LA-UR-01-6217
Approved for public release; distribution is unlimited

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

CHRISTOPHER L. BARRETT    MARTIN DROZDA    MADHAV V. MARATHE

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

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

## 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. four important parameters: *(i) number of received packets, (ii) average latency of each packet, (iii) long term fairness and (iv) throughput*. The following general qualitative conclusions were obtained; some of the conclusions reinforce the earlier claims by other researchers. Our results provide in many cases a plausible explanation of these results.

1. Typically, all protocols degrade significantly at higher packet injection rate. Moreover, the degradation often happens rather sharply.
2. In general while the performance of 802.11 was better than CSMA at lower injection rates, the performance of 802.11 is worse than that of CSMA at higher injection rate, on the other hand, CSMA assigns inequitable amount of resources; in this regard 802.11 performs quite well.
3. MACA typically was dominated either by CSMA or by 802.11 w.r.t. any of the performance measures.
4. Protocols in the higher level of the protocol stack affect the MAC layer performance.

The main general implications of our work is that *No single protocol dominated the other protocols across various measures of efficiency.* In other words the performance of protocols depends on all of the parameters mentioned above. 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.

---

[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 Informatics.

# 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 multi-hop radio networks 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.

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 intervals 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 for simulating wireless communication networks.

The rest of the report is organized as follows. In Section 2 we give an overview of the three MAC layer protocols – 802.11, CSMA, and MACA protocols used in our simulations. Section 3 summarizes our results and offers qualitative explanations. In Section 4, we describe the experimental setup. Section 5 reports the specific results for each scenario and provides a qualitative analysis. In particular Section 5.1 considers the generalized hidden terminal scenario, Section 5.2 studies the effect of connectivity and Section 5.3 the effect of sparsity. Section 6 briefly discusses two extensions (i) the performance when the number of connections is varied (Section 6.1) and (ii) the effect of mobility (Section 6.2). Finally concluding remarks are presented in Section 7.

# 2 MAC layer protocols: Issues and Description

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 transceiver $A$ wants to send a message to transceiver $C$. Following the protocol, it first checks to find and finds that the medium is busy. It thus backs off and refrains from sending any packets. 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 receive packets from $A$ without much interference. Additionally, $D$ which is receiving messages from $B$ 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.

## 2.1 CSMA

CSMA is an acronym for Carrier Sense Multiple Access. As the name suggests, this protocol exploits capability of transceivers to listen to the on-going traffic in the adjacent area. This information is used to decide whether to start its own transmission, or whether to postpone this transmission until the channel gets idle, i.e., there is no carrier sensed. By monitoring the carrier CSMA improves efficiency by lowering collisions with neighboring transceivers. However, if CSMA senses carrier, the protocol has a built-in mechanism for delaying transmission. These mechanism are referred to as *back-off* algorithms. These come in several flavors – we refer reader to [Ra96] for details. CSMA/CA, where CA stands for Collision Avoidance applies similar mechanism to those used in MACA, notably, use of RTS/CTS control packets. Unlike CSMA/CD, where CD stands for Collision Detection, in CSMA transceivers do not monitor the carrier during their own transmission. CSMA/CD usually implements this feature by interrupting the transmission (single channel), or transceivers need to have *listen-while-talk* capability.

## 2.2 MACA

MACA uses a different approach for channel reservation than CSMA. Unlike CSMA, MACA does not reserve the channel at the originator of a transmission but rather at the destination for the transmission. This is done by exchange of an RTS/CTS pair of control packets. To start a transmission to a chosen destination, originator of the transmission sends an RTS (Ready-To-Send) control packet to the destination. In case the channel at the destination is idle, a CTS (Clear-To-Send) control packet is returned, and the transmission can start. Should that not be the case, binary exponential back-off algorithm is applied. This algorithm postpones the next RTS request, and the delay period for this grows with the number of unsuccessful RTS requests. The motivation behind MACA was observation that congestion mostly occurred at the destination rather than at the origin of transmission. MACA has been considerably improved in MACAW, see [BD+94].

## 2.3 802.11

The 802.11 standards consists of not one but three basic specifications: (i) the physical layer, (ii) the medium access layer control specification and (iii) the power saving functionality that operates on both layers. An 802.11 WLAN has two basic working modes: ad hoc mode that allows only peer to peer communication and an infrastructure mode where a control infrastructure such as the base station be accessed for control purposes. 802.11 allows two control/access schemes that are allowed.

The first scheme is the centralized scheme called the *Point control scheme (PCF)* in 802.11 where in a base station like control entity controls the access to the medium. Informally speaking, whenever mobile node wants to transmit, it requests access to the medium until the base station gives it explicit signal to do so. Centralized schemes need extensive information to be stored at the base station (or a control node in general) and thus are not robust to single point failures.

The second scheme is a distributed scheme and is called the distributed coordination function (DCF) in 802.11. Under this scheme, each transceiver wishing to transmit the information does local polling and communication to obtain access to the medium. Being inherently distributed the protocol does not require large amounts of information about the medium such as currently assigned frequencies etc. to be stored. On the other hand, the scheme does not, in general, provide QoS guarantees.

Our experiments used DCF access scheme due to its suitability for ad-hoc networks. This scheme is based on a variant of the CSMA/CA (Collision avoidance) protocol called the DFWMAC (Distributed Foundation Wireless MAC). Roughly speaking in this protocol, the station that intends to transmit and senses that the channel is busy waits for the end of the ongoing transmission, then waits for a time period of DIFS length (DCF Interframe spaces) and then randomly selects a time slot within the back-off window. The basic access mechanism can optionally be extended to the RTS/CTS message exchange mechanism followed by data delivery. To justify the additional overhead caused by the control packets, their usage depends on the size of the packet to be transmitted. More information on 802.11 can be found in [802.11].

## 3 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 interval, (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; hence to the extent possible, we have chosen very simple instances to effectively argue about an issue.*

### 3.1 Scenario Specific Results

For now, view a scenario as a combination of, the injection rate, spatial connection locations and network topology. We have considered three basic scenarios – each scenario consists of a number of sub-scenarios. Alternatively, each scenario can be viewed as an experimental design set up to verify/test a certain hypothesis. The experimental designs are discussed in detail in Section 4. We first discuss these results.

1. The first scenario was created to verify performance of MAC layer protocols under hidden terminal situations. For this we designed the basic hidden terminal sub-scenario as well as extensions of this idea to the case of multi-hop networks. Results show that CSMA inequitably assigned resources to the two connections over individual runs. On the other hand CSMA performed quite well in terms of latency and in fact had 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. See Section 5.1 for more details.

2. The second scenario was to test behavior under low and high connectivity. We can see that the protocols fail to perform in extreme situations. Here we have a case where communication depends on a few isolated nodes that convey packets between clusters of nodes. CSMA dominates in this scenario. 802.11 failed in both the latency and packets received. MACA's performance is quite poor under high injection rate. It shows up that exchange of RTS/CTS pair is not very efficient in See Section 5.2 for more details.

3. The third scenario was to convey results on the effect of grid width and sparsity. CSMA shows domination in terms of latency. MACA shows low latency but only for low injection rates. The reason is the overhead caused by exchange of RTS/CTS pairs. 802.11 performs well in terms of packets received but at extremely high injection rates the protocol performance plummets. CSMA on the other hand shows improved performance with increased injection rate. See Section 5.3 for more details.

4. All the protocols do an inequitable assignment of channel resources for low injection interval. 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.

5. For high injection interval 802.11 assigned resources equitably. On the other hand CSMA and MACA had a wide variation.

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

7. 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.[3]

### 3.2 Explanation and conclusions

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 and inequitable resource allocation. The main conclusions of our study include the following:

1. The network connectivity, spatial location of connections, injection rate and packet size all play a crucial role in determining the performance of a media access protocol. While, the effect of last two parameters has been studied earlier to some extent [WS+97, BD+94], the effect of first two parameters has not been extensively studied to the best of our knowledge.

2. Mobility deteriorates the performance of all the protocols. Although not surprising, the extent of its effect is certainly worth noting. Due to the RTS/CTS/ACK protocol of 802.11, this protocol has substantially superior performance compared to the other two. See Section 6.2 for additional explanation of plausible reasons.

---

[3] 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.

3. In general the following broad conclusions can be drawn: (i) higher injection rates, (ii) smaller packets and (iii) increased density of network affect the protocol performance adversely. Section 3 discusses this in more detail and provides qualitative reasons for this.

4. *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* (PARADYCE) and as a first step suggest key design requirements for such a class of protocols.

# 4 Experimental Setup

The experimental set up consists of a description of (i) the scenarios used, (ii) simulation setup, (iii) input and output variables,

## 4.1 The scenarios.

We studied the performance of the three protocols under three different basic scenarios. Each scenario consisted of a number of sub-scenarios. Each scenario was designed to test a distinct hypothesis. Unlike most of the earlier studies, our scenarios were designed to understand the performance of the MAC protocols at the "network level" rather than at "link level", i.e. most of our scenarios consisted of source sink pairs that were at least 2 links apart. We briefly describe the scenarios below; additional details for each scenario are given in the section describing the results for that scenario.

1. **Scenario 1: Effect of General Hidden terminal.** This scenario is motivated by the well known hidden terminal problem. It has been well documented that hidden terminal configuration causes CSMA to assign inequitable resources to connections. 802.11 overcomes this problem using the RTS/CTS/ACK mechanism. We wanted to see if the random delays introduced by the network can mitigate the hidden terminal to some extent. We call this the *generalized hidden terminal scenario.* Section 5.1 describes the scenario setup in more detail.

2. **Scenario 2: Effect of Network Connectivity.** In this scenario, our goal was to investigate the effect of network connectivity on MAC layer protocols. We consider successively denser network keeping the set of nodes constant. Another motivation for this scenario was to provide insights into optimal power settings for power aware MAC protocols. Intuitively, increasing the network density has two conflicting effects. On one hand, increasing the power range implies that paths between source destination pair tend to be shorter (the packets make faster progress towards their destination); this reduces the number of collisions that a packet might participate in. On the other hand, the network becomes dense (the node and edge connectivity); this implies that one is likely to encounter more spatial interference from adjacent radios. The second issue has been studied analytically by a number of authors for CSMA and ALOHA like protocols, most notable by Nelson, Kleinrock, Takagi and Tobagi [NK83, NK84, TK84, KT75, KT75a]. But no such analytical results are known for 802.11;

moreover, the analysis in [NK83, NK84, TK84, KT75, KT75a] is done only on randomly distributed set of points.

3. **Scenario 3: Effect of Separator size and sparsity.** In the final scenario, we aim to understand the effect of network sparsity and separator size on the performance of MAC protocols. Intuitively, it is obvious, that smaller separator imply higher probability of collisions and thus reduced performance. Again, as mentioned earlier, our broad goal is to look for network level effects as opposed to link level effects. The importance of separators has been well established in the study of circuit switched networks.

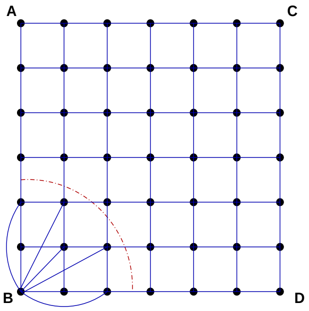

Figure 1: The figure schematically illustrates the connectivity of the graph. For clarity only the edges incident on the node $(0, 0)$ are shown. The dotted arc shows the transceiver's radio range.

## 4.2 Simulation Setup Characteristics

We now describe the details of the parameters used.

1. **Network Characteristics:** In each scenario we have kept the following parameter constant:
    - **Network Topology:** Although specific scenarios use specific network topologies, one particular topology is used frequently. We call it the *grid-squared* topology. It consists of $7 \times 7$ node grids with the radio radius of 2.5 grid units (1 grid unit = 100m). The name comes from the fact that it can be viewed as constructing $G^2(V_1, E_1)$, where $G(V, E)$ denotes the grid. In the graph $G^2$, there is an edge between $u$ and $v$ iff $u$ and $v$ were no more than a distance 2 apart. A vertical connection, i.e. source being $(x, 0)$ and destination $(x, 6)$ required at least three hops for a packet to reach its destination, whereas for a diagonal connection at least four hops were required. By vertical connection we mean those running up-down (A–B, C–D), and by diagonal connection we mean those running from upper right corner to lower left corner (C–B), or from upper left corner to lower right corner (A–D) as it is shown in Figure 1. Finding out the number of hops required to reach a destination from a source is an easy task and is omitted. Most of our topologies are derived from this basic structure.
    - **Number of connections:** We used two connections, except for the experiment in which we studied influence of number of connection on quality of service.

- **Routing protocol** : AODV.

2. **Mobility Parameters.** There was no movement of nodes except for the last experiment in which we studied influence of mobility on quality of service.

3. **Traffic Characteristics.**

   - The initial packet size was 512 bytes, the initial number of packets was 1,000, and the initial injection interval was 0.1 second. We 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 0.05 second 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.
   - The bandwidth for each channel was set to 1Mbit. The propagation path-loss model was two-ray.
   - Other radio propagation model details are as follows: (i) Propagation path-loss model: two ray (ii) Channel bandwidth: 1 Mb (iii) Channel frequency: 2.4 GHz (iv) Topography: Line-of-sight (v) Radio type: Accnoise (vi) Network protocol: IP (vii) Connection type: TCP.

4. **Simulation Characteristics.**

   - 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.
   - Unless otherwise stated, we used two connections (source destination pairs) per run. For each protocol and each value of injection interval 30 independent runs using a new random seed were carried out for each sub-scenario. In a few cases the number of runs were reduced to 10; but in all such cases the basic trend is evident.
   - Hardware used in all cases was a Linux PC with 512MB of RAM memory, and Pentium III 500MHz microprocessor.

## 4.3 Input Variables and Measured Quantities

The MAC layer protocols studied are 802.11, CSMA and MACA. The independent (input) variables for a given scenario were: the network topology and the injection interval for packets. The following three pieces of information were collected: (i) Latency: Average end to end delay for each packet as measured in seconds, (ii) Total number of packets received, (iii) Throughput: number of in bits/second received. Note that while calculating latency, we only consider packets that were successfully received.

Apart from latency and packets received that are plotted for each connection (recall for most part we deal with two connections), we also report the average behavior of the protocols. We briefly describe the method used to calculate these parameters. Average throughput and average latency is simply the average over 60 runs of each protocol over the two connections (30 for each connection). An example of a plot showing average fairness (discussed below), throughput and latency is seen in Figure 6.

**Measuring and Plotting Fairness.** Informally speaking, a fair assignment of resources means that all the agents participating in the game/process get equal access to the resources. This informal notion can be extended in several ways and indeed formalizing the concept is beyond the scope of this paper. Here as discussed earlier, we only consider long term fairness of protocols as opposed to short term fairness.

To measure long term fairness, let $r = p_1/p_2$ denote the ratio of packets received for a given run of the protocol for the two connections. Then $r$ denotes the fairness index of the protocol. Note that in case of perfectly equitable allocation the fairness index is 1. Average fairness is $\frac{1}{30}\sum_{i=1}^{i=30}\max\{r_i, \frac{1}{r_i}\}$, where $r_i$ is the above stated ratio for the $i^{th}$ run of the protocol[4]. One way to see the behavior of a protocol w.r.t. its fairness characteristics is to plot the fairness index for each run. For example consider any one of the 6 plots in Figure 5. The X-axis has the run number and the Y-axis displays the fairness index of the protocol for each run. These points are connected by a straight line so that differences in the height of the points is better reflected. Additionally, in order to depict smaller changes, all ratios above 10 are scaled down to 10. Thus the maximum $y$-value is always 10; the maximum $x$-value is 30 denoting the number of independent runs. Note that the runs are just independent invocations of the simulator with exactly the same parameters but with a different random seed. The dotted line is a horizontal line with y-value equal to 1. This denotes the expected behavior of a perfect protocol that assigns equitable resources. Any deviation of the curve above or below this line implies that the protocol was not fair. Also note that a fairness index such as .2 is in essence the same as 5 – in both cases one connection got 5 times more resources that the other connection.

## 5 Results and Analysis

We now summarize the experimental results. The results are summarized in form of graphs. The graphs show the data as a function of varying injection interval. The dependence of (i) latency and (ii) the number of packets received are plotted against the injection interval[5].

### 5.1 Scenario 1: Generalized Hidden Terminal Effect

We now discuss the experimental setup for the first experiment: effect of the generalized hidden terminal. The experimental design consists of three sub-scenarios and is depicted in Figure 2(a–c). Figure 2(a) depicts the base case; the classical hidden terminal setting. We have two connections: one from $A$ to $B$ and the other from $C$ to $B$. The setting is such that $B$ can hear both $A$ and $C$ but $A$ and $C$ cannot hear each other. Figure 2(b) depicts the first form of generalized hidden terminal setting. We have a grid-squared network and two connections shown by arrows from source to the destination. The arrows represent the rough flight path of packets: the path is not deterministic in general. As in the hidden terminal scenario, the connections have the same destination but different sources. Moreover, in contrast to the classical scenarios, the shortest path from source to destination for both connections is 4. This is the only difference. The rationale is the following: although the destinations are the same, the packets are likely to encounter random delays as they traverse the network and hence it is likely that the inequitable resource assignment problem for CSMA is mitigated to some degree. Figure 2(c) considers another variant. Here the destinations are not the same

---

[4] This is done so that all the summed quantities are at least 1.

[5] Note that we have used log-scale for y axis in all graphs.

but very closely located spatially. Again, one would expect the inequitable resource assignment problem is mitigated to a degree.

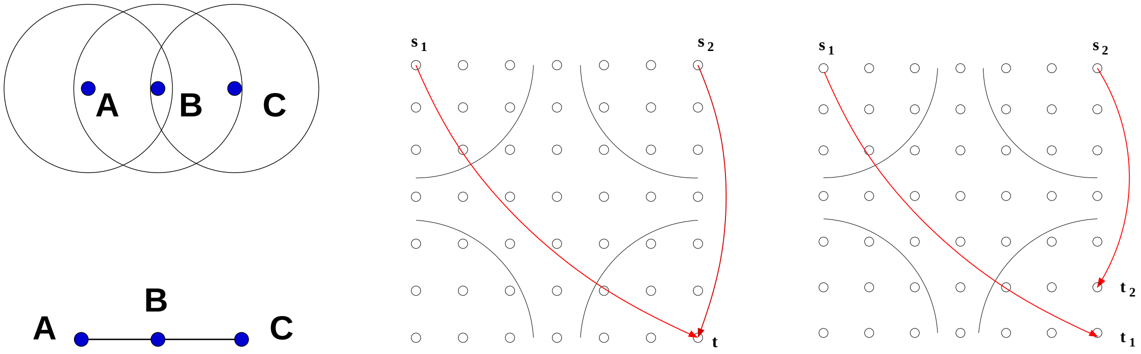

Figure 2: Distinct sources. (a) Three-node hidden terminal, $B$ can hear $A$ and $C$, but $A$ and $C$ cannot hear each other; (b) Identical sinks; (c) Closely positioned sinks. Minimum connectivity for (b) and (c) is 8, and maximum connectivity is 21. The four quarter circles denote the radio range of the corner radios. Each radio has the same range. The basic connectivity is the same as in a grid-squared graph.

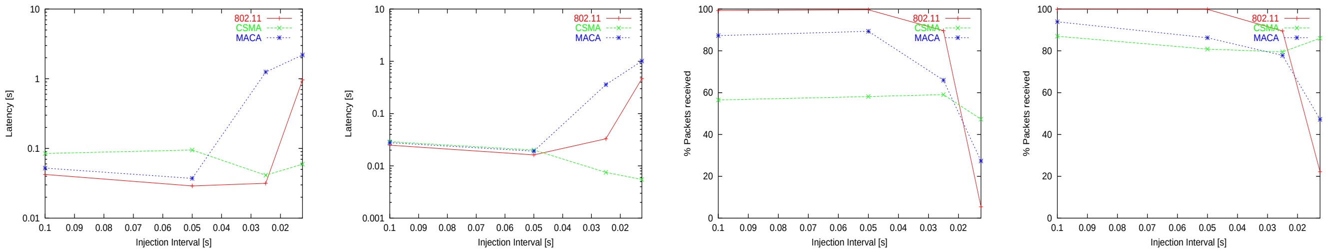

Figure 3: Distinct sources, identical sinks. The two leftmost figures represent data for latency (connection 1 and 2) and the other two show data for packets received (connection 1 and 2). The graphs shows dependency of these parameters on injection interval. These results correspond to scenario in Figure 2(b).

*Broad Conclusions for Experiment 1:* Results are shown in Figures 3, 4, 5, 6, and in Table 2. The plotted values are averaged over 30 runs with different random seed for each run of the simulator.

1. Looking at the numbers in Table 2 (results for Figure 2(a)), we see that CSMA essentially did not assign any resources when the connections started at the same time. In contrast, when the connections were started 1 millisecond apart, resources were assigned equitably. 802.11 did very well for both connections with and without any delays; in fact its performance was essentially indistinguishable. MACA's performance was somewhere in between the performance of CSMA and 802.11. The poor performance of CSMA is not obvious by looking at Figure 3 since the plots are obtained by averaging over 30 independent runs.

2. Results for the generalized hidden terminal scenario (Figure 3 and 4 for scenarios shown in Figures 2(b),(c)) show that latency is low for all protocols at low injection rates. However at high injection rates, both 802.11 and MACA exhibit much higher latency. On the other hand number of packets

| Protocol | 802.11 | 802.11 | CSMA | CSMA | MACA | MACA |
|---|---|---|---|---|---|---|
| Case | a | b | a | b | a | b |
| **Connection 1** | | | | | | |
| Injection interval [s] | | | | | | |
| 0.1 | 0.0097 | 0.0105 | 0.0037 | 0.0083 | 0.0258 | 0.0095 |
| 0.05 | 0.0067 | 0.0065 | 0.0028 | 0.0042 | 0.0218 | 0.0055 |
| 0.025 | 0.0051 | 0.0043 | 0.0024 | 0.0022 | 0.0200 | 0.0035 |
| 0.0125 | 0.0032 | 0.0032 | 0.0020 | 0.0011 | 0.0610 | 0.0610 |
| **Connection 2** | | | | | | |
| Injection interval [s] | | | | | | |
| 0.1 | 0.0097 | 0.0055 | 0.0019 | 0.0046 | 0.0262 | 0.0057 |
| 0.05 | 0.0067 | 0.0034 | 0.0016 | 0.0026 | 0.0217 | 0.0035 |
| 0.025 | 0.0051 | 0.0025 | 0.0015 | 0.0016 | 0.0200 | 0.0025 |
| 0.0125 | 0.0021 | 0.0021 | 0.0016 | 0.0010 | 0.0578 | 0.0578 |

Table 1: Three-node hidden terminal – latency. Case (a) The connections started at the same time. Case (b) The connections started with a difference of 1ms. Results correspond to Figure 2(a).

| Protocol | 802.11 | 802.11 | CSMA | CSMA | MACA | MACA |
|---|---|---|---|---|---|---|
| Case | a | b | a | b | a | b |
| **Connection 1** | | | | | | |
| Injection interval [s] | | | | | | |
| 0.1 | 999 | 999 | 0 | 998 | 494 | 998 |
| 0.05 | 1998 | 1998 | 1 | 1998 | 998 | 1998 |
| 0.025 | 3997 | 3997 | 1 | 3997 | 1973 | 3996 |
| 0.0125 | 7995 | 7995 | 2 | 7995 | 7188 | 7188 |
| **Connection 2** | | | | | | |
| Injection interval [s] | | | | | | |
| 0.1 | 999 | 999 | 0 | 999 | 506 | 998 |
| 0.05 | 1998 | 1998 | 1 | 1998 | 1001 | 1997 |
| 0.025 | 3997 | 3997 | 1 | 3997 | 1969 | 3996 |
| 0.0125 | 7995 | 7995 | 2 | 7995 | 7184 | 7184 |

Table 2: Three-node hidden terminal – packets received. Case (a) The connections started at the same time. Case (b) The connections started with a difference of 1ms. Results correspond to Figure 2(a).

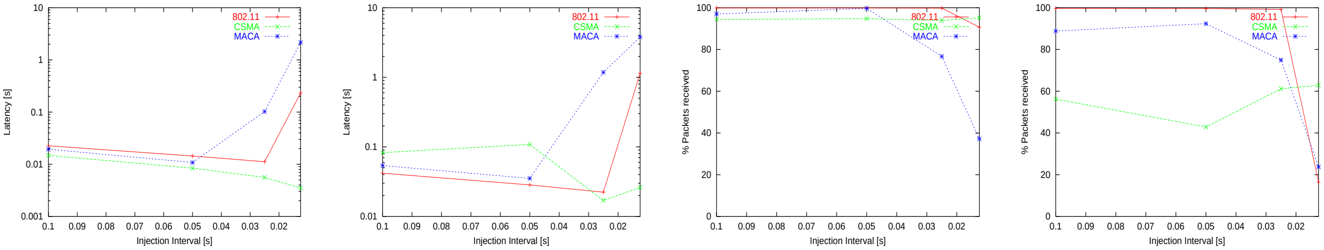

Figure 4: Distinct sources, closely positioned sinks. The two leftmost figures represent data for latency (connection 1 and 2) and the other two show data for packets received (connection 1 and 2). The graphs shows dependency of these parameters on injection interval. These results correspond to Figure 2(c).

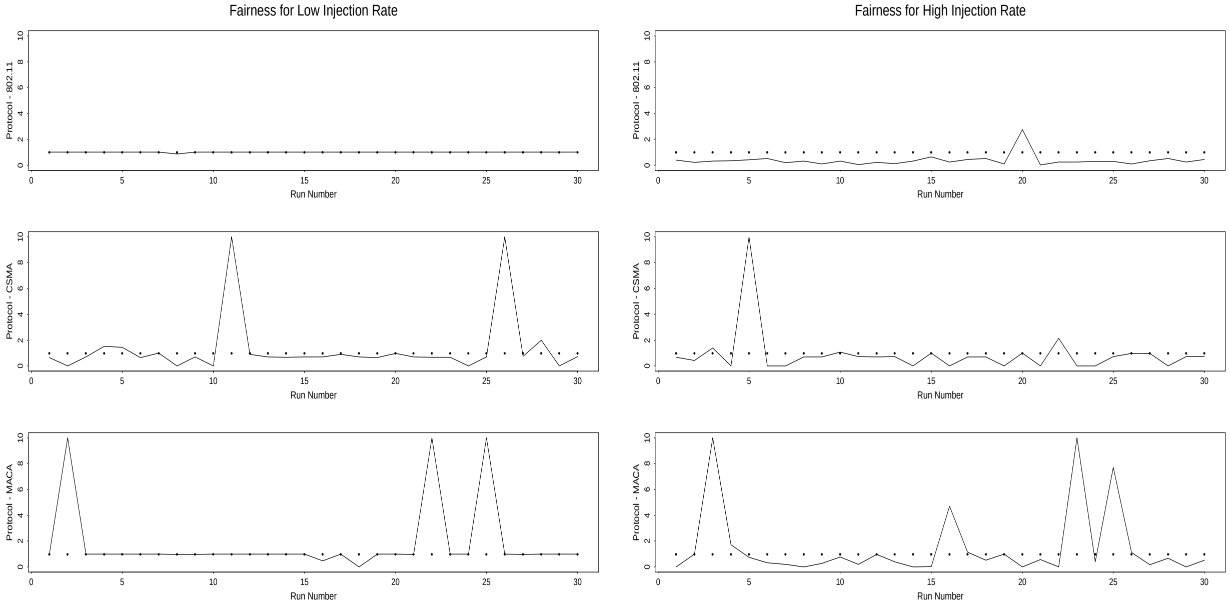

Figure 5: Fairness over a set of 30 runs for the three protocols. The $x$-axis shows 30 runs with different simulation seeds. The $y$-axis shows the fairness as a ratio of packets received for connection 1 to packets received for connection 2. The dotted line shows a ratio of 1; for fair protocols plot should coincide with this line. Case (a) - low injection rate are shown in the left column, Case (b) - high injection rate are shown in the right column. See Section 4.3 for more detail on the fairness measure. These results correspond to Figure 2(b).

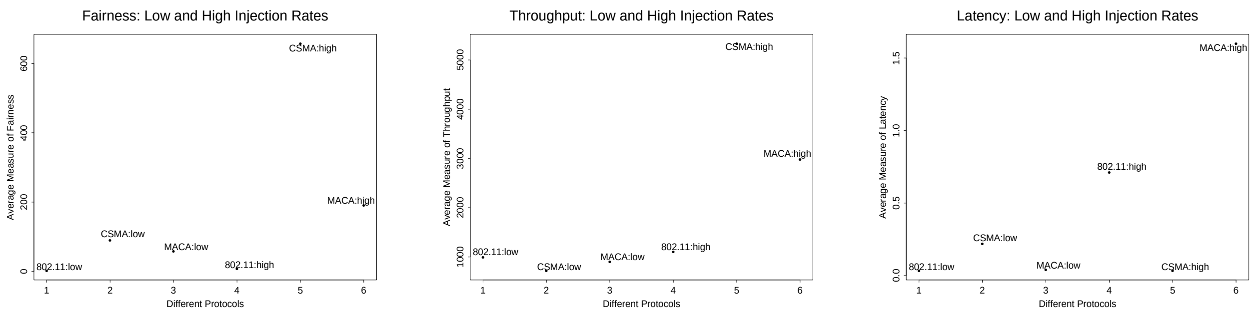

Fairness: Low and High Injection Rates

Throughput: Low and High Injection Rates

Latency: Low and High Injection Rates

Figure 6: Average (Un)Fairness, throughput and latency of the three MAC protocols under low and high injection rates. Note that at high injection rate we have correspondingly reduced the packet sizes. These results correspond to Figure 2(b).

received falls steeply for 802.11 as one increases the injection rates. On the other hand CSMA shows a steady increase.

3. The results for the two variant hidden terminal scenarios (Figure 3 and 4 for scenarios shown in Figures 2(b),(c)) exhibit similar performance characteristics. In particular, as expected the random delays introduced by the network improved the fairness characteristics of CSMA considerably over its performance for scenario Figure 2(a).

4. Figure 5 shows the behavior of the three protocols w.r.t. fairness ratio discussed in the earlier Section. It shows that almost every run of the CSMA and MACA protocol produce inequitable assignment of resources to the two connections. CSMA assigns inequitable resources more frequently than MACA but MACA has much high levels of inequitable resource assignment when they are so assigned. 802.11 behaves quite well across low as well as high injection rates.

5. Figure 6 shows that no single protocol dominates the other protocols across the three different performance metrics (fairness, throughput and latency) and over range of injection rates. This is an important conclusion and will be reinforced as we alter the scenarios.

*Qualitative Explanations for Experiment 1:* We provide plausible qualitative explanation for the above conclusions. First consider the relative behavior of 802.11 and CSMA. The RTS/CTS/ACK mechanism in conjunction with IFS (Interframe spaces) of 802.11 reduces the probability of collisions. On the other hand, it sometimes (unnecessarily) reserves media space thus disallowing other transmitters to use the space even if they could have probably used it without causing collisions. Additionally the control packets (RTS/CTS/ACK) imply additional overhead on the system which increases latency and decreases the good throughput (also known as goodput). These opposing aspects of the control packets used in 802.11 makes the analysis of 802.11 complicated. Nevertheless note the following: at high injection rates we use smaller packets and thus the relative overhead of the control packets in 802.11 exceeds the gain obtained by decreasing the number of collisions. Furthermore, the paths used by the two connections are by and large distinct (except near the destination). Thus the collisions we are avoiding are primarily those that occur between packets belonging to the same connection (collisions that occur while transmitting packets over three consecutive links of a routing path). At low injection rates, the number of control packets are significantly smaller and we have larger packet sizes: thus implying a higher bandwidth utilization. Moreover, although the collision probability is low, recovering from collisions at link level as done in 802.11 using the ACK part helps its overall performance. Thus 802.11 does quite well at low injection rates but deteriorate substantially at higher injection rates. It appears that the time for a packet to travel over one link together with the time it takes to move the packet from input buffer to the output buffer is less than the time it takes to generate the next packet at the source. Thus packets transmitted using CSMA do not typically experience collisions in this case. CSMA on the other hand does not assign equitable resources to the connections. This fact is clearer on inspecting Figure 5, rather than Figure 3 that reports the average over 30 runs. The reason for this is clear: once one connection gets access to the channel, it prevents the other connection from acquiring any resources. More surprisingly, MACA in spite of using RTS/CTS control packets, also exhibits inequitable resource assignment. Thus it appears that the random delays used in 802.11 play an important role in improving the fairness characteristics of 802.11. CSMA and MACA on the other hand rely on the transport layer to recover from collisions and thus pay a high price when collisions do occur. The qualitative difference between 802.11 and MACA at high injection rates is due to the ACK and IFS mechanism present in 802.11.

### 5.2 Scenario 2: Effects of Connectivity

We now discuss the set up for the second experiment. It aims to understand the effect of graph connectivity on the performance of the MAC protocols. As in the case of first experiment we have three sub-scenarios. The first scenario consists of a grid graph. The second and third scenarios are obtained by progressively increasing the radio range of all transceivers. More formally: (i) first we set the radio range of transceivers to one grid unit, (ii) in the second case the radio range was set to 2.5 grid units, and (iii) in the last case the radio range was set to 5 grid units. This gave us increasing minimum and maximum connectivity. The minimum node degrees were 3, 8, and 26 respectively, for corner nodes, and the maximum node degrees were 5, 21, and 49, respectively, for centrally positioned nodes. The edge connectivity of the graph is 2, 7 and 26 respectively[6] The topology of these experiments is shown in Figure 7.

*Broad Conclusions for Experiment 2:* The results are depicted in Figures 8, 9, 10, 11, 12, 13, and 14. The graphs are again averaged over 30 runs with different random seeds for each run.

1. First note that increasing connectivity has a mild effect on latency and packets received for both 802.11 and CSMA. In particular the number of packets received dropped somewhat at high connectivity and at high injection rate. Although the latency shows a slight drop, this is due to the fact that it is reported only for good packets. The performance of MACA dropped considerably at higher injection rates. As in experiment 1 this was mainly due to increased control packets.

2. Comparing the results for low connectivity Figure 8 with results for scenario in Figure 2(b)(shown in Figure 3) we see that all the protocols in general do better. This is because the two connections do not interfere with each other at all in Figure 7. Thus the latency for each protocol in this case is lower than

---

[6] The edge connectivity $\kappa$ is the minimum number of edge disjoint paths between any two pair of vertices.

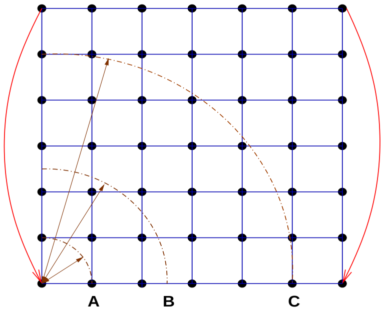

Figure 7: Two parallel connections, dense grid network with low, medium, and high connectivity. Circles show the radio range – (A) range is 1 grid unit: low connectivity, minimum and maximum degree is 2 and 5 respectively. (B) range is 2.5 grid unit : medium connectivity, minimum and maximum node degree is 7 and 20 respectively. (C) range equals 5 grid units: high connectivity, minimum and maximum degree is 25 and 48 respectively.

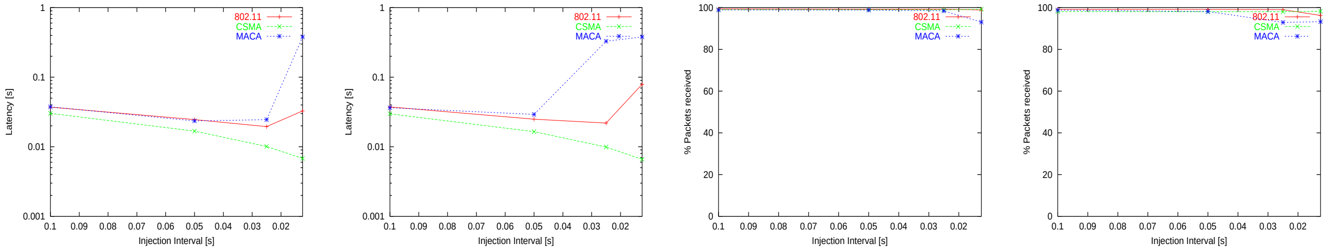

Figure 8: Grid network, low connectivity. The two leftmost figures represent data for latency (connection 1 and 2) and the other two show data for packets received (connection 1 and 2). The graphs shows dependency of these parameters on injection interval. The plots correspond to Figure 7(A).

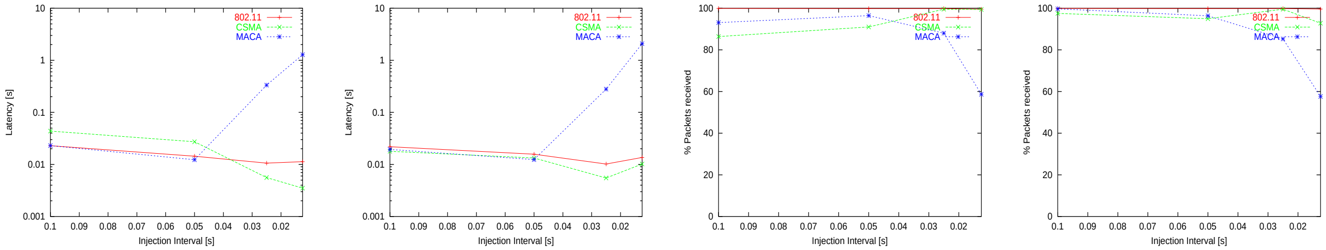

Figure 9: Grid network, medium connectivity. The two leftmost figures represent data for latency (connection 1 and 2) and the other two show data for packets received (connection 1 and 2). The graphs shows dependency of these parameters on injection interval. The plots correspond to scenario described in Figure 7(B).

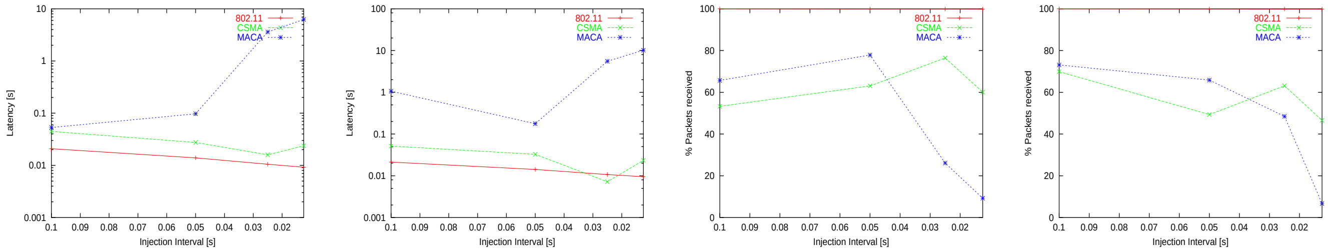

Figure 10: Grid network, high connectivity. The two leftmost figures represent data for latency (connection 1 and 2) and the other two show data for packets received (connection 1 and 2). The graphs shows dependency of these parameters on injection interval. The plots correspond to scenario in Figure 7(C).

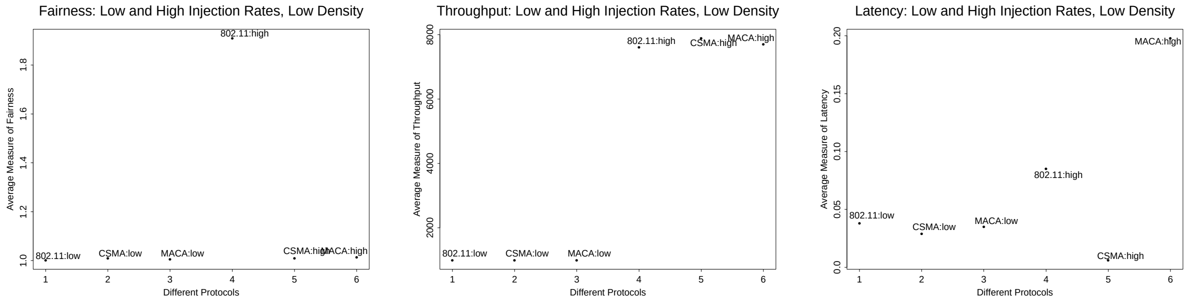

Figure 11: Average Fairness, throughput and latency of the three MAC protocols under low and high injection rate, and low density. Corresponds to the scenario described in Figure 7(A).

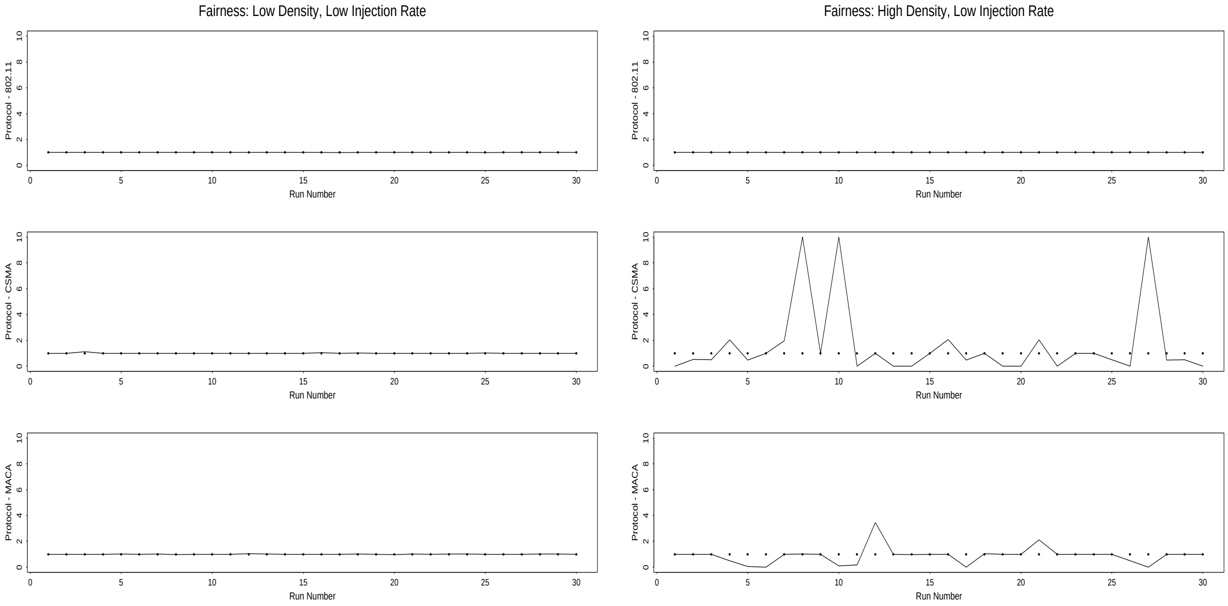

Figure 12: Fairness over a set of 30 runs for the three protocols. The $x$-axis shows 30 runs with different simulation seeds. The $y$-axis shows the fairness as a ratio of packets received for connection 1. and 2. Case (a) - low density and low injection rate (left), Case (b) - high density and low injection rate (right). See Section 4.3 for more detail on the fairness measure. These results correspond to Figure 7(A and C).

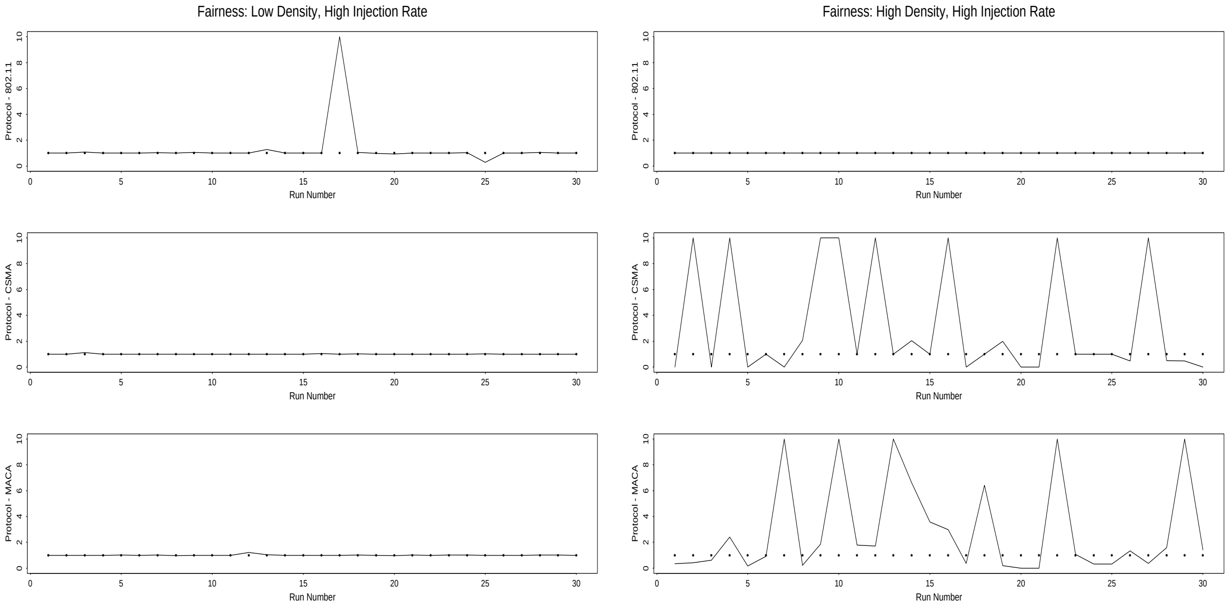

Figure 13: Fairness over a set of 30 runs for the three protocols. The $x$-axis shows 30 runs with different simulation seeds. The $y$-axis shows the fairness as a ratio of packets received for connection 1. and 2. Case (a) - low density and high injection rate (left), Case (b) - high density and high injection rate (right). See Section 4.3 for more detail on the fairness measure. These results correspond to Figure 7(A and C).

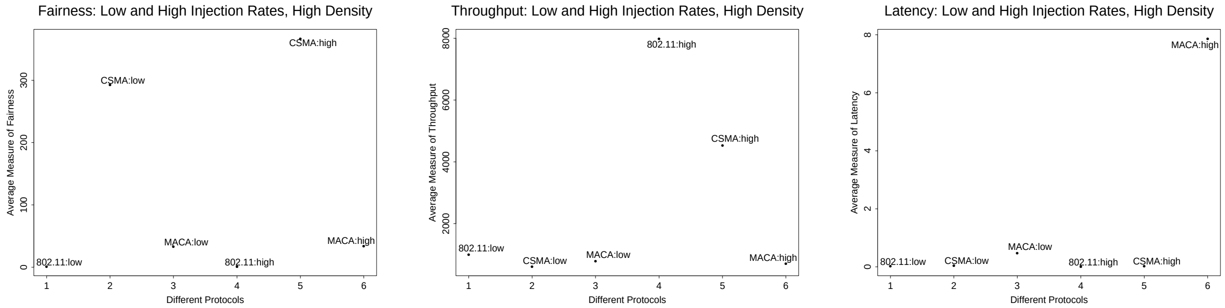

Figure 14: Average (Un)Fairness, throughput and latency of the three MAC protocols under low and high injection rate, and high density. Note that at high injection rate we have correspondingly reduced the packet sizes. These results correspond to Figure 7(C).

the latency experienced in Figure 2(b). The number of packets received are also substantially higher in this experiment.

3. The fairness characteristics also exhibit an intuitively expected behavior. At high connectivity both CSMA and MACA perform poorly; interestingly the performance was poor even at low injection rates. This can be seen by inspecting the right side of Figures 12 and 13. In contrast, comparing the results in the left column in the same figure (for low connectivity) with the results in Figure 5, we see that the protocols exhibit much better fairness behavior. The main reason is simple – there is hardly any integration between the two connections in this case (Figure 7 (A)).

In average for the three levels of connectivity 802.11 performs very well, but CSMA dominates in case of lower connectivity. MACA's performance decreases with increasing connectivity.

*Qualitative Explanations for Experiment 2:* We provide plausible qualitative explanation for the above conclusions. As we have observed 802.11 and CSMA had a fairly uniform performance at all connectivity levels. The main reason is that even at highest connectivity, there was one path for each connection that was not affected at all by the other connection. These paths are the sequence of nodes on the left and right edge of the graph. Thus if this path was indeed used then one would not expect any performance drop. On the other hand if a slightly different path was used by either connection, we have an interaction. The interaction can cause a performance drop if the injection rate was high enough. The reason is simple: as we have discussed in Experiment 1, even at high injection rates, the probability of interaction between packets on consecutive links was small if the connections did not interact. This is no longer true if the connections interact.

We note that although the routing paths need not have common nodes, they might be close enough so as to cause MAC layer interaction. In particular, consider the following setting illustrated in Figure 15. We have shown three paths from 1 to 2 and similarly $3$ paths from $3$ to $4$. The paths $1-6-2$ and $3-5-4$ are completely non-interfering. Paths $1-x-2$ and $3-x-4$ share the node $x$ and thus clearly interfere. The paths $1-y-2$ and $3-z-4$ are interesting. These paths do not share nodes but influence each other in that $y$ and $z$ cannot simultaneously transmit. This is because although they do not share nodes these paths influence each other. This holds since under the radio propagation model, nodes $y$ and $z$ can not simultaneously transmit.

### 5.3 Scenario 3: Effects of Separator Size and Sparsity

We discuss the final experiment for static configuration. The experiments aim to understand the effect of sparsity and minimum cut in the network on the performance of MAC layer protocols. Like the first two cases, we have three sub scenarios. In the first scenario, we have our usual *grid-squared graph* and two connections that are going diagonally across. In the second scenario the grid size is $3 \times 15$ nodes, and in the third experiment we use a sparse near-grid of 53 nodes. The topologies are depicted in Figures 16, 17, and 18, respectively. The basic qualitative difference between the sub-scenarios is obvious. In the first case, the minimum cut of the graph is roughly $O(\sqrt{n})$ where $n$ is the number of nodes. The minimum path length is $4$ for both connections. The situation is close to the generalized hidden terminal scenario considered in Experiment 1. The difference is that the paths for the two connections may not intersect at all. In the second case, the minimum cut is a constant and thus independent of the size of the graph. The length of the paths

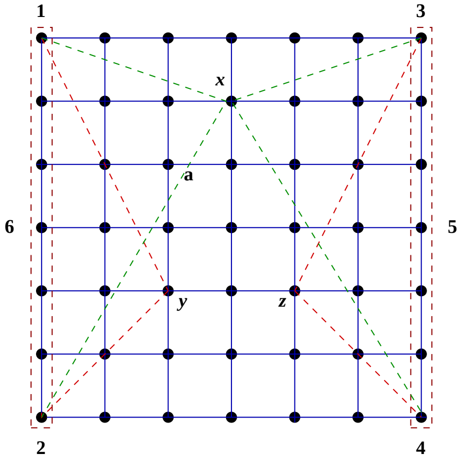

Figure 15: Figure illustrating that the routing paths need not intersect to be interfering.

on the other hand are $O(n)$, where $n$ again is the number of nodes in the graph. The second topology can be thought of as "stretching" the first topology in one direction. As a result, although the "vertical" cut in the $x$-direction is small the horizontal cut is $O(n)$. The third topology is somewhat different. Here both the horizontal and vertical cuts are small (a constant) as compared to graph size.[7] As a result, the situation portrayed can be viewed as a study of the tradeoff between connectivity (and thus multiple paths) on one hand and the increased interaction at the MAC layer. Another rationale for this study was to study the exposed and hidden terminal problems when the nodes being affected are not the end points but intermediate nodes.

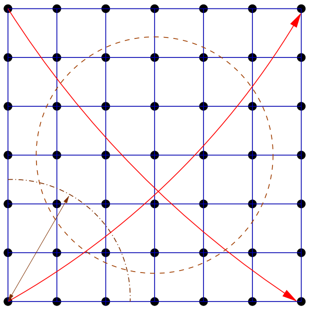

Figure 16: Grid-squared graph. Quarter-circle show the radio range of corner transceivers. The complete circle shows the range of transceiver that is at the canter of the grid. Minimum degree of the graph is 7 and the maximum degree is 20.

*Broad Conclusions for Experiment 3:* The results are depicted in Figures 19, 20, and 21, respectively.

CSMA again performed quite well w.r.t. latency and packets received, especially at high injection rates. But its fairness dropped significantly at high injection rates. Interestingly, the performance of the protocols on scenarios given in Figures 17 and 18 is *qualitatively similar*. In both cases, 802.11 and MACA had a drop in performance at high injection rate. In general the performance of the protocols for scenarios in Figures 17 and 18 is worse than their respective performance for scenario is Figure 16.

---

[7] Although all the grids are small as used in the experiments, it is easy to see how an infinite family of such grids can be created.

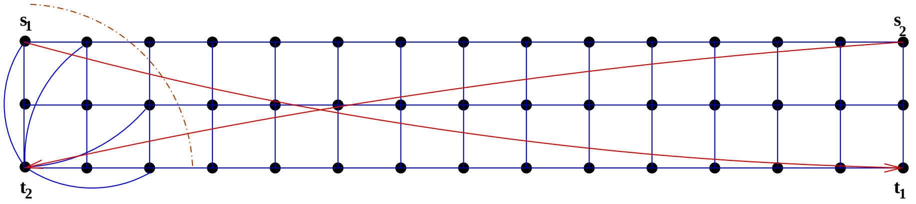

Figure 17: Long dense grid. The quarter-circle shows radio range from the lower left node. The node is effectively connected to seven other nodes. Minimum connectivity is 8, maximum connectivity is 16.

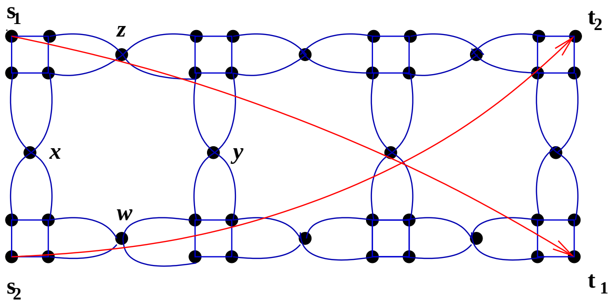

Figure 18: Long sparse graph. The figure shows the radio range for one of the nodes. The node has direct connection to four other nodes. Minimum connectivity is 4, maximum connectivity is 6.

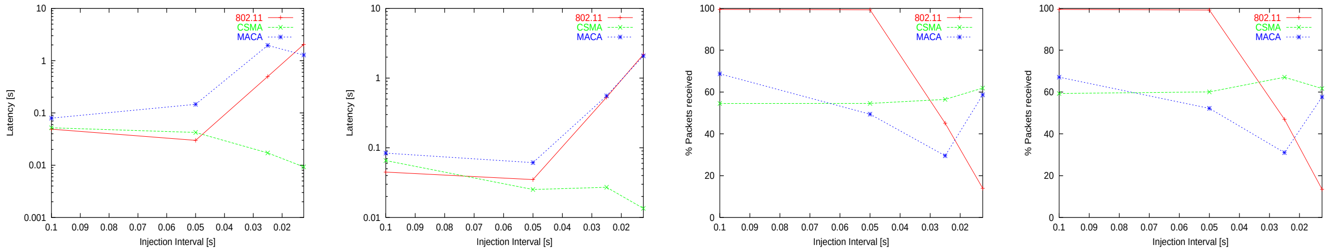

Figure 19: Grid-Squared graph. The two leftmost figures represent data for latency (connection 1 and 2) and the other two show data for packets received (connection 1 and 2). The graphs shows dependency of these parameters on injection interval. The plots correspond to scenario in Figure 16.

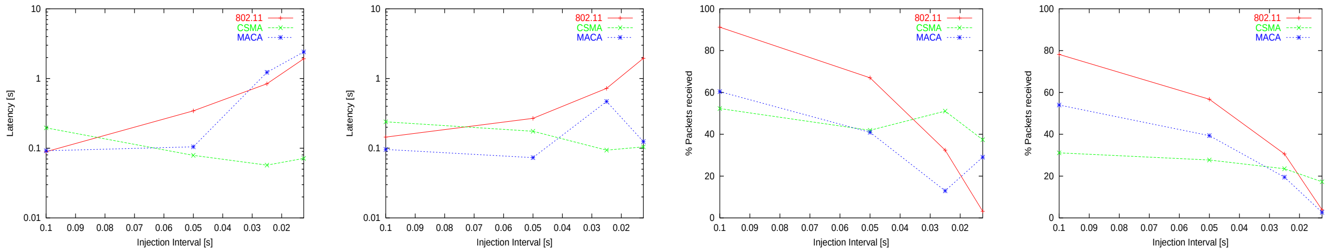

Figure 20: Long dense grid. The two leftmost figures represent data for latency (connection 1 and 2) and the other two show data for packets received (connection 1 and 2). The graphs shows dependency of these parameters on injection interval. The plots correspond to scenario described in Figure 17.

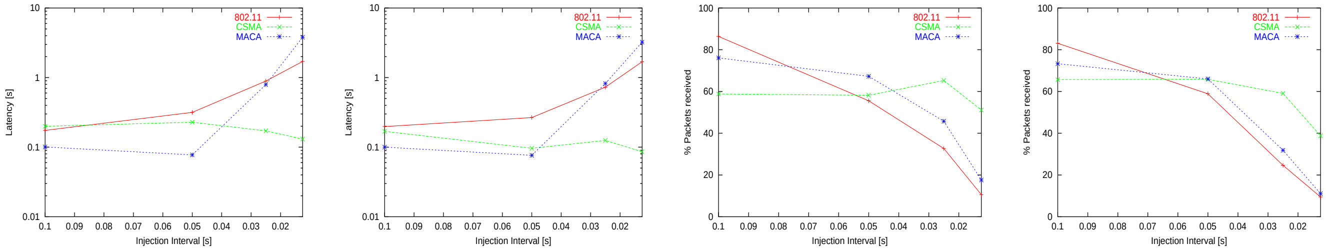

Figure 21: Long sparse near-grid. The two leftmost figures represent data for latency (connection 1 and 2) and the other two show data for packets received (connection 1 and 2). The graphs shows dependency of these parameters on injection interval. The plots correspond to scenario in Figure 18.

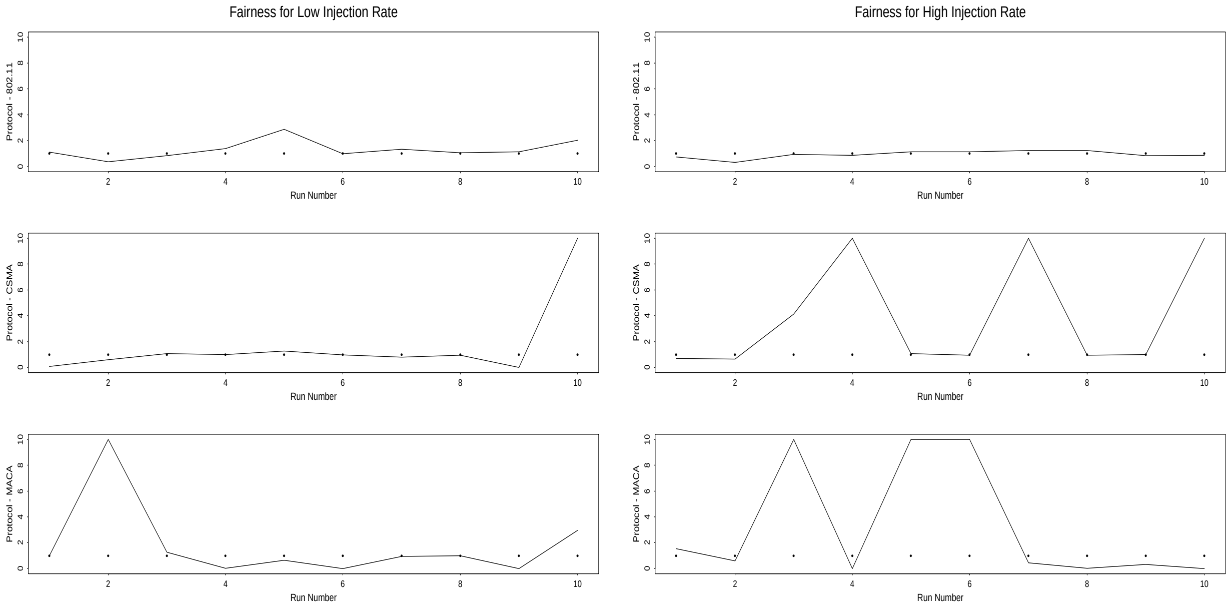

Figure 22: Results for scenario in Figure 16. Fairness over a set of 10 runs for the three protocols. The $x$-axis shows 10 runs with different simulation seeds. The $y$-axis shows the fairness as a ratio of packets received for connection 1. and 2. Case (a) - low injection rate (left), Case (b) - high injection rate (right). See Section 4.3 for more detail on the fairness measure.

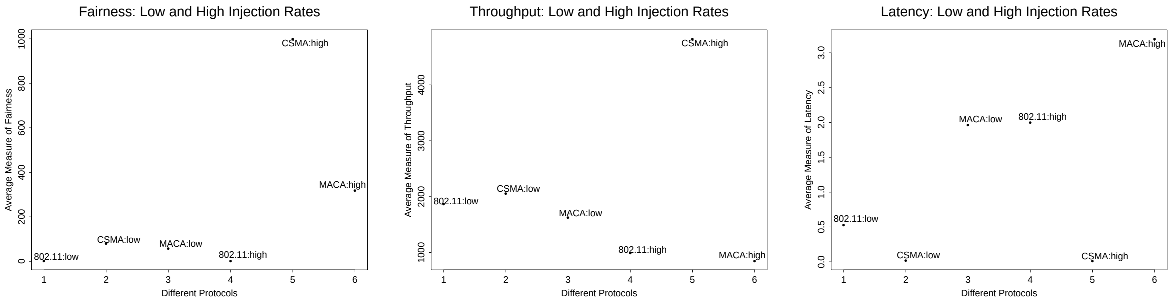

Figure 23: Average (Un)Fairness, throughput and latency of the three MAC protocols under low and high injection rate, and high density. Note that at high injection rate we have correspondingly reduced the packet sizes. These results correspond to Figure 16.

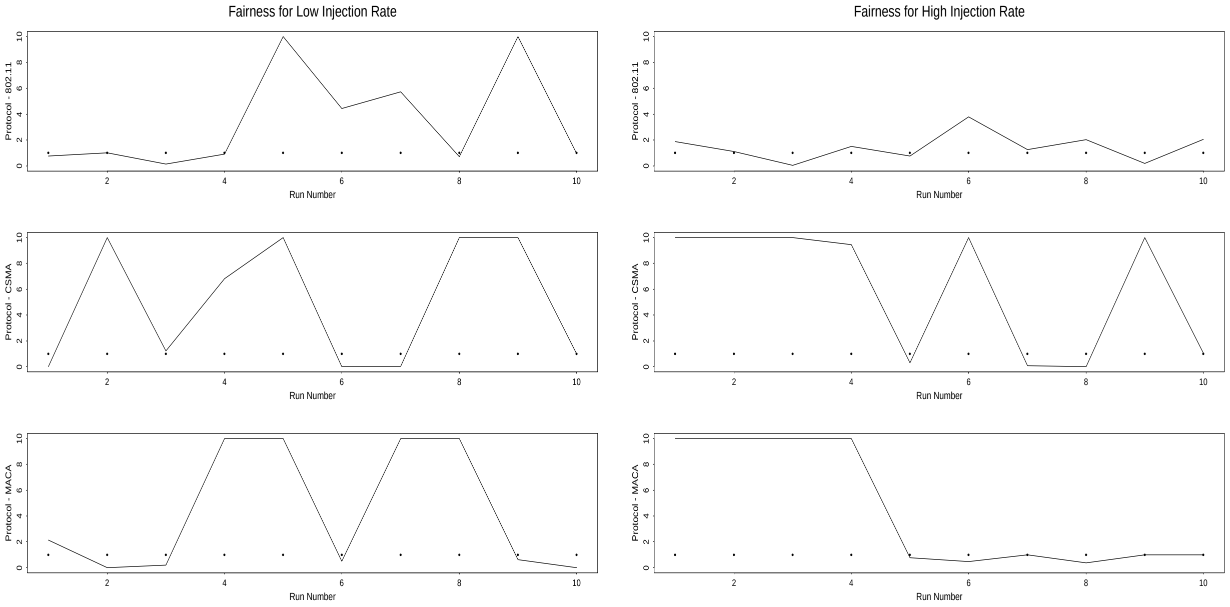

Figure 24: Fairness over a set of 10 runs for the three protocols. The $x$-axis shows 10 runs with different simulation seeds. The $y$-axis shows the fairness as a ratio of packets received for connection 1. and 2. Case (a) - low injection rate (left), Case (b) - high injection rate (right). See Section 4.3 for more detail on the fairness measure. These results correspond to Figure 17 (Long Dense Grid).

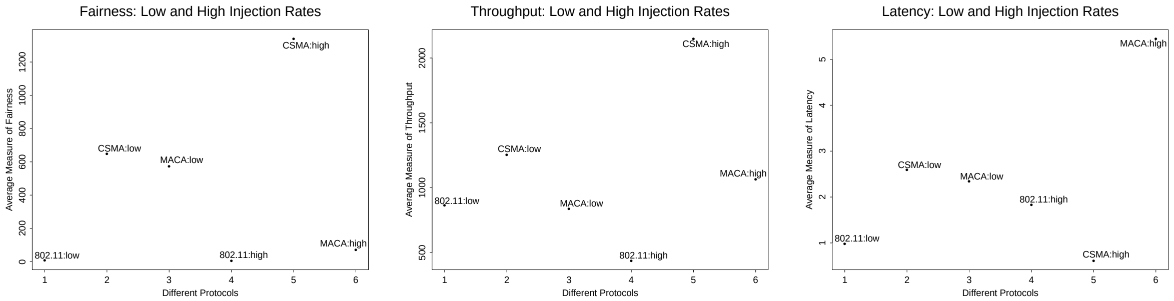

Figure 25: Results for long dense grid. Average (Un)Fairness, throughput and latency of the three MAC protocols under low and high injection rate, and high density. Note that at high injection rate we have correspondingly reduced the packet sizes. These results correspond to Figure 17.

It is instructive to compare the results for this scenario with the results for the previous scenario (scenario shown in Figure 7 and results shown Figures 8, 9, 10). First note that the medium connectivity case in Figure 7 is essentially the same as the grid-squared graph shown in Figure 16. The main difference is that in one case the connections have crossing paths while in the other case the connections do not have crossing paths. As expected the crossing paths scenario has a slightly worse performance.

*Qualitative Explanations for Experiment 3:* The main reasons for the observed behavior of the protocols is again related to (i) the control packets, (ii) the cuts in the graph and the ensuing probability of collision. As can be observed 802.11 sacrifices the packets received to get a better per run fair behavior. CSMA's performance appears quite good on the average but is quite poor when one notes the per run fairness characteristics as depicted in Figure 22 and 24.

The reason for poorer performance of the protocols for scenarios given in Figures 17 and 18 as compared to their performance for scenario is Figure 16 is quite simple: sparse connectivity and long paths. Both these factors increase the spatial contention for the media.

We finally discuss our results in light of the theoretical results by Nelson, Kleinrock, Takagi and Tobagi [NK83, NK84, TK84, KT75, KT75a]. The authors obtain analytical bounds on the "best possible degree" of a node when a greedy routing algorithm is used along with a CSMA/ALOHA class of MAC protocol. The results are obtained for a random set of points distributed according to a Poisson point process such that expected number of transceivers per unit area is $\lambda$. The authors also use the notion of *capture* for the radio propagation model. Specifically, the authors pose the question of calculating the optimal power transmission for maximizing the expected progress of a packet towards its destination. The general conclusion is that the expected number of transceivers that should be in the neighborhood of a given transceiver should be between $5$ and $8$. All the topologies considered in Experiments 2 and 3 use fairly uniform graphs (excepting the boundary nodes). Our experimental results show that indeed for cases with low constant degree networks, the performance of protocols is quite good while in cases when the degree of the network is high ($\Omega(n)$) the performance falls significantly. Obtaining exact numerical bounds on the vertex degrees is not very meaningful until we can simulate very large systems. Thus our experiments provide additional insights into the work of [NK83, NK84, TK84, KT75, KT75a]. As stated earlier, analytical results for 802.11 have not been carried out to our knowledge.

# 6 Extensions

We consider two important extensions of our basic study. The first is to study the performance of the MAC protocols when we increase the number of connections. The second study to evaluate their performance in ad-hoc networks.

## 6.1 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 26.

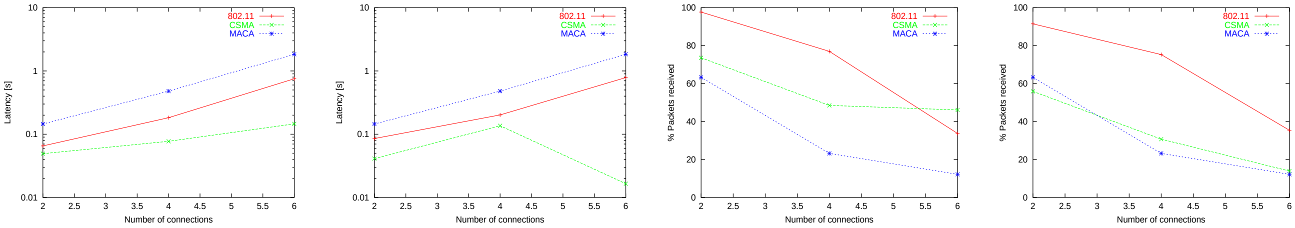

Figure 26: Effect of number of connections on Quality of Service. Two basic parameters are plotted: Latency and Packets received. The $x$-axis denotes the number of connections and the $y$-axis shows the latency/packets received.

## 6.2 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[8]. 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 27. 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.

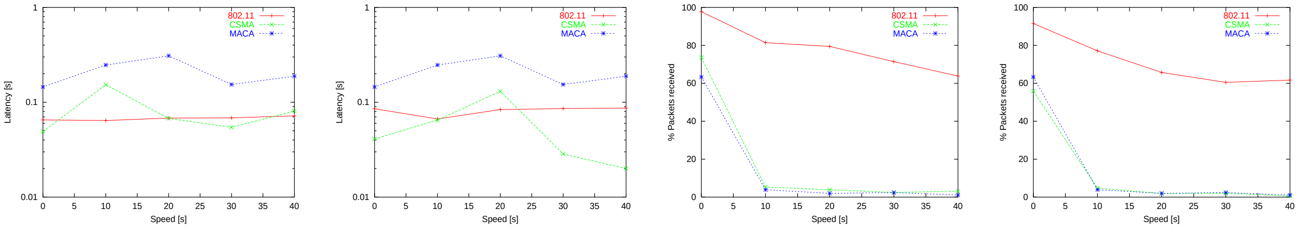

Figure 27: Effect of Mobility. Latency (first two figures) and Packets received, connection 1 and 2. All the graphs plot the change w.r.t. to speed; the first two graphs plot latency while the third and fourth plot packets received.

---

[8] 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.

## 7 Conclusions

We experimentally analyzed the performance of three MAC layer protocols: (i) CSMA/CA, (ii) MACA and (iii) 802.11. The performance of the protocols was measured in terms of (i) latency, (ii) throughput, (iii) number of data packets received and (iv) equitable resource assignment. The study was carried out by varying (i) the rate at which packets were injected in the network, (ii) the network topology, (iii) the spatial layout of the connections. The main conclusions are two folds:

1. No protocol dominated the other protocols over all the performance measures **even** for a given combination of all the input parameters (injection rate, topology and spatial location of connections). Although, the conclusion in itself is not surprising, the frequency of its occurrence and the variation displayed by the protocols was certainly surprising. The conclusion is important when service providers are likely to guarantee a given level of quality of service.

2. MACA was by and large dominated over the entire range of combinations by either CSMA/CA or 802.11. Interestingly, it appears that CSMA/CA might indeed be a good protocol for lightly loaded systems (in terms of number of connections). It is also seen from our experiments that the routing layer can affect the performance of the underlying MAC protocols.

As discussed earlier, this motivates a new class of protocols we refer to as **PARADYCE**. Although designing such protocols is non-trivial, the results do suggest key design requirements. They include: (i) ability to shut of the RTS/CTS/ACK mechanism when the traffic streams are non-interfering, (ii) use adaptive back off mechanisms that change with traffic conditions. Both these changes are likely to improve the performance of 802.11 significantly. MACAW designers have also suggested per-channel priority queues. Although this is possible, its overall effect on the network throughput remains to be understood. Additionally, the next generation radio units can likely control their power. This will give rise to interesting questions about simultaneously adjusting power levels for routing and MAC layer for the best utilization of resources.

# References

[802.11] Wireless LAN Medium Access Control (MAC) and Physical (PHY) Layer Specification. IEEE Standard 802.11, IEEE, June 1999.

[Al93] D. Allen. *Hidden terminal Problems in Wireless LANs*, IEEE 802.11 Working group paper.

[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 (MOBICOM)*, Berkeley, CA, November 1995.

[Ba+99] L. Bajaj et al. Improving simulation for network research. Technical Report 99-702, USC. March 1999.

[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

[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, (MOBICOM)* 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.

[Ce] Francisco Cercas. Simulation Tools. citeseer.nj.nec.com/317360.html

[JM96] D. Johnson and D. Maltz. Dynamic Source Routing in Ad Hoc Wireless Networks. *Mobile Computing,* Tomasz Imielinski and Hank Korth, Eds. Chapter 5, pages 153-181, Kluwer Academic Publishers, 1996.

[JT87] J. Jubin and J. D. Tornow. The DARPA Packet Radio Network Protocols. *Proc. 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.

[KT75a] F. Tobagi, and L. Kleinrock, Packet Switching in Radio Channels: Part II–the Hidden Terminal Problem in Carrier Sense Multiple-Access and the Busy-Tone Solution. *IEEE Transactions on Communications*, Vol. COM-23(12), pp. 1417-1433, 1975.

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

[KV98] Y. Ko and N. Vaidya. Location-Aided Routing(LAR) in Mobile Ad Hoc Networks. *4th Annual International Conference on Mobile Computing and Networking (MOBICOM'98),* Dallas, TX. October 1998.

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

[LSG00] S. Lee and W. Su and M. Gerla. Demand Multicast Routing Protocol in Multi-hop Wireless Mobile Networks. *ACM/Baltzer Mobile Networks and Applications, special issue on Multipoint Communication in Wireless Mobile Networks,* 2000.

[LNB98] S. Lu, T. Nandagopal, and V. Bharghavan. A Wireless Fair Service Algorithm for Packet Cellular Networks. *4th Annual International Conference on Mobile Computing and Networking (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

[MG96] S. Murthy and J. J. Garcia-Luna-Aceves. An Efficient Routing Protocol for Wireless Networks. *Mobile Networks and Applications,* 1(2), pages 183-197, 1996.

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

[NK83] R. Nelson and L. Kleinrock. Maximum Probability of Successful Transmission in a Random Planar Packet Radio Network. *Proc. IEEE INFOCOMM* San Diego, California, pp. 365-370, 1983.

[NK84] R. Nelson and L. Kleinrock. The Spatial Capacity of a Slotted ALOHA Multihop Packet Radio Network with Capture. *IEEE Transactions on Communications,* Vol. COM-32, No. 6, pp. 684-694, June 1984.

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

[Otcl] D. Wetherall and C. J. Lindblad. Extending Tcl for Dynamic Object-Oriented Programming. *Proceedings of the Tcl/Tk Workshop '95,* Toronto, July 1995.

[Pars] R. Bagrodia, R. Meyer, M. Takai, Y. Chen, X. Zeng, J.Martin, B. Park, H. Song. Parsec: A Parallel Simulation Environment for Complex Systems. *IEEE Computer,* 31(10), October 1998, pages 77–85.

[PB94] C. E. Perkins and P. Bhagwat. Highly dynamic Destination-Sequenced Distance-Vector routing (DSDV) for mobile computers. *Proc. SIGCOMM'94,* pp. 234–244, August 1994.

[PC97] V. Park and S. Corson. A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks. *Proc. IEEE INFOCOM '97,* Kobe, Japan 1997.

[PDNS] www.cc.gatech.edu/computing/compass/pdns/

[PGC00] G. Pei and M. Gerla and T. Chen. Fisheye State Routing: A Routing Scheme for Ad Hoc Wireless Networks. *Proc. IEEE International Conference on Communications (ICC),* pp. 70–74, New Orleans, LA, June 2000.

[PR99] C. E. Perkins and E. M. Royer. Ad-hoc On-Demand Distance Vector Routing. *Proc. 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pp. 90–100, New Orleans, LA, February 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, pp.342-353. Mobile Networks and Applications, 1-2, pp. 89-104, 1996.

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

[TK84] H. Takagi, H. and L. Kleinrock. Optimal Transmission Ranges for Randomly Distributed Packet Radio Terminals. *IEEE Transactions on Communications*, Vol. COM-32, No. 3, pp. 246-257, March 1984. Also appears in *Multiple Access Communications, Foundations for Emerging Technologies,* Norman Abramson (Ed.), IEEE Press, pp.342-353, 1992.

[Sc01] J. Schiller. *Mobile Communications,* Addison Wesley, Boston, MA, 2001.

[SK] D. Søbirk and J. Karlsson. *A Survey of Wireless ATM MAC Protocols.* citeseer.nj.nec.com/161151.html

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

[SW] N. Schacham and J. Westcott. Future directions in packet radio architectures and protocols. *Proc. 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).