99

LA-UR- 01-2879

Approved for public release; distribution is unlimited.

Title: A COMPARATIVE EXPERIMENTAL STUDY OF MEDIA ACCESS PROTOCOLS FOR WIRELESS RADIO NETWORKS

Author(s):

Christopher L. Barrett, D-2
Martin "NMI" Drozda, D-2
Madhav V. Marathe, D-2

Submitted to:

ACM Symposium on Mobile Ad Hoc Networking & Computing
October 4-5, 2001
Long Beach, CA

LOS ALAMOS NATIONAL LABORATORY
3 9338 00789 0360

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this ar[illegible] the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the pu[illegible]d form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Form 836 (10/96)

LA-UR-01- ____

A Comparative Experimental Study of Media Access Protocols for Wireless Radio Networks

CHRISTOPHER L. BARRETT MARTIN DROZDA MADHAV V. MARATHE

LOS ALAMOS
NATIONAL LABORATORY
Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

A Comparative Experimental Study of Media Access Protocols for Wireless Radio Networks

CHRISTOPHER L. BARRETT[1] MARTIN DROZDA[1,2] MADHAV V. MARATHE[1]

May 24, 2001

Abstract

We conduct a comparative experimental analysis of three well known media access protocols: 802.11, CSMA, and MACA for wireless radio networks. Both fixed and ad-hoc networks are considered. The experimental analysis was carried out using GloMoSim : a tool for simulating wireless networks. The main focus of experiments was to study how *(i) the size of the network, (ii) number of open connections, (iii) the spatial location of individual connections, (iv) speed with which individual nodes move and (v) protocols higher up in the protocol stack (e.g. routing layer)* affect the performance of the media access sublayer protocols. The performance of the protocols was measured w.r.t. three important parameters: *(i) number of received packets, (ii) average latency of each packet, and (iii) throughput.* The following general qualitative conclusions were obtained; some of the conclusions reinforce the earlier claims by other researchers.

1. Although 802.11 performs better than the other two protocols with respect to fairness of transmission, packets dropped, and latency, its performance is found to (i) show a lot of variance with changing input parameters and (ii) the overall performance still leaves a lot of room for improvement.
2. CSMA does not perform too well under the fairness criteria, however, was the best in terms of the latency criteria.
3. MACA also shows fairness problems and has poor performance at high packet injection rates.
4. Protocols in the higher level of the protocol stack affect the MAC layer performance.

The main general implications of our work is two folds:

1. *No single protocol dominated the other protocols across various measures of efficiency.* This motivates the design of a new class of parameterized protocols that adapt to changes in the network connectivity and loads. We refer to these class of protocols as *parameterized dynamically adaptive efficient protocols* and as a first step suggest key design requirements for such a class of protocols.
2. *Performance analysis of protocols at a given level in the protocol stack need to be studied not locally in isolation but as a part of the complete protocol stack.* The results suggest that in order to improve the performance of a communication network, it will be important to study the entire protocol stack as a single algorithmic construct; optimizing individual layers in the 7 layer OSI stack will not yield performance improvements beyond a point.

---

[1] Corresponding Author. Los Alamos National Laboratory, P. O. Box 1663, MS M997, Los Alamos NM 87545. The work is supported by the Department of Energy under Contract W-7405-ENG-36. Email: `barrett,drozda,marathe@lanl.gov`.

[2] Work done while the author was a graduate research student in the Basic and Applied Simulation Science Group at Los Alamos National laboratory. Affiliation: Slovak Academy of Sciences, Institute of Control Theory and Robotics.

1 Introduction and motivation

Design of MAC layer protocols for wireless mobile networks has become an important area of research in recent years (See [BD+94, Ka90] and the references therein). An extreme form of wireless mobile networks are the ad-hoc networks – networks that do not rely on any fixed infrastructure, e.g., base stations. The upsurge in a variety of mobile computing devices such as laptops, PDAs, and other portables has caused an unprecedented interest in this form of communication. Early progress on multihop radio networks was funded by DARPA. This includes the **PRNET** (Packet Radio Network) [JT87], and **SURAN** (Survivable Adaptive Networks) [SW] projects. Interest in ad-hoc networks for mobile communications has also resulted in a special interest group for Mobile, Ad-hoc Networking within the Internet Engineering Task Force (IETF) [MC].

Network protocols in general need to fulfill a multitude of design and functional requirements, including, *(i) High throughput; (ii) Low average latency; (iii) Heterogeneous traffic (e.g. data, voice, and video); (iv) Preservation of packet order; and (v) Support for priority traffic.* (See [Sa95, RS96, Ra96, Pa97a, Ba98].) As ad-hoc networks lack fixed infrastructure in the form of base stations, fulfilling the above stated functional requirements becomes all the more difficult. Many MAC layer protocols have been proposed and designed to meet one/many of these criteria; the research area continues to be very active.

There are two basic issues that distinguish traditional LANs with Wireless LANs w.r.t. media access control protocols. The first is the well known *hidden terminal problem.* In this situation two data transmission sources try to communicate with a common node. Moreover, assume that the sources are not within radio range of each other. As a result, either source is oblivious of the fact that the other is trying to transmit data to the common node. This results in packet collisions or one of the sources backing-off after sensing a busy carrier. Despite clever back-off mechanisms the situation either ends up with a lot of collisions and hence poor throughput or inequitable resource allocation to one of the connections, leaving the other connection *starved.* A related problem is the *exposed terminal problem.* Here a station A listening finds that the medium is busy and hence refrains from sending. But it is quite possible that the station C to which it wants to send the signal is far away from the currently transmitting station B and thus can C can easily hear the transmission of A. Additionally, D which is hearing B is also far from A and thus receives the signal without much noise. The second basic issue is that while transmitting, the transceiver cannot simultaneously listen and hence makes collision detection and avoidance much more difficult. This has been referred to as the mute-deaf-time problem; most protocols do not address the exposed terminal problem at all and hence are inherently inefficient from the resource utilization standpoint.

A commonly known group of MAC protocols is based on the carrier sense multiple access (CSMA) paradigm. The idea behind this paradigm is to reserve transmission channel at the originator (source) by carrier sensing. Until recently CSMA based protocols supported only single channel communication, recently, multiple channel extensions have been proposed [NZD99]. Many protocols have been proposed to avoid the hidden terminal problems. Two notable examples are the MACA [Ka90] and MACAW [BD+94] protocols. MACA introduced a reservation system achieved with exchange of an RTS/CTS (Request To Send/Clear To Send) pair of control packets. MACAW also recognizes the importance of congestion, and exchange of knowledge about congestion level among entities participating in communication. An advanced back-off mechanism was proposed to spread information about congestion. Furthermore, the basic RTS/CTS/DATA reservation schema has become an RTS/CTS/DS/DATA/ACK schema with significantly improved performance. In these protocols message originators reserve reception area at the sink by exchange of RTS/CTS control messages. This is in contrast to CSMA where reservation was done at originators. This powerful method has a drawback of introducing small control packets into the network that later collide with other data, control, or routing packets. IEEE 802.11 MAC standard [OP] was designed with a reservation system similar to MACA or MACAW in mind. 802.11 has also improved fairness characteristics, however, in [LNB98] authors point out deficiencies in the fairness of this protocol, as well.

Because of the limited bandwidth of wireless channels message complexity of both MAC layer and network layer (routing) protocols needs to be kept low. Informally speaking, we define the message complexity of a protocol as the ratio of the number of data packets successfully transmitted to the total number of packets actually sent (including control packets, duplicates etc.). Latency is defined to be the average time it takes for a packet to reach its destination. Note that as defined, the definition does not distinguish between the type of packets received. Thus it is conceivable that a connection might be deemed to have good latency but might not deliver too many new packets. Thus a good protocol should have the following characteristics: (i) high throughput as measured by the total number of good packets received in a unit time and (ii) fairness. Intuitively speaking, high throughput implies low message complexity and low latency. The interplay between message complexity and latency with dynamically changing network connectivity, and traffic load is the main focus of the study described in this document. Additionally, we study the effect of (i) spatial locations of sources and sinks, (ii) injection rates of packets, and (iii) type of network on quality of service. We have also studied the impact of mobility on performance of MAC layer protocols. We chose the following MAC layer protocols to test: 802.11, CSMA, and MACA. All simulations were done in GloMoSim, a tool specialized in wireless networks simulation.

2 Summary of Results and Implications

We experimentally evaluate the performance of three well known MAC protocols in wireless radio networks. Both static and ad-hoc radio networks are considered. The goal is to see how (i) the network topology, (ii) the traffic injection rate, (iii) the spatial location of the source destination pairs, (iv) the effect of mobility (in case of ad hoc networks), all affect the performance of the protocols. *Moreover, we want to do this in settings where the results are interpretable; as a result to the extent possible, we have chosen very simple instances to effectively argue about an issue.*

2.1 Scenario Specific Results

Each distinct spatial source destination pair location is referred to as a scenario in this paper. We consider eight basic scenarios in the current draft. The first four scenarios are shown schematically in Figure 1. The results for each of these four scenarios are briefly discussed below; the results for the remaining four scenarios are to be found in subsequent sections.

1. The first scenario was to verify performance of MAC layer protocols under hidden terminal situation. Results show that CSMA inequitably assigned resources to the two connections on short time scales. On the other hand CSMA performed quite well in terms of latency and in fact the lowest latency among all the three protocols for this case. MACA had a very high latency as well as inequitable resource assignment. 802.11 had worse latency than CSMA but was better than MACA. On the other hand, it allowed the most equitable access to the media and had the best throughput.
2. In the second scenario we tried to test the MAC layer protocols on performance in crossing traffic. What we have found out is that in this case CSMA does slightly better under lower injection rate

than in the previous case, however, with increasing injection rate fairness becomes a problem again. Performance of MACA was satisfiable under low injection rate, but plummeted under high injection rate. 802.11 did okay under this test – even latency for this protocol was very low.

3. The third scenario was chosen to test our hypothesis that CSMA should perform well if the packet flow paths of the two connections are disjoint causing no perceivable interference. Our results show that this is indeed the case.
4. The fourth scenario can be thought as a mix of Scenario 1 and 3 and thus intuitively one would expect performance that is somewhere between the performance of the two scenarios.

A qualitative explanation of many of the results can be given. For instance, CSMA has low overhead since it does not have the RTS/CTS control mechanism; this makes collisions more likely but on the other hand allows for lower latency (at least for the connections that are given access) and adequate throughput for the connections that are scheduled. 802.11 has RTS/CTS mechanism; the overhead that such a control mechanism causes for small packets is evident from the degradation of 802.11 for small packet sizes. MACA appears to be probably the worst overall: it has high latency, inequitable resource allocation.

Three basic quality of service metrics were considered in our experiments: (i) Throughput, (ii) Latency and (iii) Packet Loss. From the experiments, it can be seen that CSMA has the best performance in terms of latency while 802.11 is best in terms of throughput and packet loss characteristics.

Fairness: All the three protocols assigned resources in an inequitable fashion although 802.11 appears to be the best among the three protocols. The current experiments were not designed to study the short term fairness problem – we hope to incorporate this in subsequent studies.

2.2 Protocol Interaction

Another important goal of the experimental analysis is to see the interaction of the different layers in the protocol stack and its effect on the network performance. While intuitively it is clear that different levels in the protocol stack should affect each other in most cases; to the best of our knowledge a thorough understanding of this interaction is lacking. The only related reference in this direction that we are aware of is the work of Balakrishanan et. al. [BS+97, KKB00]. The authors were specifically considering TCP/IP protocol and have devised an elegant snoop protocol that conceptually sits between the transport layer and the network layer to overcome this problem. In [KKB00] also point out how short term fairness of the MAC can affect the TCP/IP performance which in turn can affect the overall performance of the communication system.

Here, we study the interaction between the routing layer and the MAC layer. The results show the paths chosen by routing protocol can significantly affect the MAC layer protocol especially with regards to inequitable assignment of channel resources. Combined with the results of [KKB00], the results show that discussion about the performance of a MAC layer cannot be usually carried out without putting it in context of the other protocols in the stack. Specifically, we show that the performance of CSMA depends on what paths are chosen by the routing protocols to send packets. Intuitively, it is obvious that if the the virtual paths corresponding to two connections intersect then they should yield different dynamics than when the paths do not. This is exactly what happens; moreover, it is hard to predict exactly the paths the routing protocols are going to take given the randomized local nature. As a result all the MAC protocols exhibit varying levels of performance.

2.3 Qualitative Conclusions

The work presented here has two important conclusions:

1. *No single protocol dominated the other protocols across various measures of efficiency.* This motivates the design of a new class of parameterized protocols that adapt to changes in the network connectivity and loads. We refer to these class of protocols as *parameterized adaptive efficient protocols* and as a first step suggest key design requirements for such a class of protocols.
2. *The protocol stack needs to be treated as a monolithic algorithmic construct if we wish to improve the performance of the communication network.* Optimizing individual layers might be only possible to a certain point and beyond that might cause completely unintuitive side effects.

3 Experimental Setup: Static case

In the next several experiments described in the following, we tried to obtain empirical results on the performance of three well known MAC layer protocols namely, 802.11, CSMA, and MACA.

We studied the performance of the three protocols under eight different scenarios. Each scenario was designed to test a distinct hypothesis. The eight scenarios are shown in Figures 1 and 6, and are described below.

1. **Scenario 1: Hidden terminal Scenario.** The first scenario is a hidden terminal type of scenario. We have distinct sources for the two connections, however, the sinks are the same. Motivation for this scenario was to test capture-of-the-receiver avoidance capabilities of the three MAC layer protocols.
2. **Scenario 2: Crossing Traffic Scenario.** The second scenario represents a case when sources and sinks are different but the traffic is crossing. Motivation for this scenario was to test quality of the three MAC layer protocols for crossing traffic. In case of use of a congestion unaware routing protocol, crossing traffic can lead to increased dropping of packets.
3. **Scenario 3: Independent paths scenario.** The third scenario represents distinct sources and sinks, but the paths are disjoint, i.e., there should be a minimum of interacting traffic. Motivation for this scenario was to have a base case with mutual minimum interference.
4. **Scenario 4: Closely positioned sources and sinks.** In the fourth scenario we show a case where sources are distinct, but sinks, though, distinct are closely positioned on the grid. Motivation for this scenario was to test the three protocols for quality under high congestion on the receiver side.
5. **Scenario 5: Multiple Shortest Paths.** The fifth scenario depicts a situation in which if the routing protocol chooses the shortest path there would be no interference between connection 1 and 2. Should the routing protocol choose a less than optimal routing interference between connections will arise.
6. **Scenario 6: Long Term Fairness.** The scenario is designed to study how the protocols treat various connections w.r.t. access to channel resources. The current experiment is designed for studying long term fairness; short term fairness is interesting and will be studied subsequently.
7. **Scenario 7: Effect of Number of connections on Quality of Service.** The experiment was designed to see how the quality of service is affected with increased number of connections. In this case the

graph is kept constant. A similar study is done by increasing the graph size; due to lack of space the results are reported in a complete version of this document.

8. **Scenario 8: Effect of Mobility.** The final experiment studies the effect of mobility on the protocol performance. For this we moved the transceivers at various speeds from a starting configuration and observed the latency and throughput of the ensuing system.

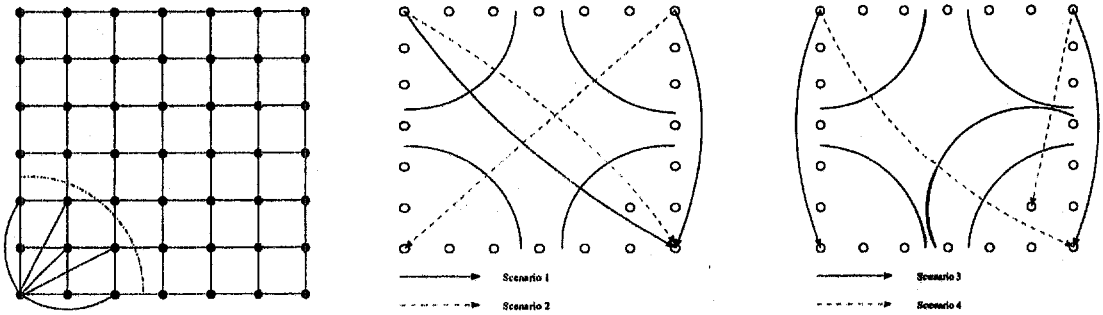

Figure 1: The first figure schematically illustrates the connectivity of the graph. For clarity only the edges incident on the node $(0,0)$ are shown. The dotted arc shows the transceiver's radio range. The next two figures show a schematic of the four scenarios used in the experimental analysis. Each scenario consists of two connections. The connections are schematically illustrated by drawing an arrow from the source to its destination. Different line types imply different scenarios. Each figure shows the source and sink position for two scenarios. The reason for drawing arrows instead of merely marking the positions is to give a rough idea of the paths the packets are likely to take when going from a source to the destination.

We now describe the details of the parameters used in these scenarios:

1. **Network Characteristics:** In each scenario we have kept the following parameter constant: (i) number of nodes 49, (ii) Network Topology: The network is schematically shown in Figure 1. The transceivers were placed on a 6×6 grid; i.e. by scaling we assume that the coordinates of transceivers are (x,y) $0 \leq x, y \leq 6$. Each transceiver is equipped with an omnidirectional antenna. The antenna range is 2.5 units. Thus a transceiver positioned at $(0,0)$ is adjacent to transceivers $(1,0),(2,0),(0,1),(0,2),(1,1),(1,2),(2,1)$. Given the symmetry, this implies a maximum degree of seven for any node. This meant that maximum possible connectivity was 20. The minimum node degree is 7 and is that of corner nodes. In Figure 1 we have drawn partial circles around the corner transceivers to schematically show this. (iii) **Number of connections:** 2, (iv) **Routing protocol** : AODV.

2. There was no movement of nodes. Nodes were positioned on a square area of 600×600 meters.

3. The initial packet size was 512 bytes, the initial number of packets was 1,000, and the initial injection interval was 0.1s. We have reduced the packet size by a factor of 2 and increased the number of packets by a factor of 2 every time the injection interval was reduced by a factor of 2. For example, if the injection interval was halved to 0.05s then the new packet size was 256 bytes and the new number of packets was 2,000. This allowed us to keep the injection at input nodes constant in terms of bits per second.

4. The bandwidth for each channel was set to 1Mbit. The propagation pathloss model was two-ray. The transmission radius for a transceiver was 250 meters (corresponds to 2.5 units).

5. For any vertical connection at least two hops were required for a packet to reach its sink, for any diagonal connection at least three hops were necessary to reach its sink; the minimum path length in terms of number of edges is 3 for vertical connections, and 4 for diagonal connections. By vertical connection we mean those running up-down, and by diagonal connection we mean those running from right upper corner to lower left corner, or from upper left corner to lower right corner as it is shown in Figure 1.

6. To keep the simulation time 100 seconds, and the total size of packets in traffic constant, we halved the size of packets with each doubling of packet injection interval.

7. Hardware used in all cases was a Linux PC with 512MB of RAM memory, and Pentium III 500MHz microprocessor.

Independent and Dependent variables The independent (input) variable for a given scenario, connection and protocol triple was the injection interval for packets. The following three pieces of information were collected: (i) Average end to end delay for each packet as measured in seconds, (ii) Total number of packets received, (iii) Throughput in bits/second.

4 Results and Analysis

We now summarize the experimental results for each of the four scenarios discussed in the preceding section. We have summarized the results of running the experiments in form of graphs. The graphs show basically the same data as a function of varying injection rate. There is one graph per connection.

4.1 Scenario 1 – hidden terminal

Results[3] [4] for this scenario are presented in Figure 2. The broad conclusions for Scenario 1 are:

1. Performance of 802.11 is fairly okay at low injection rates. CSMA and MACA do not perform very well; the main problem is that of fairness. In particular, connection one has better reception at the common sink; for all practical purposes, connection 2 never gets a turn to transmit. The result confirms the well established *hidden terminal problem* of CSMA protocol. In contrast, 802.11 treats the packet streams from both the connections fairly; the throughput and delays are very similar across connections.

2. On the other hand average end-to-end delay of CSMA is better 802.11 and MACA in some cases. This is mostly due to lower message complexity of MACA and 802.11 protocols.

3. MACA's performs deteriorates rapidly with increased injection rate. We had problems to run GloMoSim with high injection rate for CSMA. We could not find a logical reason to explain this.

---

[3] Note that we have used logscale for y axis in all figures.

[4] In some cases the number of packets that a MAC layer protocol transmitted was zero. This translates as infinite latency for packet delivery. This fact is shown as vertical line in our graphs.

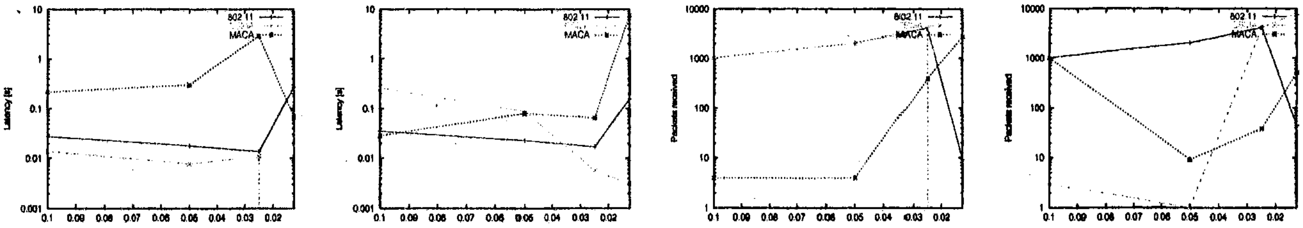

Figure 2: Scenario 1: Hidden terminal scenario. Two parameters are plotted for each connection: latency and number of packets received. The first two figures are for latency and the third and fourth are for number of packets received. On x-axis we plot the injection rates.

4.2 Scenario 2 – distinct sources and sinks with crossing traffic

Results for Scenario 2 are presented in Figure 3. The main rationale for studying this scenario along with Scenario 1 was that the hidden terminal effect is likely to occur at an intermediate destination. The broad conclusions for Scenario 2 are:

1. In contrast to Scenario 1, here 802.11 performs quite poorly at high injection rate: both in terms of latency and throughput. On the other hand CSMA does quite well at least for low injection rates.
2. MACA had problems allocating channel resources equitably. CSMA and 802.11 appear to behave similarly for this criteria.
3. On the other hand, CSMA outperformed 802.11 and MACA in terms of latency. Latency versus injection rate does not show any clear trend, except in the case of MACA for which latency progressively increases with injection rate. In contrast, for 802.11 latency actually decreased with increased injection rate. Similar situation in fact also occurs in Scenario 1.

One possible explanation for the results is that the data paths chosen by the routing protocol did not intersect.

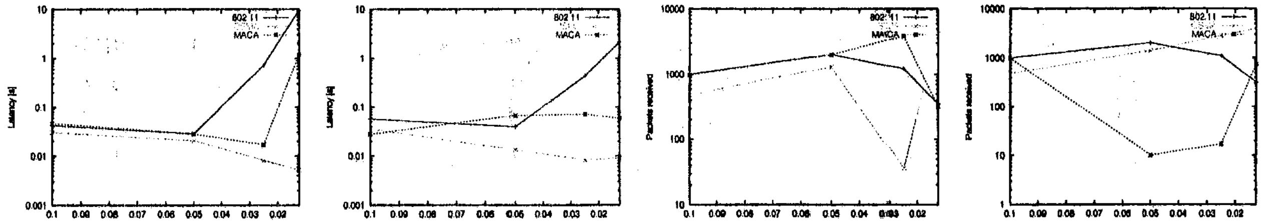

Figure 3: Scenario 2: Crossing Traffic Scenario. As in previous figure, x-axis plots injection rate and y-axis plots latency (first two plots) and number of packets received (the next two plots).

4.3 Scenario 3 – distinct sources and sinks, disjoint paths

Results for this scenario are presented in Figure 4. Intuitively, this scenario would be expected to yield the best results for protocol with least control packet overhead. The conclusions for this scenario are:

1. In this experiment we can see that 802.11 does very well in terms of fairness and handling of high injection rate.

2. Fairness still appears to be a problem for CSMA. Given that the paths share minimal number of nodes, the results appear somewhat surprising. On the other hand, CSMA was able to keep near maximum throughput for one connection.
3. Performance of MACA became poor with increased injection rate. With increased injection rate the performance of MACA in terms of latency becomes intolerable.

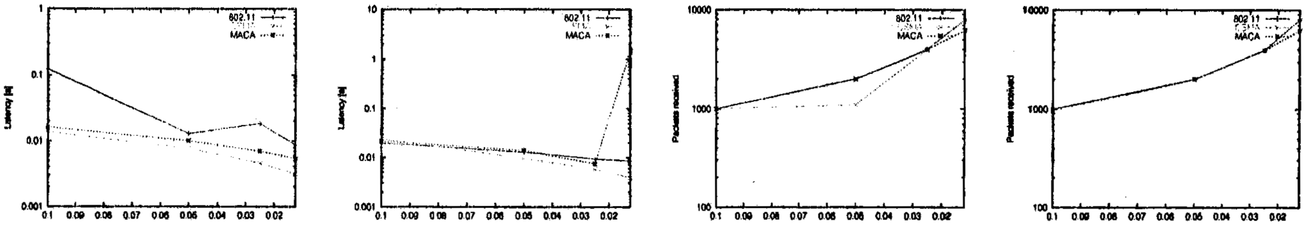

Figure 4: Scenario 3: Independent paths scenario. The description of x and y-axis is similar to the previous figures.

4.4 Scenario 4 – Distinct Sources, Closely Positioned Sinks

Results for this scenario are presented in Figure 5.

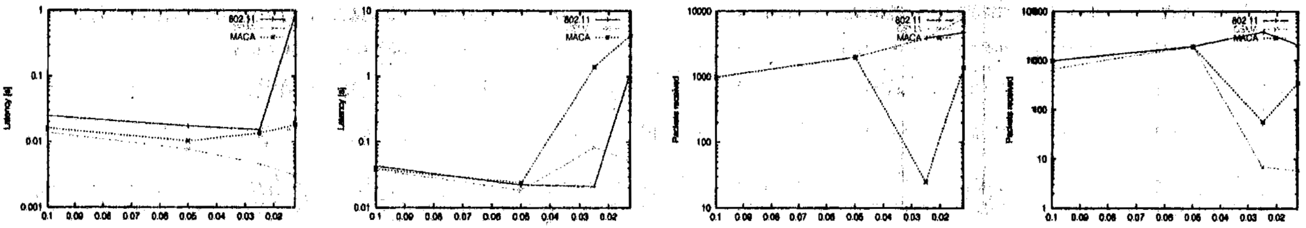

Figure 5: Scenario 4: Closely positioned sources and sinks. The description of x and y-axis is similar to the previous figures.

- In this case 802.11 performs well for lower injection rates.
- Latency for CSMA is very low. With increased injection rate the protocol trends to drop a quantity of packets for one of the connections.
- Latency for MACA increased with the injection rate. Performance of MACA plummeted in terms of dropped packets with increased injection rate.

4.5 Scenario 5 – Multiple Short Paths and Effect of Routing Layer

Recall that in this scenario, our objective was to show the effect of the routing layer on MAC layer. In order to do this, the end points of the two connections were placed in such a way that if the routing protocol chooses the shortest path there would be no interference between connection 1 and 2. If the routing protocol chooses a less than optimal routing interference between connections will arise. Parameters of the network for this experiment were slightly different from those in the previous scenarios.

Here is a short summary. We used the same number of nodes 49 to produce a grid of 7×7 nodes. The nodes, however, were positioned closer at a distance of 50 meters from each other gaining a physical size of the grid of 300×300 meters. This fact is depicted in Figure 6. Note that transmission radii from nodes in the very left column are just short to achieve direct communication with nodes in the very right column, and vice-versa. We have only used CSMA as the underlying MAC layer protocol. We have only used injection interval of 0.1s; the number of packets then was 1,000 and the packet size was 512 bytes. All other parameters were the same. We use AODV as the routing protocol, the channel bandwidth was 1Mb, the propagation pathloss model was two-ray, the transmission radius of a transceiver was 250 meters, the simulation time was 100s, and the hardware used was identical.

In this scenario we tried to capture only qualitative results as opposed of quantitative results of the previous four scenarios. The setup was as follows. We have two connections: one going from the upper left corner to the lower left corner, and the other one going from the right upper corner going to the lower right corner. From this and the network setup it is obvious that if the routing protocol finds the shortest paths between the sources and sinks there will be no interference between the two connections. Vice-versa if the routing protocol finds a less than optimal path for routing packets between sources and sinks than interference will result. For example, if the routing protocols does not find the shortest path from the upper left corner to the lower left corner, i.e., it will use some nodes not belonging to the very left column than interference with the other connection will results.

We have observed several modes of operation. One of them happened when the routing protocol found the shortest path for the connections. In this case, the number of received packets at sinks was 1,000, i.e., 100%. This fact is depicted in Figure 6 by nodes in thick squares. These nodes were the actual nodes used for routing the packets for respective connections. In the second case the routing protocol did not find the shortest path for one of the connections. The actual node that was used to route packets is in thick circle in Figure 6. This fact caused interference between the two connections and resulted in delivering only one packet for the connection for which the routing was not optimal. The number of received packets for the other connection was 1,000, i.e., all.

This scenario resulted in four basic modes of operation which we summarize in Table 1. We ran the scenario 15 times, each time with a different seed. We counted the different modes as follows:

- We counted 1,000 received packets for connection 1, and 0 received packets for connection 2 the same as 0 for connection 1 and 1,000 for connection 2, i.e., in general, we regarded symmetric results to be the same.
- If the number of packets received for a connection was e.g. 995 we counted it as 1,000, i.e., in general, we discarded small fluctuations and regarded such results as identical.

Ratio of received packets	**Number of cases**
1,000:0	6
1,000:500	5
500:500	3
1,000:1,000	1

Table 1: Scenario 5, modes of operation – summary.

From Table 1 we can see that the routing protocol (AODV) [5] managed to find the shortest path only in one case, and that resulted in 100% of packets received for both connections.

The main conclusion we can draw from this scenario is that routing layer can influence overall performance of MAC layer. In our case, had the routing protocol found the shortest path, the two connection would not interfere. We showed to what extent can the suboptimal routing path change the overall throughput. The throughput for connection one varied from 100% to almost 0%. *Admittedly, the situation used to illustrate is very simple* but designed on purpose to illustrate the basic interaction.

An important implication of this result and the associated results in [BS+97] on the design of snoop protocols suggest that optimizing the performance of the communication network by optimizing the performance of individual layers is not likely to work beyond a certain point. We need to treat the entire stack as a single algorithmic construct in order to improve the performance beyond a certain point. Specifically, optimizing a particular layer might improve the performance of that layer locally but might produce non-intuitive side effects that will degrade the overall system performance. The issue is likely to become more important in ad hoc networks where the topology is changing constantly and hence it is not even easy to discern what shortest paths mean.

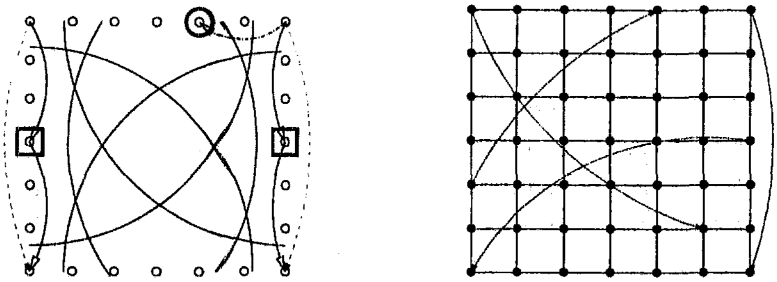

Figure 6: Figures used in two experiments – (i) effect of routing protocols (Scenario 5) and (ii) fairness of resource assignment (Scenario 6). The first figure shows the possible paths used by the protocol. The second figure shows the four connections studied for equitable resource assignment study.

4.6 Scenario 6 – Long Term Fairness

We designed a simple experiment to study the equitability of resources assignment by the three MAC protocols. Figures 7 and 8 shows 2 (for high injection rate) and 10 (for low injection rate) independent runs (with independent random seeds) and its results for the three protocols. The following qualitative conclusions can be obtained:

1. All the protocols do an inequitable assignment of channel resources for high injection rates. We have deliberately refrained from calling this unfair: what does it mean to be fair is not obvious and has been subject of a extensive research in the past in Economics and Social Science.

2. For low injection rates 802.11 assigned resources equitably. On the other hand CSMA and MACA had a wide variation.

3. At least two notions of equitable resource allocations can be formulated: one in which we see how the protocol does in a particular run and one in which measure the relative resources assigned to each

---

[5] We have run this experiment also with DSR. In this case the routing performance was worse then that of AODV. We have observed three modes of operation.

connections over a given set of runs. Using the other measure CSMA and MACA appear to have a more equitable resource assignment.

4. Many researchers have in the past designed specific algorithms and argued (heuristically or formally) about the fairness of protocols. We believe that the topic deserves more attention. For instance [VBG99] propose distributed fair scheduling algorithm. The essential idea is to assign resources to each flow in proportion to the amount that is backlogged for that particular flow. In [NK+99], the authors have discussed per-node versus per-flow fairness. We merely point out that, each such proposed mechanism can have subtle side effects; the goal is merely point out undertaking a more in-depth study.[6]

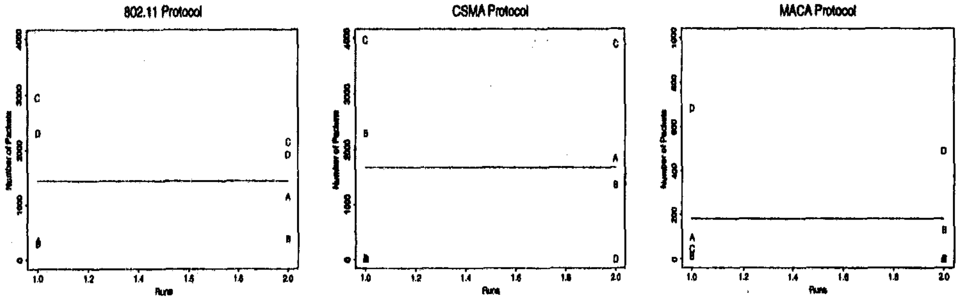

Figure 7: Scenario 6: Long term fairness. Equitability of resource assignment. A, B, C and D are the four connections. The line in the center is the grand mean. Thus the plot shows one relative way of comparing how equitable the resource assignment was for the protocols. The results are for ***high injection rates***. The first figure is for 802.11, the second one is for CSMA and the last is for MACA. Only two runs are plotted for clarity.

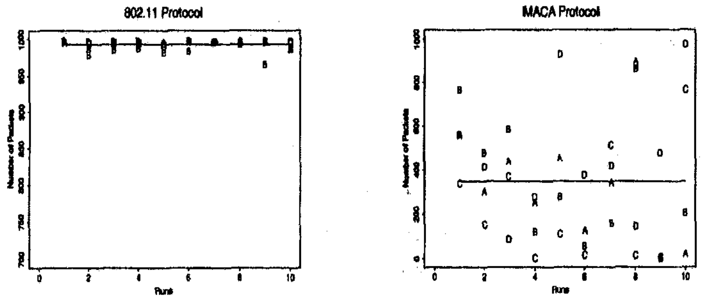

Figure 8: Scenario 6: Long term fairness. Equitability of resource assignment. 10 runs and low injection rates.

---

[6] A very simple example will make the point. Consider for instance an adversary, who wishes to slow down a network without any goal of transmitting useful information. Furthermore, imagine the adversary to have control over the protocol stack. The adversary can easily compromise the network's good throughput by not implementing a voluntary back off scheme and thus flooding the intermediate nodes. If per flow fairness is implemented this will end up giving unusually high resources to this connections making the other connections have low throughput.

4.7 Scenario 7 – Increasing Number of Connections and Quality of Service

In this experiment, we studied influence of increasing number of connections on quality of service. Quality of service was captured in terms of latency and number of packets received.

Number of connections was 2, 4, and 6. First, we ran the experiment with two base connections and gathered information on these. Then we added two other connections, but we gathered information only on the base connections. At the end, we added again two connections, bringing the number of connections to six, we ran the experiment and we gathered information on the two base connections. Results are depicted in Figure 9.

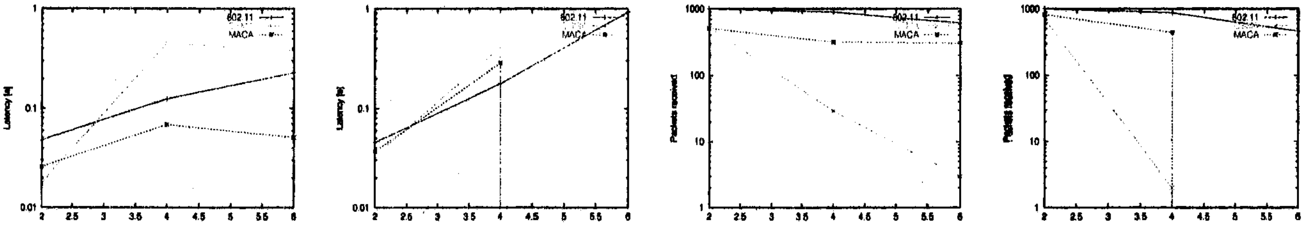

Figure 9: Scenario 7: Effect of Number of connections on Quality of Service. Two basic parameters are plotted: Latency and Packets received. The x-axis denotes the number of connections and the y-axis shows the latency/packets received.

4.8 Scenario 8 – Effect of Mobility on Performance of MAC layer protocols

In this experiment we tried to obtain some knowledge on dependence between the speed with which nodes are moving and performance of MAC layer protocols. We have uniformly increased the speed from 0 m/s to 40 m/s. In the beginning of each simulation we positioned nodes onto a grid of 10×10 nodes. Then we used the random-waypoint movement protocol to simulate mobility[7]. We have tested this scenario with 802.11, CSMA and MACA. Interval of packet injection was 0.1 second, packet size was 512 bytes, there were two connections active at each time, number of packets injected by each connection was 1,000. The physical size of the underlying area was 1,000×1,000 meters. The routing protocol used was AODV. The results are shown in Figure 10. We can see that 802.11 does very well in terms of both latency and packets received. MACA does slightly better than CSMA, however, in both cases the performance is not good.

References

[Ba98] H. Balakrishnan. Challenges in Reliable Data Transport Over Heterogeneous Wireless Networks. Ph.D. Thesis, Department of Computer Science, University of California at Berkeley, 1998.

[BS+97] H. Balakrishnan, S. Seshan, E. Amir, R. Katz. Improving TCP/IP Performance over Wireless Networks. Proc. 1st ACM Conf. on Mobile Computing and Networking, Berkeley, CA, November 1995.

[BD+94] V. Bharghavan and A. Demers and S.Shenker and L. Zhang. MACAW: A Media Access Protocol for Wireless LANs. *Proc. 1994 SIGCOMM Conference,* London, UK, pages 212–225, 1994. citeseer.nj.nec.com/bharghavan94macaw.html

---

[7] Random waypoint in GloMoSim implements movement on trajectories. A random destination is produced, mobile node moves to the destination with a preset speed, stays there preset amount of time, and then new random destination is produced.

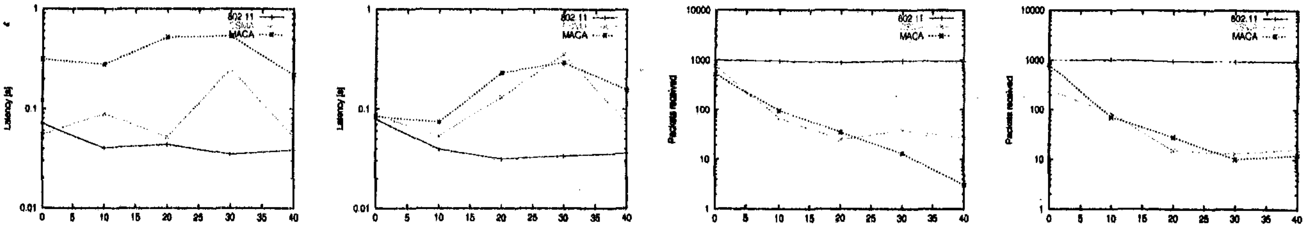

Figure 10: Scenario 8: Effect of Mobility. Latency (first two figures) and Packets received, connection 1 and 2. All the graphs plot the change w.r.t. to speed; the first two graphs plot latency while the third and fourth plot packets received.

[BM+98] J. Broch, D. Maltz, D. Johnson, Y. Hu, and J. Jetcheva. Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols. *Proc. 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking,* ACM, Dallas, TX, October 1998.

[BT+99] L. Bajaj, M. Takai, R. Ahuja, K. Tang, R. Bagrodia, and M. Gerla. GloMoSim: A Scalable Network Simulation Environment. UCLA Computer Science Department Technical Report 990027, May 1999.

[JT87] J. Jubin and J. D. Tornow. The DARPA Packet Radio Network Protocols. *Proceedings of the IEEE,* 75(1), pp. 21-32, Jan, 1987.

[Ka90] P. Karn. MACA - a new channel access method for packet radio. MACA - a new channel access method for packet radio ARRL/CRRL Amateur Radio 9th Computer Networking Conference 1990

[KT75] L. Kleinrock and F. Tobagi. Packet Switching in Radio Channels: Part I - Carrier Sense Multiple Access Modes and their Throughput-Delay Characteristics, *IEEE Transactions in Communications, COM,* 23(12), pp. 1400-1416, 1975.

[KKB00] C. Koksal, H. Kassab, and Hari Balakrishnan An Analysis of Short-Term Fairness in Wireless Media Access Protocols. *Proc. ACM SIGMETRICS,* June 2000. MIT-LCS-TR-807, May 2000.

[LC01] Y. Lin and I. Chlamtac, *Wireless and Mobile Network Architectures,* Wiley Computer Publishing, New York, 2001.

[LNB98] S. Lu, T. Nandagopal, and V. Bharghavan. A Wireless Fair Service Algorithm for Packet Cellular Networks. *ACM Mobicom'98,* Dallas, TX. October 1998.

[NK+99] T. Nandagopal, T. Kim, X. Gao and V. Bharghavan. Achieving MAC layer Fairness in Wireless Packet Networks. *ACM Mobicom'99.*

[MC] J. Macker and S. Corson. Mobile ad-hoc networks (MANET). IETF Working Group Charter. www.ietf.org/html.charters/manet-charter.html

[NZD99] A. Nasipuri, J. Zhuang and S. R. Das. A Multichannel CSMA MAC Protocol for Multihop Wireless Networks. *Proc. IEEE Wireless Communications and Networking Conference (WCNC),* New Orleans, September, 1999.

[OP] B. O'Hara and A. Petrick *802.11 Handbook, A Designer's Companion.* IEEE Press, 1999.

[Pa97a] V. Paxson. Measurements and Analysis of End-to-End Internet Dynamics. Ph.D. Thesis, Department of Computer Science, University of California at Berkeley, 1997.

[RS96] S. Ramanathan and M. Steenstrup, A survey of routing techniques for mobile communication networks, Mobile Networks and Applications, 1-2, pp. 89-104, 1996.

[Ra96] T.S. Rappaport. Wireless Communications. Prentice-Hall, 1996.

[Sa95] M. Satyanarayanan. Fundamental challenges of mobile computing. Proc. ACM Symposium on Principles of Distributed Computing, 1995 (PODC'95 invited lecture).

[SW] N. Schacham and J. Westcott. Future directions in packet radio architectures and protocols. Proceedings of the IEEE, 75(1), pp. 83–99, January 1987.

[VBG99] N. Vaidya, P. Bahl and S. Gupta. Distributed fair Scheduling in a Wireless LAN. *Proc. ACM MOBICOM 99.*

[WS+97] J. Weinmiller, M. Schlager, A. Festag and A. Wolisz Performance Study of Access Control in Wireless LANs - IEEE 802.11 DFWMAC and ETSI RES 10 Hiperlan. *Mobile Networks and Applications*, pp. 55-67, 2, (1997).