An ORBIT Testbed Study of 802.11b DCF: Throughput, Latency, and the Capture Effect

Haris Kremo, Ivan Seskar, and Predrag Spasojevic
WINLAB, ECE Department, Rutgers University
Email: {harisk, seskar, spasojev}@winlab.rutgers.edu

***Abstract*— We present an experimental study of 802.11b Distributed Coordination Function (DCF) performance conducted on the Open-Access Research Testbed for Next-Generation Wireless Networks (ORBIT). The experiments involve relatively large number of wireless nodes in a scalable and controllable testbed environment. The nodes use 802.11b wireless interfaces with rate adaptation feature disabled. Both basic access and RTS/CTS handshaking are considered for two topologies. The first one is the set of bi-directional point-to-point links and the second one is a sink topology in which the sources send packets to the same receiver. We measure the transport layer throughput and latency as a function of the offered load. The results for user datagram protocol (UDP) as the transport protocol show linear dependency between the throughput and the offered load until the saturation has been reached. The saturation throughput measurements are in good agreement with analytically predicted values. The latency changes exponentially for the same range of offered load for which the throughput changes linearly, and is limited by the finite capacity of senders' queues. In saturation, the contribution of the nodes to the aggregate throughput is unequal due to the capture effect. The received signal strength measurements corroborate our conjecture that the main reason for the capture effect is difference in the link signal-to-noise ratios (SNRs).**

I. Introduction

The dominant standard for wireless networks at present is 802.11 [1], [2]. The protocol complexity originates from the nature of the radio channel as the shared medium. Unlike for the case of wired networks, in the case of wireless medium, the link quality changes dramatically in time because of fading, multipath, and shadowing. Wireless network protocols must support mobility, and very often must provide quality of service for different types of user applications.

The previous experimental work on 802.11 DCF performance [3]–[9] involved setups with very small number of nodes and limited capabilities to set available parameters, either on physical, or on medium access control layer (MAC). On the other hand, existing analytical results cover only the throughput saturation conditions, and involve simplifying assumptions [10]–[14]. Typically, to analytically characterize the MAC layer performance, the complexity of radio environment is circumvented by considering the communication channel to be error-free. Another important motivation for experimental approach to characterize 802.11 MAC performance using testbeds like ORBIT [15]–[17], is to obtain results that are more reliable than results of computer simulations, since simulations typically trade processing effort for accuracy.

Our work is motivated by the fact that the ORBIT testbed allows us to perform a set of experiments involving different network scenarios in the controllable and repeatable environment [18]. The ORBIT software enables us to aggregate and analyze the measurements in a simple and efficient way [16].

In our experiments we measure transport layer throughput and latency. Although transport layer measurements might appear to be inconsistent with our intention to investigate the MAC performance, we are more interested in how the 802.11 MAC performs from the application point of view. Using statistics on the number of successfully transmitted packets and no delivered packets obtained from the wireless interface driver, we measure the contribution of each sender to the aggregate network throughput. Throughout the experiments we increase offered load to all senders in equal discrete steps. The throughput measurements for the case of saturation are compared to [10], and they show good agreement with analytical prediction for both tested network topologies, and both access schemes. The difference between the measurement results and the analytical prediction are caused by simplifying assumptions in [10]. The measured statistics on unsuccessful transmissions in saturation agrees with the assumption in [10] on collision probability being constant and independent on previous collisions for each node.

The rest of the paper is organized as follows. Section II is the overview of the distributed coordination function (DCF), which is the random access scheme of the 802.11 protocol. Section III is the overview of the previous work on the characterization of DCF performance, both theoretical and experimental. Section IV briefly presents the ORBIT testbed and measurements setup. Section V contains the results of the experimental study. In Section VI we summarize the results.

II. The 802.11 Distributed Coordination Function Overview

The 802.11 distributed coordination function (DCF) employs the basic two-way or the RTS/CTS four-way handshaking random access scheme.

It is based on Carrier Sense Multiple Access with Collision Avoidance (CSMA-CA) [1]. All nodes listen to the medium for the duration of the DCF interframe space (DIFS) before attempting to transmit. In the absence of a carrier, a node with a queued packet is allowed to immediately attempt transmission after DIFS. Since 802.11 MAC belongs to the class of Stop-and-Wait protocols, after transmitting a packet,

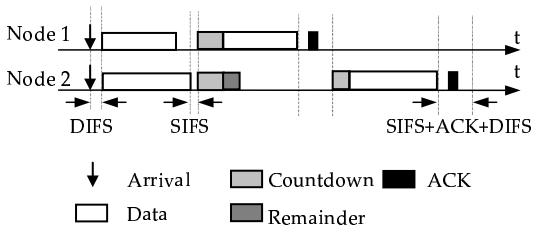

Fig. 1. An example of collision in the case of basic access

the sender waits for positive acknowledgment (ACK) from the receiver for the duration of a short interframe space (SIFS). The absence of an ACK indicates reception failure and the transmitter schedules the packet for retry. When a packet is not acknowledged the transmitter declares erroneous packet delivery and engages in a binary random backoff scheme with an exponentially growing contention window. After each successful transmission contention window is reset to its minimum value.

Fig. 1 shows the behavior of the protocol in the case of collision. Suppose two nodes observe an unoccupied channel for the DIFS duration and attempt to transmit simultaneously. The node sending shorter packet becomes aware of the collision first due to the missing ACK. Both nodes subsequently apply random backoff to resolve the collision. The figure illustrates situation in which, by chance, the nodes chose the number of backoffs that does not cause repeated collision.

The basic access handshaking does not include any additional signaling other than packet acknowledgments, thus being susceptible to packet collisions. RTS/CTS handshaking first attempts reserving the channel for the duration of the packet prior to initiating its transmission. Hidden nodes which are not able to decode either RTS or CTS may still transmit and cause a collision.

In the Table I are given the numerical values of the protocol parameters for two versions: 802.11a and 802.11b Direct Sequence Spread Spectrum (DSSS).

TABLE I

802.11A AND 802.11B DSSS MAC PROTOCOL PARAMETERS

Parameter	**802.11a**	**802.11b DSSS**
CW_{min}	16 slots	32 slots
CW_{max}	1024 slots	1024 slots
Slot time	9 μs	20 μs
SIFS	16 μs	10 μs
DIFS	34 μs	50 μs
ACK	14 bytes	14 bytes
RTS	20 bytes	20 bytes
CTS	14 bytes	14 bytes
PLCP header [a]	24 μs	192 μs
MAC header [b]	34 bytes	34 bytes
PHY rates	6, 9, 12, 18, 24, 36, 48, 54 Mbps	1, 2, 5.5, 11 Mbps

[a]PLCP header is for both protocol versions transmitted with the smallest available PHY rate to be decodable by all nodes.

[b]Includes 4 bytes of CRC-32 parity check at the end of each packet.

III. Related Work

A. Analytical Approach to 802.11 MAC Protocol Modeling

Two of the most widely quoted analytical models of IEEE 802.11 MAC [10], [11] address the saturation state of network, in which all participating nodes at every time instance have a packet available for transmission. Both models are limited to the case of *ideal channel*. The radio channel is considered to be noiseless, meaning that, in absence of a collision, all transmitted bits are correctly decoded at the receiver. All nodes in the network are assumed to be in the transmission range. Thus, both models assume absence of mobile or hidden nodes.

We extensively use the model presented in [10] to verify the results of experiments. From the two-dimensional Markov chain, the author determines two probabilities: the probability that a node will transmit, and the probability of collision in a randomly chosen *time slot*. The time slot is defined as a random fraction of time needed to decrement the backoff counter. These probabilities are assumed to be constant, equal, and independent on previous attempts to transmit for all nodes. From these two probabilities it is possible to determine the normalized throughput in saturation state. The resulting throughput formula is independent on the employed access method. Based on [10], the authors of [14] derive an expression for the average packet delay in 802.11.

In [11] the authors derive the model of a p-persistent memoryless protocol equivalent to 802.11 MAC. It is possible to model 802.11 MAC in such a way because after each successful transmission the protocol resets the contention window size to CW_{min}, regardless of the number of accumulated backoffs.

B. Prior Experimental Work

Both [3] and [4] consider simple wireless network with up to five nodes connected to an access point (AP). The AP is connected to the LAN traffic analyzer over the Ethernet. Simple file exchange is used to generate the traffic. The results show that the throughput, considering the transport layer payload as the actual payload, is roughly 50% of the nominal PHY rate.

In [5]–[8] are given measurements on point-to-point links. In [5] are presented two different setups, one with high SNR, and the other with low SNR. In [7], [8] are first given the results for a single point-to-point link with variable distance between the nodes, and then, the results for a topology in which two point-to-point links create mutual interference. Measurements in [9] are collected on a multihop topology consisting of only two hops.

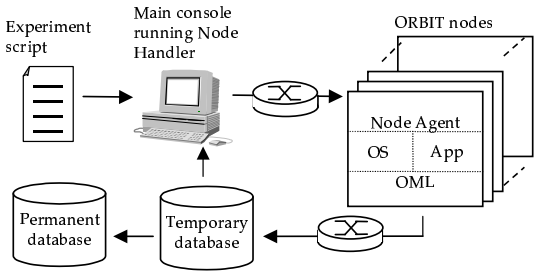

Fig. 2. The ORBIT testbed components. "OS" denotes operating system, and "App" stands for user designed application.

IV. The ORBIT Testbed Setup

Simplified to the very core, the ORBIT testbed [15] is given in Fig. 2. The front line of the testbed are 64 *ORBIT nodes*[1] designed as custom made personal computer platforms, each equipped with two wireless 802.11a/b/g interfaces. The nodes are placed in the two-dimensional rectangular grid separated by 1-meter distance. Behind this frontline is the very complex wired infrastructure that enables simple and scalable access to the grid.

The main responsibility of a user is to write a script that explains the lifecycle of experiment. During the experiment execution the program named *Node Handler* interprets the script, and sends commands to its counterparts installed on the nodes, named *Node Agents*. Node Agent is capable of passing the commands both to the operating system, as well as to user-defined applications, which are respectively denoted as "OS" and "App" in Fig. 2. This, for instance, includes the commands to set a particular Internet protocol address on an interface, to set the appropriate PHY rate on a wireless interface, or to change the packet size and rate of a network traffic generator.

A central component of the ORBIT software is the ORBIT Measurement Library (OML) [16]. It is the framework that processes measurements by means of filtering and compression, and collects the measurements to the database.

V. The Results of the Experimental Case Study

A. Experimental Setup

In our experiments we use UDP as the transport protocol. We test two network topologies with fixed number of nodes, given in Fig. 3. The first topology considers five single-hop bi-directional links. The second topology is formed of ten senders, all transmitting packets to the single receiver. The wireless interfaces used in experiments are Cisco AiroNet 350 802.11b [19]. The interfaces employ 802.11b DSSS version of the standard. PHY rate is set to constant 11Mbps. Transmit power on all nodes is set to 5mW. The experiments involve both basic access and RTS/CTS handshaking. All sources generate constant bit rate traffic (CBR). The maximal number of retries for a packet is set to $k = 8$, thus effectively setting $CW_{max} = 1024$.

[1]When fully operational the testbed will consist of 400 nodes.

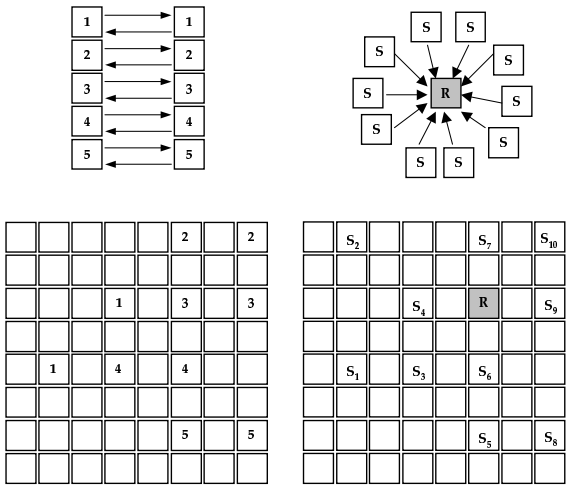

Fig. 3. Experimental network topologies. Plots on the left of the figure relate to the point-to-point topology and plots on the right relate to the sink topology.

Throughout the experiment lifecycle we use five different payload lengths (256, 512, 768, 1024, and 1280 bytes) and ten different packet rates, equally set on each node (from roughly 100kbps to 1Mbps in 100kbps steps). Hence, the aggregate offered load changes from 1Mbps to 10Mbps in 1Mbps steps. The range of aggregate offered load covers values from light load up to the network saturation. The settings are kept constant over two-minutes period. The throughput and offered load are measured every second. This means that the throughput and offered load statistics are calculated over 120 samples for a particular set of tunable parameters.

Our custom-made application was incorporating a time stamp in every packet before it was passed to the nonblocking UDP socket. This time stamp was compared with the local clock immediately after the packet was extracted from the receiver's socket. In order to synchronize the clocks all the nodes were running Network Timing Protocol (NTP) daemon process. The latency measurements are averaged over the values extracted from all successfully received packets and recorded every second.

B. Throughput Measurements

The results for the point-to-point topology and basic access over the set of different payloads are given in Fig. 4. A larger throughput is achieved with longer packets, since longer packets introduce smaller protocol overhead. As the offered load increases, the throughput increases linearly, and it saturates to a practically constant value after a certain threshold. For shorter packets, the protocol introduces larger overhead, and the throughput curve saturates at a smaller value of offered load. Because of the overhead, even in the case of the largest payload length, the maximum achieved throughput is only around a half of the nominal PHY rate. In the Fig. 5 is shown comparison of the results for both access methods and both topologies in the case of 1024 bytes payload. The results are

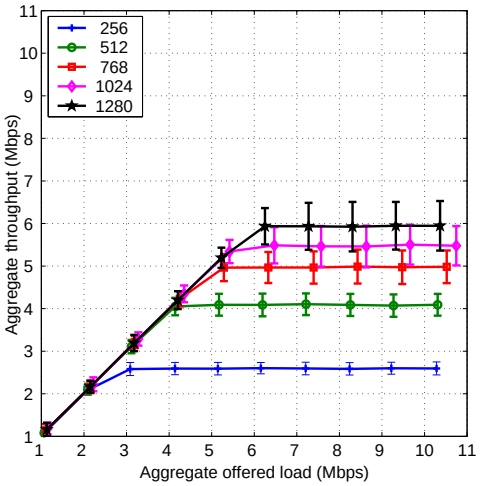

Fig. 4. Throughput for point-to-point topology and DCF basic access as a function of aggregate offered load. Payload lengths are 256, 512, 768, 1024, and 1280 bytes. The bars on the curves denote one standard deviation limits. Larger payload lengths result in a larger saturation throughput because of smaller protocol overhead.

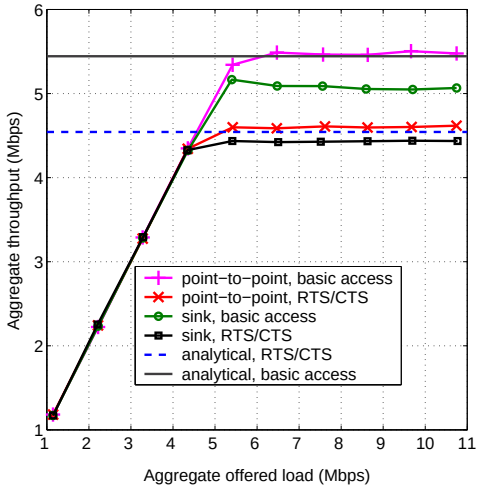

Fig. 5. Throughput for payload length 1024 bytes. Basic access shows larger throughput regardless of the topology. Point-to-point topology results in a higher saturation throughput for both access schemes. Analytically predicted saturation throughput falls between the results for two topologies.

qualitatively the same for all other payloads, as it can be seen in detail in [20].

Comparing all four throughput curves in the Fig. 5, we observe that the basic access performs better than RTS/CTS. This is in agreement with [10]. The analysis presented in [10] predicts that a larger number of active nodes than the number of nodes used in the experiments is required in order to trade between the additional overhead introduced by RTS/CTS and the smaller penalty in lost time slots due to collisions.

Regardless of the employed access scheme, the point-to-point topology performs slightly better than the theoretical prediction. A careful look at Fig. 3 reveals the reason for this. The nodes that create the point-to-point links are in almost all cases the nodes with the smallest distance. Other point-to-point links, which act as the sources of interference for a particular link, are typically farther away. Such allocation of the nodes raises aggregate throughput slightly above the theoretical prediction because collisions due to farther links do not necessarily result in dropped packets, as would be the case for randomly distributed nodes.

Table II presents the comparison of results for saturation throughput with the prediction given in [10] for the case of basic access, as a function of payload length. The second and the third column contain the values given as the points corresponding to the highest offered load setting in the experiments, respectively for point-to-point and for sink topology.

Theoretical predictions from [10] are listed in the fourth column. The saturation throughput formula from [10] is adapted in (1) to take into account the fact that the data portion of a packet is sent using a bit rate different from the rate of the physical layer header. The throughput in (1) is expressed in bits per second.

$$S = \frac{8BP_sP_{tr}}{P_sP_{tr}T_s + (1-P_{tr})\sigma + P_{tr}(1-P_s)T_c} \qquad (1)$$

The payload length in bytes is represented by B. The probability P_{tr} is the probability of at least one transmission in a time slot [10]. Similarly, the probability P_s is the probability that this transmission in a time slot is successful, or in other words, that a single node attempts to transmit in a time slot. Under the assumption of an ideal communication channel these quantities do not depend on the network size or bit rate. The quantities T_s and T_c represent, respectively, the average time utilized for successful transmission (including all overheads), and the average time wasted in collisions. They take into account that the payload, as well as MAC, IP, and UDP headers, are transmitted with a rate larger than the minimum rate. They also take into account, depending on the employed access scheme, the duration of overhead introduced by the signaling packets ACK, RTS, and CTS, and protocol synchronization time spaces SIFS and DIFS. The numerical values for these protocol parameters, as well as for the protocol slot time σ are listed in the Table I.

Table III follows the same outline for the case of RTS/CTS handshaking. The points corresponding to the highest offered load in the experiments with RTS/CTS are used to populate the second and the third column of the Table III.

Tables II and III show large throughput inefficiency of the protocol. Even for the case of large packets, the throughput close to a half of the nominal PHY rate. The reason for such inefficiency are:

- Inherent protocol overhead:
  - Signaling overhead includes DIFS, SIFS, ACK, RTS, and CTS.
  - Protocol stack overhead includes PHY, MAC, Internet, and transport layer headers for each data packet.
  - Support for multiple PHY rates means that the PHY header, known as Physical Layer Convergence Protocol (PLCP) header, must be transmitted with the

TABLE II

COMPARISON OF THE THEORETICAL PREDICTION TO THE EXPERIMENTAL RESULTS FOR BASIC ACCESS, 11MBPS PHY RATE, AND 10 ACTIVE NODES

Payload length	**Point-to-point topology**	**Sink topology**	**Analytical prediction**
bytes	Mbps	Mbps	Mbps
256	2.5952	2.3959	2.4427
512	4.0896	3.7719	3.8618
768	4.9833	4.6383	4.7892
1024	5.4762	5.0661	5.4427
1280	5.9449	5.4625	5.9281

TABLE III

COMPARISON OF THE THEORETICAL PREDICTION TO THE EXPERIMENTAL RESULTS FOR RTS/CTS ACCESS, 11MBPS PHY RATE, AND 10 ACTIVE NODES

Payload length	**Point-to-point topology**	**Sink topology**	**Analytical prediction**
bytes	Mbps	Mbps	Mbps
256	1.6931	1.6129	1.6452
512	2.9122	2.7779	2.8624
768	3.8562	3.7034	3.7992
1024	4.6174	4.4347	4.5427
1280	5.2151	5.0132	5.1470

smallest possible rate in order to be decodable by all nodes. It also must contain additional bits explaining the actual rate at which the rest of the packet is modulated.

- Possible collisions affect the throughput by creating time slots which are wasted in collisions, or not at all utilized for transmission [10], [11].

C. Latency Measurements

The average latency values vary from a few milliseconds for lightly loaded network, up to the values that may seem extraordinarily huge in the case of saturation. This is because, unlike [14], we take into account the queuing delay.

Comparing Figs. 4 and 6, we observe that the latency increases exponentially and reaches saturation for the same offered load for which the throughput reaches saturation. At that point, the queues at the senders are full, because the rate at which packets are generated is larger than the rate at which the packet are successfully sent. Excessive packets are simply rejected and discarded at the sending sockets.

As explained at the beginning of this section, only the packets that enter the queues and are successfully delivered contribute to the results presented in Figs. 6 and 7. The limited queue capacity results in bounding the average packet latency for delivered packets. The same relationship between the throughput and latency saturation is described in the Figs. 5 and 7 for both access schemes and topologies.

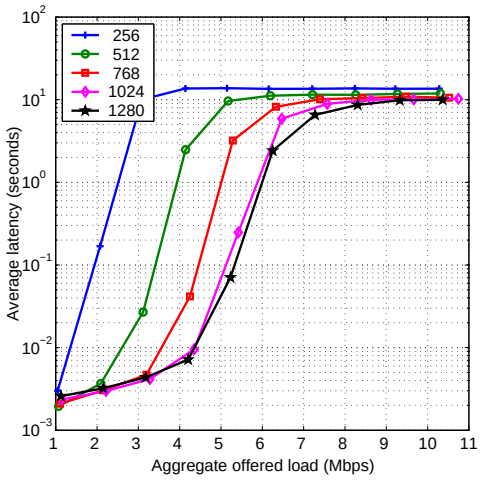

Fig. 6. Latency for point-to-point topology and DCF basic access as a function of aggregate offered load. Payload lengths are 256, 512, 768, 1024, and 1280 bytes. Smaller payload lengths result in larger latency.

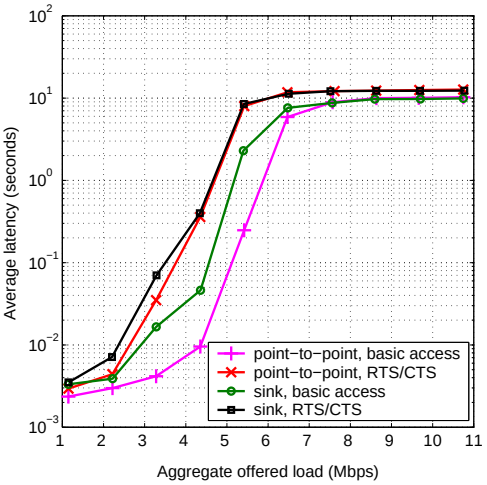

Fig. 7. Latency for payload length 1024 bytes. Basic access shows smaller latency regardless of the topology. RTS/CTS shows larger latency compared to basic access for both topologies.

D. Statistics on Failed Transmissions

The Fig. 8 shows a typical data collected from the wireless driver about the number of failed transmissions. The driver reports the number of packets which were transmitted, but the corresponding ACK was not received. Transmission failures occur either due to collisions, or due to insufficient SNR. The protocol property is that a transmitter cannot distinguish which one of these two was the reason for failure. The values presented in the Fig. 8 are collected for the point-to-point topology, basic access, and payload length of 1024 bytes. The solid line represents an average over the measurements collected from each of ten nodes. These ten measurements are the time average calculated over two-minute runs during which we keep constant offered load. These time averages fall into the limits denoted by the bars in the Fig. 8. Therefore, the solid line represents typical statistics on unsuccessful transmissions

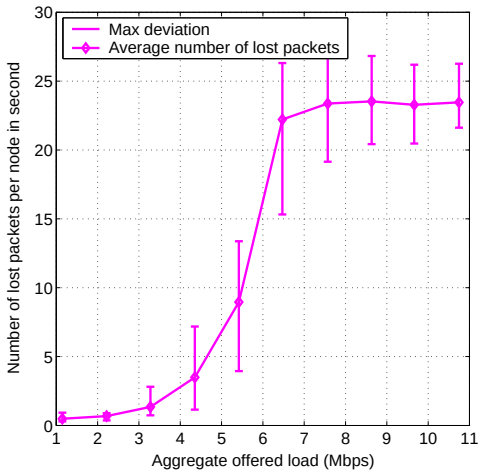

Fig. 8. Number of unsuccessful transmissions as a function of offered load for point-to-point topology, basic access, and payload 1024 bytes. Bars on the graph represent the maximum deviation over the measurements collected from ten nodes.

per node in second. Detailed results for different payload lengths, topologies, and access methods are presented in [20].

The Fig. 8 shows that there exists a particular point at which further increase in the offered load does not cause equivalent increase in the average number of failed transmissions. At that point the network has reached the saturation state, and the average number of failed transmissions remains constant, regardless of the increase in offered load. Assuming that, due to the proximity of the nodes on the testbed, the majority of failures are caused by collisions, this result is in coherence with the key assumption in [10] on the collision probability in saturation being constant and independent on previous collisions.

E. Fairness of the Protocol and Capture Effect

We consider the protocol as being fair if the nodes with equal offered load contribute the same to the total throughput. To get an insight into the issue of fairness we use the data collected for the sink topology. Presented results consider both basic access and RTS/CTS access.

Fig. 9 shows the throughput contribution for each of ten nodes and for the payload length of 1024 bytes, depending on the offered load per node. This figure shows that, under the load that is not large enough to saturate the channel, all nodes contribute to the throughput by the same amount, since all ten measurement instances per node overlap and lay on the curve presenting the average. This is because even if the access to the medium is not equally likely for all nodes, there is a sufficient number of time slots for all transmitters to acquire the channel.

On the other hand, under saturation conditions, there is a clear separation in the contribution of each node. We observe that the nodes show certain regularity in the contribution to the total throughput regardless of the access scheme. In [20] it is shown that the same regularity exists regardless of payload length. For instance, node 7 always contributes the most,

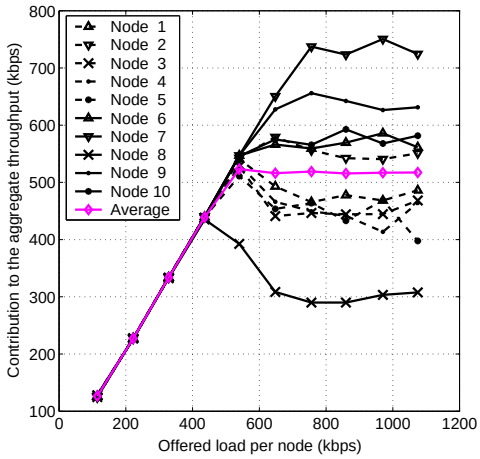

(a) Basic access

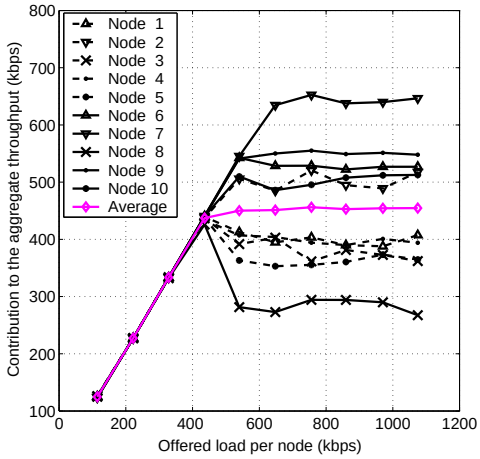

(b) RTS/CTS

Fig. 9. The contribution of particular nodes to the aggregate throughput for sink topology. In saturation the nodes show differences in contribution to the aggregate throughput. This stands for both access schemes.

node 8 always has the smallest contribution, node 5 contributes slightly more than the average, etc. We can observe similar regularities for the rest of nodes.

Fig. 10 illustrates how successful the nodes are when attempting to transmit. The figure shows the fraction of no delivered packets in all attempted transmissions. We see that the curves in Fig. 10 follow the curves from Fig. 9 in reverse order.

Fig. 11 reveals a part of the reason for such behavior. It gives the signal strength at the senders derived from measured received strength signal index (RSSI), averaged over the whole course of the experiment. RSSI values are converted to dBm according to the table from [21]. The numbers placed above the points that present signal power classify the nodes by their participation in the aggregate throughput. Node 8 in Fig. 11 has the smallest received signal power compared to all other nodes and the smallest contribution to the total throughput. Since the carrier sensing mechanism declares the channel

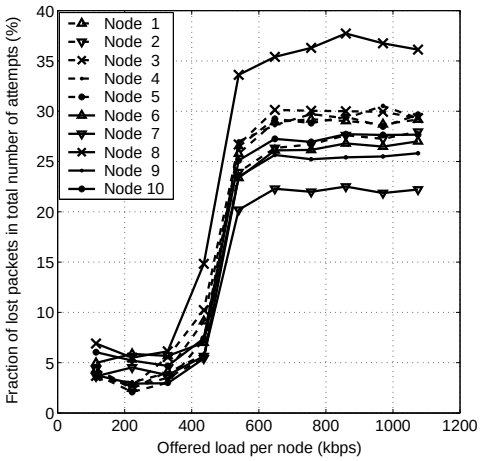

(a) Basic access

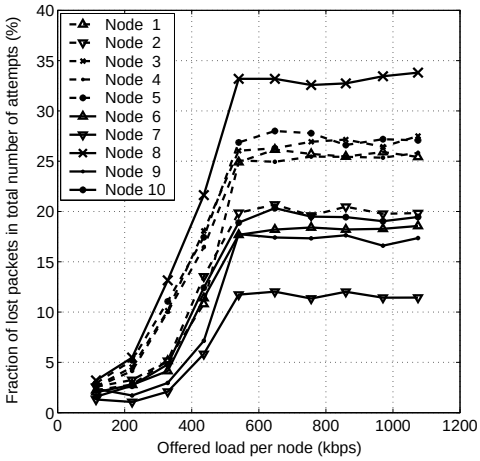

(b) RTS/CTS

Fig. 10. Average ratio of transmission failures for payload length 1024 bytes. The nodes which contribute less to the total throughput exercise larger ratio of unsuccessful transmissions. The same stands for both access schemes.

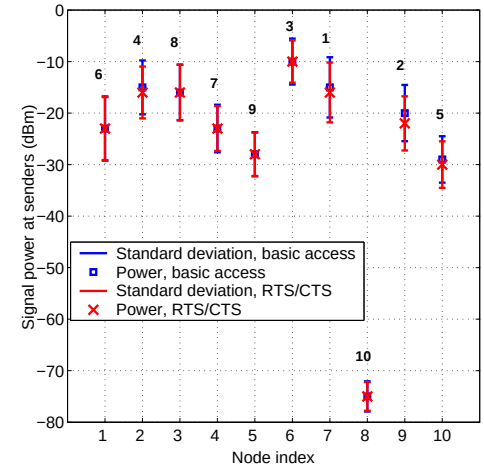

Fig. 11. Signal power at the senders for sink topology. The numbers above the measurements order the nodes according to the contribution to aggregate throughput.

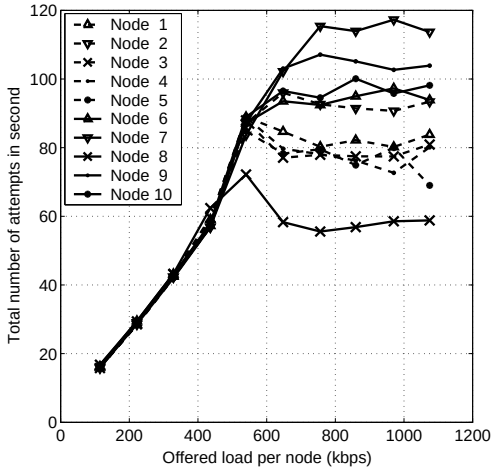

(a) Basic access

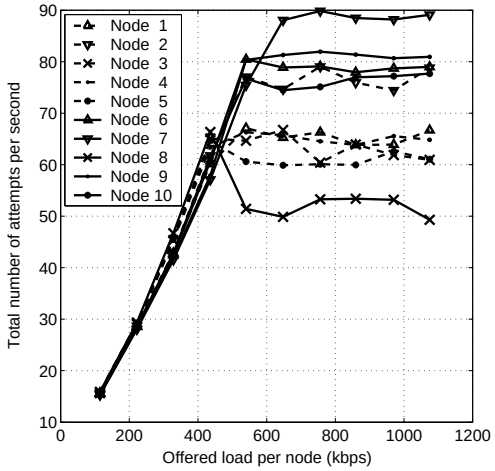

(b) RTS/CTS

Fig. 12. The total number of attempts to transmit, both successful and collisions, in second, for payload length 1024 bytes. The nodes which contribute less to the total throughput have smaller average attempt rate. The same stands for both access schemes.

to be free if the instant value of RSSI is smaller than the carrier sensing threshold, our conjecture is that this node has a reduced carrier sensing ability, and hence, collides with the highest probability. The consequence of higher collision probability is the increased average contention window size and longer average backoff period (proved in [11] to be a half of the average contention window size). Therefore, this node is the most hesitant in attempting to send, as is shown in the Fig. 12. This does not exclude the possibility that the node experiences delivery failure because of the transmission power which is insufficient to provide correct reception.

The situation in which the nodes with better SNR capture the channel [22] is caused by:

- The receiver observes different power levels from different senders because of unequal distance to the senders and tolerance in transmit power.
- The senders observe different power level of the receiver's signal due to the different distances and tolerance in

sensitivity.

The work presented in [18] emphasizes the tolerance of wireless card electronics as the main reason for link asymmetry on the ORBIT testbed.

VI. Conclusions

In this paper we have presented a set of experimental studies of the 802.11b Distributed Coordination Function performance conducted on the ORBIT testbed.

The experiments illustrate the throughput inefficiency of the protocol relative to the nominal physical layer bit rate since, at best, (for relatively large packets) only close to a half of the available bandwidth is used for data transmission.

Our results are in agreement with [10] when predicting that, for a network of ten nodes, RTS/CTS, due to its overhead, does not benefit from a smaller fraction of time spent in collisions relative to the basic access scheme, even for large packet lengths.

The throughput measurements show that, as the aggregate offered load increases, the aggregate throughput also increases linearly until saturation is reached. The latency, limited by the finite capacity of senders' queues, changes exponentially for the same range of offered load, for which the throughput changes linearly.

The sink topology experiments illustrate that the differences in receiver SNRs do not affect the per-node throughput contribution in the linear regime in which a sufficient number of available time slots accommodates for most attempted transmissions. On the other hand, in saturation, our measurements of RSSI at the senders suggest that links with better SNR capture the multiple access medium and, hence, contribute more to the aggregate throughput.

Acknowledgment

We would like to thank to the whole ORBIT team for their help, in particular to Manpreet Singh, Sachin Ganu, Kishore Ramachandran, and Faiyaz Ahmed. Special thanks are due to Richard Howard.

References

[1] *Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE-SA Standards Board ANSI/IEEE Std 802.11, Rev. June 12 2003, August 1999.

[2] *Supplement to IEEE Standard for Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: High-speed Physical Layer in the 5 GHZ Band*, IEEE-SA Standards Board IEEE Std 802.11a-1999 (Supplement to IEEE Std 802.11-1999), September 1999.

[3] A. Kamerman and G. Aben, "Throughput performance of wireless LANs operating at 2.4 and 5 GHz," The 11th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, vol. 1, London, UK, September 2000, pp. 190–195.

[4] ——, "Net throughput with IEEE 802.11 wireless LANs," IEEE Wireless Communications and Networking Conference, vol. 2, Chicago, IL USA, September 2000, pp. 747–752.

[5] M. G. Arranz, R. Aguero, L. Munoz, and P. Mahonen, "Behavior of UDP-based applications over IEEE 802.11 wireless networks," 12th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, vol. 2, San Diego, CA USA, October 2001, pp. F72–F77.

[6] B. Bing, "Measured performance of the IEEE 802.11 wireless LAN," Conference on Local Computer Networks, Lowell, MA USA, October 1999, pp. 34–42.

[7] G. Anastasi, E. Borgia, M. Conti, and E. Gregori, "IEEE 802.11 ad hoc networks: Performance measurements," 23rd International Conference on Distributed Computing Systems Workshops, May 2003, pp. 758–763.

[8] ——, "Wi-fi in ad hoc mode: a measurement study," Second IEEE Annual Conference on Pervasive Computing and Communications, March 2004, pp. 145–154.

[9] K. Benekos, N. Pogkas, G. Kalivas, G. Papadopulos, and A. Tzes, "TCP performance measurements in IEEE 802.11b-based wireless LANs," 12th IEEE Mediterranean Electrotechnical Conference, vol. 2, Dubrovnik, Croatia, May 2004, pp. 575–578.

[10] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, May 2000.

[11] F. Cali, M. Conti, and E. Gregori, "Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit," *IEEE/ACM Transactions on Networking*, vol. 8, no. 6, pp. 785–799, December 2000.

[12] L. Bononi, M. Conti, and E. Gregori, "Design and performance evaluation of an asymptotically optimal backoff algorithm for IEEE 802.11 wireless LANs," IEEE Proceedings of the 33rd Annual Hawaii International Conference on System Sciences, vol. 2, January 2000, pp. 1–10.

[13] F. Cali, M. Conti, and E. Gregori, "IEEE 802.11 protocol: Design and performance evaluation of an adaptive backoff mechanism," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 9, pp. 1774–1786, September 2000.

[14] P. Chatzimizios, A. C. Boucouvalas, and V. Vitsas, "IEEE 802.11 packet delay - A finite retry limit analysis," IEEE Global Telecommunications Conference GLOBECOM '03, vol. 2, December 2003, pp. 950–954.

[15] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, "Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols," IEEE Wireless Communications and Networking Conference, March 2005. [Online]. Available: http://www.orbit-lab.org/download/publications/Orbit_Overview.pdf

[16] M. Singh, M. Ott, I. Seskar, and P. Kamat, "ORBIT measurements framework and library (OML): Motivations, design, implementation, and features," First International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, Trento, Italy, February 2005. [Online]. Available: http://www.orbit-lab.org/download/publications/Orbit_OML.pdf

[17] M. Ott, I. Seskar, R. Siracusa, and M. Singh, "ORBIT testbed software architecture: Supporting experiments as a service," First International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, Trento, Italy, February 2005. [Online]. Available: http://www.orbit-lab.org/download/publications/Orbit_Software.pdf

[18] S. Ganu, H. Kremo, R. Howard, and I. Seskar, "Addressing repeatability in wireless experiments using ORBIT testbed," First International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, Trento, Italy, February 2005. [Online]. Available: http://www.orbit-lab.org/download/publications/Orbit_Repeatability.pdf

[19] *Cisco Aironet Wireless LAN Client Adapters Installation and Configuration Guide for Linux 340 and 350 Series*, Cisco Systems, Inc., 170 West Tasman Drive, San Jose CA 95134-1706 USA, 2002. [Online]. Available: http://www.cisco.com/univercd/cc/td/doc/product/wireless/airo_350/350cards/linux/instlcfg/icglttl.pdf

[20] H. Kremo, P. Spasojevic, and I. Seskar, "IEEE 802.11 medium access protocol: An experimental case study," WINLAB, Rutgers University, Piscataway, NJ, Tech. Rep. TR-276, May 2005.

[21] J. Bardwell, "Converting signal strength percentage to dBm values," WildPackets, Inc., November 2002.

[22] Z. Hadzi-Velkov and B. Spasenovski, "An analysis of CSMA/CA protocol with capture in wireless LANs," IEEE Wireless Communications and Networking, March 2003, pp. 1303–1307.