

Wi-Fi in Ad Hoc Mode: A Measurement Study

Giuseppe Anastasi
Dept. of Information Engineering
University of Pisa
Via Diotisalvi 2 - 56122 Pisa, Italy
Email: g.anastasi@iet.unipi.it

Eleonora Borgia, Marco Conti, Enrico Gregori
Istituto IIT
National Research Council (CNR)
Via G. Moruzzi, 1 - 56124 PISA, Italy
Email: {firstname.lastname}@iit.cnr.it

Abstract

In this paper we investigate the performance of IEEE 802.11b ad hoc networks by means of an experimental study. This analysis reveals several aspects that are usually neglected in simulation studies. Firstly, since different transmission rates are used for control and data frames, different transmission ranges and carrier-sensing ranges may exist at the same time in the network. In addition, the transmission ranges are in practice much shorter than usually assumed in simulation analysis, not constant but highly variable (even in the same session) and depends on several factors (i.e., mobile height, interference condition, etc.). Finally, exploiting our performance measurements, we present a channel model for an 802.11 network that indicates virtual carrier sensing is generally not necessary and the RTS/CTS mechanism only introduces additional overhead.

1. Introduction

The IEEE 802.11 technology [1] is a good platform to implement single-hop ad hoc networks because of its extreme simplicity. Single-hop means that stations must be within the same transmission radius (say 100-200 meters) to be able to communicate. This limitation can be overcome by multi-hop ad hoc networking. This requires the addition of routing mechanisms at stations so that they can forward packets towards the intended destination, thus extending the range of the ad hoc network beyond the transmission radius of the source station. Routing solutions designed for wired networks (e.g., the Internet) are not suitable for the ad hoc environment, primarily due to the dynamic topology of ad hoc networks. Even though large-scale multi-hop ad hoc networks will not be

available in the near future, on smaller scales, mobile ad hoc networks are starting to appear thus extending the range of the IEEE 802.11 technology over multiple radio hops. Most of the existing IEEE 802.11-based ad hoc networks have been developed in the academic environment, but recently even commercial solutions have been proposed (see, e.g., MeshNetworks¹ and SPANworks²).

The characteristics of the wireless medium and the dynamic nature of ad hoc networks make (IEEE 802.11) multi-hop networks fundamentally different from wired networks. Furthermore, the behavior of an ad hoc network that relies upon a carrier-sensing random access protocol, such as the IEEE 802.11, is further complicated by the presence of hidden stations, exposed stations, “capturing” phenomena [2, 3], and so on. The interactions between all these phenomena make the behavior of IEEE 802.11 ad hoc networks very complex to predict. Recently, this has generated an extensive literature related to the performance analysis of the 802.11 MAC protocol in the ad hoc environment. Most of these studies have been done through simulation. To the best of our knowledge, only very few experimental analysis have been conducted. For this reason, in the paper we extend the 802.11 performance analysis with an extensive set of measurements that have been conducted on a real testbed. The measurements were done in an outdoor environment, by considering different traffic types (i.e., TCP and UDP traffics). For the sake of comparison with the previous studies, our analysis is mostly related to the basic IEEE 802.11 MAC protocol (i.e., we consider a data rate of 2 Mbps). However, some results related to IEEE 802.11b are also included. Our experimental results indicate that transmission ranges are much shorter than assumed in simulation studies. In addition, we also observed a dependency of transmission range on transmission speed and mobile device’s height. Using our experimental results we present an innovative channel model for 802.11

[‡] Work carried out under the financial support of the FET IST Mobile MAN Project and the Italian Ministry for Education and Scientific Research (MIUR) in the framework of the Projects VICOM and PERF.

¹ <http://www.meshnetworks.com>

² <http://www.spanworks.com>

networks. This model indicates that virtual carrier sensing is not necessary and the RTS/CTS mechanism only introduces additional overhead. In addition this model is used to identify new hidden stations or capture phenomenon that are not solved by current 802.11 mechanisms. The paper is organized as follows. In Section 2 the main problems of a wireless ad hoc network are briefly discussed. The experimental analysis of 802.11b network is presented in Section 3 and 4, respectively. Our conclusions are summarized in Section 5.

2. Common Problems in Wireless Ad Hoc Networks

In this section we shortly discuss the main problems that arise in wireless ad hoc networks. A detailed discussion can be found in [4].

The characteristics of the wireless medium make wireless networks fundamentally different from wired networks. Specifically, as indicated in [5]:

- the wireless medium has neither absolute nor readily observable boundaries outside of which stations are known to be unable to receive network frames;
- the channel is unprotected from outside signals;
- the wireless medium is significantly less reliable than wired media;
- the channel has time-varying and asymmetric propagation properties.

In wireless (ad hoc) networks that rely upon a carrier-sensing random access protocol, like the IEEE 802.11, the wireless medium characteristics generate complex phenomena such as the hidden station and the exposed station problems.

Figure 1 shows a typical “hidden station” scenario. Let us assume that station B is in the transmitting range of both A and C, but A and C cannot hear each other. Let us also assume that A is transmitting to B. If C has a frame to be transmitted to B, according to the *Distributed Coordination Function (DCF)* protocol, it senses the medium and finds it free because it is not able to hear A’s transmissions. Therefore, it starts transmitting the frame but this transmission will result in a collision at the destination Station B.

The hidden station problem can be alleviated by extending the basic mechanism by a **virtual carrier sensing** mechanism (also referred to as floor acquisition mechanism) that is based on two control frames: *Request To Send (RTS)* and *Clear To Send (CTS)*, respectively.




Figure 1. The “hidden station” problem

According to this mechanism, before transmitting a data frame, the source station sends a short control frame, named RTS, to the receiving station announcing the upcoming frame transmission (see Figure 2). Upon receiving the RTS frame, the destination station replies by a CTS frame to indicate that it is ready to receive the data frame. Both the RTS and CTS frames contain the total duration of the transmission, i.e., the overall time interval needed to transmit the data frame and the related ACK. This information can be read by any station within the transmission range of either the source or the destination station. Such a station uses this information to set up a timer called *Network Allocation Vector (NAV)*. While the NAV timer is greater than zero the station must refrain from accessing the wireless medium. By using the RTS/CTS mechanism, stations may become aware of transmissions from hidden stations and on how long the channel will be used for these transmissions.




Figure 2. Virtual Carrier Sensing mechanism

Figure 3 depicts a typical scenario where the “exposed station” problem may occur. Let us assume that Station A and Station C can hear transmissions from B, but Station A can not hear transmissions from C. Let us also assume that Station B is transmitting to Station A and Station C receives a frame to be transmitted to D. According to the DCF protocol, C senses the medium and finds it busy because of B’s transmission. Therefore, it refrains from transmitting to D although this transmission would not cause a collision at A. The “exposed station” problem may thus result in a throughput reduction.




Figure 3. The “exposed station” problem

2.1. Simulation Analysis of IEEE 802.11 Ad Hoc Networks

The performance provided by the 802.11 MAC protocol in an ad hoc environment have been extensively analyzed via simulation. The studies presented in the literature have pointed out several performance problems. They can be summarized as follows. In a dynamic environment, mobility may have a severe impact on the performance of the TCP protocol [6, 7, 8, 9, 10, 11, 12]. However, even when stations are static, the performance of an ad hoc network may be quite far from ideal. It is highly influenced by the operating conditions, i.e., TCP parameter values (primarily the congestion window size) and network topology [13, 14]. In addition, the interaction of the 802.11 MAC protocol (hidden and exposed station problems, exponential back-off scheme, etc.) with TCP mechanisms (congestion control and time-out) may lead to unexpected phenomena in a multi-hop environment. For example, in the case of simultaneous TCP flows, severe unfairness problems and - in extreme cases - capture of the channel by few flows [2, 3, 15, 16] may occur. Even in the case of a single TCP connection, the instantaneous throughput may be very unstable [2, 3]. Such phenomena do not appear, or appear with less intensity, when the UDP protocol is used.

3. Experimental Analysis of IEEE 802.11 Ad Hoc Networks

All these previous analysis were carried out using simulation tools (GloMosim [17], *ns-2* [18], Qualnet [19] etc.), and thus the results observed are highly dependent on the physical layer model implemented in the simulation tool. Hereafter, we extend these results by presenting analyses carried on a real testbed. The simulation results presented in the literature were obtained by considering WaveLAN IEEE 802.11 network cards (operating at the nominal bit rate of 2Mbps). Currently, however, the Wi-Fi network interfaces are becoming more and more popular. Wi-Fi cards implement the IEEE 802.11b standard. It is therefore important to

extend the previous studies to IEEE 802.11b ad hoc networks.

The 802.11b standard extends the 802.11 standard by introducing a higher-speed Physical Layer in the 2.4 GHz frequency band still guaranteeing the interoperability with 802.11 cards. Specifically, 802.11b enables transmissions at 5.5 Mbps and 11 Mbps, in addition to 1 Mbps and 2 Mbps. 802.11b cards may implement a dynamic rate switching with the objective of improving performance. To ensure coexistence and interoperability among multirate-capable stations, and with 802.11 cards, the standard defines a set of rules that must be followed by all stations in a WLAN. Specifically, for each WLAN is defined a *basic rate set* that contains the data transfer rates that all stations within the WLAN will be capable of using to receive and transmit.

To support the proper operation of a WLAN, all stations must be able to detect control frames. Hence, RTS, CTS, and ACK frames must be transmitted at a rate included in the basic rate set. In addition, also frames with multicast or broadcast destination addresses must be transmitted at a rate belonging to the basic rate set. These differences in the rates used for transmitting (unicast) data and control frames have a big impact on the system behavior as clearly pointed out in [20].

Actually, since 802.11 cards transmit at a constant power, lowering the transmission rate permits the packaging of more energy per symbol, and this leads to an increased transmission range. In the next subsections we investigate, by a set of experimental measurements,

- i) the relationship between the transmission rate of the wireless network interface card (NIC) and the maximum throughput (two-stations experiments);
- ii) the relationship between the transmission and carrier sensing range and the transmission rate (two-stations experiments);
- iii) Hidden and/or exposed station situations (four-stations experiments).

To better understand the results presented below, it is useful to provide a model of the relationships existing among stations when they transmit or receive. In particular, it is useful to make a distinction between the transmission range, the interference range and the carrier sensing range. The following definitions can be given.

- The *Transmission Range* (TX_range) is the range (with respect to the transmitting station) within which a transmitted frame can be successfully received. The transmission range is mainly determined by the transmission power and the radio propagation properties.

- The *Physical Carrier Sensing Range* (PCS_range) is the range (with respect to the transmitting station) within which the other stations detect a transmission. It mainly depends on the sensitivity of the receiver (the receive threshold) and the radio propagation properties.
- The *Interference Range* (IF_range) is the range within which stations in receive mode will be "interfered with" by a transmitter, and thus suffer a loss. The interference range is usually larger than the transmission range, and it is function of the distance between the sender and receiver, and of the path loss model.

In the previous simulation studies the following relationship has been generally assumed:

$$TX_range \leq IF_range \leq PCS_range$$

For example, in the ns-2 simulation tool [18] the following values are used to model the characteristics of the physical layer:

$$TX_range = 250m, \\ IF_range = PCS_range = 550m$$

In addition, the relationship between TX_range , PCS_range , and IF_range are assumed to be constant throughout a simulation experiment. On the other hand, from our measurements we have observed that the physical channel has time-varying and asymmetric propagation properties and, hence, the value of TX_range , PCS_range , and IF_range may be highly variable in practice.

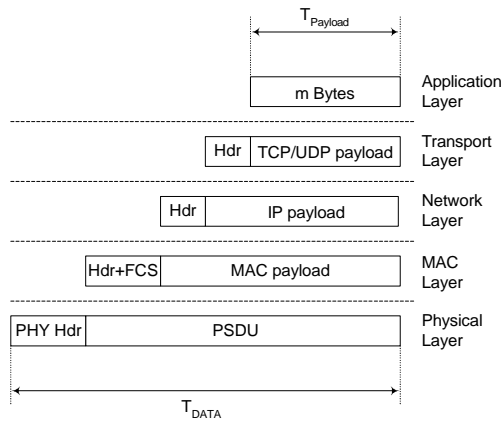


Figure 4. Encapsulation overheads

3.1. Available Bandwidth

In this section we will show that only a fraction of the 11 Mbps nominal bandwidth of the IEEE 802.11b cards can be used for data transmission. To this end we need to

carefully analyze the overheads associated with the transmission of each packet (see Figure 4). Specifically, each stream of m bytes generated by a legacy Internet application is encapsulated in the TCP/UDP and IP protocols that add their headers before delivering the resulting IP datagram to the MAC layer for transmission over the wireless medium. Each MAC data frame is made up of: *i*) a *MAC header*, say MAC_{hdr} , containing MAC addresses and control information³, and *ii*) a variable length *data payload*, containing the upper layers data information. Finally, to support the physical procedures of transmission (carrier sense and reception) a *physical layer preamble* (PLCP preamble) and a *physical layer header* (PLCP header) have to be added to both data and control frames. Hereafter, we will refer to the sum of PLCP preamble and PLCP header as PHY_{hdr} .

It is worth noting that these different headers and data fields are transmitted at different data rates to ensure the interoperability between 802.11 and 802.11b cards. Specifically, the standard defines two different formats for the PLCP: Long PLCP and Short PLCP. Hereafter, we assume a Long PLCP that includes a 144-bit preamble and a 48-bit header both transmitted at 1 Mbps while the MAC_{hdr} and the $MAC_{payload}$ can be transmitted at one of the NIC data rates: 1, 2, 5.5, and 11 Mbps. In particular, control frames (RTS, CTS and ACK) can be transmitted at 1 or 2 Mbps, while data frame can be transmitted at any of the NIC data rates.

By taking into considerations the above quantities, a detailed analysis of the maximum expected throughput for a single active session (i.e., only a sender-receiver couple active) with UDP and TCP traffic can be found in [21]. The obtained results will be presented in the next sections. They depend on the specific setting of the IEEE 802.11b protocol parameters. Table 1 gives the values for the protocol parameters used hereafter.

Table 1. IEEE 802.11b parameter values

$Slot_Time$	τ	PHY_{hdr}	MAC_{hdr}	$Bit\ Rate$ (Mbps)
20 μsec	$\leq 1 \mu sec$	192 bits ($9.6 t_{slot}$)	272 bits	1, 2, 5.5, 11
$DIFS$	$SIFS$	ACK	CW_{MIN}	CW_{MAX}
50 μsec	10 μsec	112 bits + PHY_{hdr}	$32 t_{slot}$	$1024 t_{slot}$

In Table 2 and Table 3 we report the expected throughputs (with and without the RTS/CTS mechanism) by assuming that the NIC is transmitting at a constant data rate equal to 1, 2, 5.5 or 11 Mbps, respectively for a UDP

³ Without any loss of generality we have considered the *frame error sequence* (FCS), for error detection, as belonging to the MAC header.

and TCP connection, and assuming a data packet size at the application level equal to $m=512$ and $m=1024$ bytes.

Table 2. Maximum throughput in Mbit/sec (Mbps) at different data rates for a UDP connection

	m= 512 Bytes		m= 1024 Bytes	
	No RTS/CTS	RTS/CTS	No RTS/CTS	RTS/CTS
	11 Mbps	3.337 Mbps	2.739 Mbps	5.120 Mbps
5.5 Mbps	2.490 Mbps	2.141 Mbps	3.428 Mbps	3.082 Mbps
2 Mbps	1.319 Mbps	1.214 Mbps	1.589 Mbps	1.511 Mbps
1 Mbps	0.758 Mbps	0.738 Mbps	0.862 Mbps	0.839 Mbps

Table 3. Maximum throughput in Mbit/sec (Mbps) at different data rates for a TCP connection

	m= 512 Bytes		m= 1024 Bytes	
	No RTS/CTS	RTS/CTS	No RTS/CTS	RTS/CTS
	11 Mbps	2.456 Mbps	2.917 Mbps	2.739 Mbps
5.5 Mbps	1.931 Mbps	0.840 Mbps	1.979 Mbps	0.621 Mbps
2 Mbps	1.105 Mbps	0.997 Mbps	1.423 Mbps	1.330 Mbps
1 Mbps	0.661 Mbps	0.620 Mbps	0.796 Mbps	0.766 Mbps

As shown in Table 2, only a small percentage of the 11 Mbps nominal bandwidth can be really used for data transmission. This percentage increases with the payload size. However, even with large packets sizes (e.g., $m=1024$ bytes) the bandwidth utilization is lower than 44%.

The above theoretical analysis has been complemented with the measurements of the actual throughput at the application level. Specifically, we have considered two types of applications: ftp and CBR. In the former case the TCP protocol is used at the transport layer, while in the latter case the UDP is adopted. In both cases the applications operate in asymptotic conditions (i.e., they always have packets ready for transmission) with constant size packets of 512 bytes.

The results obtained from this experimental analysis are reported in Table 4. The experimental results related to the UDP traffic are very close to the maximum throughput computed analytically. As expected, in the presence of TCP traffic the measured throughput is much lower than the theoretical maximum throughput.

Similar results have been also obtained by comparing the maximum throughput derived from the theoretical analysis and the real throughputs measured when the NIC data rate is set to 1, 2 or 5.5 Mbps.

Table 4. Comparison between the theoretical maximum throughput and the actual throughput achieved by TCP/UDP applications at 11Mbps

	No RTS/CTS		RTS/CTS	
	Ideal	Real	Ideal	Real
UDP	3.337 Mbps	2.917 Mbps	2.739 Mbps	2.221 Mbps
TCP	2.456 Mbps	0.840 Mbps	1.979 Mbps	0.621 Mbps

3.2. Transmission Ranges

The dependency between the data rate and the transmission range was investigated by measuring the packet loss rate experienced by two communicating stations whose network interfaces transmit at a constant (preset) data rate. Specifically, four sets of measurements were performed corresponding to the different data rates: 1, 2, 5.5, and 11 Mbps. In each set of experiments the packet loss rate was recorded as a function of the distance between the communicating stations. The resulting curves are presented in Figure 5. We also experienced a high variability in the transmission range depending on the weather conditions: for example measuring the transmission ranges at 1 Mbps in two different days the results differ of about 20 meters.

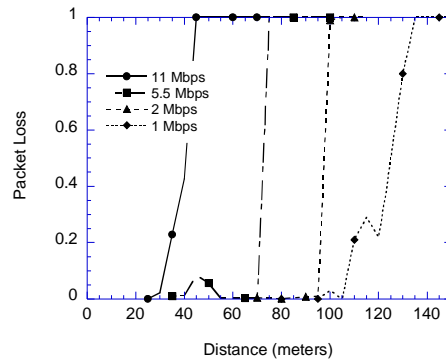


Figure 5. Packet loss rate as a function of the distance between communicating stations for different data rates

The results presented in Figure 5 are summarized in Table 5 where the estimates of the transmission ranges at different data rates are reported. These estimates point out that, when using the highest bit rate for the data transmission, there is a significant difference in the transmission range of control and data frames, respectively. For example, assuming that the RTS/CTS mechanism is active, if a station transmits a frame at 11 Mbps to another station within its transmission range (i.e., less than 30 m apart) it reserves the channel for a radius of

approximately 90 (120) m around itself. The RTS frame is transmitted at 2 Mbps (or 1 Mbps), and hence, it is correctly received by all stations within station transmitting range, i.e., 90 (120) meters.

Again, it is interesting to compare the transmission range used in the most popular simulation tools, like ns-2 and Glomosim, with the transmission ranges measured in our experiments. In these simulation tools it is assumed $TX_range = 250m$. Since the above simulation tools only consider a 2 Mbps bit rate we make reference to the transmission range estimated with a NIC data rate of 2 Mbps. As it clearly appears, the value used in the simulation tools (and, hence, in the simulation studies based on them) is 2-3 times higher than the values measured in practice. This difference is very important for example when studying the behavior of routing protocols: the shorter is the TX_range , the higher is the frequency of route re-calculation when the network stations are mobile.

Table 5. Estimates of the transmission ranges at different data rates

	11 Mbps	5.5 Mbps	2 Mbps	1 Mbps
Data	30 meters	70 meters	90-100 meters	110-130 meters
TX_range				
Control			≈ 90 meters	≈ 120 meters
TX_range				

3.3. Transmission Ranges and the Mobile devices' Height

During the experiments we performed to analyze the transmission ranges at various data rates, we observed a dependence of the transmission ranges from the mobile devices' height from the ground. Specifically, in some cases we observed that while the devices were not able to communicate when located on the stools, they started to exchange packets by lifting them up. In this section we present the results obtained by a careful investigation of this phenomenon. Specifically, we studied the dependency of the transmission ranges from the devices' height from the ground. To this end we measured the throughput between two stations⁴ as a function of their height from the ground: four different heights were considered: 0.40 m, 0.80 m, 1.2 m and 1.6 m. The experiments were performed with the Wi-Fi card set at two different transmission rates: 2 and 11 Mbps. In each set of experiments the distance between the two devices was set close to guarantee that the receiver is always inside the sender transmission range. Specifically, the sender-receiver distance was equal to 30 and 70 m. when the cards operated at 11 and 2 Mbps, respectively.

⁴ In these experiments UDP is used as the transport protocol.

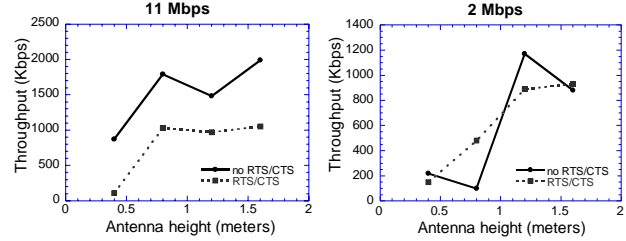


Figure 6. Relationship between throughput and devices' height

As it clearly appears in Figure 6, the ground height may have a big impact on the quality of the communications between the mobile devices. For example, at 11 Mbps, by lifting up the devices from 0.40 meters to 0.80 meters the throughput doubles, while further increasing the height does not produce significant throughput gains. A similar behavior is observed with a 2 Mbps transmission rate, however in this case the major throughput gain is obtained lifting up the devices from 0.80 meters to 1.20 meters. A possible explanation of this difference is related to the distances, in the two cases, between the communicating devices. This intuition is confirmed by the work presented in [22] that provides a theoretical framework to explain the height impact on IEEE 802.11 channel quality. Specifically, the channel power loss depends on the contact between the Fresnel zone and the ground. The Fresnel zone for a radio beam is an elliptical area with foci located in the sender and the receiver. Objects in the Fresnel zone cause diffraction and hence reduce the signal energy. In particular, most of the radio-wave energy is within the First Fresnel Zone, i.e., the inner 60% of the Fresnel zone. Hence, if this inner part contacts the ground (or other objects) the energy loss is significant. Figure 7 shows the Fresnel zone (and its inner 60%) for a sender-receiver couple at a distance D . In the figure, R_1 denotes the height of the First Fresnel Zone. As shown in [22] R_1 is highly dependent on the stations distance. For example, when the sender and the receiver are at an height of 1 meter from the ground, the First Fresnel Zone has a contact with the ground only if $D > 33$ meters. While at heights of 1.5 and 2 meters the First Fresnel Zone contacts the ground only if D is greater than 73 and 131 meters, respectively. These theoretical computations are aligned with our experimental results.




Figure 7. The Fresnel Zone

3.4. Four-Stations Network Configurations

The results presented in the previous sections show that the IEEE 802.11b behavior is more complex than the behavior of the IEEE 802.11 standard. Indeed the availability of different transmission rates may cause the presence of several transmission ranges inside the network. In particular, inside the same data transfer session there may be different transmission ranges for data and control frame (e.g., RTS, CTS, ACK). Hereafter, we show that the superposition of these different phenomena makes very difficult to understand the behavior of IEEE 802.11b ad hoc networks. To reduce this complexity, in the experiments presented below the NIC data rate is set to a constant value equal to 11 Mbps for the entire duration of the experiment.⁵

The network configuration is shown in Figure 8 and the related results are presented in Figure 9.




Figure 8. Network configuration at 11 Mbps

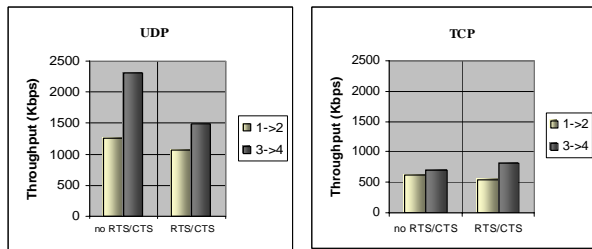


Figure 9. Throughputs at 11 Mbps

The results show that dependencies exist between the two connections even though the transmission range is smaller than the distance between stations S1 and S3. In detail, the throughput experienced by each session is much smaller than the throughput obtained by a session in isolation, e.g., about 3.3 Mbps with UDP (see Table 4).

Furthermore, the dependency exists also when the basic mechanism (i.e., no RTS/CTS) is adopted.⁶ To summarize, these experiments show that *i*) interdependencies among the stations extends beyond the

transmission range; *ii*) our hypothesis is that the physical carrier sensing range, including all the four stations, produces a correlation between active connections and its effect is similar to that achieved with the RTS/CTS mechanism (virtual carrier sensing). The difference in the throughputs achieved by the two sessions when using the UDP protocol (with or without RTS/CTS) can be explained by considering the asymmetric condition that exists on the channel: station S2 is exposed to transmissions of station S3 and, hence, when station S1 sends a frame to S2 this station is not able to send back the MAC ACK. Therefore, S1 reacts as in the collision cases (thus re-scheduling the transmission with a larger backoff). It is worth pointing out that also S3 is exposed to S2 transmissions but the S2's effect on S3 is less marked given the different role of the two stations. When using the basic access mechanism, the S2's effect on S3 is limited to short intervals (i.e., the transmission of ACK frames). When adopting the RTS/CTS mechanism, the S2 CTS forces S3 to defer the transmission of RTS frames (i.e., simply a delay in the transmission), while RTS frames sent by S3 forces S2 to not reply with a CTS frame to S1's RTS. In the latter case, S1 increases the back off and reschedules the transmission. Finally, when the TCP protocol is used the differences between the throughputs achieved by the two connections still exist but are reduced. The analysis of this case is very complex because we must also take into consideration the impact of the TCP mechanisms that: *i*) reduces the transmission rate of the first connection, and *ii*) introduces the transmission of TCP-ACK frames (from S2 and S4) thus contributing to make the system less asymmetric.

3.5. Physical Carrier Sensing Range

The results presented in the previous section seem to indicate that dependencies among the stations extend far beyond the transmission range. For example, taking as a reference the scenario presented in Figure 8, the distance between the two couples of transmitting stations is about three times the transmission range. The hypothesis is that dependencies are due to a large physical carrier sensing range that includes all the stations. To validate this hypothesis and to better understand the system behavior we designed some experiments to estimate the physical carrier sensing range. A direct measure of this quantity seems difficult to achieve because the 802.11b cards we utilized do not provide to the higher layers information about the channel carrier sensing. Therefore, we defined an indirect way to perform these measurements. We utilized the scenario shown in Figure 10 with fixed distance between each couple of communicating stations ($d(1,2)=d(3,4)=10$ meters), and variable distance between the two couples (i.e., $d(2,3)$ is variable).

⁵ It is worth pointing out that we experienced a high variability in the channel conditions thus making a comparison between the results difficult.

⁶ A similar behavior is observed (but with different values) by adopting the RTS/CTS mechanism.




Figure 10. Reference network scenario

The idea is to increase $d(2,3)$ until no correlation is measured between the two sessions. To quantify the correlation degree between the two sessions we measured the throughput of each session in isolation, i.e., when the other session is not active. Then we measured the throughput achieved by each session when both sessions are active. Obviously, no correlation exists when the aggregate throughput is equal to the sum of the throughput of the two sessions in isolation. By varying the distance $d(2,3)$ we performed several experiments. The results were obtained with the cards transmission rates set to 11 and 2 Mbps, and they are summarized in Table 6. As it clearly appears from the table, there are two steps in the aggregate throughput: one after 180 m and the other after 250 m. This behavior can be explained as follows. Taken a session as a reference, the presence of the other session may have two possible effects on the performance of the reference session: 1) if the two sessions are within the same physical carrier sensing range, they share the same physical channel; 2) if they are outside the physical carrier sensing range the radiated energy from one session may still affect the quality of the channel observed by the other session. As the radiated energy may travel over unlimited distances, we can expect that the second effect completely disappears only for very large distances between the sessions [20]. Hence, we can assume that the first step coincides with the end of the physical carrier sensing range, while the second one occurs even when the second effect becomes almost negligible.

Table 6. Aggregate Throughput vs. Distance between sessions at 11 Mbps and 2 Mbps (No RTS/CTS, payload size = 512 byte)

Distance (meters)	11 Mbps		2 Mbps	
	Independent sessions	Aggregate throughput	Independent sessions	Aggregate throughput
150	5800 Kbps	3750 Kbps	2600 Kbps	1660 Kbps
180	5800 Kbps	3790 Kbps	2600 Kbps	1750 Kbps
200	5800 Kbps	4590 Kbps	2600 Kbps	2160 Kbps
250	5800 Kbps	4460 Kbps	2600 Kbps	1920 Kbps
300	5800 Kbps	5260 Kbps	2600 Kbps	2290 Kbps
350	5800 Kbps	5520 Kbps	2600 Kbps	2370 Kbps

It is worth noting that the physical carrier sensing range is almost the same for the two different transmission rates. Indeed the physical carrier sensing range mainly depends only on two parameters: the

stations' transmitting power and the distance between transmitting stations. The rates at which data are transmitted have no significant effect on these parameters.

The results obtained confirm the hypotheses we made above to justify the apparent dependencies existing between the two couples of transmitting stations even if the distance among them is about three times greater than the transmission range (see, for example, Figure 9).

3.6. Channel Model for an IEEE 802.11b Ad Hoc Network

The results presented in this paper indicate that to correctly understand the behavior of an 802.11b network operating in ad hoc mode, several different ranges must be considered.

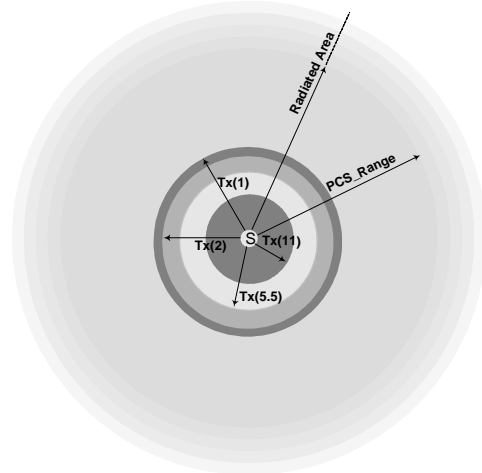


Figure 11. Channel model for an 802.11 ad hoc network

Specifically, as shown in Figure 11, given a transmitting station S , the stations around will be affected by the station S transmissions in a different way depending on the distance from S and the rate used by S for its transmissions.

Specifically, assuming that S is transmitting with a rate x ($x \in \{1, 2, 5.5, 11\}$) stations around it can be partitioned into three classes depending on their distance, d , from S :

- i. Stations at a distance $d < TX_Range(x)$ are able to correctly receive data from S , if S is transmitting at a rate lower or equal to x ;
- ii. Stations at a distance d , where $TX_Range(x) < d < PCS_Range$, are not able to receive data correctly from station S . However, as they are in the S physical carrier sensing range, when S is transmitting they observe the channel busy and thus they defer their transmissions;

iii. Stations at a distance $d > \text{PCS_Range}$ do not measure any significant energy on the channel when S is transmitting, therefore they can start transmitting contemporarily to S; however, the quality of the channel they observe may be affected by the energy radiated by S. In addition, if $d < \text{PCS_Range} + \text{TX_Range}(x)$ some interference phenomena may occur (see below). This interference depends on the IF_Range value. This value is difficult to model and evaluate as it depends on several factors (mainly the power at the receiving site) but as explained before $\text{TX_Range}(1) < \text{IF_Range} < \text{PCS_Range}$.

Several interesting observations can be derived by taking into consideration points i-iii above. Firstly, the hidden station phenomenon, as it is usually defined in the literature (see Section 2), is almost impossible with the ranges measured in our experiments. Indeed, the PCS_Range is more than twice $\text{TX_Range}(1)$, i.e., the larger transmission range. Furthermore, two stations, say S1 and S2, that can start transmitting towards the same receiver, R, must be at a distance $\leq 2 \cdot \text{TX_Range}(1)$, and thus they are inside the physical carrier sensing range of each other. Hence, if S1 has an ongoing transmission with R, S2 will observe a busy channel and thus will defer its own transmission. This means that, in this scenario, virtual carrier sensing is not necessary and the RTS/CTS mechanism only introduces additional overhead.

While the hidden station phenomenon, as defined in the literature, seems not relevant for this environment point iii above highlights that packets cannot be correctly received due to the interference caused by a station that is “hidden” to the sending station. An example of this type of *hidden station phenomenon* is presented in Figure 12. In this figure we have two transmitting stations, S and S1 that are outside their respectively PCS_Range and hence they are hidden to each other. In addition we assume that the receiver of station S (denoted by R in the figure) is inside the interference range (IF_Range) of station S1. In this scenario S and S1 can be simultaneously transmitting and, if this occurs, station R cannot receive data from S correctly. Also in this case the RTS/CTS mechanism does not provide any help and new coordination mechanisms need to be designed to extend the coordination in the channel access beyond the PCS_Range . It is worth noting that, in our channel model, the exposed station definition (see Figure 3) must be modified too. In this scenario, exposed stations are those stations at a distance $\text{PCS_Range} - \text{TX_Range}(1) < d < \text{PCS_Range}$. Indeed, these stations are exposed to station S transmissions, while they are in the transmission range of stations with $d > \text{PCS_Range}$. The following example outline problems that may occur in this case.

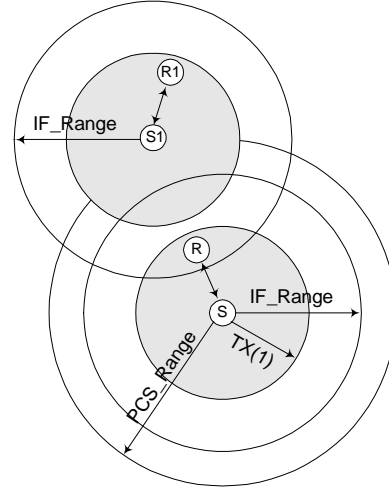


Figure 12. Interference-based hidden station phenomenon

Let us denote with S1 a station at a distance d from S: $\text{PCS_Range} < d < \text{PCS_Range} + \text{TX_Range}(x)$. Station S1 can start transmitting, with a rate x , towards a station E that is inside the physical carrier sensing of S; station E cannot reply because it observes a busy channel due to the ongoing station S transmissions, i.e., E is exposed to station S. Since station S1 does not receive any reply (802.11 ACK) from E, it assumes an error condition (collision or CRC error condition), hence it backoffs and then tries again. If this situation repeats for several times (up to 7), S1 assumes that E is not anymore in its transmission range, gives up the transmission attempt and (wrongly) signals to the higher layer a link breakage condition, thus forcing higher layers to attempt a recovery action (e.g., new route discovery, etc. – see Section 2.1).

To summarize, results obtained in the configuration we analyzed indicate that the hidden station and exposed station definitions must be extended. These new hidden-station and exposed-station phenomena may produce undesirable effects that may degrade the performance of an ad hoc network, mainly if the TCP protocol is used. Extending the coordination in the channel access beyond the PCS_Range seems to be the correct direction for solving the above problems.

4. Conclusions

In this paper we have investigated the performance of IEEE 802.11b ad hoc networks. Previous studies in this framework have pointed out that the behavior of IEEE 802.11b ad hoc networks are complicated by the presence of hidden stations, exposed stations, “capturing” phenomena, and so on. Most of these studies have been done through simulation. In this paper we have extended the 802.11 performance analysis with an extensive set of

measurements performed on a real testbed by considering IEEE 802.11b cards. This analysis has pointed out several aspects that are commonly neglected in simulation studies. Specifically, the system behavior is very complex as several transmission and carrier-sensing ranges exist at the same time on the channel. For example, the physical level preamble is always transmitted at 1 Mbps, the signaling frames can be transmitted up to 2 Mbps, while the data frames can be transmitted up to 11 Mbps. As a consequence of this complex behavior, with IEEE 802.11b, we never observed capture phenomena (i.e., it never happened that a couple of stations monopolized the channel). Finally, in simulation studies the transmission and carrier sensing ranges are quite large and constant for the entire duration of the experiment. On the other hand, in our testbed we have observed that the transmission and physical sensing ranges are much shorter than assumed in simulation studies, and highly variable even in the same session in time and space, depending on several factors (weather condition, place and time of the experiment, etc.). From these experimental results, we have developed an innovative channel model for 802.11b networks. This model indicates that the RTS/CTS mechanism does not solve the hidden and exposed station problems. The conditions for the existence of these new hidden and exposed stations are identified in the paper and depend on the current values of the TX_range, IF_range and PCS_range. We are currently working on the identification of mechanisms to solve the new phenomena. The most promising direction is to extend the coordination in the channel access beyond the PCS_Range.

References

- [1] Official Homepage of the IEEE 802.11 Working Group, <http://grouper.ieee.org/groups/802/11/>.
- [2] S. Xu and T. Saadawi, "Does the IEEE 802.11 MAC protocol Work Well in Multihop Wireless Ad Hoc Networks?", *IEEE Comm Magazine*, Volume 39, N. 6, June 2001, pp. 130-137.
- [3] S. Xu and T. Saadawi, "Revealing the Problems with 802.11 MAC Protocol in Multi-hop Wireless Networks", *Computer Networks*, Volume 38, N. 4, March 2002.
- [4] G. Anastasi, M. Conti, E. Gregori, "IEEE 802.11 Ad Hoc Networks: Protocols, Performance and Open Issues", in *Mobile Ad hoc networking*, S. Basagni, M. Conti, S. Giordano, I. Stojmenovic (Editors), IEEE Press and John Wiley and Sons, Inc., New York, 2003.
- [5] IEEE standard 802.11, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", August 1999.
- [6] G. Holland, N. Vaidya, "Analysis of the TCP Performance over Mobile Ad Hoc Networks", Proceedings of the *ACM MobiCom'99*, Seattle (WA), August 1999, pp. 207-218.
- [7] Gavin Holland, Nitin H. Vaidya "Analysis of TCP Performance over Mobile Ad Hoc Networks", *ACM/Kluwer WINET 8(2-3)*, (2002) pp. 275-288.
- [8] K. Chandran, S. Raghunathan, S. Venkatesan, R. Prakash, "A Feedback Based Scheme for Improving TCP Performance in Ad Hoc Wireless Networks", *IEEE Personal Communication Magazine*, Vol. 8, N. 1, pp. 34-39, February 2001.
- [9] J. Liu, S. Singh, "ATCP: TCP for mobile ad hoc networks", *IEEE JSAC*, 19(7):1300-1315, July 2001.
- [10] Z. Fu, X. Meng, S. Lu, "How Bad TCP Can Perform in Mobile Ad Hoc Networks", Proceedings of the IEEE ISCC, Italy, July 2002, pp. 298-303.
- [11] A. Ahuja et al., "Performance of TCP over different routing protocols in mobile ad-hoc networks," Proceedings of IEEE VTC 2000, Tokyo, Japan, May 2000.
- [12] T.D. Dyer, R.V. Boppana "A Comparison of TCP Performance over Three Routing Protocols for Mobile Ad Hoc Networks", *Proc. ACM MobiHoc*, October 2001.
- [13] Z. Fu, P. Zerfos, K. Xu, H. Luo, S. Lu, L. Zhang, M. Gerla, "On TCP Performance in Multihop Wireless Networks", *UCLA Technical Report*, 2002 (available at <http://www.cs.ucla.edu/~hluo/publications/WINGTR0203.pdf>).
- [14] J. Li, C. Blake, D. De Couto, H. Lee, R. Morris, "Capacity of Wireless Ad Hoc Wireless Networks", Proceedings of *ACM MobiCom'01*, Rome (I), July 2002, pp. 61-69.
- [15] K. Tang, M. Gerla, "Fair Sharing of MAC under TCP in Wireless Ad Hoc Networks", Proceedings of IEEE MMT'99, Venice (I), October 1999.
- [16] K. Xu, S. Bae, S. Lee, M. Gerla, "TCP Behavior across Multihop Wireless Networks and the Wired Networks", Proceedings of *ACM WoWMoM 2002*, Atlanta (GA), September 28, 2002, pp. 41-48.
- [17] GloMoSim, Global Mobile Information Systems Simulation Library, <http://pcl.cs.ucla.edu/projects/glomosim/>.
- [18] The Network Simulator - ns-2, <http://www.isi.edu/nsnam/ns/index.html>.
- [19] Qualnet simulator, <http://www.qualnet.com/>.
- [20] T. Ephremides, "A Wireless Link Perspective in Mobile Networking", *ACM Mobicom 2002 keynote speech*, available at <http://www.acm.org/sigmobile/mobicom/2002/program/>
- [21] G. Anastasi, E. Borgia, M. Conti, E. Gregori, "IEEE 802.11 Ad Hoc Networks: Performance Measurements", *Cluster Computing Journal*, (to appear).
- [22] D. B. Green, M. S. Obaidat, "An Accurate Line of Sight Propagation Performance Model for Ad-Hoc 802.11 Wireless LAN (WLAN) Devices", Proceedings of IEEE ICC 2002, New York, April 2002.