Rate Avalanche: The Performance Degradation in Multi-rate 802.11 WLANs

Liqiang Zhang and Yu-Jen Cheng
Dept. of Computer & Information Sciences
Indiana University South Bend
South Bend, IN 46615, USA

Xiaobo Zhou
Dept. of Computer Science
University of Colorado at Colorado Springs
Colorado Springs, CO 80918, USA

Abstract

The Request-to-Send/Clear-to-Send (RTS/CTS) exchange was defined as an optional mechanism in DCF (Distributed Coordination Function) access method in IEEE 802.11 standard to deal with the hidden node problem. However, in most infrastructure-based WLANs, it is turned off with the belief that the benefit it brings might not even be able to pay off the transmission overhead it introduces. While this is often true for networks using fixed transmission rate, our investigation leads to the opposite conclusion when multiple transmission rates are exploited in WLANs. In particular, through extensive simulations using realistic channel propagation and reception models, we found out that in a heavily loaded multi-rate WLAN, a situation that we call rate avalanche *often happens if RTS/CTS is turned off. The rate avalanche effect could significantly degrade the network performance even if no hidden node presents. Our investigation also reveals that, in the absence of effective and practical loss-differentiation mechanisms, simply turning on the RTS/CTS could dramatically enhance the network performance in most cases. Various scenarios/conditions are extensively examined to study their impact on the network performance for RTS/CTS on and off respectively. Our study provides some important insights about using the RTS/CTS exchange in mutlirate 802.11 WLANs.*

1 Introduction

With its great success in the last decade, the prevailing IEEE 802.11 Wireless Local Network (WLAN) technology still attracts significant research efforts today to further enhance its performance. Dynamically adjusting transmission rates according to the time-varying and location-dependent link quality is one example of such efforts.

Many current and proposed wireless networking standards support multiple transmission rates at the PHY layer that use different modulation and coding schemes. For example, the 802.11a PHY supports eight transmission rates (6∼54Mbps) . Recently, a number of rate adaptation schemes have been proposed in the literature [1, 2, 5, 6, 7, 8, 12, 17, 18]. These schemes could be roughly divided into two categories: SINR (Signal to Interference plus Noise Ratio)-based [2, 6, 17] and statistics-based [1, 5, 7, 8, 12, 18]. The former usually leads to a higher performance gain than the later, because channel conditions are measured more accurately and timely in the former thus optimal rates are more likely to be selected to match the channel conditions in a timely manner. However, several factors have made SINR-based approaches hard to implement in practice (refer to section 2 for details). On the other hand, statistics-based rate adaptation algorithms ,such as ARF (Auto Rate Fallback) [7], still remains the most widely deployed rate control schemes in today's 802.11 networks because of their simplicity. Therefore, we limit our efforts on enhancing the present statistics-based schemes and keeping them compatible with currently-deployed 802.11 WLANs, which we believe is of practical importance.

During our study on the performance of rate adaptation mechanisms, we noticed that in a heavily loaded multi-rate WLAN using ARF, a situation that we call *rate avalanche* often happens if RTS/CTS exchange is not used. That is, high collision rates not only lead to retransmissions but also drive the nodes to switch to lower date rates; the retransmissions and the longer channel occupation caused by lower rates will further deteriorate the channel contention, which yields more collisions. This vicious circle could significantly degrade the network performance if no hidden node presents. There are two important reasons behind this phenomenon: (1) Statistics-based rate adaptation schemes, like ARF, are lack of the ability to differentiate frame losses that are caused by collisions and link errors. All the packet losses, even caused by collisions, are counted to reduce the transmission rate; Lower transmission rates deteriorate channel contentions which leads to further rate reducing and more nodes involvement in the cycle. (2) Even if a node is luckily less involved in the rate reducing cycle, its performance is still dragged down to the same level as those transmitting at lower rates, which is a verification to the so-called performance anomaly effect [4]. Interestingly, however, we found that the rate avalanche could be effectively ameliorated through turning on the RTS/CTS mechanism.

The above mentioned discovery has motivated us to thoroughly investigate the effect of RTS/CTS exchange on the performance of multi-rate WLANs. The investigation itself, however, is more challenging than it appears. The complexity of time-varying channel characteristics, node mobility, various traffic patterns, node interactions, and their effects on the behavior of ARF algorithm make it difficult to build a complete analytical model to study the network performance. On the other hand, a real testbed that contains a large number of wireless nodes is expensive. Therefore, we have chosen to use the Network Simulator – ns-2 [13] to conduct the study. However, the current implementation of ns-2, e.g., the recent all-in-one version 2.29, does not satisfy our requirements – while it has a rather complete, detailed, and accurate simulation of the 802.11 MAC protocol, the support to the PHY layer is much less substantial. For a simulation study where link dynamics and frame errors need to be carefully addressed, a concrete and realistic PHY implementation is particularly important. To facilitate the investigation, we have made significant changes in the implementation of 802.11 MAC and PHY. Besides the support to the multi-rate and ARF algorithm, the implementation also includes two important features: (1) it has a realistic channel propagation model that is suitable for mobile nodes deployed in urban area; and (2) it uses a FER (Frame Error Rate) - based reception model that reflects the error performance of 802.11a PHY modulation and coding schemes.

Our study leads to following conclusions about the impact of RTS/CTS exchange on multi-rate 802.11 WLANs:

- Different than in single-rate networks, simply turning off RTS/CTS exchange in heavily-loaded multi-rate 802.11 WLANs could lead to severe degradation of network performance. Particularly, the rate avalanche effect contributes significantly to the performance degradation.

- With all factors considered, keeping RTS/CTS on will have a much higher chance to gain a better network performance than keeping RTS/CTS off.

- If dynamic RTS/CTS exchange is to be employed, simply using a pre-configured RTS threshold will only yield sub-optimal performance. This is because the optimal RTS threshold depends on several factors, such as, the number of competing nodes, the geographic distribution of nodes, and node mobility, etc. All these factors can vary over time in real networks.

The rest of the paper is organized as follows. Section II presents the related work. Some related background issues are presented in Section III. We describe the detailed simulation modeling in section IV. Performance analysis is presented and discussed in section V. Finally, section VI concludes the paper.

2 Related Work

Wireless channel conditions are time-varying and can be influenced by several factors, such as path loss, shadowing, multi-path fading, and interference, etc. The idea of rate adaptation is to adapt the data rate over time according to the fluctuating channel conditions so that optimal throughput is achieved.

2.1 SINR-based Rate Adaptation

Ideally, if both of the following could be achieved: (A) accurately modeling the relation between SINR (Signal to Interference plus Noise Ratio) and the optimal rate for throughput; and (B) senders being able to precisely estimate the SINR that will be observed by receivers at the moment of transmission; then an optimal rate adaptation algorithm could be implemented simply by the means of a lookup table. Unfortunately, in reality, neither (A) nor (B) is easy to achieve. A complete model for the relation between SINR and the optimal rate for throughput not only includes a radio channel model describing the relation among FER, SINR, modulation and coding schemes, but also a MAC model capturing the interactions of multiple contending stations. On the other hand, estimating time-varying channel conditions is a very challenging task, especially when dynamic interferences are considered; not to mention that the SINR observed at the sender side could be significantly different than that at the receiver side. Despite these difficulties, the optimality of such SINR-based approach has attracted significant research efforts, such as RBAR [6], OAR [17], RAF [2], and others [11]. Assumptions or approximations are often used in these schemes to simplify the problem. However, these SINR-based approaches have not been applied in practice so far.

2.2 Statistics-based Rate Adaptation

An alternative to SINR-based approaches is to estimate link conditions through maintaining statistics about the transmitted data like the achieved throughput [1], consecutive transmission successes/losses [7, 12, 8, 5], and short-term lose ratio [18].

ARF [7], the first documented rate adaptation algorithm which was originally designed for Lucent Technologies' WaveLAN-II WLAN devices, belongs to this category. ARF incrementally increases or decreases the rate by keeping the track of acknowledged/unacknowledged transmissions as well as a timing function. A sender will switch its transmission rate to the next lower level if it experiences two consecutive failed (unacknowledged) transmissions, a timer will also be started. When either the timer expires or it sees 10 consecutive successful (acknowledged) transmissions, the sender will raise its rate to the next higher level and the timer is cancelled. However, if the first transmission at the higher rate fails, the rate will be lowered down again instantly and the timer is restarted. Some deficiencies of ARF have been revealed by previous research efforts: it is unable to adapt effectively to fast-changing channel conditions [6]; On the other hand, if the channel conditions do not change at all, or change very slowly, ARF will try to use a higher rate every 10 consecutive transmission successes or after a timer expiration, which results in increased retransmission attempts and thus do harmful to the throughput [12]. However, due to its simplicity, ARF is still the most widely implemented rate adaption scheme in the 802.11 market [8].

3 Preliminaries

3.1 IEEE 802.11 DCF Access Method

The IEEE 802.11 WLAN operates with two modes at MAC layer a mandatory Distributed Coordination Function (DCF) and an optional Point Coordination Function (PCF). DCF is a contention-based protocol, derived from CSMA/CA (Carrier Sense Multiple Access / Collision Avoidance). In DCF, each wireless station has to sense the channel before sending its packet. Only if the channel has been idle for a certain period of time – DCF Inter Frame Space (DIFS), the transmission is allowed. Between any two consecutive transmissions of the same station, there must be a back-off operation with a randomly-picked back-off count value between [0, CW-1] (CW standards for Contention Window). This backoff count decreases by one for each idle time slot. This backoff procedure is frozen every time the station finds the channel busy and is resumed every time the channel has been idle longer than DIFS. The transmission is triggered when the back-off counter reaches zero. Such a CSMA/CA based access method cannot totally avoid collisions. If two or more stations happen to transmit at the same time, a retransmission will be triggered at all involved stations repeating the above process. However the CW will be doubled. The CW is limited by a minimum value (CWmin) and a maximum value (CWmax).

3.2 RTS/CTS Exchange

Besides the basic access mechanism, the 802.11 DCF also includes an optional RTS/CTS based access scheme

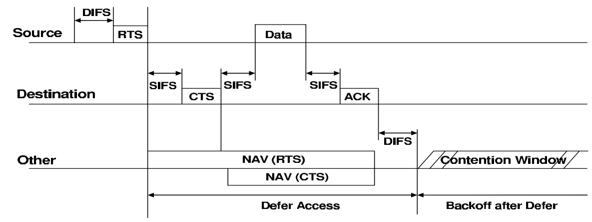

Figure 1. RTS/CTS Exchange Scheme of IEEE 802.11

which was defined to ameliorate the hidden node problem. A hidden node is one that is in the transmission range of the receiver, but out of the range of the sender. Because of their unawareness of each other, the hidden node and the sender unintentionally create a high-collision area in the vicinity of the receiver, which could significantly degrade the network performance. This effect could be effectively mitigated through the using of RTS/CTS exchange which helps to reserve time intervals in both the sender's and the receiver's neighborhood. The RTS/CTS based access scheme is shown in Fig. 1 [9], where NAV (Network Allocation Vector), contained in the MAC header of the RTS/CTS frame to indicate the duration of the transmission, are used to reserve the channel in the vicinity of the sender/receiver. Albeit effective in dealing with hidden node problem, the RTS/CTS exchange itself introduce transmission overhead (although RTS and CTS frames are short, only 20 and 14 octets respectively, they are defined to be transmitted using low rates, i.e. Basic Data Rates [9]). The tranmission overhead of RTS/CTS exchange could be a significant factor to the network performance, especially when data frames are short. With this considered, the RTS/CTS exchange is defined in the 802.11 standard as an optional mechanism, which is turned on only if the length of the data frame is larger than a configurable value called RTS threshold.

In infrastructure-based WLANs, since the hidden node problem is often much less severe than in ad hoc networks, it is a common belief that it might not worth to employ the RTS/CTS exchange due to its transmission overhead. This has led to the fact that the RTS/CTS exchange has been effectively disabled, by setting a large RTS threshold value, e.g., 3000 octets, in most today's 802.11 WLANs.

There is a lack of a thorough study on the effects of RTS/CTS exchange in 802.11 WLANs, especially when multiple transmission rates are enabled and channel conditions are time-varying and non-ideal.

4 Simulation Modeling

To facilitate our investigation, we have made significant changes in the implementation of 802.11 MAC and PHY. Besides the support to the multi-rate and ARF algorithm,

Mode	Modulation	Code Rate	Data Rate
1	BPSK	1/2	6 Mbps
2	BPSK	3/4	9 Mbps
3	QPSK	1/2	12 Mbps
4	QPSK	3/4	18 Mbps
5	16-QAM	1/2	24 Mbps
6	16-QAM	3/4	36 Mbps
7	64-QAM	2/3	48 Mbps
8	64-QAM	3/4	54 Mbps

BPSK: Binary Phase-Shift Keying
QPSK: Quadrature Phase-Shift Keying
QAM: Quadrature Amplitude Modulation

Table 1. Eight PHY Modes of IEEE 802.11a

Parameters	Values
Preamble Length	16μs
PLCP Header Length	4μs
MAC Header Size	28 Bytes
Slot Time	9μs
Short Inter Frame Space (SIFS)	16μs
DCF Inter Frame Space (DIFS)	34μs
Minimum Contention Window Size (CWmin)	31
Maximum Contention Window Size (CWmax)	1023
Clear Channel Assessment(CCA) Time	3μs
RxTxTurnaround Time	1μs
RTS Threshold	0 or 3000 Octets
Fragmentation Threshold	2100 Octets
LongRetryLimit	7
ShortRetryLimit	7

Table 2. IEEE 802.11a MAC/PHY Channel Parameters Used in Simulations

the implementation also includes two important features: (1) it has a realistic channel propagation model that is suitable for mobile nodes deployed in urban area; and (2) it uses a FER (Frame Error Rate) - based reception model that reflects the error performance of 802.11a PHY modulation techniques and coding schemes.

We present the simulation modeling in details from the following four aspects: the multi-rate network model and parameters, the channel propagation model, the frame reception model, and the traffic model.

4.1 Multi-rate Network Model and Parameters

Sharing the same protocol at the MAC layer, currently deployed 802.11 wireless devices follow different standards at the PHY layer, for example, 802.11a, 802.11b, 802.11g, 802.11n, etc. All these standards support multi-rate transmissions. For the simulation purpose, however, we have chosen to use 802.11a at the PHY layer, which supports up to eight transmission rates (6$\sim$54Mbps), yet it is less complex than 802.11g which must support the four transmission rates of 802.11b to keep the compatibility.

IEEE 802.11a employs OFDM (Orthogonal Frequency Division Multiplexing) modulation technology in PHY layer. As shown in Table 1, by using different modulation schemes, such as BPSK, QPSK, 16-QAM, 64-QAM, and different coding rates, 802.11a supports 8 transmission rates. To guide against data loss, FEC (Forward Error Correction) coding/decoding is added in 802.11a PHY, which is performed by bit interleaving and convolutional coding.

Table 2 summarizes some major MAC/PHY parameters that we used in simulations. Just a comment on RTS threshold: when it is set as 0, RTS/CTS exchange will be used for all data frames. On the contrary, to disable RTS/CTS exchange sequence, we simply set the RTS threshold to 3000 octets, which is larger than any legal MAC data frame.

4.2 Channel Propagation Model

In ns-2, Friis free-space model is used to simulate short distance propagations and two-ray-ground model is used for long distances. Although widely used, these two models are inherently limited – they assume ideal propagation conditions without considering multipath fading effects, which often dominates the propagation effects in mobile networks. An enhanced model – the shadowing model was introduced in ns-2 to characterize the multipath fading effects. However, as pointed out in [3], using an added Gaussian random variable to simulate the fluctuations of the channel is lack of experimental validation. On the other hand, extensive measurement campaign has shown that the signal deviation can be better modeled using a Rayleigh or Ricean distribution. Ricean distribution is used when there is a dominant stationary signal component present, such as a line-of-sight propagation path; while Rayleigh distribution is more suitable if such a component does not present [16].

To simulate a mobile network that is deployed in an urban area, we developed a new implementation of the 802.11 PHY layer in ns-2, which combines the log-distance path loss model [16] and the Ricean propagation model [14].

With log-distance path loss model, and assuming the transmitter-receiver separation distance is d, the received power (in dBm) of a frame is computed as

$$Pr_d[dBm] = Pr_{d_0}[dBm] + 10 \cdot n \cdot \log(\frac{d}{d_0}), \quad (1)$$

where n is the path loss exponent which indicates the rate at which the path loss increases with distance, d_0 is the close-in reference distance which is determined from measurements close to the transmitter, and Pr_{d_0} (in $watt$) can be calculated using the Friis free-space model:

$$Pr_{d_0}[watt] = \frac{P_t G_t G_r \lambda^2}{(4\pi)^2 d_0^2 L}, \quad (2)$$

where P_t is the transmitted power, G_t and G_r are the antenna gains for the transmitter and receiver respectively, L is the system loss factor not related to propagation ($L \leq 1$), and λ is the radio wavelength in meters. In our simulations, the path loss exponent n is set as 3 to simulate an urban area environment, and d_0 is set as 1 meter.

With Ricean propagation model, the envelop of the received signal r has the following probability density function (pdf):

$$p(r) = \begin{cases} \frac{r}{\sigma^2} e^{-\frac{r^2+A^2}{2\sigma^2}} I_0(\frac{Ar}{\sigma^2}), \text{for}(A \leq 0, r \leq 0), \\ 0, \text{for}(r < 0), \end{cases} \quad (3)$$

where the parameter A denotes the peak amplitude of the dominant signal component and $I_0(\cdot)$ is the modified Bessel function of the first kind and zero-order. The Ricean distribution is often described in terms of a parameter K (called Ricean factor) which is defined as the ratio between the deterministic signal power and the variance of the multipath. It is given by $K = A^2/(2\sigma^2)$ or, in terms of dB, $K(dB) = 10\log(A^2/2\sigma^2)$. As $A \to 0$, $K \to -\infty\ dB$, the Ricean distribution degenerates to a Rayleigh distribution [16].

Punnoose et al. in [14] proposed an efficient implementation of Ricean propagation model based on a simple table lookup, which has been integrated into the implementation of our combined model. With the Pr_d calculated in (1), the final received power after Ricean fading effect will be

$$Pr_d^*[dBm] = Pr_d[dBm] + 10\lg(F_{Ricean}), \quad (4)$$

where F_{Ricean} is the envelope multiplicative factor resulted by the Ricean propagation model. In our simulations, the transmitted power and and the Reicean factor K were set as $15dBm$ and $6dB$ respectively.

4.3 Frame Reception Model

The frame reception model used in ns-2 is a simple threshold-based model: if the received power of the frame falls below the carrier sense threshold – $CS_{Threshold}$, the frame is discarded as noise to simulate not being sensed. If the received power is above $CS_{Threshold}$, however lower than the receive threshold – $Rx_{Threshold}$, the frame is marked as error and discarded. Otherwise, if the received power is above $Rx_{Threshold}$, the frame is received without errors. This reception model is obviously over-simplified. In reality, the error performance of 802.11 devices is much more complicated: besides platform-dependent factors (i.e., variations in transmitter power, antenna gain, reception sensitivity, and frequency control, etc.) and external interferences, it also highly depends on the underlying PHY modulation modes and coding schemes.

Towards a more realistic reception model, a FER(Frame Error Rate)-based model was set up in our simulations. If the received power of a frame is above $Rx_{Threshold}$ and not discarded because of the capture effect, instead of being delivered to MAC layer instantly, it has to go through the following steps:

1. Calculate the SINR for the incoming frame using the following formula:

$$SINR = \frac{Pr_i}{NoiseFloor + \sum_{j=1, j\neq i}^{n} Pr_j}, \quad (5)$$

   where Pr_i is the received power for the frame that we care, and $\sum_{j=1,j\neq i}^{n} Pr_j$ is the sum of the received power of all other ongoing frames, taken as interferences to the incoming frame.

2. Calculate BER (Bit Error Rate) based on the SINR and the transmission mode (see Table 1) with which the frame was sent. An AWGN (Additive White Gaussian Noise) channel noise model is assumed in the calculation.

3. Calculate FER (Frame Error Rate) using the upper bond probability of error that is given in [15] under the assumption of binary convolutional coding and hard-decision Viterbi decoding.

4. Generate a random number uniformly in the range of [0,1). If the number is larger than FER, the frame is correctly received (to be delivered to the MAC layer); otherwise, the frame is marked as error and then discarded.

Due to the space limitation, please refer to [19] for details about the calculation of the BER and FER and the frame reception model.

4.4 Traffic Model

We focus our simulations on infrastructure-based 802.11a WLANs. In the simulation, traffic flows are symmetric in the sense that all of them share the same traffic load in each scenario. There is a UDP traffic flow between any non-AP node to the AP with the non-AP node as the sender. The aggregated traffic load in the network could simply be expressed as

$$E(n, l, \tau) = n \times l \times (1/\tau), \quad (6)$$

where n represents the number of nodes, l represents the average data packet size [1], and τ represents the average packet interval. We vary the traffic load and patterns through changing the n, l, and τ respectively.

[1] The data packet size referred in this work includes the length of TCP/UDP and IP headers, however, not the MAC header.

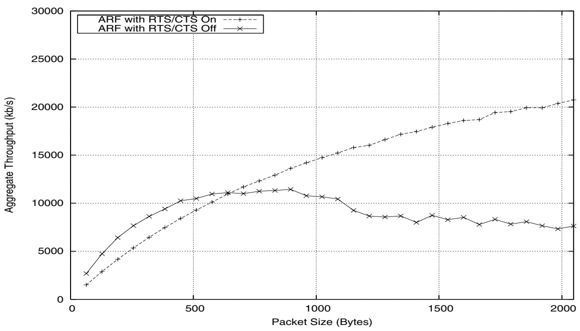

Deployment area: 80meters × 80meters

Traffic load: ($n = 40$, $\tau = 5ms$, and l varies)

Figure 2. Aggregated Throughput: RTS/CTS-on vs. RTS/CTS-off

5 Performance Analysis

Different from most work on rate adaptation, we chose to use network-wise aggregated throughput, instead of per-flow throughput, as the major performance metric. Two reasons are behind: (1) the performance anomaly of multi-rate networks – a per-flow optimal strategy might lead to the performance degradation of the whole network; and (2) the network revenue is better represented by the network-wise aggregated throughput than the throughput of some individual flows.

5.1 The Rate Avalanche Effect

First, let us use one example scenario to demonstrate the effect that we call rate avalanche. In this scenario, 40 802.11a wireless nodes are randomly deployed in a square area of 80meters × 80meters and one AP is installed in the center of the area (i.e. (40, 40)). There is one UDP-based traffic flow from each node to the AP. Each flow has the packet arrival rate of 200 packets/second (or a packet interval of 5ms). All the flows start at $0s$ and stop at $30s$. Nodes are kept static during the whole simulation procedure. We first turn RTS/CTS on and repeat the experiment with different packet sizes, from 64 Bytes to 2048 Bytes with the step of 64 Bytes. Using the exactly same settings, the same group of experiments are repeated with RTS/CTS turned off.

Fig. 2 compares the aggregated throughput for the cases with RTS/CTS turned on and the cases with RTS/CTS turned off (referred as RTS/CTS-on and RTS/CTS-off respectively in the rest of the paper). As we can see from the figure, between 64 Bytes and 640 Bytes, RTS/CTS-on delivers slightly lower throughput than RTS/CTS-off. However, starting from 640 Bytes, the performance of RTS/CTS-off are much lower than that of RTS/CTS-on.

With no hidden node presenting, it was expected the RTS/CTS-off would always deliver higher throughput than RTS/CTS-on due to the transmission overhead of RTS/CTS exchanges. So what causes the performance degradation of RTS/CTS-off? The question can be answered by Figs. 3 (a) and (b), which compare the number of collisions and the cumulative distribution of transmission rates used for RTS/CTS-on and RTS/CTS-off respectively (for the experiment with packet size of 1024 Bytes): (1) Fig. 3 (a) demonstrates that RTS/CTS-on leads to fewer collisions. This is because RTS frame is much shorter compared to most data frames, thus the collision probability of RTS frames is much lower. And once a RTS frame is successfully transmitted, the channel access is then reserved for the CTS and data frames that follow. Less collisions lead to less retransmissions which save the channel resources. (2) More importantly, Fig. 3 (b), with rates 1 to 8 corresponding to 6 to 54Mbps respectively, clearly shows that RTS/CTS-off leads to a much higher percentage of lower-rate transmissions than RTS/CTS-on. ARF behaves differently because the only feedback that ARF used to adapt transmission rates is the statistics of acknowledged/unacknowledged data transmissions. When RTS/CTS is off, every collision will cause an unacknowledged transmission, which will lead to rate dropping; on the other hand, when RTS/CTS is on, most collisions will be RTS frame collisions, which only lead to retransmissions of RTS frames, instead of unacknowledged data transmissions.

It is worthwhile noting, that as the RTS/CTS transmission overhead is always incurred for all the cases when RTS/CTS is turned on, so does the rate avalanche so RTS/CTS-off; however, they are differently influenced by the (average) size of the data packets. While the effect of RTS/CTS transmission overhead becomes less severe when packet size increases, the rate avalanche effect goes to the opposite direction – it is amplified with relatively large data packets. This explains the crossing of the two curves in Fig. 2. *The packet size corresponding to the* **cross-point** *could be used as the optimal RTS Threshold*.

5.2 The Impact of Node Numbers

To study the impact of node numbers, in this group of experiments, we vary the number of nodes from 20 to 60. A fixed packet interval of 5ms is used for the UDP applications. As shown in Fig. 4, the cross-points highly depend on the number of nodes. *When the number of nodes gets larger, the cross points (or the optimal RTS thresholds) are shifting toward smaller packet size, which means the chance that RTS/CTS-on outperforms RTS/CTS-off is getting larger*. This could be easily understood: higher number of nodes yield higher level of contentions and more collisions, which lead to more retransmissions and deeper de-

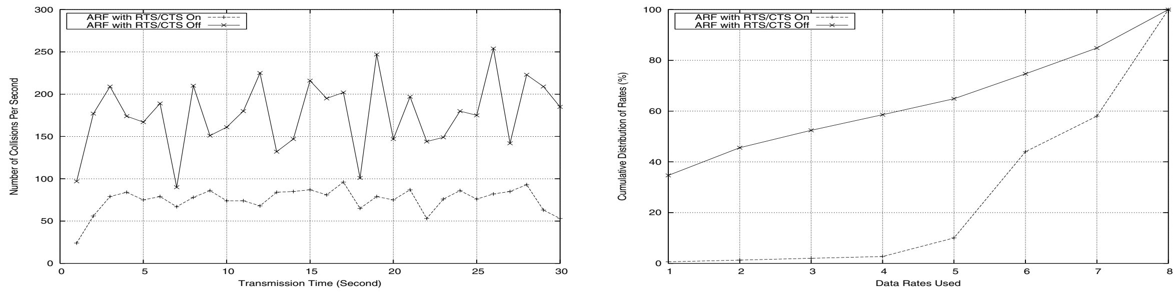

(a) Number of Collisions Per Second (b) Cumulative Distributions of Data Rates Used

Deployment area: 80meters × 80meters

Traffic load: ($n = 40$, $\tau = 5ms$, and $l = 1024 Bytes$)

Figure 3. Collisions and Data Rates

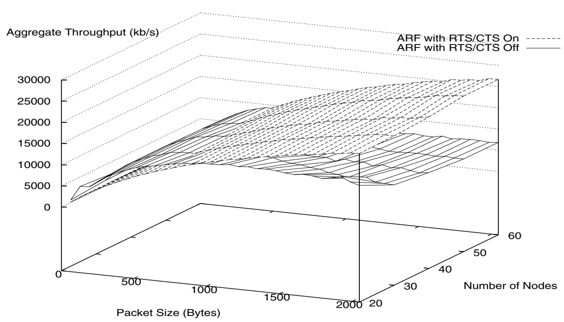

Deployment area: 80meters × 80meters

Traffic load: ($\tau = 5ms$, n and l vary)

Figure 4. The Impact of Node Numbers

gree of rate avalanche.

5.3 The Impact of Node Geographical Distributions

To study how geographical distribution could influence the performance of RTS/CTS-on and RTS/CTS-off, we varied the size of the node deployment area. Figs. 5 (a) and (b) show the results for the cases of 60meters × 60meters and 100meters × 100meters respectively. Comparing these two figures with Fig. 4, we can see that *as the deployment area is larger (i.e. nodes are more spreading out), the cross-points (i.e. the optimal thresholds) are shifting toward smaller packet sizes*. This could be explained as follow. When nodes are deployed in a wider area, the chances that they have to use lower rates for transmission are getting higher. The side-effect is that the relative overhead of transmitting RTS/CTS exchanges is getting lower since the difference between the transmission rates for data frames and RTS/CTS exchanges are smaller.

Beside that, we also notice that when the deployment area is larger, the performance gain of RTS/CTS-off over

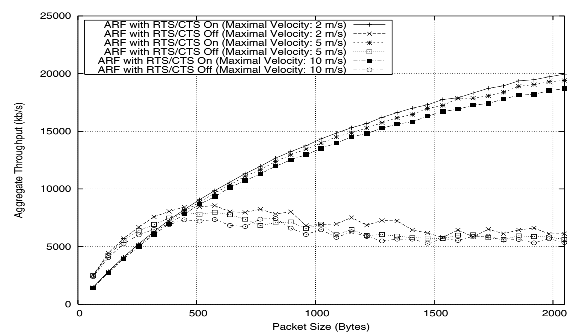

Deployment area: 80meters × 80meters

Traffic load: ($n = 40$, $\tau = 5ms$, and l varies)

Figure 6. The Impact of Node Mobility

RTS/CTS-on before the cross-points in the curves, diminishes quickly.

5.4 The Impact of Node Mobility

Next we investigate how node mobility influences the performance of RTS/CTS-on and RTS/CTS-off. Node mobility is simulated using the widely adopted Random Waypoint Mobility Model (RWPM). We compare the results for experiments that use three different Maximal Velocity values, namely, 2, 5 and 10 meters/second. In this group of experiments, 40 nodes are first randomly deployed in the area of 80meters × 80meters, then move individually following the RWPM model. Results are shown in Fig. 6. As we can see from the figure, a higher moving speed causes a slight drop in the performance of both RTS/CTS-on and RTS/CTS-off. The cross-point also shifts toward smaller packet sizes when the moving speed increases. This is because, overall, a faster moving speed of the sender, the receiver and other surrounding nodes will also yield faster fluctuations of the received signal strength at the receiver due to the Doppler shifts and multipath propagations. Reacting to a faster changing channel is surely more challeng-

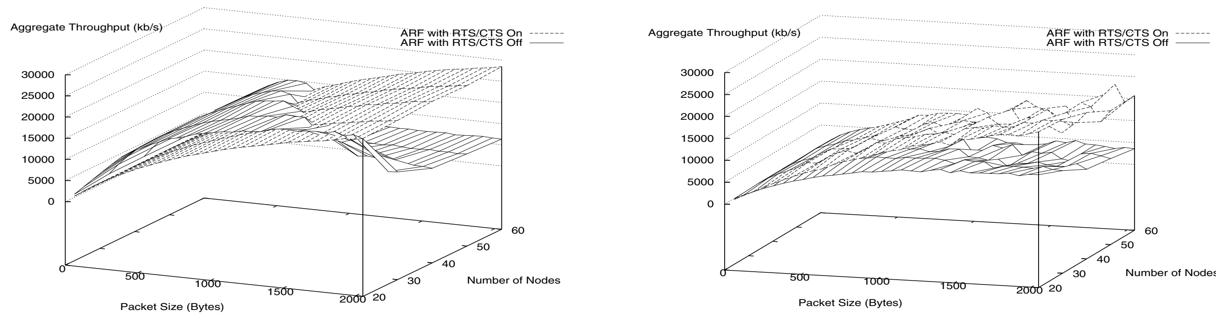

(a) Deployment area: 60meters $\times$ 60meters

(b) Deployment area: 100meters $\times$ 100meters

Traffic load: ($\tau = 5ms$, n *and* l *vary)*

Figure 5. The Impact of Node Geographical Distributions

ing to the ARF scheme.

6 Conclusion

In this paper, we report an important phenomenon called rate avalanche that happens in heavily-loaded multi-rate 802.11 WLANs when RTS/CTS exchange is turned off. Motivated by this discovery, we conducted an extensive investigation on the effects of the RTS/CTS exchange on the performance of multi-rate 802.11 WLANs. Our investigation reveals that, with all factors considered, keeping RTS/CTS on will have a much higher chance to gain a better network performance than keeping RTS/CTS off. If dynamic RTS/CTS exchange is to be employed, simply using a pre-configured RTS threshold will only yield sub-optimal performance.

Acknowledgement

Liqiang Zhang's research was supported in part by an Indiana University Faculty Research Grant. Xiaobo Zhou's research was supported in part by US National Science Foundation grant CNS-0720524.

References

[1] J. Bicket, "Bit-rate Selection in Wireless Networks," *MIT Master Thesis*, 2005.

[2] C. Chen et al., "Rate-adaptive Framing for Interfered Wireless Networks," *in Proc. of IEEE INFOCOM'07*, Anchorage, Alaska, May 2007.

[3] J. Dricot and P. Doncker, "High-accuracy Physical Layer Model for Wireless Network Simulations in ns-2," *in Proc. of International Workshop on Wireless Ad Hoc Networks (IWWAN 2004)*, Finland, May-June 2004.

[4] M. Heusse et al., "Performance Anomaly of 802.11b," *in Proc. of IEEE INFOCOM 2003*, San Francisco, California, April 2003.

[5] I. Haratcherev et al. "Hybrid Rate Control for IEEE 802.11", *in the proc. of ACM MOBIWAC'04*, Philadelphia, PA, Sept. 2004.

[6] G. Holland et al., A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks," *in Proc. of ACM MOBICOM'01*, Rome, Italy, July 2001.

[7] 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.

[8] J. Kim et al., "CARA: Collision-Aware Rate Adaptation for IEEE 802.11 WLANs," *in Proc. of IEEE INFOCOM'06*, Barcelona, Spain, April 2006.

[9] IEEE 802.11, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std. 802.11-1999, 1999.

[10] IEEE 802.11a, *802.11a, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-speed Physical Layer in 5 GHz Band*, Sept. 1999.

[11] D. Qiao et al., "Goodput Analysis and Link Adaptation for IEEE 802.11a Wireless LANs", *IEEE Transactions on Mobile Computing*, Vol. 1, No. 4, Oct.-Dec. 2002.

[12] M. Lacage et al., "IEEE 802.11 Rate Adaptation: A Practical Approach," *in Proc. of ACM MSWiM'04*, Venezia, Italy, Oct. 2004.

[13] The Network Simulator – ns-2, http://www.isi.edu/nsnam/ns.

[14] R. Punnoose et al., "Efficient simulation of Ricean Fading within a Packet Simulator," *in Proc. of IEEE VTC 2000 Fall*, Sept. 2000.

[15] M. Pursley and D. Taipale, "Error Probabilities for Spread-Spectrum Packet Radio with Convolutional Codes and Viterbi Decoding," *IEEE Trans. Comm.*, vol. 35, no.1, pp. 1-12, Jan. 1987.

[16] T. S. Rappaport, "Wireless Communications: Principles and Practice," second edition, Prentice Hall, 2002.

[17] B. Sadeghi et al., "Opportunistic Media Access for Multirate Ad Hoc Networks," *in Proc. of ACM MOBICOM'02*, Atlanta, Georiga, Sept. 2002.

[18] S. Wong et al., "Robust Rate Adaptation for 802.11 Wireless Networks," *in Proc. of MOBICOM'06*, Los Angeles, CA, Sept. 2006.

[19] L. Zhang, Y. Cheng, and X. Zhou, "Effects of RTS/CTS Exchange on the Performance of Multi-rate 802.11 WLANs," *Technical Report, available at http://www.cs.iusb.edu/~liqzhang/RTSCTS-Multirate.pdf*, Dec. 2007.