IEEE 802.11 Ad Hoc Networks: Performance Measurements

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

E. Borgia, M. Conti, E. 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 802.11 ad hoc networks by means of an experimental study. Measurements on IEEE 802.11 ad hoc networks confirm previous simulative results (e.g., TCP connections may actually experience significant throughput unfairness). The analysis of IEEE 802.11b reveals several aspects that are usually neglected in simulative 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 simulative analysis, and are not constant but highly variable, even in the same session.

1. Introduction

The IEEE 802.11 technology [IEE02] 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 radium (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 radium 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., MeshNetworks[1] and SPANworks[2]).

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 [XuS01, XuS02], 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 [Ana03]. To the best of our knowledge, only very few experimental analysis have been conducted. For this reason, in Section 3 we present an extensive set of measurements that have been conducted on a real test-bed. The measurements were performed in an outdoor environment, by considering different traffic types (i.e., TCP and UDP traffics).

2. IEEE 802.11b

Currently, the Wi-Fi network interfaces are becoming more and more popular. Wi-Fi cards implement the IEEE 802.11b standard.

The 802.11b standard extends the 802.11 standard [IEE99] 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

---

‡ 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

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 [Eph02].

IEEE 802.11 cards transmit at a constant power, nence lowering the transmission rate permits the packaging of more energy per symbol. For this reason a reduction of the transmission rate results in an higer transmission range.

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, defined as:

- The *Transmission Range* (*TX_range*): 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*): 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*): 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 simulative studies the following relationship was been generally assumed:

TX_range≤ IF_range ≤ PCS_range.

For example, in the *ns*-2 simulative tool [Ns02] 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*, *IF_range* and are assumed to be constant throughout a simulative 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.

3. Experimental Analysis

The measurement test-bed is based on laptops running the Linux-Mandrake 8.2 operating system and equipped with D-Link*Air* DWL-650 cards using the DSSS physical layer operating at the nominal bit rate of 1,2,5.5, and 11 Mbps. The target of our study is the analysis of the TCP/UDP performance over an IEEE 802.11b ad hoc network. Since we are interested in investigating the impact of the CSMA/CA protocol on the TCP/UDP performance, we have considered static, single-hop ad hoc networks, i.e., communicating stations are within their transmission range and stations do not change their position during the experiment. This allows to remove other possible causes that may interfere with the TCP behavior, e.g., link breakage, route re-computation, etc.

The experiments were performed in an outdoor space. Each station was located in an open environment (a field without buildings) and the distance among them was changed to generate scenarios in which hidden and/or exposed stations may be present.

Specifically 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-nodes experiments);
ii) the relationship between the transmission range and the transmission rate (two-nodes experiments);
iii) hidden and/or exposed node situations (four-nodes experiments).

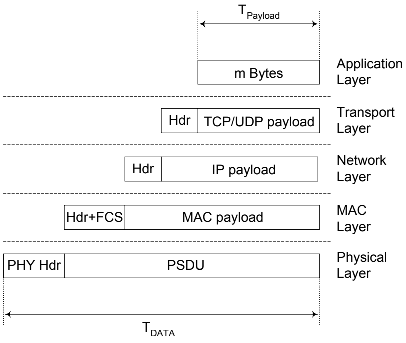

Figure 1. Encapsulation overheads.

3.1. Maximum Throughput

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 1). 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 information,[3] 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 Equation (1) defines the maximum expected throughput for a single active session (i.e., only a sender-receiver couple active) when the basic access scheme (i.e., DCF and no RTS-CTS) is used. Specifically, Equation (1) is the ratio between the time required to transmit the user data and the overall time the channel is busy due to this transmission:

$$Th_{noRTS/CTS} = \frac{T_{payload}}{DIFS + T_{DATA} + SIFS + T_{ACK} + \frac{CW\min}{2} * Slot\_Time} \qquad (1)$$

where

$T_{payload}$ is the time required to transmit only the *m* bytes generated by the application; $T_{payload}$ is therefore equal to *m / data_rate*, where *data_rate* is the data rate used by the NIC to transmit data, i.e., 1, 2, 5.5, or 11 Mbps.

T_{DATA} is the time required to transmit a MAC data frame; this includes the PHY_{hdr}, MAC_{hdr}, $MAC_{payload}$ and FCS bits for error detection.

T_{ACK} is the time required to transmit a MAC ACK frame; this includes the PHY_{hdr} and MAC_{hdr}.

$\frac{CW\min}{2} * Slot\_Time$ is the average back off time

When the RTS/CTS mechanism is used, the overheads associated with the transmission of the RTS and CTS frames must be added to the denominator of (1). Hence, in this case, the maximum throughput $Th_{RTS/CTS}$, is defined as

$$Th_{RTS/CTS} = \frac{T_{payload}}{DIFS + T_{RTS} + T_{CTS} + T_{DATA} + T_{ACK} + 3 * SIFS + \frac{CW\min}{2} * Slot\_Time} \qquad (2)$$

where T_{RTS} and T_{CTS} indicate the time required to transmit the RTS and CTS frames, respectively.

The numerical results presented in the next sections depend on the specific setting of the IEEE 802.11b protocol parameters. Table 2 gives the values for the protocol parameters used hereafter.

Table 1. IEEE 802.11b parameter values.

$Slot\_Time$	τ	PHY_{hdr}	MAC_{hdr}	FCS	*Bit Rate(Mbps)*
20 sec	1 sec	192 bits (2.56 t_{slot})	240 bits (2.4 t_{slot})	32 bits (0.32 t_{slot})	1, 2, 5.5, 11
$DIFS$	$SIFS$	ACK		CW_{MIN}	CW_{MAX}
50 sec	10 sec	112 bits + PHY_{hdr}		32 t_{slot}	1024 t_{slot}

In Table 2 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. These results are computed by applying Equations (1) and (2), and assuming a data packet size at the application level equal to m=512 and m=1024 bytes.

Table 2. Maximum throughputs in Mbit/sec (Mbps) at different data rates.

	m= 512 Bytes		m=1024 Bytes	
	No RTS/CTS	RTS/CTS	No RTS/CTS	RTS/CTS
11 Mbps	3.06 Mbps	2.549 Mbps	4.788 Mbps	4.139 Mbps
5,5 Mbps	2.366 Mbps	2.049 Mbps	3.308 Mbps	2.985 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

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

---

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

always have packets ready for transmission) with constant size packets of 512 bytes.

The results obtained from this experimental analysis are reported in the Figure 2.

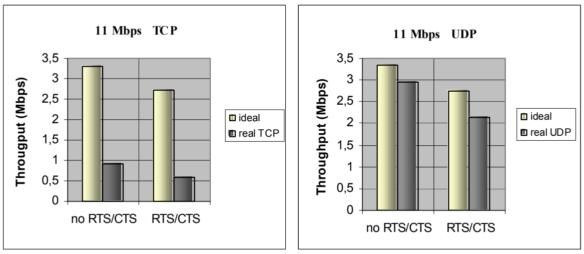

Figure 2. Comparison between the theoretical maximum throughput and the actual throughput achieved by TCP/UDP applications.

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. Indeed, when using the TCP protocol overheads related to the TCP-ACK transmission must be taken into account.

Similar results have been also obtained by comparing the maximum throughput derived according to (1) and (2), and the real throughputs measured when the NIC data rate is set to 1, 2 or 5.5 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 shown in Figure 3 In Figure 4 we report the transmission-range curves (when the data rate is equal to 1 Mbps) estimated in two different days. The graph highlights the variability of the transmission ranges depending on the weather conditions[4].

The results presented in Figure 3 are summarized in Table 3 where the estimates of the transmission ranges at the different data rates are reported.

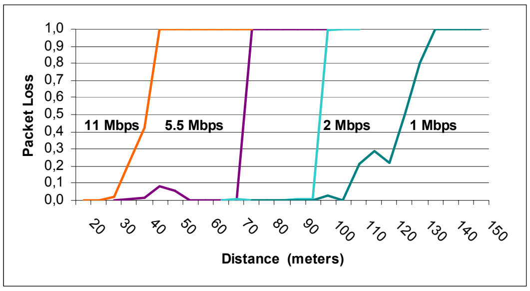

Figure 3. Packet loss rate as a function of the distance between communicating stations for different data rates.

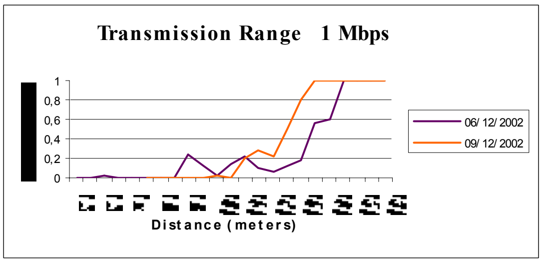

Figure 4. 1 Mbps transmission ranger in different days.

These estimates are very important since they 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 11Mbps to another station within its transmission range (i.e., less then 30m apart) it reserves the channel for a radius of approximately 90 (120) m around itself. In fact, the RTS frame is transmitted at 2Mbps (or 1Mbps), and, hence, it is correctly received by all stations that are less then 90 (120) meters away from the transmitting station.

Table 3. Estimates of the transmission ranges at different data rates.

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

Again it is interesting to compare the transmission range used in the most popular simulative tools, like ns-2 and Glomosim, with the transmission range measured in our experiments. In this simulative tools it is assumed *TX_range*=250m. Since the above simulative tools only consider a 2-Mbps bit rate we only consider the transmission range related to 2 Mbps. As it clearly appears, the values of the transmission range used in the simulative tools (and, hence, in the simulative studies

---

[4] It is worth pointing out that we experienced a high variability in the channel conditions during the same experiment.

based on them) are 2-3 times higher that 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.

3.3 Four-Stations Network Configurations

The results presented in the previous sections show that the IEEE 802.11b behavior is very complex. 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.

The reference network scenario for the experiments is shown in Figure 5. In this scenario, we have two contemporary active sessions. Specifically, Station S1 communicates with Station S2 (Session 1), while Station S3 communicates with Station S4 (Session 2). In the figure, the arrows represent the direction of the data flow (e.g., Station 1 is delivering data to Station 2), and *d*(*i,j*) is the distance between Station *i* and Station *j*. Data to be delivered are generated by either an ftp application, or a Continuous Bit Rate (CBR) application. In the former case the TCP protocol is used, while in the latter case UDP is the transport protocol.

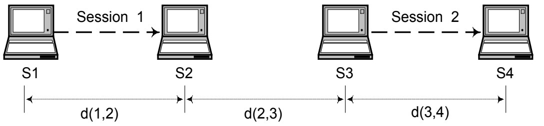

Figure 5. Reference network scenario.

Two sets of experiments were performed by varying the NIC data rate. In the each set of experiments the data rate is constant, and equal to 11Mbps, and 2Mbps, respectively, and the distance between the two couples of stations is different to take into account the different transmission range. The network configurations are shown in Figure 6 and Figure 8, while the related results are presented in Figure 7 and Figure 9, respectively. These results are the superposition of several factors that make the system behavior similar in the two cases (even though numerical values differ due to the different transmission rates).

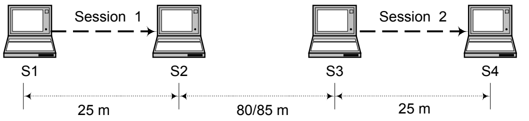

Figure 6. Network configuration at 11 Mbps.

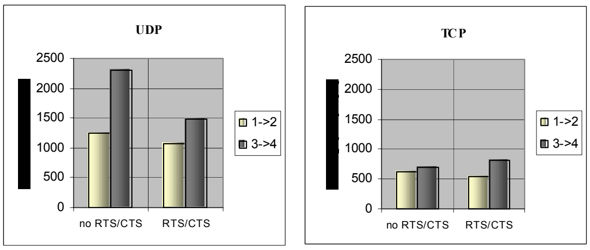

Figure 7. Throughputs at 11 Mbps.

In the first set of experiments (11Mbps) dependencies exist between the two connections even though the transmission range is smaller than the distance between stations S1 and S3. Furthermore, the dependency exists also when the basic mechanism (i.e., no RTS/CTS) is adopted.[5] To summarize, this set of experiments show that i) interdependencies among the stations extends beyond the transmission range; ii) the physical carrier sensing range often produces an effect that 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 S4, 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 S4 is exposed to S2 transmissions but the S2's effect on S4 is less marked given the different role of the two stations. When using the basic access mechanism, the S2's effect on S4 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 throughput 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.

---

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

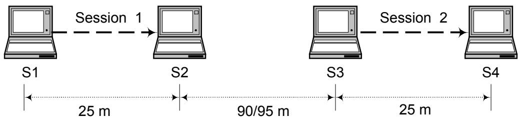

Figure 8. Network configuration at 2 Mbps.

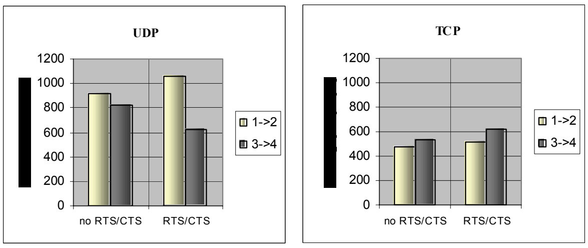

Figure 9. Throughputs at 2 Mbps.

In the second set of experiments (data rate equal to 2Mbps), whose results are presented in Figure 9, S1 and S3 are within the same transmission range and, in addition, it can assumed that all stations are within the same physical carrier sensing range. It is also worth noting that in this case the system is more balanced from the throughput standpoint. This can be expected as by transmitting with a 2-Mbps data rate the transmission range is significantly larger than with the 11-Mbps transmission rate, and hence the stations have a more uniform view of the channel status.

We have performed several other experiments by considering the symmetric scenario shown in Figure 10. The results obtained with 11 Mbps and 2 Mbps are reported in Figure 11 and Figure 12, respectively. These results are aligned with the previous observations.

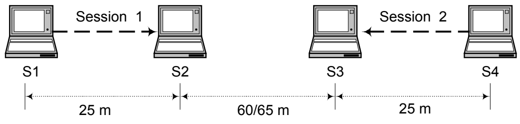

Figure 10. Simmetric Scenario.

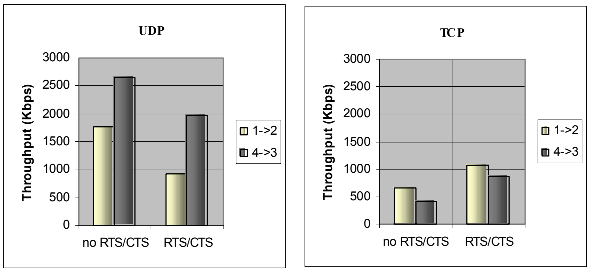

Figure 11. Throughputs at 11 Mbps.

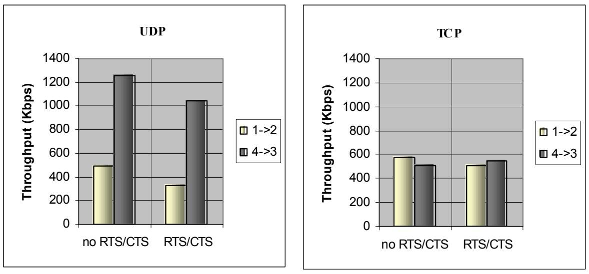

Figure 12. Throughputs at 2 Mbps.

References

[Ana03] G. Anastasi, E. Borgia, M. Conti, E. Gregori, “IEEE 802.11 Ad Hoc Networks: Performance Measurements”, *IIT Internal Report*, Gennaio2003.

[Eph02] T. Ephremides, “A Wireless Link Perspective in Mobile Networking", ACM Mobicom 2002 keynote speech, available at http://www.acm.org/sigmobile/mobicom/2002/program/

[Fu02] 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) .

[IEE02] Official Homepage of the IEEE 802.11 Working Group, http://grouper.ieee.org/groups/802/11/.

[IEE99] IEEE standard 802.11, “Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, August 1999.

[Ns02] The Network Simulator - ns-2, http://www.isi.edu/nsnam/ns/index.html.

[XuS01] S. Xu and T. Saadawi, “Does the IEEE 802.11 MAC protocol Work Well in Multihop Wireless Ad Hoc Networks?”, *IEEE Communication Magazine*, Volume 39, N. 6, June 2001, pp. 130-137.

[XuS02] 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.