Performance Analysis of the IEEE 802.11 Wireless LAN Standard[1]

Craig Sweet and Deepinder Sidhu

Maryland Center for Telecommunications Research
Department of Computer Science and Electrical Engineering
University of Maryland Baltimore County
1000 Hilltop Circle
Baltimore, MD 21250

{sweet, sidhu}@cs.umbc.edu

1. Abstract

IEEE 802.11 is a relatively new standard for communication in a wireless LAN. Its need arose from the many differences between traditional wired and wireless LANs and the increased need for interoperability among different vendors. To date, detailed performance measures for this CSMA/CA protocol are not known. We describe the results of our Discrete-Event Simulation of the Distributed Coordination Function (DCF) within the MAC sublayer. We model an ideal LAN and describe the best case performance. Our results show the relationship between the protocol options and total system throughput.

2. Introduction

Over the last several years, we have witnessed widespread deployment of Wireless LANs in virtually every industry. Until recently, there has been no agreed upon standard by which wireless stations communicate. This lack of standardization usually results in decreased interoperability. The Industry for Electrical and Electronics Engineers (IEEE) has been working with leaders from industry to develop a standard to which wireless stations from different vendors can communicate. In 1997, the IEEE finally ratified their standard 802.11, the Physical and MAC specification for Wireless LANs [IEE97].

Since traditional Ethernet has been in existence for quite some time, much research has been done studying its attributes under various conditions [BUX81, GON83, and GON87]. A detailed study of the Carrier Sense Multiple Access with Collision Detection scheme used in Ethernet can be found in [TOB80].

Since wireless networking is a recent development, not much is known about how the protocol performs. The goal of this research is to uncover some of the hidden performance issues in this new direction in networking.

3. Modeling and Simulation

In our experiments, the goal was to explore the efficiency of the MAC protocol under ideal conditions. While many of these conditions may be unrealistic, the end result is useful in telling us the highest performance that can be expected from the protocol. This section describes some of the assumptions and limitations assumed in our system. Also, the simulation model and computation variables are described.

Assumptions

All stations are assumed to be using a Direct Sequence Spread Spectrum (DSSS) radio. The operation of Frequency Hopping Spread Spectrum (FHSS) and Infrared (IR) radios had too much of an impact on a given transmission to study the aspects of the protocol itself. Additionally, it is assumed that there are no power considerations for either the radios or the wireless stations that could interfere with the operation of the protocol.

A significant aspect of any transmission protocol is how it handles transmission errors. In order to focus on the core MAC protocol, we assumed error-free channels. Additionally, all stations have unobstructed access to all other stations and thus can hear all transmissions.

To minimize complexity, we chose to model our wireless LAN as an ad-hoc network, also known as an Independent Basic Service Set (IBSS). This is the simplest type of wireless LAN defined in the standard. There is no Access Point and therefore no tie to a wired LAN.

1 This research was supported in part by the Department of Defense at the Maryland Center for Telecommunications Research, University of Maryland Baltimore County. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Department of Defense or the U.S. Government.

Finally, we decided to only model the DCF portion of the protocol. The DCF forms the heart of the MAC protocol. The PCF has been omitted in this simulation, as it is an optional portion of the MAC protocol that works on top of the DCF and can significantly alter its results.

Description of Simulation Model

To perform this analysis, we constructed a discrete-event simulation of the MAC portion of the IEEE 802.11 protocol. A complete description of simulation techniques can be found in [BAN84]. For all experiments, each station is assumed to have a one MAC Service Data Unit (MSDU) buffer. An MSDU is the basic unit delivered between two compatible MAC sub-layers. For uniformity all MSDUs transmitted are of equal size. Initially, each station is given one MSDU to transmit. Upon completing the transmission attempt, another MSDU is assigned for transmission after some exponentially distributed inter-arrival time. In this manner, changing the mean inter-arrival time between MSDUs can be used to alter system load.

Offered Load Computation

Upon transmitting a message, the station generates the next message with an inter-arrival time exponentially distributed with mean θ. Additionally, each station is sending the same size packets, in bytes P, for the duration of a run. The offered load of station i, G_i, is defined as in [GON87] to be the throughput of station i if the network had infinite capacity, i.e.,

$$G_i = T_p / \theta_I$$

where $T_p = P/C$ and C is the transmission speed in Mbit/s. The total offered load can thus be computed to be

$$\sum_{i=1}^{N} G_i$$

4. Performance Analysis

In this analysis, we performed four experiments measuring various aspects of the MAC protocol. Each of these experiments was conducted at several transmission speeds. 1 and 2 Mbit/s were selected because they are explicitly supported in the specification. 10 Mbit/s was selected to provide a comparison at traditional LAN speeds. The results of these experiments are the topic of this section. Current research is aimed at providing 802.11 operation at 10 and 20 Mbit/s.

Experiment 1: Variable Load

In our first experiment we wanted to see what effect the total load on the system played on performance. This experiment is similar to one found in [GON87]. Figure 1 shows the variation of total throughput with total offered load G for various message sizes P at 1 Mbit/s. In this experiment, the fragmentation threshold has been set to 2346 bytes and the RTS threshold has been set to 3000 bytes.

We can see that with an offered load of about 80% or less virtually no collisions occur and throughput and load are approximately equal. Once the system load increases beyond 90-100% we see the impact of collisions. As can be expected, greater throughput is achieved via a greater packet size. Due to the overhead present in the protocol, acceptable throughput was not seen with packet sizes below 2000 bytes.

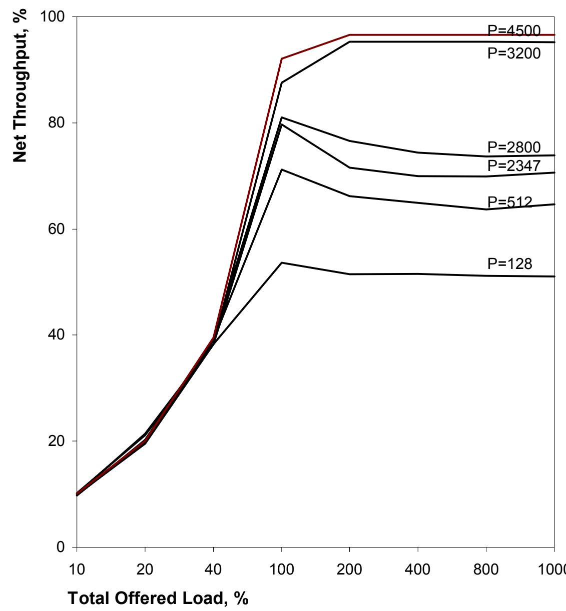

Fig. 1. Throughput vs. Offered Load at 1 Mbit/s, 32 stations, Parameter, P.

Packet sizes above and below the fragmentation threshold did not yield much difference. Even then, it all but disappeared with loads in excess of 200%. While increasing the number of packets per message produces more overhead, it also reduces the collision probability.

In this example, the RTS threshold played a crucial role in the performance of the protocol. The throughput peaked out at approx. 80% for all packet sizes below 3000 bytes. For packet sizes above the RTS threshold, noticeable performance gains were seen and throughput peaked at 96%.

The RTS threshold acts as a medium reservation mechanism. Collisions, and subsequent retransmissions, can occur on the smaller RTS frames but not normally on the longer data frames. The result is a better utilization of the bandwidth.

Table I
Simulation Results at 200% Offered Load for Various Packet Sizes and Transmission Speeds

Mbit/s	Packet Size	Throughput %
1	4500	96.61
	2800	76.54
	2347	71.52
2	4500	96.11
	2800	76.07
	2347	73.08
10	4500	91.80
	2800	73.17
	2347	68.87

Our results were similar for transmission speeds of 2 and 10 Mbit/s. Table I summarizes some of these results. What we saw was that as the transmission speed increased, the throughput dropped. This can be attributed to the fact that the inter-frame spaces are independent of transmission speed. At higher speeds, since it takes less time to send the same packet, an Inter-Frame Space (IFS) of 50 µs has more of an impact than at lower speeds.

Experiment 2: Variable Stations

In our second experiment, our goal was to determine how many stations would overload a wireless network. Certainly the performance characteristics for 10 stations would be different than for 20 stations, all contending for access to the medium. Figure 2 shows the effect on throughput with an increasing number of stations and a constant Offered Load of 100%.

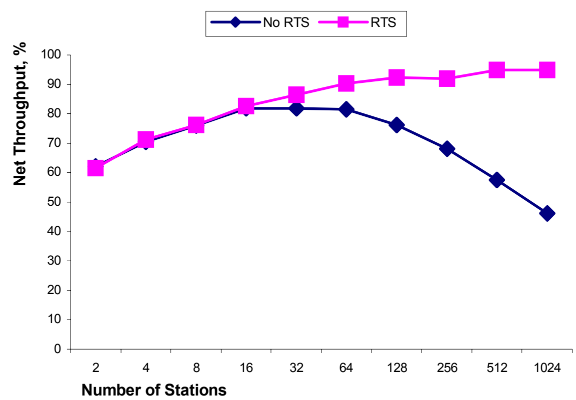

Fig. 2. Throughput vs. Number of stations with packet size above and below RTS Threshold (1 Mbit/s)

Results are shown both with and without the RTS mechanism implemented. For all runs, the message size was set to 3000 bytes and the fragmentation threshold was set to 2346 bytes.

Without RTS enabled, we can see that the maximum throughput reached was approx. 82% with 16 participating stations. In fact, with few stations (below 16), we see that there is not much difference in performance with and without RTS enabled.

As more stations are added to the simulation the probability that two or more stations will calculate the same backoff window is increased. Thus, the chance for collision increases. This can be seen by the large differences between the RTS and No-RTS runs with higher station counts, above 64.

Since IEEE 802.11 uses CSMA/CA, collisions are expensive. The transmitting station must continue to transmit the entire message and wait a minimum amount of time before determining that the transmission was in error. With RTS enabled, the collisions occur on smaller RTS frames, allowing for a quicker turn-around time. We can see that with RTS enabled, the system stabilized to approx. 92% or higher with 128 or more stations.

As in the previous experiment, we saw similar results in our 2 and 10 Mbit/s experiments. As the speed of the medium increased there was still the same pattern between RTS and No RTS results. We can see that higher transmission speeds yielded lower average throughput results. Table II summarizes some of the results from these experiments with and without RTS enabled.

Table II
Simulation Results at 100% Offered Load with Variable Number of Stations

Mbit/s	# of Stations	Throughput with RTS	Throughput without RTS
1	16	82.59	81.93
	128	92.38	76.22
	1024	94.81	46.13
2	16	82.83	82.98
	128	93.04	73.28
	1024	94.07	53.91
10	16	80.34	78.01
	128	88.6	70.04
	1024	89.95	57.84

Experiment 3: Variable Fragmentation

In our third experiment, our goal was to determine what effect the fragment size played on system performance. The simulation was run with 32 stations at 200% load with varying fragmentation thresholds. Each message sent was 3000 bytes long. Therefore, the fragmentation threshold merely determined how many fragments the 3000 byte messages were broken up into.

Intuitively, advantages can be gained by both increasing and decreasing the fragmentation threshold. Smaller thresholds limit the loss of performance due to retransmissions but come with an increase in overhead. This is important because the 802.11 protocol has considerable overhead [IEE97]. On the other hand large fragmentation thresholds, while limiting the overhead, become expensive in the event of a collision.

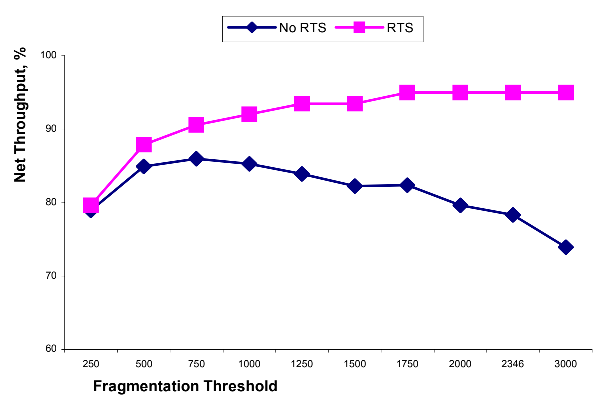

Fig. 3 Throughput vs. Fragmentation Threshold with packet size above and below RTS Threshold (1 Mbit/s)

Figure 3 shows the results of this experiment run at 1 Mbit/s. As predicted, the RTS mechanism does a great deal to improve the performance of this aspect of the protocol. The reason can be attributed to the reduction in collisions that it provides. At smaller thresholds, there is little difference between the RTS and No RTS figures. There is nearly a balance between three factors: the overhead provided by the RTS mechanism, the smaller fragment sizes that are retransmitted in the event of a collision, and the overhead provided by multiple smaller fragments.

It is not until the fragmentation threshold increases that we see the largest variation in performance. As was expected, with larger fragments comes a decrease in performance. Each collision requires retransmission of a much larger fragment. Since 802.11 does not have a collision detection mechanism the entire fragment must be transmitted before success or failure of that fragment can be determined.

This experiment has also shown that, in this specific case, little improvement can be seen with fragments above 1000 bytes when the RTS mechanism is used. While this may be true in this experiment, note that we are assuming that all fragments are transmitted error-free. This assumption will certainly not hold in a real-world case. In fact, performance may decrease as the bit-error rate increases. The probability of each fragment being successfully delivered will decrease as the fragment size increases and results will most certainly differ.

The results for 1, 2, and 10 Mbit/s experiments are summarized in Table III. We can see that the same pattern is exhibited regardless of the transmission speed. As we have seen in the previous experiments, the constant inter-frame space times effectively reduce the system performance at higher speeds.

Table III
Simulation Results at 200% Offered Load with Variable Fragmentation Threshold and Transmission Speed

Mbit/s	Frag. Threshold	Throughput with RTS	Throughput without RTS
1	250	79.61	78.93
	1250	93.46	83.90
	3000	94.96	73.89
2	250	78.44	77.72
	1250	92.68	83.49
	3000	94.27	73.38
10	250	70.47	70.09
	1250	88.79	80.56
	3000	89.45	70.50

Experiment 4: Variable Propagation Delay

In our previous experiments, we assumed a constant delay of 1 μs between stations. This allowed us to measure the protocol performance without respect to the interoperability in a real-life situation. In our fourth experiment, our goal was to determine how far apart stations can be from one another, in terms of propagation delay, before system throughput degrades. In a real-world wireless network, some stations may be constantly moving while others are stationary for periods of time.

In this experiment, we set the fragmentation threshold to 2346 bytes and the message size to 3000 bytes. The system is run at 100% Offered Load. Figure 4 shows the results of increasing the propagation delay between any two wireless stations operating at 1 Mbit/s.

We can see that, with the current fragmentation threshold and a 50 μs IFS, throughput drops when the propagation delay between stations exceeds 50 μs.

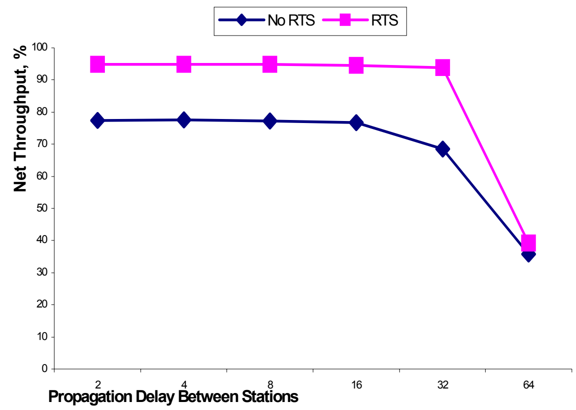

Fig. 4. Throughput vs. Propagation Delay with packet size above and below RTS Threshold (1 Mbit/s)

When a station transmits a message, it waits only a finite amount of time for the response. If this response does not arrive in time, it will retransmit the message. This timer begins immediately after the sender finishes transmitting the message. If the receiver is sufficiently far away from the sender, much of this time is taken up by twice the delay between the stations, once for the message to reach the recipient and once for the response to arrive at the source.

If the distance between two stations becomes too large, it will be impossible for the sender to hear the acknowledgement from the receiver. In this case, it becomes increasingly difficult for messages to be received correctly. The result is increased retransmissions and decreased throughput.

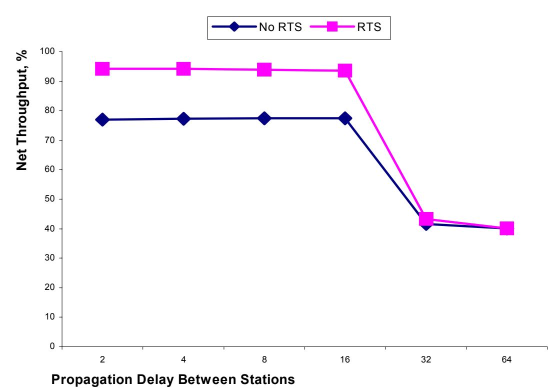

Fig. 5. Throughput vs. Propagation Delay with packet size above and below RTS Threshold (2 Mbit/s)

Unfortunately the problem only compounds itself as the transmission speed increases. Figure 5 shows the same experiment run at 2 Mbit/s. Here we see that the same drop off in throughput occurs with stations only 32 μs apart. The reason is that the initial transmission is shorter at the higher speed, which forces the station to begin its waiting period earlier. Therefore, this timer can expire with a shorter propagation delay.

As can be expected, the results are even worse for transmissions at 10 Mbit/s. These results are shown in figure 6. An interesting point in all three graphs is that the RTS mechanism can do little to improve this performance. This assures us that the loss in throughput is not attributed to collisions but rather to too much distance between stations. In fact, the added overhead of the RTS mechanism slightly reduces the performance once this problem occurs.

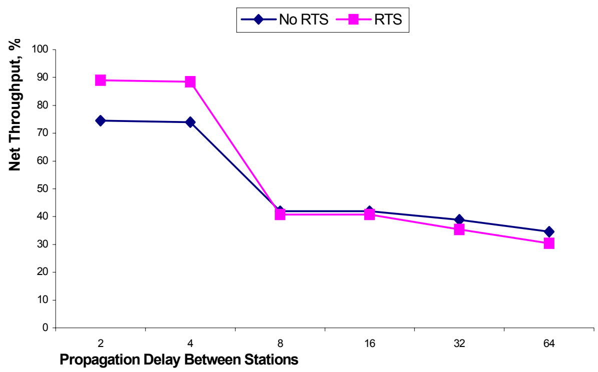

Fig. 6. Throughput vs. Propagation Delay with packet size above and below RTS Threshold (10 Mbit/s)

It a reasonable assumption that there is a limit to the distance that any two communicating stations can be from one another before system performance suffers. This limit is based upon the attributes of the communication medium and the protocol. From this experiment we can see that the transmission speed also plays a crucial role.

5. Conclusion

While the experiments described in this paper do not reflect any real-life scenario, they are useful in determining the maximum system performance under a variety of conditions. Our goal has been to see what the maximum performance we can expect out of the protocol is and what it takes to reach it.

We see from our experiments that Ethernet speeds are possible but only with the RTS mechanism that is built into the 802.11 MAC protocol. This mechanism, while adding some overhead, offers considerable improvement in most highly loaded systems.

We found that the best performance can only be achieved in systems with relatively slow transmission speeds. Transmission speed and throughput were inversely proportional. This is due to the constant delays and timers used in the protocol, which are not altered as transmission speed increases.

Future Work

Currently our research does not take into account the transmission errors that are inherent in all forms of communication. One area of research will be to incorporate a bit-error rate into the simulation, based upon the transmission device, and see how the system performance is affected.

Our system did not allow for a subset of stations to be hidden from the others. We assumed that all stations can hear all transmissions from all others. With this medium, stations can be obstructed from some other stations in the network. This would prevent them from reading all of the Network Allocation Vector (NAV) values that are transmitted. Future research could take this into account.

The aim of our research was focused on the DCF but completely ignored the optional PCF. It is quite possible that some of the inefficiencies found in our experiments can be overcome by the PCF. Our current research aim is to explore possible performance gains by exploiting the PCF functionality.

References

[BAN84] Banks, J. and J. S. Carson, "Discrete-Event System Simulation," Prentice-Hall, Englewood, NJ, 1984.

[BUX81] W. Bux, "Local-area subnetworks: A performance comparison," *IEEE Trans. Commun.*, vol COM-29, pp. 1465-1473, 1981.

[GON83] T. A. Gonsalves, "Performance characteristics of 2 Ethernets: An experimental study," ACM SIGCOMM *Symp. On Commun. Architectures and Protocols*, Austin, TX, Mar. 1983, pp. 178-185.

[GON87] T. A. Gonsalves, "Measured Performance of the Ethernet," in *Advances in Local Area Networks*, Kummerle, K., Tobagi, F., and Limb, J.O. (Eds.), New York: IEEE Press, 1987.

[IEE97] IEEE Std 802.11-1997, "IEEE Standard for Local and Metropolitan Area Networks: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification."

[TOB80] F. A. Tobagi and V. B. Hunt, "Performance analysis of carrier sense multiple access with collision detection," *Comput. Networks*, vol. 4, Oct./Nov. 1980.