds for rate choices based on the frame error rate. Another hybrid algorithm proposed in [13] utilizes RSSI to clamp frame-error based bit-rate changes.

The Opportunistic Auto Rate (OAR) protocol [14], which can be layered on top of any of the above rate adaptation mechanisms can optimize individual, as well as network throughput, by sending multiple back-to-back frames under favorable channel conditions.

2.2 Collisions and Their Effect on Rate Adaptation

A collision occurs if two simultaneously transmitted frames interfere at the receiver, so that frames are lost. To avoid collisions, the IEEE 802.11 distributed coordination function (DCF) employs a CSMA/CA mechanism [15]. Carrier sensing prevents transmissions that start while another transmission is in progress. A random backoff mechanism is used to reduce the probability of two stations simultaneously starting a transmission. Specifically, on detecting the wireless medium to be idle for a DCF interframe space (DIFS) duration, each station initializes a counter to a random number selected uniformly from the interval $[0, CW(\text{retransmission})-1]$ and starts counting down.

Time is slotted and the countdown halts when the medium becomes busy, resuming only after the medium is idle again for a period DIFS. Given its half duplex nature, 802.11 transmitters require the receiver to send an acknowledgment (ACK) after a short interframe space (SIFS) duration. The absence of an ACK is interpreted as a collision, following which, CW is doubled (uptil a maximum value of CW_{max}) and the process repeated. CW is reset to its minimum value, CW_{min} (16 for 802.11a and 32 for 802.11b) after successful transmissions, as well as when the maximum retry limit is reached. Note that, for a CSMA/CA MAC, simultaneous transmission of frames can occur either because the two senders (a) happen to select the same time slot for transmission or (b) cannot hear each other's transmissions (hidden terminals).

In the absence of hidden terminals, the probability that two nodes select the same time slot increases with the number of stations and the load on the network [16], since nodes reset their contention window to CW_{min} after every

(HEC)



Figure 1. Preliminary ORBIT testbed setup

successful transmission. This increase in collision-based packet errors leads the auto rate fallback (ARF) algorithm to unnecessarily decrease bitrates as observed in [17], even though the interference from collisions is usually strong enough to prevent decoding even at the lowest rate. Moreover, transmissions at lower rate consume more time, decreasing the overall network throughput [18]. This anomaly occurs due to bit-rate diversity — hosts using lower bit-rates limit the throughput of hosts using higher bit-rates. These shortcomings are addressed in the design of the Collision-Aware Rate Adaptation Algorithm [7] based on ARF. This algorithm [7] uses RTS packets to probe the state of the channel—the loss of an RTS frame is interpreted as a collision loss rather than being due to low SNR. To avoid the overhead, the sender invokes the RTS/CTS exchange only after a DATA frame transmission failure at the current bit-rate. If an ACK is not received (after an RTS/CTS exchange), the algorithm interprets this event as being due to poor channel conditions and drops the bit rate.²

For throughput-based algorithms, one might expect that the collision probability remains independent of the rate choice and that the collisions should cancel each other out when comparing different rates. Based on this assumption, these algorithms should be resilient to congestion. SNR-based algorithms are expected to perform optimally in congested environments. However, it is unclear whether the RSSI provided by a majority of 802.11 implementations reflects the SNR or the Signal-to-Interference-and-Noise ratio (SINR).

Overall, the collision resiliency of many of the above-mentioned rate adaptation techniques (apart from ARF) remains, to our knowledge, experimentally unexplored. This motivates our experimental study of these algorithms in a controlled, high-density setting.

²The authors also present an enhancement using the Clear Channel Assessment (CCA) feature of 802.11 – if the channel is not idle immediately after the reception of a DATA frame for SIFS period, it is interpreted as a frame collision. However, this functionality will be difficult to implement without requiring firmware changes for current wireless NICs.

3 Experimental Methodology

In this section, we discuss key aspects of our experimental platform, setup and methodology.

Hardware and Software: Our study is based on systematic experiments on a preliminary version of the ORBIT indoor testbed [8]. This testbed consists of 64 nodes (standard Linux PCs), each of which is equipped with two wireless 802.11a/b/g interfaces. Half of these nodes use the Atheros 5212 chipset-based wireless NICs and the remaining use Intel 2915 chipset-based wireless NICs. The nodes are placed in a two-dimensional rectangular grid separated by 1-meter distance (see Figure 1) and the antennas are mounted on the sides of the benches in 135 and 215 degree positions (viewed from the top). The testbed nodes run Linux and we utilize only those nodes with Atheros chipset-based wireless NICs. Several features of this platform facilitate our research on rate adaptation. First, since the rate adaptation mechanisms are implemented in the open-source MadWiFi driver [11] (rather than in the firmware), we can develop new algorithms and modify existing ones. Second, by offloading most of the MAC protocol processing to the node’s CPU, these cards are more open to protocol modifications. MadWiFi allows for the configuration of a number of MAC parameters, including the transmission rate, on a per-frame basis. Third, the platform provides a controlled and repeatable environment, where surrounding objects are stationary. Shielding in the walls of the room, housing this testbed, limit the effect that outside interference could have on experimental results. Also, we are not aware of any other adjacent 802.11 networks operating in the 5GHz band (confirmed using an additional sniffer to ensure that no background traffic exists on the channel in question).

In addition, we have instrumented the MadWiFi driver to report both successful and failed transmissions at the sender, as well as successful frame reception at the receiver (every 100ms). Given a constant packet size this allows for goodput calculations. The driver was also modified to report the source MAC address, RSSI, bit-rate and hardware timestamp (microsecond resolution) for each successfully received frame.

Experiment setup: We focus on an infrastructure-based 802.11a system in which, all nodes are within communication range of each other emulating future very high density

Parameter	Default Setting
Mode	802.11a
Channel	36
Transmit Power	18 dbm
Packet Size	1350 bytes

Table 1. Default configuration parameters

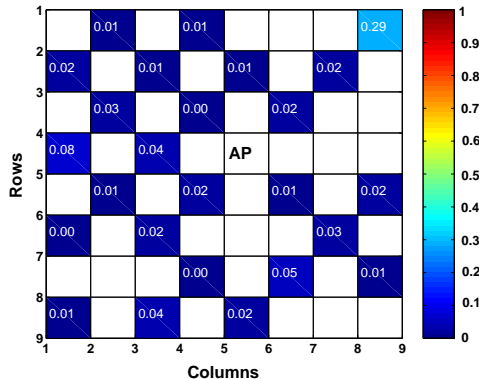
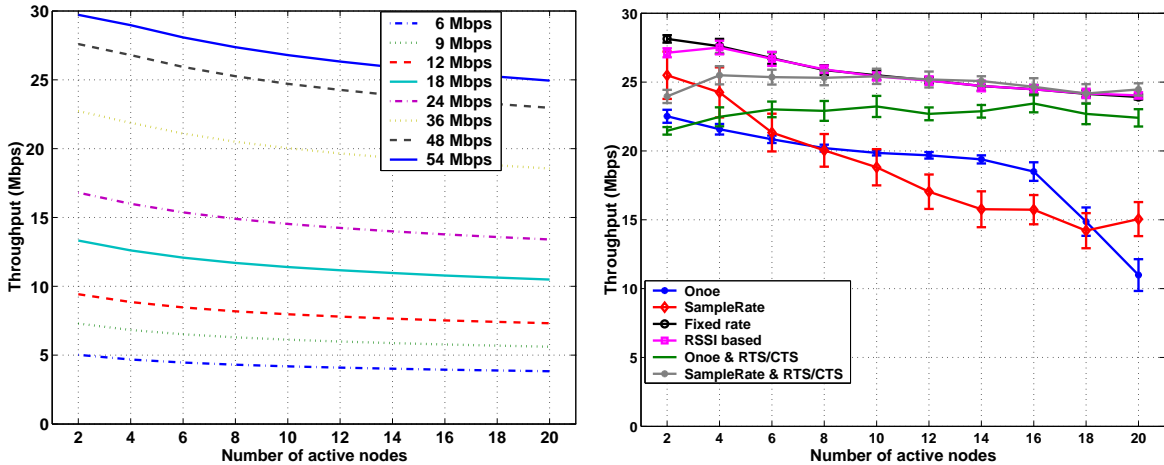


Figure 2. Single link (no collision) mean packet error rate from each sender node to the AP at 54 Mbps (across five six-minute experimental runs).

deployments. We vary the number of clients from 2 up to 20 (we could not use the remaining 11 Atheros-based nodes due to hardware issues). To characterize the radio links, we rely on single-link received signal strength indicator (RSSI) values and packet error rate (PER). RSSI is an estimate of the signal energy at the receiver and is reported by all commodity wireless NICs on proprietary scales (Atheros cards report RSSI in dB relative to the noise floor). In our setup, RSSI measurements serve to approximate “true” SNR values, which would require a calibrated comparison with an accurate RF measurement device. In essence, they represent SNR as measured by actual 802.11 radios and although our results may not apply to future radios with improved measurement accuracy, we believe that our findings have significant implications with regard to practical mechanisms which must depend on similar measurements in real deployments. From our single-link experiments, we observe that link RSSI values (not shown here due to space constraints) range between approximately 30 and 60, which translates to an SNR of -65dBm to -35dBm (assuming constant noise floor of -95dBm), indicating good to excellent connectivity [19]. Figure 2 confirms that all links support the highest bitrate (54Mbps) with near-zero packet error rate (in the absence of contention).³

Note that, although our experiments use a single AP, we believe that our results serve to highlight the significant issues, related to rate adaptation, in congested environments. We believe that the same issues will assume significance in networks consisting of multiple APs on the same channel, albeit with fewer clients per AP (for e.g. home wireless networks with 3-4 clients per AP).

³Note that we do not use the nodes with poor PER (due to defective NICs).



(a) Analytical saturation throughput for fixed bit-rate as a function of the number of active nodes. (b) Experimental saturation throughput for the different rate adaptation algorithms vs a fixed 54Mbps rate approach.

Figure 3. Comparison of theoretical and empirical throughput.

Methodology: In our experiments, clients generate constant bit-rate UDP traffic (using the ORBIT traffic generator [20]) to emulate streaming media applications. Further, recent IETF measurement studies [2], which show that highly congested environments represent realistic scenarios, motivate our study of these algorithms under network saturation. Other advantages of this approach are that it enables comparisons with prior theoretical work [16] and provides an estimation of the worst-case performance.

We carry out multiple runs of each experiment, and the results presented are the average over all runs. We empirically choose the experiment durations so as to provide low variance in results. We also vary the packet size in our experiments. Due to space constraints, we only present the results with 1350-byte packets in this paper. Unless otherwise specified, for the results presented in this paper, the default configuration parameters are specified in table 1.

We evaluate and compare ONOE [11], SampleRate [6], and Pavon and Choi’s algorithm [4] as representatives of the packet-error-, throughput-, and SNR-based categories. We attempt to measure the performance of CARA [7] by approximating its behavior using a combination of ONOE, which is similar to ARF [9], with RTS/CTS enabled for all frames. We also report results for SampleRate with RTS/CTS enabled, since this configuration can serve as an indicator of how throughput-based approaches, in conjunction with RTS/CTS, will perform. Note that we use vanilla versions of SampleRate and ONOE and implement Pavon and Choi’s RSSI-based algorithm. We believe that these selected algorithms provide a good sample of representative designs in literature.


4 Scalability in Dense Congested Environments

Our experimental analysis begins by comparing the rate adaptation algorithms by how well they maximize cumulative throughput under congestion. Note that we concentrate on the congestion case, since adaptation to changing SNR has been studied elsewhere [10].


4.1 Cumulative Throughput

Figure 3(b) shows an experimental comparison of the rate adaptation algorithms in the access point scenario, with a fixed rate of 54 Mbps (i.e., deactivated rate adaptation). For reference, Figure 3(a) also depicts analytical saturation throughput curves for the same scenario, obtained using Bianchi’s model [21] with 802.11a parameters. These curves assume fixed (no adaptation) PHY rates and predict a graceful degradation in cumulative throughput. The analytical results show about 14% reduction in throughput when the number of transmitting nodes increases from 2 to 20. The experimental results for fixed rate (deactivated rate adaptation) closely track this performance. While deactivated rate adaptation cannot represent a useful approach in general, it illustrates that the basic MAC protocol scales as expected.


With both ONOE and SampleRate, the cumulative throughput drops with an increase in the number of transmitters. As the number of transmitting nodes increases from 2 to 20 nodes, the throughput falls by more than half, compared to a drop of less than 15% corresponding to the single bitrate analytical results. SampleRate performs slightly bet-



(a) Individual flow throughputs for ONOE (b) Individual flow throughputs for SampleRate



(c) Individual flow throughputs for fixed rate (54Mbps, no adaptation) (d) Individual flow throughputs for ONOE with RTS/CTS



(e) Individual flow throughputs for SampleRate with RTS/CTS

Figure 4. Throughput fairness characteristics of rate adaptation algorithms from an experimental run. We observed similar trends across multiple runs of this experiment.

ter than ONOE when the network size approaches 20 nodes, whereas, ONOE maintains a higher throughput when the number of users is between 6 to 16 nodes. Given that the average frame transmission time would be minimal at the highest bit-rate, even in congested environments, we would expect SampleRate to perform much better.

RSSI-based rate adaptation appears resistant to collisions and shows excellent performance, in terms of cumulative throughput.

Note that ONOE shows significant throughput improvement when RTS/CTS is enabled. Additionally, RTS/CTS benefits SampleRate as well, with cumulative throughput approaching the performance of fixed rate. For both algorithms, the throughput improvements can be explained, in part, by the smaller time spent in collisions through the use of RTS (which is smaller than the frame header for 802.11 DATA frames), even though it is sent at the lowest bit-rate.




Figure 5. Mean packet error rate (PER) from each sending node to the AP when the network is saturated (across five six-minute experimental runs).

More significantly, this implies that the reduced time spent in collisions outweighs the overhead of using the RTS/CTS exchange in such environments.

In summary, we observed improved throughput for RSSI-based adaptation and through the use of RTS/CTS. We also notice lower than expected throughput for SampleRate. Since throughput gains can be easily achieved at the expense of fairness, let us now look at the throughput fairness characteristics of these algorithms.

4.2 Fairness

Table 2 reports the mean and std. dev. in Jain’s fairness index (JFI) [22]⁴ for the RTS/CTS-based approaches with

⁴The index, F, is calculated as $F = \frac{(\sum_i x_i)^2}{n \times \sum_i x_i^2}$ where x_i is the individual flow throughput and n is the total number of flows. An index value equal to one is considered to be perfectly fair.

Table 2. Fairness comparison for the 20-sender case. Mean and Std. Dev. in JFI across 5 runs is reported.

Rate adaptation scheme	Avg. JFI	Std. Dev. in JFI
ONOE	0.822	0.032
SampleRate	0.819	0.024
Fixed Rate (54Mbps)	0.917	0.027
ONOE w/ RTS/CTS	0.491	0.020
SampleRate w/ RTS/CTS	0.709	0.034

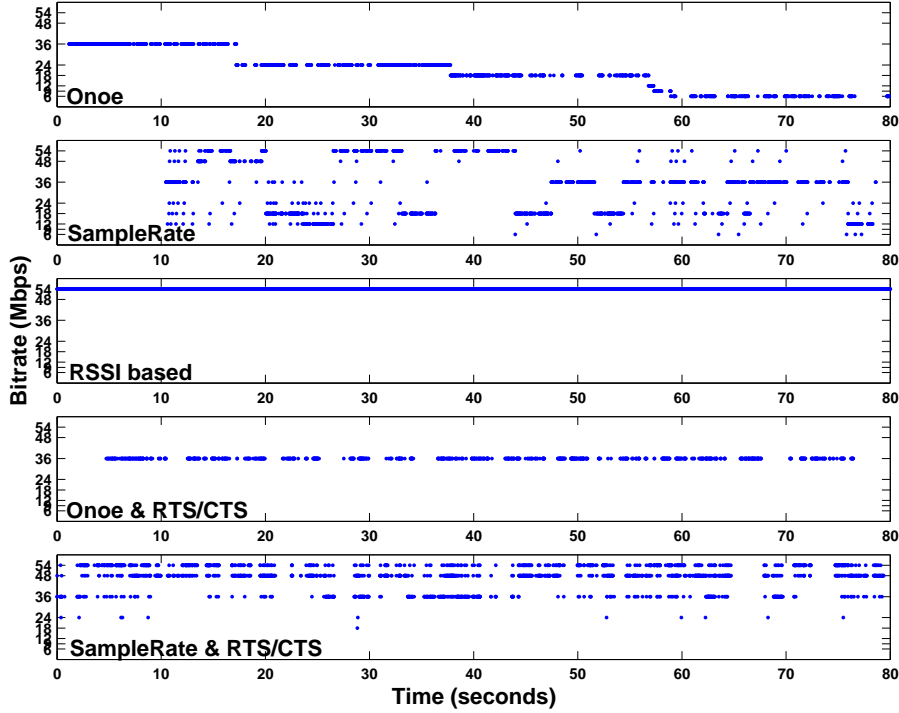


Figure 6. Rate vs time for rate adaptation algorithms from an experimental run. Transmission rate for both SampleRate and ONOE decreases with the addition of nodes with SampleRate showing higher variance.

that for ONOE, SampleRate, and fixed rate for the standard 20 sender experiment. ONOE with RTS/CTS stands out with the lowest fairness index. To analyze fairness in more detail, figures 4(a) through 4(e) compare the throughput distribution across senders for each of the algorithms. Fixed rate shows slight imbalances that, as we will see, is due to the physical layer capture (PLC) effect.

The presence of PLC is illustrated in the PER imbalances observed in a 26 node setup in Figure 5 (we also confirmed PLC by looking at packet traces from multiple sniffers). PER for each link in saturation ranges from approximately 50% to 90%. Since all other parameters in this experiment were the same as that for Figure 2, we can attribute these PERs solely to collisions. Note that the PER of nodes with lower RSSI (relative to the stronger sender) at the access point tends to be higher, a typical result under PLC.


SampleRate and ONOE both show more pronounced throughput variations, most likely because rate diversity increases the probability of capture as shown in the previous subsection. SampleRate with RTS/CTS shows significant throughput imbalances also reflected by its relatively low Jain fairness index. ONOE with RTS/CTS clearly shows the largest throughput imbalances. A closer inspection of the packet error traces in the experiments involving ONOE

with RTS/CTS reveals that for a majority of senders, PER was higher than 10% (ONOE’s threshold for rate increase) and below 50% (threshold for rate decrease). This indicates that significant DATA frame losses occur even though the channel is reserved and links are not SNR-limited. To investigate this anomaly, we carried out single-link measurements, with and without RTS, and observed that the PER, unexpectedly, shows a relative increase of approx. 4% when RTS/CTS is enabled. We speculate that this indicates an incorrect implementation of this mechanism on Atheros 5212 NICs (also discovered by [23] on different hardware) and we hypothesize that the throughput gains will be higher with more accurate implementations. In addition, the observed fairness reductions may be also caused by this issue. We plan to investigate this issue further as part of our future work.

The surprisingly low throughput of SampleRate motivates us to look more closely at the bit-rate choices of the individual algorithms.

4.3 Bit-rate choices

Figure 6 depicts transmission rate changes over time for one of the senders. Note that in this 80 second experi-



(a) Estimated average transmission time for two (b) Individual goodput vs network goodput for (c) Number of transmission attempts for different senders - the “strong” sender is able to capture different rate choices of one node, under PLC. nodes depending on the rate choice of one node. transmissions from “weak” senders.

Figure 7. Performance in the presence of Physical Layer Capture (PLC).

ment, all 20 senders start simultaneously. The bit-rate (obtained from the PLCP header) for each received packet is logged at the access point—each dot in the plot represents one packet. The results show that ONOE starts off at its default initial rate of 36Mbps and steadily decreases the bitrate until it reaches 6Mbps. SampleRate shows a similar trend even though it is not as pronounced because the algorithm tends to change rates more frequently. For example, around 30s into the experiment, transmissions occur at nearly all the rates between 12 and 54Mbps. This confirms, however, that poor bitrate decisions are also the cause for the reduced throughput obtained with SampleRate.

In comparison, RSSI-based adaptation shows nearly perfect rate choices, with all nodes choosing high bitrates. This indicates that RSSI-based adaptation is not affected by collisions. ONOE with RTS/CTS remains at 36Mbps for the entire experiment. This is contrary to what we expect – the use of higher bit-rates by transmitters, provided only DATA frame losses are taken into account (while calculating PER). However, as mentioned before, we do see DATA frame losses, even when RTS is turned ON. SampleRate with RTS/CTS shows bit-rate fluctuations, but critically, we can see that it is more prone to select the higher bit-rates (48Mbps and 54Mbps). We believe that more accurate bit-rate choices are the primary reason for the difference in cumulative throughput gains between ONOE and SampleRate, when RTS/CTS is enabled. The adaptation stability of both algorithms is significantly improved with RTS/CTS, nevertheless, SampleRate still lacks stability due to a reduction in the number of measurement samples.

We highlight some of the more interesting performance details characterizing the selected rate adaptation algorithms in the sections that follow.

4.4 SampleRate with, and without, RTS/CTS

SampleRate’s rate decisions compare the expected transmission time of different rates relative to each other. One might expect, based on arguments offered in section 2.2, that this algorithm is resilient in high collision environments. Surprisingly, SampleRate’s performance degrades with increasing node density, similar to ONOE. We identified two reasons:

1. In highly congested environments, few samples (packets) per node are available to accurately estimate the transmission time.
2. Due to the PLC effect [24], some nodes can decrease their collision probability by decreasing their rate while maximizing their individual throughput.

The frequent rate changes observed in Fig. 6 support that the algorithm bases its rate choice on too few samples. To confirm that nodes can maximize throughput by lowering rates even on the high SNR channels in our setup, Figure 7(b) shows the goodput for different rate choices for a station that is closer to the access point as compared to the other 10 competing stations. The competing stations use a fixed rate of 54 Mbps and the average RSSI observed by the access point for the close and far stations are 60 and 47, respectively. Evidently, the closer station obtains maximum throughput while operating between 18Mbps and 36Mbps. We can explain this result with a capture probability that depends on the bit rate choice of the stronger sender—capture becomes more likely when the stronger sender reduces its rate. Each capturing transmission in turn causes the capturing node to reset its contention window to CW_{min} while the other colliding nodes double their current contention windows. This leads to a larger number of transmission opportunities for the capturing node, at the expense of reduced

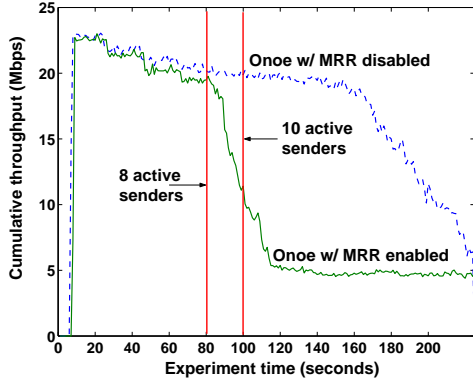


Figure 8. Cumulative throughput with and without multi-rate retry (MRR) for ONOE. Behavior is similar for other algorithms.

transmission opportunities for the other nodes, as shown in Figure 7(c). In effect, the station sacrifices total network throughput for a small gain in individual throughput. While Tan and Guttag [25] have noted the existence of such inefficient equilibria through simulations of mobile nodes under a Rayleigh fast fading model, these results show that these inefficiencies also exist in relatively stationary multiple-client single-access point scenarios.

To further validate this hypothesis, Figure 7(a) translates the observed packet error rates into the expected packet transmission time, on which SampleRate bases its rate decision⁵. Indeed, the minimal expected transmission time for the stronger sender occurs at 36Mbps. This explains why stations choose lower rates. Moreover, the differences in expected packet transmission times between the rates of 18Mbps to 54Mbps are small. This explains the oscillatory behavior in Figure 6.

As mentioned earlier, SampleRate with RTS/CTS selects higher bit-rates with greater frequency and shows corresponding increase in cumulative throughput (relative to when RTS is disabled).

4.5 ONOE with, and without, RTS/CTS

Auto rate fallback is known to lead to degraded performance with less than 10 senders. Since vanilla ONOE also bases its rate decisions on packet errors, one might expect similar performance. Instead, the cumulative throughput with ONOE remains more stable until a significant reduction occurs with 18 active senders. The exact number of senders tolerated is, however, very sensitive to the detailed algorithm configuration. When the multi-rate retry (MRR)

⁵The theoretical curve derivation is outlined in 7.

feature in Atheros cards is enabled, throughput collapse occurs with just 10 senders as depicted in Figure 8. Since this mechanism is configured to pick lower rates for retransmissions, we hypothesize that the pathological effects of collision on packet-error-based adaptation are amplified by MRR. All ONOE results in this paper were obtained with MRR disabled.

Results from ONOE with RTS/CTS indicate that the performance of packet-error-based adaptation can be stabilized through channel reservations, as proposed in CARA [7]. Here, ONOE is modified to only consider packets that did not contend with other stations (e.g., the data frame following a CTS), thus avoiding unnecessary rate decreases (due to RTS losses). However, as mentioned before, ONOE does not increase the bit-rate, as would be expected when RTS/CTS is used, in near-perfect channel conditions, likely due to implementation issues.

Now that we have looked at the performance of ONOE, SampleRate and RTS/CTS-based rate adaptation in detail, we proceed to highlight some practical issues with current implementations.

5 Implementation Experiences

In this section, we first discuss precision issues associated with the reporting of SNR in existing wireless NICs. This is followed by a discussion on how the use of RTS/CTS enables the accurate estimation of channel quality.

5.1 RSSI-based rate adaptation

From the comparative evaluation, RSSI-based algorithms proved to be more resistant to collisions in a congested scenario. However, given that the RSSI thresholds for all rates lie in a small interval of the total RSSI measurement range, there is a low margin of error w.r.t. comparison with thresholds to increase (or decrease) the bit-rate. Hence, we expect that these algorithms will fail to perform optimally in SNR-limited environments.

Table 3 lists the RSSI threshold values for which the frame error rate (FER) approaches 1.0 for any of the 802.11a rates. We measured these thresholds by placing an (additive white gaussian noise) AWGN source [8] near the receiver, fixing the sender's bit-rate and steadily increasing noise power until the receiver did not decode any frames. The RSSI values for the frames decoded last were noted as the approximate RSSI thresholds. These thresholds are specific to the Atheros 5212 card because the absolute interpretation of RSSI values is not defined in the standard. However, for convenience, many manufacturers use a similar scale where each step in RSSI signals an increase of approximately one dB in signal strength. According to simulations of the modulation schemes, they cover a range of

20 dB [26]. Thus, we would expect the thresholds to lie within an interval of 20 RSSI values in most implementations. With thresholds compressed into such small intervals, slight measurement errors might have a large effect on RSSI-based rate adaptation. To make RSSI useful as a primary indicator for rate selection, wireless NIC’s should provide more fine-grained RSSI measurement differentiation in the range relevant to bitrate selection. In addition, protocols for exchanging receiver RSSI information must be defined to allow a pure RSSI-based approach, as opposed to Pavon’s hybrid RSSI/frame-error-based algorithm.

5.2 RTS/CTS-based rate adaptation

In congested environments, both SampleRate and ONOE make rate adaptation decisions based on “collision-tainted” measurements obtained during a fixed time interval (ONOE uses frame retry measurements and SampleRate measures average frame transmission time). Measurements are tainted since, an increase in congestion results in a reduction in the number of available samples and simultaneously, these samples are more likely to be “affected” by collisions. For ONOE, the number of frame retry and success samples that are available to infer PER will be lower, thereby affecting the accuracy of PER estimation. SampleRate will also suffer from the same issue — reduction in the number of samples, which in turn, will affect accuracy in estimating the average frame transmission time.

The use of RTS/CTS ensures that the rate adaptation decisions are made solely on measurements not tainted by collisions (provided RTS transmission errors are not taken into account). Both, packet error-based mechanisms (as proposed by CARA [7]) and throughput-based mechanisms infer channel quality on DATA frames. Thus, they exhibit a stable behavior with an increase in the number of nodes. However, in practice, we observed reduced gains due to the

likely implementation issues with RTS/CTS, which results in DATA frame losses even when the channel is reserved and not SNR-limited.

An alternative to RTS/CTS-based collision detection includes the *passive estimation* of PER due to collisions [27]. However, although the proposed PER estimation technique appears promising, it is yet to be experimentally evaluated. Moreover, the proposed technique requires precise information regarding CSMA/CA slot usage at each IEEE 802.11 transmitter and this information, to our knowledge, is not exposed by the majority of existing open-source device drivers. Similarly, the dynamic tuning of MAC contention windows based on achieving a balance between the time spent in collisions and the time spent waiting in idle slots [28] requires a number of changes to the MAC, which may not be possible to implement on existing wireless NICs.

6 Conclusions

We have experimentally analyzed rate adaptation in dense congested IEEE 802.11 networks and highlighted the substantial effect that rate adaptation algorithms have on its performance. Specifically, we conclude that:

- Contrary to expectations, throughput-based adaptation shows a throughput degradation similar to packet-error-based mechanisms. This occurs primarily due to physical layer capture, which we show to be a function of the bit-rate and whose occurrence is highly likely in congested networks. More specifically, we illustrated that some nodes may experience higher throughput with lower bit-rates, even if the channel could sustain a higher bit-rate in the absence of collisions. However, bit-rates which maximize individual throughput may reduce network throughput and fairness. This has implications for current deployments and experimental studies—the throughput-based SampleRate algorithm is currently the default algorithm adopted by the widely used MadWiFi driver (used in both Linux and BSD).
- Similar to Auto Rate Fallback, a default ONOE implementation does not scale gracefully beyond 10 simultaneous senders. Scalability of such packet-error-based algorithms appears very sensitive to parameter choices and exact implementation details. By disabling MRR, which reduces rates more aggressively on retries, the network could accommodate up to 18 simultaneous senders.
- RSSI measurements provided by the Atheros 5212 chipset remained unaffected by interference from packet collisions. Thus, SNR-based rate adaptation performed well in the static high-density scenario.

Table 3. RSSI thresholds for 802.11a PHY bit-rates

Rate (Mbps)	RSSI Threshold
6	9
9	10
12	11
18	12
24	13
36	15
48	19
54	23

However, our experiments indicate that the RSSI readings do not allow for a fine-grained differentiation in the range relevant to bit-rate selection (especially at bit-rates below 36Mbps). We expect that this lack of precision will have serious performance implications for SNR-based approaches in dynamic, real-world scenarios.

- RTS/CTS with throughput-based adaptation appears most promising, providing both relatively acceptable fairness and throughput stability. Similarly, RTS with packet error-based adaptation also appears effective—although it trades throughput fairness against higher cumulative throughput on our hardware platform.

Note that collision-aware algorithms based on reservations using RTS/CTS could be incrementally deployed on new client devices and provide benefits even in congested networks with legacy devices.

This work raises several questions for future investigation. First, these results need to be combined with results from experiments in dynamic channel environments to inform the design of rate adaptation algorithms. Second, the effect of mixing several different rate adaptation schemes remains an open problem. Third, the effect of bit-rate on physical layer capture warrants a more detailed study to improve simulation models.

7 Appendix: SampleRate Expected Transmission Time

The transmission time estimates for different rates in Fig. 7(a) are calculated from the equations given in [6].

$$t_t = \text{DIFS} + B(k) + (k + 1) \left(\text{SIFS} + \text{ACK} + h + \frac{8L}{R} \right) \quad (1)$$

$$B(k) = \frac{2^k CW_{min} - 1}{2} \sigma$$

$B(k)$ is an estimate of the backoff duration, σ is the system time slot, k is the retransmission index, h is the overhead introduced by the packet header, L is the (average) payload length in bytes, R is the physical rate, and CW_{min} is the minimum contention window.

Under assumption of ideal communication channel and constant collision probability p [21], packet transmissions can be modelled as a geometric random variable. By taking the expectation of (1) we obtain

$$E[t_t] = \text{DIFS} + E[B(k)] + (E[k] + 1) \left(\text{SIFS} + \text{ACK} + h + \frac{8L}{R} \right). \quad (2)$$

Let \hat{k} be the order of retransmission after which the contention window does not grow, and it stays fixed at CW_{max} . Also, let k_{max} be the maximum allowed number of retransmissions after which a packet is being dropped. Without loss of generality we can assume $\hat{k} < k_{max} = \infty$.

$$\hat{k} = \log_2 \frac{CW_{max}}{CW_{min}}$$

$$E[B(k)] = \frac{\sigma}{2} \left[CW_{min}(1-p) \left(\sum_{k=0}^{\hat{k}} 2^k p^k + \sum_{k=\hat{k}+1}^{k_{max}} p^k \frac{CW_{max}}{CW_{min}} \right) - 1 \right] \quad (3)$$

$$E[k] = \frac{p}{1-p}$$

If $k_{max} < \infty$, the collision probability derived in [29] should be used instead the one outlined in [21].

8 Acknowledgments

This material is based upon work supported by the National Science Foundation under Grants No. CNS 0335244 and CNS 0627032.

References

- [1] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting," in *Proc. of ACM E-WIND workshop (held with SIGCOMM)*, 2005, pp. 5–10.
- [2] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer, "Understanding congestion in IEEE 802.11b wireless networks," in *Proc. of USENIX Internet Measurement Conference (IMC)*, 2005, pp. 279–292.
- [3] A. Akella, G. Judd, S. Seshan, and P. Steenkiste, "Self-management in chaotic wireless deployments," in *Proc. of ACM MOBICOM*. New York, NY, USA: ACM Press, 2005, pp. 185–199.
- [4] J. Pavon and S. Choi, "Link Adaptation Strategy for IEEE 802.11 WLAN via Received Signal Strength Measurement," in *Proc. of IEEE ICC*, May 2003.
- [5] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *Proc. of ACM MOBICOM*, Rome, Italy, 2001.

- [6] J. C. Bicket, "Bit-rate selection in wireless networks," Master's thesis, M.I.T., Cambridge, MA, February 2005.
- [7] J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs," in *Proc. of IEEE INFOCOM*, Barcelona, Spain, March 2006.
- [8] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, "Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols," *Proc. of IEEE WCNC*, March 2005.
- [9] A. Kamerman and L. Monteban, "WaveLAN-II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Technical Journal*, pp. 118–133, Summer 1997.
- [10] M. Lacage, M. H. Manshaei, and T. Turletti, "IEEE 802.11 rate adaptation: a practical approach," in *Proc. of ACM MSWiM*, 2004, pp. 126–134.
- [11] MadWifi, "Multiband Atheros Driver for WiFi," <http://www.madwifi.org/>.
- [12] S. H. Y. Wong, S. Lu, H. Yang, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *Proc. of ACM MOBICOM*. New York, NY, USA: ACM Press, 2006, pp. 146–157.
- [13] I. Haratcherev, K. Langendoen, R. Lagendijk, and H. Sips, "Hybrid rate control for IEEE 802.11," *ACM International Workshop on Mobility Management and Wireless Access*, October 2004, pp. 10–18.
- [14] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," *Proc. of ACM MOBICOM*, 2002, pp. 24–35.
- [15] IEEE, *Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE-SA Standards Board, 1999.
- [16] G. Bianchi, L. Fratta, and M. Olivieri, "Performance evaluation and enhancement of the CSMA/CA protocol for 802.11 wireless lans," *IEEE PIMRC*, vol. 2, Taipei, Taiwan, October 1996, pp. 392–396.
- [17] S. Choi, K. Park, and C. Kim, "On the performance characteristics of WLANs: Revisited," in *Proc. of SIGMETRICS*, 2005.
- [18] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. of IEEE INFOCOM*, vol. 2, 2003, pp. 836–843.
- [19] MADWiFi, "RSSI in MadWifi," <http://madwifi.org/wiki/UserDocs/RSSI>.
- [20] Z. Wu, "Orbit traffic generator," <http://www.orbit-lab.org/wiki/Documentation/OTG>, WINLAB, Rutgers University.
- [21] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, May 2000.
- [22] R. Jain, D. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," *DEC Tech Report TR-301*, Tech. Rep., September 1984.
- [23] J. Bellardo and S. Savage, "802.11 denial-of-service attacks: Real vulnerabilities and practical solutions," in *Proc. of USENIX Security Symposium*, Washington D.C., 2003.
- [24] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. of IEEE ICNP*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 252–261.
- [25] G. Tan and J. Guttag, "The 802.11 MAC Protocol Leads to Inefficient Equilibria," in *Proc. of IEEE INFOCOM*, Miami, FL, March 2005.
- [26] P. Gopalakrishnan, P. Spasojevic, L. Greenstein, and I. Seskar, "A method for predicting the performance of rate-adaptive WLANs," *WINLAB, Rutgers University, Piscataway, NJ, Tech. Rep. TR-259*, September 2004.
- [27] D. Malone, P. Clifford, and D. J. Leith, "MAC Layer Channel Quality Measurement in 802.11," in *IEEE Communication Letters*, February 2007, vol. 11.
- [28] M. Heusse, F. Rousseau, R. Guillier, and A. Duda, "Idle sense: an optimal access method for high throughput and fairness in rate diverse wireless lans," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 4, pp. 121–132, 2005.
- [29] P. Chatzimizios, A. C. Boucouvalas, and V. Vitsas, "IEEE 802.11 packet delay - A finite retry limit analysis," in *Proc. of IEEE GLOBECOM*, vol. 2, December 2003, pp. 950–954.