Analyzing the Short-Term Fairness of IEEE 802.11 in Wireless Multi-hop Radio Networks

Christopher L. Barrett¹

Madhav V. Marathe¹

Los Alamos National Laboratory

Los Alamos NM 87545

Email: barrett,marathe@lanl.gov

D. Charles Engelhart^{1,2}

Anand Sivasubramaniam

Penn. State University

University Park PA 16802

Email: engelhar,anand@cse.psu.edu

Abstract

We extensively analyze the short (and long) term fairness performance of wireless ad hoc networks. The primary goal of our study was to investigate how multi-hop routing in an ad hoc network would affect the fairness of the MAC layer. Recently the fairness of the MAC layer of such systems has been studied. We demonstrate that higher-level protocols must also be considered when studying fairness. We analyze the behavior of a network using DSR and DSDV as an ad hoc routing protocol. The MAC layer used was IEEE 802.11.

The study is carried out for a wide range of independent parameters. These include: (i) network topology, (ii) the number of connections and their spatial location and (iii) the injection rate and burstiness of individual connections. We also studied the correlation between fairness and reception ratio.

1 Introduction and Summary of Results

Multi-hop ad hoc wireless networks (a.k.a packet radio networks) consist of a set of transceivers communicating with each other using a set of radio frequencies. There is no fixed infrastructure, network topology can change rapidly as a function of time and packets are transmitted by relaying them over a series of intermediate transceivers before reaching their final destination. Each transceiver has a fixed range (or geographic region) within which it can communicate with other transceivers. This allows spatial reuse of frequencies in transmitting data. Over the recent years,

wireless ad hoc radio networks have been a subject of active research. Early work on this topic includes the **PR-NET** (Packet Radio Network) [6], and **SURAN** (Survivable Adaptive Networks) [11] projects. Recently a special interest group for mobile ad hoc Networking within the Internet Engineering Task Force (IETF) [8] has been organized.

As wireless ad hoc radio networks are deployed in military, commercial and emergency management settings, they need to fulfill a number of functional and design requirements, including, (i) *High throughput*; (ii) *Low average latency*; (iii) *Heterogeneous traffic* (e.g. data, voice, and video), (iv) *preservation of packet order*, (v) *support for priority traffic*, and (vi) *priority of service or differentiated service* [7, 9, 12].

IEEE 802.11 is known to exhibit both short and long term unfairness [9]. Traditionally the MAC layer fairness has been studied in isolation. It was stated in a recent paper [10] by Santivez, McDonald, Stavrakakis, and Ramanathan that “*Greater understanding is required of cross-layer interactions and the impact of more general mobility models and traffic workloads*”. Here we show that the fairness characteristics of a MAC layer protocol are significantly affected by higher layer protocols as well as a variety of traffic workloads.

For this we consider two well known and widely used ad hoc routing protocols namely, DSDV and DSR. These network layer protocols use very different philosophies to handle the dynamic nature of ad hoc systems. The activity of a MAC protocol at a given node depends on if node is part of the path chosen by the routing protocol to deliver the packet to the destination. Thus, it is intuitively clear that we cannot ignore the effect of routing protocols on the *fairness* and *throughput* of the underlying MAC protocol.

¹The work is supported by the Department of Energy under Contract W-7405-ENG-36.

²Work done while the author was a graduate research student in the Basic and Applied Simulation Science Group at Los Alamos National laboratory.

1.1 Summary of Results

Motivated by the above discussion, we carry out an extensive simulation based study to understand the individual and combined effects of (i) the various protocols at the routing layer, (ii) the spatial location and the number of the active connections, (iii) the behavior of the individual connections and (iv) the mobility of the individual units on the fairness and throughput characteristics the network.

To achieve our goal we devised a number of experimental scenarios (7 in all). Each scenario can be viewed as a combination of specific network topology and traffic characteristics. The first six scenarios concentrate on fixed networks. The final scenario is based on mobile nodes. Fixed scenarios allow us to understand the network performance on invariants such as connectivity and node degree. These parameters change continuously in mobile networks and thus is harder to capture the effect of these network parameters on system performance. Before presenting results from each experiment, we briefly summarize some overall observations:

- **Remark 1:** We observed that DSDV usually performs better in terms of the packet reception ratio, than DSR for the workloads and configurations that we considered. Our results agree with those presented in [4], which points out that the set up period of on-demand protocols such as DSR causes a higher loss rate than proactive protocols such as DSDV.
- **Remark 2:** We also found that less load on the network usually improved fairness. Incidentally, if the network is able to handle the injected traffic without any degradation, then the system is perfectly fair. Consequently, all our experiments operate at loads that are beyond this point (to capture some unfairness).
- **Remark 3:** For the same reasons, in a less heavily loaded system, DSDV provides better fairness than DSR (which as pointed out above has a higher loss rate). When the system is heavily loaded the extra overheads of DSDV make DSR preferable in terms of fairness.
- **Remark 4:** Regardless of the load, we usually find DSDV to be less variable in terms of fairness from one run to the next for an experiment. Since within each run, DSDV allows different routes, the results do not deviate as much from one run to the next compared to DSR, which fixes the route during each run.
- **Remark 5:** Another interesting point to note is that while overall system load has an adverse effect on fairness, specific points of contention within the system can improve fairness. This is because the queue built

up at those points is likely to be more heterogeneous (i.e. successive packets are more likely to be from different connections), thereby improving short term fairness, compared to queues which are more isolated. The underlying 802.11 layer uses remainder of back off window to determine priorities rather than generating a new back off window at each collision (which is done in CSMA/CA).

- **Remark 6:** Fairness shows a significant amount of variation over independent runs. This poses a challenge for guaranteeing quality of service. Moreover, as noted by earlier researchers, each protocol can be seen as providing a trade-off between fairness and throughput (or latency). But this trade-off is not static, it varies with the network and traffic parameters.
- **Remark 7:** The network connectivity, spatial location of connections, injection rate and packet size all play a crucial role in determining the extent of the effect a routing protocol can have on the performance of 802.11. To the best of our knowledge such a detailed study has not been carried out in the literature.

The rest of the paper is organized as follows: In Section 2, we discuss the fairness indices used in our measurements. Section 3 discusses the distinguishing features of ad hoc networks that affect MAC layer fairness characteristics. Section 4 describes the layout of our experiments. Section 5 describes the results obtained and their possible qualitative explanation. Finally, Section 6 contains concluding remarks.

2 Fairness in Ad hoc Networks

Our study of fairness is motivated by a recent paper by Koksai, Kassab and Balakrishnan [7]. The authors studied the fairness characteristics of CSMA/CA protocol, propose analytical models and a new measure of fairness based on the Kullback Leibler distance metric. The authors seek further study of other more promising protocols such as MACAW. We choose to study IEEE 802.11 because of its similarities to MACAW and its widespread acceptance.

A study of fair medium access in IEEE 802.11 was taken up by Bensaou, Wang and Ko [1].

2.1 Meaning of Fairness

In its simplest form, the fairness of a protocol refers to its ability to allocate the channel bandwidth *equitably*. The notion of fair allocation is an active area of research in social sciences and is beyond the scope of the present article [2, 13, 14]. Here, for most part, we will be concerned with the *equal claims* situation, where an equal division of resources is the preferred solution.

2.2 Short and long term Fairness and their Impact

In communication networks the quality of a resource assignment protocol is measured in terms of how resources are assigned over a period of time. Researchers have considered two kinds of fairness in this context namely, *short term fairness* and *long term fairness*. Intuitively, short term fairness of a protocol refers to its ability to provide equitable access to resources to all the open connections over short time scales; long term fairness, in contrast, measures the amount of resources assigned over a longer time scale. In our context, short time scales usually correspond to the time required to transfer roughly a single moderate-size (about 25-50kB) message. As noted in [7], short-term fairness implies long-term fairness but not vice versa.

Short-term fairness has significant impact on the performance and Quality of Service of applications, especially those with stringent real time requirements. As discussed in [7], one way this happens is due to a degradation of the transport layer protocols such as TCP. TCP is based on the use of sliding windows, acknowledgments and time outs for flow and congestion control. If the ACKs received at the sending node are bursty in nature (due to inequitable assignment of channel at the lower layers), then we get idle periods that can be mistaken for losses by TCP. See [7] for a more thorough discussion on this subject.

2.3 Fairness Indices

For our calculations of fairness we use a sliding window method, as in [7], over the stream of packet arrivals. A sliding window method is useful because it gives a quantitative measure of fairness over a wide range of time frames. This way both short term and long term fairness can be illustrated together. A sliding window method computes, for every window size, the average fairness for every window of a given size. This can be quite computationally intensive.

In order to calculate the fairness for each window we must have some measure of QoS for each connection. We then judge the fairness of those measures of QoS for that window. Intuitively if each connection has very similar QoS measurements we would say that the system is fair, and if the measurements are radically different then we would say the system is unfair. We label each measure of QoS γ_i for each connection i .

To calculate the fairness of a given window size, w , we compute the ratios of packet arrivals from each connection over that window. Let γ_i be the fraction of packets from connection i that arrived during the window. Let N be the total number of connections competing for network resources.

Jain's Fairness Index: Jain's fairness index is the standard traditional measure of network fairness. Jain's index is

defined[5] as follows:

$$F_J = \frac{\left(\sum_{i=1}^N \gamma_i\right)^2}{N \sum_{i=1}^N \gamma_i^2}$$

Absolute fairness is achieved when $F_J = 1$ and absolute unfairness is achieved when $F_J = \frac{1}{N}$.

To calculate the fairness index for a given window size w , we begin with the left edge of the window matched with the left edge of the packet streams. We then successively move the window to the right and calculate the fairness index for the ensuing window position. The index value for a given window size is then taken as the average over all the valid window positions. Repeating this procedure for each window size yields both short and long term fairness index of the protocols.

Network QoS: Most research has used data such as packet reception ratio or routing overhead to measure network QoS[3]. We introduce a new measure, the fairness-reception ratio product, to judge network QoS.

Since Jain's fairness index is normalized to the interval $(0, 1]$ it is the most natural fairness index to apply to this product. The Kullback Leibler fairness index can also be normalized in order to be included in this measure.

3 Radio Network and Protocol Characteristics Affecting Fairness

There are several key features that intuitively affect the performance (especially with respect to fairness) of ad hoc networks, distinguishing them from wired networks. The first three characteristics have been discussed in[9] (for the most part we merely recall the discussion here) and the fourth and fifth have been discussed in [3] and [4].

Spatial Contention for Resources: Link layer protocols for wired networks have to consider contention on an individual node. In contrast, MAC protocols in radio networks have to take into account possible contention over the neighborhood of the sender as well as the receiver.

Channel Utilization and Fairness: In a shared channel wireless network, spatial reuse of the channel bandwidth may be obtained by simultaneously scheduling transmissions whose regions of contention are not in conflict. While partial reuse is very useful for improving channel utilization it introduces a fundamental conflict between optimizing aggregate allocated bandwidth and achieving fairness.

Inaccurate State and Decentralized Control: Lack of centralized control and inaccurate state information (due to information propagation delay) complicate the issue of achieving end-to-end connection-level fairness even further.

Mobility of Individual Units: Since a node in an ad hoc network can move continuously, the network connectivity

also evolves continuously. This implies that you will likely have a more bursty (unfair) resource allocation.

Effect of Routing Protocols: The above point also implies that routing protocols now would be expected to have a significant impact on the fairness characteristics of the network protocols. This is because every time a link is broken, due to mobility, the routing protocols have to initiate a route discovery process. This implies a large gap between certain successive packets at the receiving end and, in turn, implies an unfair resource allocation. Furthermore, the routing protocol also affects the TCP/IP protocol that in turn dynamically adjusts the window size and sending rate. This again implies variation in the fairness characteristics of the protocol.

Multiple paths & path caching One possible effect on fairness is that reactive algorithms will keep a route as long as it is valid where a proactive algorithm may use different routes each time a periodic update occurs. Because of this individual runs of a reactive algorithm may vary greatly since only one path is used if possible, where individual runs of a proactive algorithm should be more similar.

4 Experimental Setup

We studied the effect of several issues on the fairness of IEEE 802.11 by simulating various ad hoc networks in NS-2 (version 2.1b8a) [16]. Some of the common parameters and assumptions we make (unless explicitly varied/mentioned in later discussions) include:

1. Channel bandwidth was set to 1Mb/s.
2. There were 49 nodes.
3. We used constant bit rate UDP traffic generators for each source. Each source produced 512 byte packets at a fixed interval. No random jitter was used in packet departure times. The interval was set so that at least some packets would be dropped.
4. When the number of connections is kept constant, each source sends a total of 5000 packets to its destination. When the experiment uses varying number of connections, the total number of packets is fixed at 4200.
5. Each connection started after 100 seconds of simulated time, this wait time was used so routing protocols such as DSDV could fully establish their routing tables before any packets were sent.
6. The simulation was allowed to run until every packet that could be received was actually received, i.e. every node's link layer queue became empty.

7. Each simulation was on a flat 10 unit by 10 unit square and each node had a transmission range of 2.5 units. For node movement, the random waypoint algorithm was used [3] with a pause time of 0 seconds. The maximum speed was set to 1 and 0.125 grid units for very high and medium mobility respectively.
8. All of the experiments used only two connections. This was done in part since the fairness results in [7] showed two connections to be the most unfair.

4.1 Network Topology

For all simulations not involving node mobility we used a *grid-squared* topology. It consists of 7×7 nodes in a grid with the radio radius of 2.5 grid units. Figure 1 shows (a) The edges incident on node (0,0) and (b) The possible routes from source (2,6) to destination (2,0) are shown. The dotted arc delimits the maximum possible region of radio interference due to this connection.

All of the simulations were meant to cause some sort of contention between the connections.

Figure 1. The network topology.

4.2 Experimental Configurations

We ran several experiments to understand the role played by each issue on the resulting performance/fairness. We started out testing DSDV, DSR and AODV in each experiment. After a few simulations it was clear that the largest difference was between DSDV and DSR, with AODV performance usually somewhere between the two (since it is a mix of the other two). Consequently, in the interest of clarity, we include only the DSR and DSDV results in this paper.

Experiment 1: Generalized Hidden Terminal. The first experiment seeks to understand if the well known hidden terminal setting and its generalizations affect the fairness characteristics of a multi-hop network. In the hidden terminal scenario two transmitters send packets to a common receiver. We constructed three scenarios to study this effect: (a) the two transmitters are within each other's range; (b) the two transmitters are not within range of each other directly, but with the receiver is within range; and (c) the two transmitters are not within range of each other, nor with the destination, and packets require more hops (at least 3)

to get there. The last variant allows for more skew in packet arrivals.

Experiment 2: Crossing versus Parallel Connections. The second experiment tries to understand the inherent effect of interference between connection paths on the fairness characteristics of the MAC protocol. Three cases are considered: (a) we have two parallel connections that are reasonably far apart (minimizing interference); (b) we have two parallel connections that are close together (maximizing interference); and (c) the sources and destinations are placed in such a way that the two connections paths must cross each other at some point.

Unless explicitly stated all of the remaining simulations use parallel near connections.

Experiment 3: Number of Connections. The aim of this experiment is to find out the effect of increasing the number of connections on performance and fairness. The rationale is that increasing the number of connections will increase the resource contention due to path interference, perhaps causing a drop in fairness. We used parallel near connections as described earlier. Further connections were added parallel along side the other connections. We varied the number of connections from two to seven (results are reported for two to four connections), and the injection rates are decreased correspondingly to keep the total injected load on the network constant.

Experiment 4: Effect of Injection Rate. It has been shown that injection rate can affect MAC layer performance. Effect of injection rate (specifically burstiness) on fairness was undertaken by [7] for CSMA, and here we examine if similar conclusions can be drawn for 802.11. It is possible that because of the RTS/CTS/ACK protocol, one could expect a lesser impact of injection rate for 802.11. We conducted experiments with one connection pumping data at a faster rate than the other, and one connection pumping data more burstily (bursts of 5 packets) than the other.

Experiment 5: Effect of Path (Hop) Length. Path length can also affect fairness characteristics, and it would be interesting to find out how sensitive the results can be with varying number of hops.

We changed the hop length of one of the paths (from 1 to 3), keeping the other fixed at 3 hops.

Experiment 6: Effect of Connectivity. This experiment studies the effect of network density on the fairness.

We vary the connectivity by changing the radio range of transceivers from 1.25 to 2.5 to 5 grid units. The minimum node degrees were 2, 7, and 25 respectively, for corner nodes, and the maximum node degrees were 4, 20, and 48 respectively, for centrally positioned nodes.

Experiment 7: Effect of Mobility. While the previous experiments fix the position of the transmitters/receivers, this experiment incorporates mobility in the system using the random waypoint algorithm for capturing node movement.

For each of the runs a different random seed was used both in the simulation and also in the random waypoint node movement generator. The sources and destinations were both different in each experiment.

4.3 Metrics and Plots

For each experiment we present the results in a similar fashion, the invariants for each set of results include :

1. The results are presented over 25 different runs (with different random seeds). We tested initially with up to 50 experiments, but after 10 runs the standard deviation does not decrease. Although the standard deviation is not always less than the difference between the two protocols they are nonetheless distinguishable.
2. In general, we present the (a) *average Jain's fairness* (F_j) over these runs, (b) *the standard deviation of this index* (σ_{F_j}), (c) *the reception ratio* (RR), and (d) *the reception ratio * fairness product* ($RR * F_j$).
3. In all the performance graphs that follow, the window size is on the x -axis. For each experiment, we give two graphs:
 - The first graph shows F_j on the left y -axis, and σ_{F_j} on the right y -axis.
 - The second graph plots both RR (remember that this is independent of window size and thus a horizontal line) and $RR * F_j$.
4. We say that a network is heavily loaded, or has heavy traffic, when the total traffic is significantly greater than the available bandwidth. When there are network bottlenecks we refer to the situation as contention or interaction, rather than call it a heavily loaded network.

5 Summary of Results

In this section we discuss scenario specific results and provide plausible explanation for the observed behavior.

To conserve space many of the experimental results have been omitted, an unabridged version of this paper is available upon request.

5.1 Experiment 2: Crossing versus Parallel Connections

We find that the parallel far scenario is less fair than the parallel channels that are close to each other. The results are given in Figures 2 and 3 respectively. In the latter, there are more points of contention, and as described in Remark 5 this can improve fairness.

Figure 2. Parallel near channels.

Figure 3. Parallel far channels.

Also note that the reception ratio for DSR is lower for the case of parallel far connections than parallel near connections. This can possibly be explained by DSR's route caching and snooping strategy.

Remarks 1, 3 and 4 again hold in all these results. In fact, they hold for most of the following experiments, and henceforth we will not mention explicit comparisons between DSR and DSDV unless the results differ from this behavior.

5.2 Experiment 3. Varying the number of connections

The results with different number of connections for DSDV are given in figure 4. With higher number of connections, it is difficult to find representative packets from all of them within small time windows to make the system fair in the short term. Further, in these runs, the overall induced load does not change (i.e. the experiments with 4 connections are run with 1/2 the injection rate of those run with 2 connections, and so on). Consequently, this leads to a lower number for Jain's index as illustrated in these graphs. We would like to point out that this is a little different from results presented in [7] that uses CSMA/CA and a different network topology. The reception ratio also decreases with more connections because of increasing contention within the network as reflected in the graphs.

Figure 4. Varying the number of connections, using DSDV.

5.3 Experiment 4: Effect of Injection Rate

Figure 5. One bursty connection.

We evaluated several issues with reference to injection rate and channel behavior. We conducted experiments with (a) one channel pumping data at twice the rate of the other (i.e. same packet size but inter-injection times are half of the other), and (b) one channel pumping data in a bursty manner even though the overall average rate is the same as the other channel (i.e. it sends bursts of 5 packets, and then has a gap before sending the next burst). We also conducted experiments with both channels pumping data at the higher rate, and in the bursty mode.

In the interest of space, we present results only the case where one channel is pumping at bursts of 5 packets with the other behaving normally. These results are given in figure 5.

In general, we note that fairness and reception ratio decrease with the higher injection rates (Remark 2) and with burstiness. Similar observations hold for the other scenarios that were constructed for this experiment.

The difference we find between this and the earlier results are in the observation that DSR does better than DSDV for fairness but worse in terms of reception ratio when one channel pumps at a higher rate or in a more bursty mode.

Remember that when one channel misbehaves, completely abiding by the workload is not a good idea for fairness. DSR, which drops packets more often (as can be seen in the reception ratio graphs), makes the system more fair, since the dropped packets are more likely to belong to the connection that is pumping at a faster rate. *This result points out that performance and fairness can be different criteria for optimization, and may not necessarily go hand-in-hand, and it may sometimes be better to drop packets to make things more fair.*

5.4 Experiment 6: Effect of Radio Range/Connectivity

Figure 6. Radio range = 1.25 units.

Higher node degree indicates shorter number of hops to get to the destination, thereby reducing the points of contention along the path. At the same time, higher node degree also indicates that there is higher chance of interference of radio ranges of two different transmitters. However, in the parallel near channels case, for which the results are shown in figure 6 with maximum node degree of 4 the paths of the two connections are already colliding quite frequently (even with the smaller node degree). Consequently, we find that lower node degree leads to less fairness for this experiment. As before, when the radio range covers a lot of ground (to make the hop length one), the differences between DSR and DSDV are insignificant, while DSR turns out to be less fair for smaller radio ranges.

5.5 Experiment 7: Effect of Mobility

In the static node configurations that have been considered until now, we generally notice that fairness improves considerably usually within a window size of around 75, i.e. the system becomes fair in the long term. On the other hand, when we examine the results for the simulation when the nodes are mobile in figure 7 we notice that we do not get good fairness even for window sizes of around 1000 packets (this also holds for the case of medium mobility). The

Figure 7. Mobile nodes, very high mobility.

higher packet losses and variations in node positions result in much worse fairness properties than static networks.

With higher mobility [3], we find that DSR performs better than DSDV in terms of reception ratio. This performance has an effect on fairness as well, making DSR provide better fairness than DSDV. With high mobility, not only does DSR turn out more fair, we also find that this is the only experiment where we find that DSR provides lesser variability across experimental runs. The higher reception ratio of DSR also provides a more stable execution, leading to a smaller deviation in fairness across experiments.

6 Summary and Concluding Remarks

Performance characteristics of wireless networks have come under much scrutiny with the increasing importance of mobile computing. While throughput and packet loss are important considerations in wireless network design, *fairness* of the network to different communicating channels is an equally critical metric for optimization [7]. This paper has conducted an in-depth investigation of the fairness (and packet loss) properties of ad hoc networks, an important class of wireless networks.

The previous related study [7] has focused on analyzing the fairness of media access protocols (CSMA/CA and ALOHA in particular), demonstrating that CSMA/CA though better in terms of throughput, does poorly in terms of short-term fairness. While these results are useful to media access protocol design, it is not clear how they manifest and affect higher level protocols, particularly those that have been developed for ad hoc networks. Our results show that in ad hoc networks, the location of mobile nodes and route determination, together with numerous other factors, play as important a role as media access control in determining the overall fairness. Consequently, it is important to take into account all such factors in relative fairness comparisons.

To illustrate our point, we have conducted a detailed analysis of the short (and long) term fairness characteristics of two ad hoc network routing protocols, DSDV and DSR,

together with the IEEE 802.11 MAC layer. Some of the issues studied include the effect of hidden terminals, location of source and destination of a connections, the contending paths of connections, the number of connections, the load on the network and the injection characteristics of this load, the radio ranges of the transmitters, and the mobility of the nodes. Some specific results that were observed include the following:

- Consistent with previous results, we observed that DSDV performs better in many situations in terms of packet reception ratio for static networks. DSR however, came out in front when mobility was considered.
- In general, less load on the network usually increases fairness. However, sometimes, specific points of contention in the network can actually improve fairness.
- In a less heavily loaded system DSDV usually provides better fairness than DSR. When the system is heavily loaded the extra overheads of DSDV make DSR more preferable in terms of fairness.
- We usually found DSDV to be less variable in terms of fairness from one run to the next for an experiment, except under high mobility.
- Mobility has a considerable effect on fairness in several ways. First, it not only makes short-term fairness much worse, but the system is actually not even long-term fair. Second, mobility also changes many of the results observed with static networks. DSR turns out much better in this case compared to DSDV.

Most importantly, all these results reiterate our point that *it is not enough to study just one particular layer of the network stack in isolation when considering fairness* (the same is true for throughput as well). Fairness not only depends on the MAC protocol but also on the choice of routing protocol, and the various network parameters and traffic conditions.

In previous work different layers of the network protocol stack have typically been studied independently of one another. While such isolation makes it easier to study one issue at a time, our results here and in [2] suggest that the entire protocol stack needs to be considered as a whole.

This study has also illustrated that a sufficient measure of QoS in an ad hoc network must contain a combination of different factors such as fairness and reception ratio, rather than one single performance measure. Possible additions to this might incorporate measures of arrival-order (i.e. how many packets did not arrive in the order they were sent) and latency.

While this study has given us some key insights in this regard, there are certain points that still warrant a closer look.

In addition, we also need to take this study right up to the application level, incorporating transport and application level layers to really understand the fairness and performance implications in real systems.

References

- [1] B. BENSOU, Y. WANG, and C. C. KO. Fair medium access in 802.11 based wireless ad hoc networks. *Workshop on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2000.
- [2] R. BHARGAVA, A. GOEL and A. MEYERSON. Using approximate majorization to characterize protocol fairness. *Proc. ACM SIGMETRICS*, 2001.
- [3] J. BROCH, D. MALTZ, D. JOHNSON, Y. HU and J. JETCHEVA. Performance comparison of multi-Hop wireless ad hoc network routing protocols. *Proc. MOBICOM '98 ACM*, Dallas, TX, October 1998.
- [4] S. DAS, R. CASTANEDA, J. YAN. Simulation based performance evaluation of mobile, ad hoc network routing protocols. *Proc. 7th Intl. Conference on Computer Communication and Networks*, 1998.
- [5] R. JAIN. *The Art of Computer Systems Performance Analysis*. John Wiley & Sons, 1991.
- [6] J. JUBIN and J. D. TORNOW. The DARPA packet radio network protocols. *Proceedings of the IEEE*, 75(1), pp. 21-32, Jan, 1987.
- [7] C. KOKSAL, H. KASSAB and H. BALAKRISHNAN. An analysis of short term fairness in wireless media access protocols, *Proc. ACM SIGMETRICS*, Santa Clara, CA, June 2000.
- [8] J. MACKER and S. CORSON. Mobile ad hoc networks (MANET). IETF working group charter. www.ietf.org/html.charters/manet-charter.html
- [9] T. NANDAGOPAL, T. KIM, X. GOA and V. BHARGHAVAN, Achieving MAC layer fairness in wireless packet networks. *Proc. MOBICOM 2000*, Boston, MA, 2000.
- [10] C. SANTIVÁÑEZ, B. McDONNOLD, I. STAVRAKAKIS, AND R. RAMANATHAN. On the scalability of ad hoc routing protocols, *to appear in Proc. INFOCOM'02*, New York, June, 2002.
- [11] N. SCHACHAM and J. WESTCOTT. Future directions in packet radio architectures and protocols. *Proceedings of the IEEE*, 75(1), pp. 83-99, January 1987.
- [12] J. SCHILLER. *Mobile Communications*, Addison Wesley, Boston, MA, 2001.
- [13] H. VARIAN. Equity, envy and efficiency. *J. Economic Theory*, pp. 63-91, 1974.
- [14] L. ZHOU. Strictly fair allocations in large exchange economics. *J. Economic Theory*, pp. 158-175, 1992.
- [15] Wireless LAN Medium Access Control (MAC) and Physical (PHY) Layer Specification. IEEE Standard 802.11, IEEE, June 1999.
- [16] ns-2 Network Simulator. www.isi.edu/nsnam/ns, 2000 (version 2.1b8a)