802.11b PERFORMANCE EVALUATION

Yasir Zahur, Murtaza Doctor, Sadegh Davari, T. Andrew Yang
University of Houston Clear Lake
2700 Bay Area Blvd., Houston, TX 77058
USA

Abstract

In this paper, the impact of various key parameters on the actual performance of 802.11b wireless LAN protocol is verified by a series of controlled experiments. Overall, four sets of independent experiments were conducted to test the respective effect of one or more parameters on the performance of the 802.11b network. The parameters considered in our experiments included distance and power, Wired Equivalent Privacy (WEP), Fragmentation Threshold Setting (FTS), and Request To Send / Clear To Send (RTS/CTS). Lessons that we have learned from the series of experiments are discussed in the paper.

Key Words

Wireless LAN, IEEE 802.11b, Performance Evaluation

1. Introduction

High-speed wireless local area networks (WLAN) can provide the benefits of network connectivity without the restrictions of being tied to a location or restrained by wires. Despite the convenience of mobility, in order for the WLAN to be adopted as part of an enterprise network, two primary issues must be addressed: performance and security. The most widely implemented WLAN protocol, IEEE 802.11b, is claimed to have transfer rates of up to 11Mbps. However the actual performance demonstrated has been much lower than what is stated by the standard.

Our study has focused on evaluating the performance of the 802.11b WLAN, by studying the impact of parameters such as Wired Equivalent Privacy (WEP), physical distance, Request to Send/ Clear to Send (RTS/CTS), Fragmentation Threshold (FTS) and power variations. The rest of the paper includes a brief survey of the medium access control (MAC) and physical layers of the 802.11b standard, network architectures, configurations and the result of the performance evaluation experiments.

2. 802.11b Architecture

WLANs allow wireless stations to communicate with each other and to access the network using radio waves as the conduction medium. A WLAN, in Infrastructure mode, consists of a central connection point called the *Access Point* (AP), analogous to a hub or switch in a wired LAN. The AP transmits data between various nodes of a WLAN and, in most cases, serves as the only link between the WLAN and the wired network.

The 1999 version of the IEEE *802.11* WLAN Standards defines three types of wireless networks [5] An *Independent Basic Service Set (IBSS)* is commonly referred to as an *Ad Hoc Network*. An IBSS consists of end nodes communicating without any AP. The IBSS mode is useful for quickly setting up a wireless network, such as for a group meeting, at a convention center, or at an airport, etc. A *Basic Service Set (BSS)* is commonly referred to as an *Infrastructure Network*. A BSS consists of a single AP and all the communications between any two nodes must pass through the AP. The coverage area is greatly increased as compared to an IBSS. An *Extended Service Set (ESS)* consists of multiple BSSs, each having a single AP. The APs are linked together to form a LAN.

2.1 802.11b MAC Architecture

The latest 802.11b standard [1] is designed using *Direct Sequence Spread Spectrum (DSSS)* WLAN system which in turn uses a *complementary code keying (CCK)* modulation scheme. Overall, the symbol rate for CCK is 1.375 Mega samples/sec with a chipping rate of 11 Mega chips/sec for 8 bits per symbol. This translates into a data rate of 11 Mbps. CCK was chosen over other modulations for its superior performance when combating multi-path. CCK is a form of vector modulation and a variation of *M-ary orthogonal keying (MOK)* modulation which uses *In-Phase and Quadrature (I&Q)* modulation with complex symbol structures.

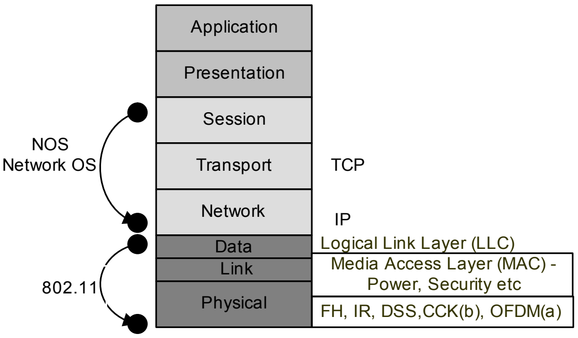

Figure 1: 802.11 Protocol Stack

Figure 1 shows the 802.11 protocol stack. As with IEEE 802.3 standard, the MAC layer defined by IEEE 802.11 standard is the lower part of the data link layer and is placed between the physical layer and Logical Link Control (LLC) sub layer of the data link layer.

The MAC layer consists of two coordination functions: *Distributed Coordination Function (DCF)* and *Point Coordination Function (PCF)* as shown in Figure 2. *DCF* is implemented in individual stations and is used in IBSS and other wireless network configurations as the *Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)*. For a station to transmit, it shall sense the medium to determine if another station is transmitting. If the medium is not determined to be busy, the transmission may advance. The CSMA/CA distributed algorithm mandates carefully designed waiting periods and medium reservation using a special timer called *Network Allocation Vector (NAV)*. In the 802.11b MAC layer, *PCF* is located above DCF and the access algorithm for this level is based on circular polling from an access point, that is, deterministic access. This mechanism allows transmission of traffic that does not tolerate random and unbounded delays or contention-free asynchronous traffic.

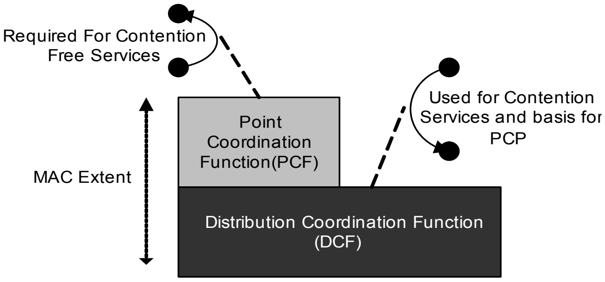
Required For Contention Free Services
Used for Contention Services and basis for PCP
MAC Extent
Point Coordination Function(PCF)
Distribution Coordination Function (DCF)

Figure 2: 802.11 MAC Architecture

The 802.11 standard further embodies the RTS/CTS feature to control the station's access to the radio medium. The primary reason for implementing RTS/CTS is to minimize collisions among hidden stations. An example of this is the *hidden node problem*, which can disrupt communications in a highly loaded WLAN. Take a scenario where stations A and B can communicate, however, station C is unable to receive from station A because of an obstruction and thus cannot determine when the channel is busy. Both stations A and C may simultaneously attempt to transmit to station B, which causes lost packets and subsequent retransmissions. Use of RTS, CTS, DATA and ACK sequences can be used to prevent this type of problem. More details about the hidden node problem and the effects of RTS/CTS are discussed in section 3.4.

3. Performance Evaluation Experiments and Results

Configuration parameters used for various test machines in the performance experiments are listed in Table 1.

	Desktop1	**Desktop2**	**Laptop1**	**Laptop2**
CPU	Intel based PII 400 MHz	Intel based PII 400 MHz	Intel based P III 750 MHz	Intel based P III 750 MHz
OS	Windows 2000 Professional	Windows 2000 Server	Windows XP	Windows 2000 Professional

Table 1: Stations used in the Experiments

All the four machines are equipped with 256 MB RAM, Cisco Aironet 350 WLAN adapters, and the Cisco Aironet client utility software. All machines run windows operating systems and J2SDK v 1.4.1. The two desktops are Pentium II machines running Windows 2000 Professional or Server, while the two laptops are Pentium III machines running Windows 2000 Professional or XP.

The access point used in the experiments was a Cisco Aironet 350 series access point with 802.11b as the network standard. The frequency band used was 2.4 to 2.497GHz; wireless medium used was DSSS and access mechanism being CSMA/CA. All experiments were performed with 100mW as transmit power (except in the power test where both 100mW and 30mW were used) and antenna at 9Dbi.

Terminologies Used

Round trip time for a particular data size is the time required for a packet carrying that size of data to travel from one host to the other and come back to the original host. It is measured in milliseconds. *Throughput* is the average rate at which the data travels between two users and is usually measured in kbps or Mbps. It should also be noted that the throughput is measured at the Application level to reflect as accurately as possible the performance that is actually experienced by the end user. This, however, implies that throughput does depend on the underlying transmission protocol (TCP or UDP) and data type being sent (i.e., HTTP, FTP, VoIP, etc). *Data bandwidth* is the maximum theoretical throughput or data rate at which data can be transmitted over the network.

3.1. Effect of Distance (Range Test) and Power

Distance and power are vital factors in evaluating the performance of 802.11b protocol. *Distance* can be tested by moving the wireless clients to locations with different distances from the AP. *Power* is a controllable attribute on the AP as well as on the clients' WLAN adapters.

Test Setup

As shown in Figure 3, the vital part of the testing was the use of a hybrid network which consisted of both an enterprise network and a wireless LAN. All machines were configured for dynamic IP addresses (DHCP). The access point was plugged into the networks switch and also configured for DHCP.

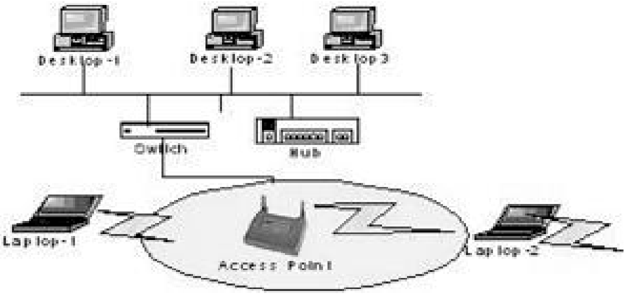

Figure 3: Configuration for Range and Power Level Tests

The WLAN consisted of approximately 40 machines. The access point was configured as an entity by itself and connected to the WLAN using a Netgear Fast Ethernet switch. All desktops as shown in the figure were connected to the wired LAN. Laptop1 and Laptop2 were wireless clients associated with the access point.

Testing Software

Network monitoring software, including *Netperf and Qcheck*, were used to perform the range and power tests. They establish two connections before actually testing the throughput. The first connection is responsible for exchanging control information and the second connection is actually used to test throughput using sample-sized packets. Regardless of the type of test being run, the control connection will be a TCP connection using BSD sockets. The software are claimed to conduct tests with 99% confidence level and have been used as testing software in various test laboratories.

Testing Cases

Case 1: Hybrid Transmissions
This case was tested using the hybrid network. Test data was sent from Laptop1 to Desktop1 in this case. The data transfer occurred partly on the WLAN and partly on the wired LAN. This is a distance test where Laptop1 was tested at various distances from the AP. The test was conducted for TCP and UDP protocols at 5, 15 and 50 feet away from the AP, making the 802.11b radio signals go through walls and around corners. The results of these tests are shown in Figure 4. Solid curves represent the throughput at different distances for TCP, and dotted curves represent the results for UDP at different distances.

Case 2: Wireless Transmissions
This test was conducted between Laptop1 and Laptop2. Data transmission occurred in a purely wireless environment. The results of these tests are shown in Figure 5.

Case 3: Power Test
This case was used to test throughput at different power levels at 50 feet away from the AP in the wireless environment between Laptop1 and Laptop2. The lines in Figure 6 indicate the performance results respectively at 30mW (dotted line) and 100mW (solid line) power levels.

Observation

The range tests were conducted on hybrid and wireless models. Looking at the results shown in Figures 4 and 5, we can easily conclude that as the distance was increased the throughput significantly dropped. In the case of hybrid model, the throughput was comparatively higher since half of the transmission was carried over a wired network. This proves that, as a client station moves away from the access point, the WLAN performance deteriorates due to decrease in signal strength.

In the case of power test, which was conducted at two power levels (the default 100mW and 30mW), the experiment indicated the higher the power level is the better the throughput. The default power on the access point and the wireless adapter cards is 100mW, which is the maximum possible power.

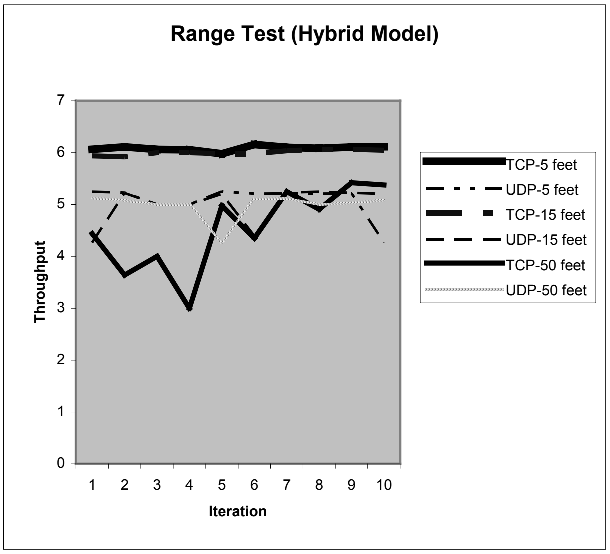

Figure 4: Result of Range Tests

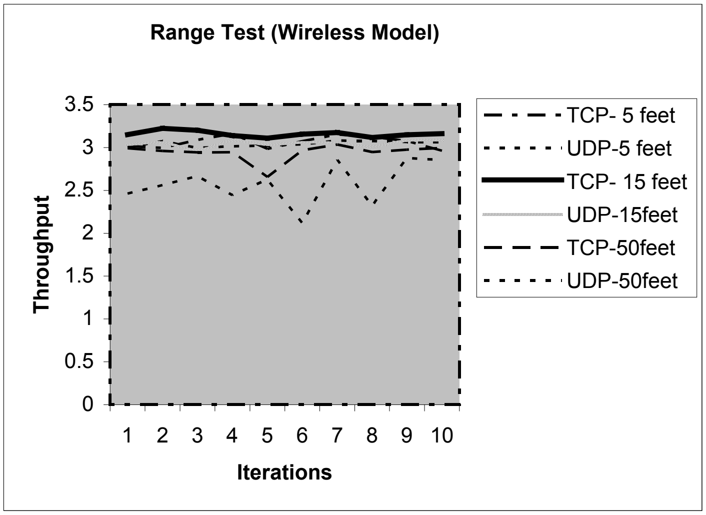

Figure 5: Result of Range Tests (wireless)

3.2. Effect of WEP

WEP is defined in the 802.11b standard as the optional encryption standard for WLANs. It is implemented in the 802.11 MAC layer and is based on RC4 encryption algorithm, which is a symmetric stream cipher where both the client station and the target station share the same key for both encryption and decryption. An Initialization vector (IV) is used to avoid encrypting two cipher texts with the same key stream and to produce a different RC4 key for each packet. However a lot of concerns were raised later regarding the effectiveness of WEP. Fluhrer, Mantin and Shamir described a passive ciphertext-only attack against RC4 as used in WEP [2]. Moreover, according to various other research publications, the vulnerability of WEP roots from its Initialization Vector and not from its smaller key size [3]. Later on an actual WEP connection was successfully attacked and the key was retrieved [4].

Test Setup

Configuration of the experiments to study the effect of WEP consisted of two wireless desktops and an access point (AP) in infrastructure mode. A socket based client-server Java program was used to measure the *round trip time* and the *throughput*. Each test sends the data back and forth 50 times (for every data size) and then the average value of *round trip time* is calculated for determining the *throughput*. Data sizes used were 1, 2, 5, 8, 11, 22 Mbits. Cisco Aironet 350 series client adapters support 40-bit and 128-bit static WEP keys. We used the 128-bit key while conducting the experiments. The tests were performed using different values of FTS and RTS with WEP disabled and then the same set of values with WEP enabled.

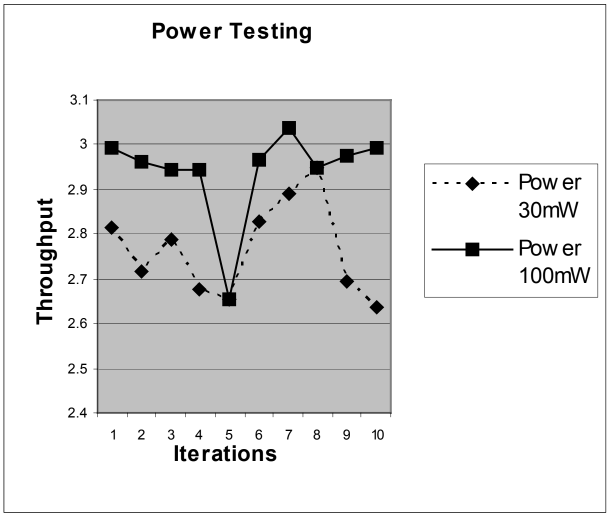

Figure 6: Result of Power Tests

Test Results

Figure 7 shows the effect of WEP on performance, when FTS = RTS = 2,312 bytes (the default). Similar results were obtained with different combinations of values of RTS and FTS

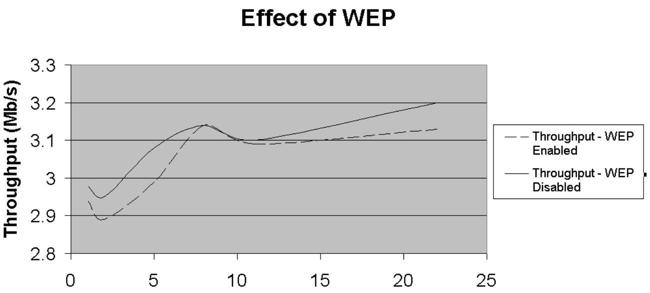

Figure 7. Effect of WEP on WLAN Performance

In Figure 7, the dashed line signifies the variations in throughput with WEP enabled, while the solid line indicates the throughput in the case without WEP being enabled. As shown in the figure, throughput essentially remains constant and the difference is less than 0.1 MB for the given data sizes and test conditions. Considering the volatile nature of wireless networks, this difference can be attributed towards occasional RF based interferences.

Observation

Encryption and decryption are typically resource-intensive and thus time-consuming processes. However as is clear from the above results, we can safely proclaim that WEP has virtually no effect on the overall throughput of WLANs. WEP is implemented in the 802.11 MAC layer. It is implemented as a firmware in the Cisco adapter. In addition to high performance, implementing the WEP algorithm as a firmware brings forth the advantage of easy updates. For example, improved security features for WEP, like Temporal Key Integrity Protocol (TKIP) and WPA (Wi-Fi Protected Access), can be applied simply by upgrading the firmware.

3.3. Effect of FTS

The 802.11 wireless stations can use the optional feature of *fragmentation* to divide a large data frame into smaller fragments, which are then sent independently to the destination. Fragmentation allows a network operator to define a MAC Service Data Unit (MSDU) across networks of varying MAC protocol Data Units (MPDU). Fragmentation is determined by the Fragmentation Threshold Parameter (FTS). FTS exists in the MAC layer. Unlike wired networks, most wireless network adapter cards allow us to change the FTS parameter. Any frames larger than the FTS value will be divided into smaller fragments. 802.11 networks suffer more interferences than traditional wired networks. Most of these are RF based interferences of short duration but of high-energy bursts. Thus sending greater number of smaller frames instead of smaller number of larger frames helps to lessen the data loss and retransmission delays especially in a wireless network experiencing heavy interference.

Test Setup

Configuration of the experiments testing the effect of FTP is the same as the WEP experiments, including the configuration of the devices, the performance calculation software, and the data sizes. The tests were performed by varying the value of FTS for the whole set of data sizes. The values of FTS included 256, 1000, 1500 and 2312 bytes. The value of RTS was kept at its default i.e., 2312 bytes and the WEP was disabled.

Test Results

As shown in Figure 8, for the given data sizes and test conditions, FTS is directly proportional to the throughput. The greater the value of FTS, the larger the throughput is.

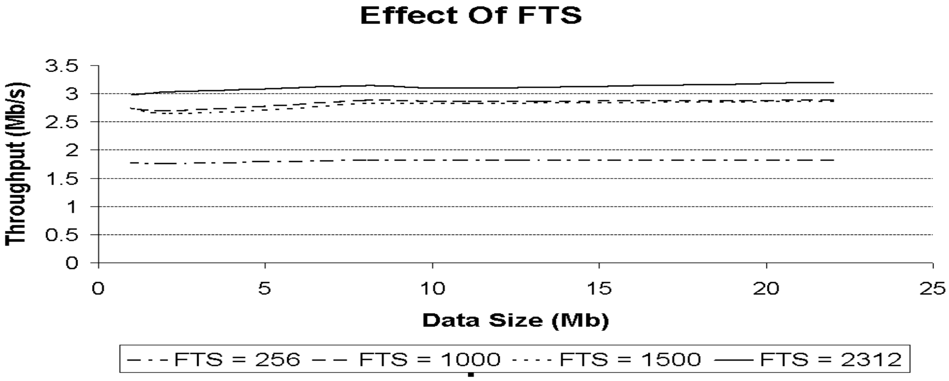

Figure 8. Effect of FTS on WLAN Performance

Observation

The default FTS value for Cisco Aironet 350 series client adapters is 2312 bytes. However depending on the network conditions and magnitude of interference, this value should be fine-tuned to achieve high performance. Higher value of FTS means there would be lesser overhead (since there would be lesser fragments); however its frame retransmission is more expensive. Lower value of FTS imparts greater fragmentation overhead; however it proves to be more efficient if the network is facing more interference, which may cause more retransmissions.

In our case the tests were performed in a controlled environment with limited RF interferences including interferences from other access points and wireless networks. The result was a lesser number of damaged frames and thus very low frame retransmission rate. This claim is validated by the obtained results. The throughput is much greater with a higher FTS value because the overhead of sending multiple smaller frames was not present. As we reduced the value of FTS, the throughput started to drop significantly. This is because for the same data size, the client adapter was sending greater number of frames. Thus the overhead was greatly increased. On the other hand, due to better network conditions, there were only few losses due to collisions and interference and although the retransmission delays in such cases were small, the increased overhead was far greater than retransmission savings.

3.4. Effect of RTS-CTS

This experiment tests the RTS/CTS mechanism which is part of the MAC layer specification. The RTS/CTS mechanism involves exchange of frames before the actual data is sent. The RTS and CTS frames contain a Duration ID field that defines the period of time that the medium is to be reserved to transmit the actual data frame and the returning ACK frame. All stations within the reception range of either the originating station (which transmits the RTS) or the destination station (which transmits the CTS) shall learn of the medium reservation. Thus a station may not be able to receive from the originating station, yet still know about the impending use of the medium to transmit a data frame.

Stations receiving the RTS or CTS will set their virtual Carrier Sense indicator NAV for the given duration, and will use this information together with the PHY Carrier Sense when sensing the radio medium. RTS/CTS attribute may be set on a per-station basis and on the AP.

Test Setup

Figure 9 shows the test setup used in this experiment.

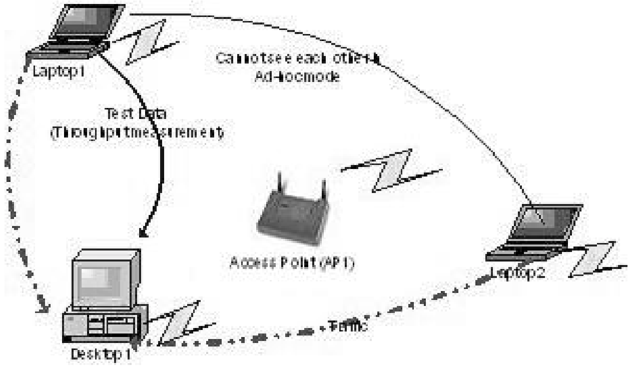
Laptop1
Cannot see each other in Ad-hoc mode
Test Data (Throughput measurement)
Access Point (AP1)
Laptop2
Traffic
Desktop1

Figure 9. Configuration for RTS/CTS Tests

In this experiment the motive was to simulate the *hidden node problem* in the infrastructure mode. Hidden node is the one which is outside the transmission range of the sender but within the range of the receiver. When a transmission has already begun, this node has no way to find this out and senses the medium to be idle and transmits its own data. However, as the node which was receiving the ongoing data is within the range, there will be a collision at the receiver.

Hidden node problem was mimicked using multiple machines, L1, L2 and D1 each capable of generating network traffic. L1 and L2 were placed in a manner where they weren't able to sense each other in the Ad-hoc mode; however they could communicate with each other via the AP1 in the Infrastructure mode as depicted in Figure 9.

In this experiment a sample client server java program was used to measure throughput between the stations by sending configurable size data packets from one station to another. This controlled piece of software was written to solely test the effect of RTS/CTS in the simulated hidden node problem.

Case A:

Base case with basic throughput measurement and without traffic generation. Setup involved L1 running the client java program, D1 running the server java program and L2 just transmitting management frames.

Case B:

In this case L2 was generating traffic in the form of file transfer aimed towards D1. D1 and L1 were running the client server programs similar to case-1.

Case C:

In this case both L1 and L2 were generating traffic towards D1. D1 and L1 assumed the same roles as the above two cases.

Test Result

Data size	**Case A**	**Case B**	**Case C**
512	120.15	6.42	3.57
1024	118.73	17.59	7.16
2048	126.54	25.33	10.75
4096	178.09	30.93	27.77
8192	157.07	58.73	47.05
Data size		**% of decrease from A to B**	**% of decrease from B to C**
512		95%	44%
1024		85%	59%
2048		80%	58%
4096		83%	10%
8192		63%	20%

Table 2: Result from RTS-CTS Experiments, with RTS Threshold = 2,312 bytes

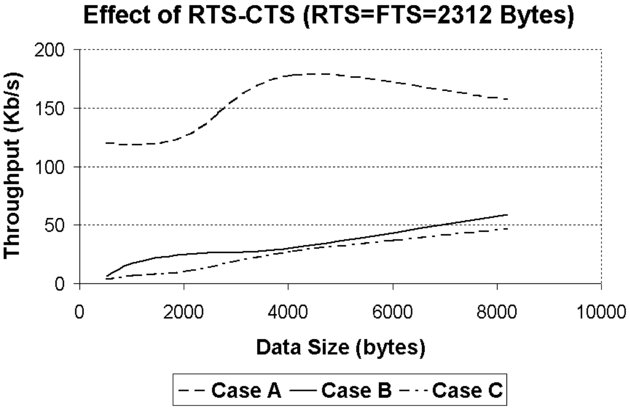

Figure 10. Effect of RTS-CTS at 2,312 Bytes

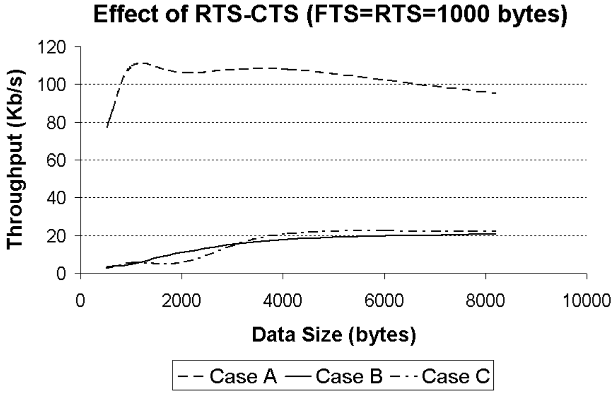

Figure 11. Effect of RTS-CTS at 1,000 Bytes

Observation

The RTS/CTS mechanism need not be used for every data frame transmission. Because the additional RTS and CTS frames add overhead inefficiency, the mechanism is not always justified, especially for short data frames.

It is clear from Figures 10 and 11 that RTS-CTS mechanism does have a significant effect upon the throughput, especially when the wireless network has higher probability of collisions. Consider Figure 10 where FTS = RTS = 2312 bytes; the data frames of size less than 2312 bytes (i.e. 512, 1024, 2048) were not proceeded by RTS packets in all the three cases. Therefore in table 1, if we consider the difference in throughput for case B (where only laptop2 was generating traffic towards desktop1) and case C (where both laptops were generating traffic thus effectively increasing the chances of collision), there is much significant drop in throughput for data sizes less than RTS threshold than the data sizes greater than RTS threshold. For data sizes greater than RTS threshold, throughput remained quite similar and the small drop in throughput can be attributed more towards the increase in traffic and thus resulting network congestion. Same arguments apply to Figure 11 where RTS = FTS = 1000 bytes.

Figure 12 shows the decrease of performance respectively from case A to B and from case B to C. In both scenarios, when data size is greater than the RTS Threshold, the drop in performance are less significant than when data size is less than the RTS Threshold. Another observation is that the performance drop is significant when traffic is introduced into the network (that is, from case A to B), but once traffic is introduced, the addition of more traffic (e.g., from case B to C) does not have as significant negative impact on the performance. The above two observations lead us to the conclusion that the performance of WLAN may not be satisfying when heavy traffic is present in the network.

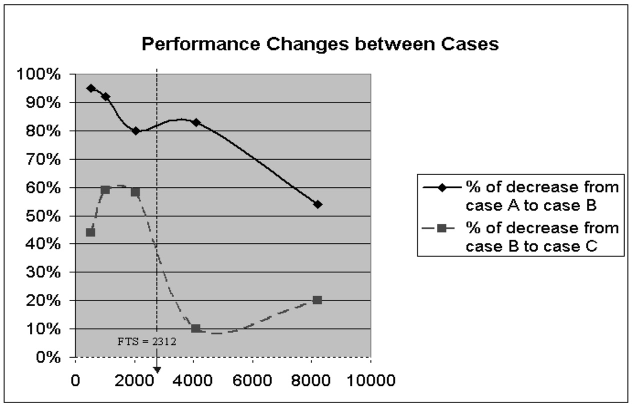

Figure 12. Change of Performance between Cases

4. Conclusion

In this paper we have studied the effect of various key parameters of IEEE 802.11b WLAN protocol on its performance in particular. The parameters we considered included: Distance and Power, Wired Equivalent Privacy (WEP), Fragmentation Threshold (FTS), and Request to Send / Clear to Send (RTS/CTS). For the Distance and Power we utilized two existing testing software on a hybrid network of WLAN and LAN. For the rest, we applied our own testing software, written in Java. A summary of our observation follows: For the Distance, the farther we are away from the AP, the lower is the performance. For the Power, the higher the power the better the performance is. The effect of WEP on the overall throughput was minimal. For FTS, the higher the value the smaller the overhead is, but in high interference environment with large re-transmissions, smaller value of FTS has advantage over large values. For RTS/CTS, we concluded that the overhead is high and it should not be used for transmission of every frame, unless the additional overhead is justified by the requirement of the application.

Acknowledgement

Zahur and Yang are partially supported by the Faculty Research Support Fund (#786, #817) of University of Houston Clear Lake (UHCL). Davari, Doctor and Yang were partially supported by the High Technology Laboratory (High Tech Lab) and by the Research Institute for Computing and Information Systems (RICIS) at UHCL. The High Tech Lab also provided some of the equipment used in the experiments.

References

[1] ANSI/IEEE, *802.11: Wireless LAN Medium Access Control (MAC) Physical Layer (PHY) Specifications*, 2000.

[2] S. Fluhrer, I. Mantin, and A. Shamir. Weaknesses in the key scheduling algorithm of RC4. In *Eighth Annual Workshop on Selected Areas in Cryptography*, Toronto, Canada, Aug. 2001.

[3] Jesse R. Walker, "Unsafe at any key size - an analysis of the WEP encapsulation", 802.11 Security Papers at NetSys.com, Oct. 27, 2000. http://www.netsys.com/library/papers/walker-2000-10-27.pdf

[4] Stubblefield, Ioannidis, Rubin, "Using the Fluhrer, Mantin, and Shamir attack to break WEP" http://philby.ucsd.edu/~bsy/ndss/2002/html/2002/papers/stubbl.pdf

[5] J. Vollbrecht, D. Rago, and R. Moskowitz. Wireless LAN Access Control and Authentication. A white paper from *Interlink Networks Resource Library*, 2001. http://www.interlinknetworks.com/images/resource/WLAN_Access_Control.pdf