

Improving Fairness and Throughput in Multi-hop Wireless Networks

Hung-Yun Hsieh and Raghupathy Sivakumar

School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, GA 30332, USA
{hyhsieh, siva}@ece.gatech.edu
<http://www.ece.gatech.edu/research/GNAN>

Abstract. In this paper we study the impact of the medium access control (MAC) layer and the routing layer on the performance of a multi-hop wireless network. At the medium access control layer, we argue that the notion of per-node fairness employed by the IEEE 802.11 standard is not suitable for a multi-hop wireless network where flows traverse multiple hops. We propose a new MAC protocol that supports prioritized per-node fairness and significantly improves performance in terms of both throughput and fairness. At the routing layer, we show that load balanced routing improves performance regardless of the nature of the underlying MAC protocol. Moreover, we show that an ideal load balanced routing protocol should take into account both the hop counts and the capacities when computing the optimal path. We propose a new routing protocol that improves performance over the conventional shortest-widest path routing.

1 Introduction

Ad-hoc networks are multi-hop wireless networks that lack the services of an established backbone infrastructure. They are typically formed by a collection of mobile stations cooperatively establishing a multi-hop wireless network. In recent years, numerous approaches have been proposed for routing [6, 7, 11–13], and medium access control (MAC) [1, 4, 10] in ad-hoc networks. While a majority of the routing protocols are similar to shortest path routing in that they use hop count as the optimization metric, the MAC schemes are mainly based on the CSMA/CA protocol. In this paper, we revisit the throughput and fairness properties of shortest path routing and CSMA/CA based MAC protocols in ad-hoc networks. We show through simulations that the end-to-end throughput and fairness properties of these routing and medium access control schemes are poor. We present simple algorithms at the two layers that significantly improve the throughput and fairness.

We make two key contributions in this paper: (i) We demonstrate that existing MAC protocols for ad-hoc networks (e.g. IEEE 802.11 [2]), based on the per-node fairness paradigm of CSMA/CA, do not provide end-to-end throughput fairness. We argue for a departure from the notion of per-node fairness to that of per-flow fairness. *We then present a new MAC protocol that has a per-flow notion of fairness for channel access and achieves improved end-to-end throughput fairness.* (ii) We show that load balanced routing not only can improve the end-to-end throughput observed by flows, but also can

have a positive impact on the fairness observed by flows. We argue that a conventional load balanced scheme such as shortest-widest path algorithm will not provide optimal results in ad-hoc networks. Finally, *we present a new load balanced routing algorithm that is suitable for the target environment.*

The rest of the paper is organized as follow: Section 2 presents the protocols and algorithms that we use in the rest of the paper. Section 3 describes the simulation model including the topology and traffic generation. Section 4 presents the simulation results. Section 5 discusses some issues and concludes the paper.

2 Algorithms

2.1 Medium Access Control

We use the IEEE 802.11 MAC protocol as the reference protocol. In order to alleviate any unfairness that the implementation of IEEE 802.11 protocol might contribute [8], we have implemented an ideal, per-node-fairness based MAC protocol (ILP) similar to the one presented in [9]. The ILP algorithm attempts to provide ideal, per-node fairness, and given a certain fairness level tries to maximize the throughput. Finally, we use an ideal per-flow-fairness based MAC protocol (IFP) that incorporates priorities in the ILP algorithm, where the priority of a node is set proportional to the number of flows traversing the node. Figure 1 presents a pseudo-code for the IFP protocol. Section 4 will present the simulation results comparing the three protocols.

2.2 Routing

We use a simple shortest path routing algorithm as the reference protocol. Initially, we show that the shortest-widest path algorithm is not suited to the ad-hoc network environment. For the rest of the simulations, we adopt a new load balanced routing algorithm that takes into account both the capacity (width) and the hop count (length) along a path. We assign a weight w to each “link” in the network, where w is proportional to the amount of contention at that link due to existing flows in the network. The shortest-widest path algorithm would then translate into finding the path with the minimum maximum-weight (MMW), while the new algorithm would involve finding the path with the minimum aggregate-weight (MAW). Figure 2 presents the algorithm for the MAW protocol. Note that a variation of Dijkstra’s algorithm (minimum maximum-weight instead of minimum aggregate-weight) can be used to achieve MMW routing with the same algorithm as shown in Figure 2. We show that the MAW algorithm performs better than the MMW algorithm in terms of the mean and variance of the end-to-end throughput. Finally, we demonstrate that the load balanced algorithm improves the fairness irrespective of whether the underlying MAC protocol is fair or unfair.

3 Simulation Model

We use the *ns2* network simulator for our simulations [3]. While we have used topologies of varying sizes (50, 100, and 200 nodes respectively) for our simulations, we

Input:

Set F of source-destination pairs (s_i, d_i)

Vector $Degree$

where $Degree(s_i)$ is the degree of node s_i

Vector $NumberOfFlows$

where $NumberOfFlows(s_i)$ is the number of flows traversing node s_i

Vector $Priority$

where $Priority(s_i)$ is the priority associated with node s_i

Vector $Allocation$

where $Allocation(s_i)$ is the number of time slots allocated to node s_i

(Both $Priority(k)$ and $Allocation(k)$ are set to 0 for all k during network initialization. The values carry over across iterations of the algorithm presented below.)

Output:

Set T of source-destination pairs allowed to transmit in the current time slot

Updated vector $Priority$

Updated vector $Allocation$

Algorithm:

Initialize set T to an empty set

While F is not empty

Find (s_i, d_i)

such that s_i has the maximum value in the lexicographic ordering of $(Priority(s_j), -Degree(s_j))$ for all s_j in F

Remove (s_i, d_i) from F

Add (s_i, d_i) to T

For each pair (s_j, d_j) in F

If node s_j is adjacent to node d_i

Remove pair (s_j, d_j) from F

If node d_j is adjacent to node s_i

Remove pair (s_j, d_j) from F

For each pair (s_i, d_i) in T

Increment $Allocation(s_i)$ by 1

$Priority(s_i) \leftarrow -Allocation(s_i)/NumberOfFlows(s_i)$

Fig. 1. *Ideal Per-Flow-Fairness Based MAC Protocol (IFP)*

present only the results for the 100 node topology in this paper. The nodes are uniformly distributed in a 1500m x 1500m grid. The simulation scenarios presented in this paper do not have any mobility. We will revisit the issue of mobility later in Section 5. The data rate of the underlying channel is set to 2 Mbps, and the transmission range

Input:

Set F of source-destination pairs (s_i, d_i)

Output:

Set R of routes for all source-destinations pairs in F

Algorithm:

Initialize R to an empty set

Initialize $weight(s_j)$ to 1 for all s_j

For each pair (s_i, d_i) in F

 Use Dijkstra's shortest path algorithm to obtain route r_i

 For each node m on route r_i except for d_i

 Increment $weight(m)$ by 1

 Increment $weight(q)$ by 1 for all q that is adjacent to m

 Insert r_i in R

Fig. 2. *Load Balanced Routing*

is set to 250m. The traffic in the network consists of 25 bi-directional TCP flows between 25 pairs of randomly (uniformly distributed) chosen sources and destinations. The simulations are run for a period of 100 seconds. Each data point is an average over 10 simulations run with different seeds for the random distribution. *We use the mean and the deviation as the metrics to compare the throughput and fairness respectively.* Unless otherwise specified, the routing protocol used is shortest path routing (SPR).

4 Simulations

4.1 MAC and Fairness

In Figure 3, we present the normalized deviation of the end-to-end throughput for the three MAC protocols. We define normalized deviation for a scenario as the standard deviation normalized by the mean throughput achieved for that scenario. As seen, IEEE 802.11 exhibits a high degree of unfairness. Note that in addition to the reasons given shortly, IEEE 802.11 has been shown to exhibit unfairness even when providing per-node fairness, and this accounts for the difference in its performance when compared to the ILP algorithm. The difference in performance between ILP and IFP can be briefly explained as follows: In ILP, nodes are given “equal” access to the channel irrespective of the number of flows traversing them. This results in lowered throughput for flows that traverse nodes handling more number of flows. However, in IFP, nodes are given access to the channel in proportion to the number of flows for which they act as relays (routers). Hence, flows are not penalized for traversing “congested” nodes. This results in the improved fairness for IFP.




Fig. 3. MAC and Fairness

4.2 Load Balanced Routing

In Figure 4, we present a comparison between the mean throughput achieved by the MMW (minimum maximum-weight, or shortest-widest path), and the MAW (minimum aggregate-weight) algorithms respectively. As observed, the MAW algorithm offers significantly more throughput than the MMW algorithm irrespective of the MAC protocol used. The reason behind the improvement is the fact that the network is moderately to heavily loaded (16 kbps to 256 kbps), and in such scenarios the longer hop counts (8.86 hops) of the MMW algorithm results in the network being overloaded sooner than in the case of the MAW algorithm (5.02 hops). Briefly, the larger number of hop counts results in more usage of the underlying network capacity:

$$Usage \approx NumberOfFlows * AverageHopCount * AverageFlowRate$$

As long as the total usage is less than the network capacity [5], the impact of larger hop counts is not noticed. However, when the network is heavily loaded, it is more likely that the larger hop count will result in the network becoming overloaded sooner, resulting in poor performance.

While the MAW algorithm is better in terms of the mean throughput, it can be seen from Figure 5 that the algorithm performs better in terms of the fairness also. Recall that the normalized deviation, and not the absolute deviation, is used as the fairness index.

4.3 Routing and Fairness

In Figure 6, we present the impact of the routing algorithm on the end-to-end throughput fairness. We again use the normalized deviation as the metric for fairness. When the




Fig. 4. Load Balanced Routing: Mean

underlying MAC is unfair, it is obvious that having a load balanced routing algorithm will improve fairness. This is because of the fact that load balancing reduces the average degree of multiplexing of flows on a single link, and hence bounds the unfairness introduced by the MAC protocol. This improvement in fairness is evident in Figure 6. However, it is interesting to note that *load balancing improves fairness even when the underlying MAC is fair with respect to flows*. Briefly, the reasons for this improvement are twofold: (i) The transport protocol used is TCP, and TCP is unfair to flows with larger RTTs. Hence, when flows with different RTTs share a single link, the mechanics of TCP will result in the flow with the smaller RTT getting a greater portion of the link capacity. Load balanced routing reduces the overlapping of flow paths, and hence reduces such effects. (ii) Although the underlying MAC protocol is fair, the variance in the degree of path overlapping (due to the existence of flows that have no or minimal link sharing along their paths, along-with flows that share links with a large number of flows) will induce unfairness in the network. Load balanced routing reduces the variance in the degree of path overlapping, and hence improves fairness.

4.4 Routing and Throughput Distribution

In our simulations, we observe that shortest path routing occasionally exhibits higher average throughput than load balanced routing. While superficially this indicates better performance, a closer look at the average throughput distribution between the different flows reveal that shortest path routing, although exhibiting higher average throughput, punishes a large number of flows (very low throughput in relation to the mean) in favor of a few flows that enjoy throughputs significantly higher than the mean throughput. Figure 7 shows the distribution of the number of flows observing different end-to-end




Fig. 5. Load Balanced Routing: Normalized Deviation

throughputs. The distribution is a consolidation of the results of 10 simulations, and hence has a total of 500 flows. As seen in the figure, the peak of the distribution for load balanced routing is closer to the mean than that of shortest path routing. Moreover, load balanced routing has a consistently better distribution curve about the mean throughput value. Finally, it can be seen that the peak of the distribution for the shortest path algorithm at the right end of the graph (high throughput) is higher than that of load balanced routing, substantiating our earlier claims that SPR greatly favors a few flows.

5 Issues and Summary

5.1 Issues

(i) *Mobility*: Due to lack of space, we do not consider mobility in the results presented thus far. However, the following observation can be made about the probable impact of mobility: While the shortest path and the MAW algorithms will suffer throughput degradation (possibly by the same amount) due to mobility induced losses, MMW can be expected to suffer significantly more losses. This is because of the fact that MMW paths, by virtue of their longer hop counts are more likely to break because of link failures. (ii) *Distributed Algorithms*: The new algorithms presented in this paper are centralized in nature. The scope of the paper is limited to highlighting the drawbacks of existing protocols and suggesting better approaches, and hence we do not present distributed versions of the algorithms. However, we believe that developing distributed versions of the algorithms introduced will not be a difficult task, and we hope to develop the distributed algorithms as part of our future work.

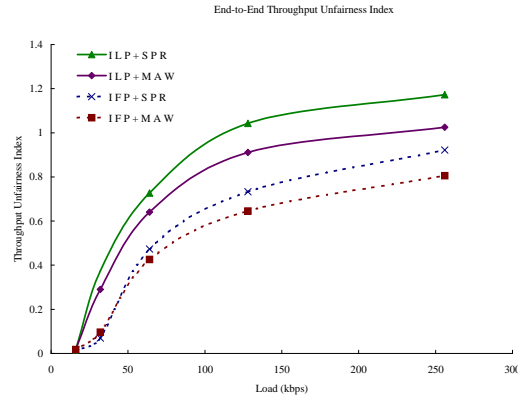


Fig. 6. Routing and Fairness

5.2 Summary

In this paper we have studied the performance of existing MAC and routing schemes in terms of their fairness and throughput characteristics. While we agree that the per-node fairness model adopted for packet cellular networks is apt for that environment, *we argue that such a model is not suitable for an ad-hoc network where nodes cooperatively act as routers or relays for flows belonging to other nodes in the network.* We propose a new MAC protocol that supports a per-flow fairness model, and in the process achieves significantly better end-to-end throughput fairness. At the routing layer, we show that *a load balanced routing scheme that takes into account both the capacity of paths and their hop counts is more suitable for ad-hoc networks* than a conventional shortest-widest approach. We demonstrate through simulations that the new routing algorithm does better than shortest path routing both in terms of throughput distribution and fairness.

6 Acknowledgments

We thank the Yamacraw organization (<http://www.yamacraw.org>) for their generous support and funding part of this work.

References

1. V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: A Medium Access Protocol for Wireless LANs," in *Proceedings of ACM SIGCOMM*, London, England, August 1994.

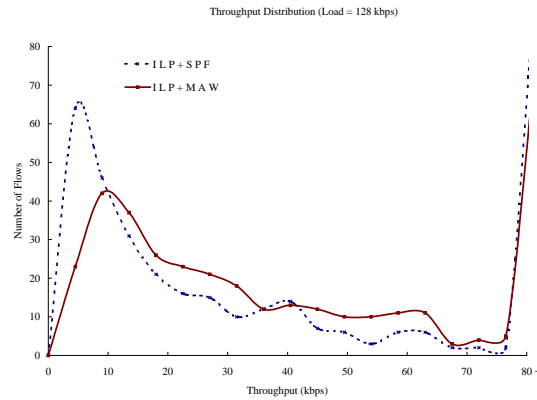


Fig. 7. Routing and Throughput Distribution

2. B. P. Crow, I. Widjaja, J. G. Kim, and P. T. Sakai, "IEEE 802.11 Wireless Local Area Networks," *IEEE Communications Magazine*, vol. 35, no. 9, pp. 116-126, September 1997.
3. K. Fall and K. Varadhan, "ns Notes and Documentation," available from <http://www.isi.edu/nsnam/ns/>, 1999.
4. C. L. Fullmer and J. J. Garcia-Luna-Aceves, "Floor Acquisition Multiple Access (FAMA) for Packet Radio Networks," in *Proceedings of ACM SIGCOMM*, Cambridge, MA USA, August 1995.
5. P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388-404, March 2000.
6. Z. J. Haas and M. R. Pearlman, "The Zone Routing Protocol (ZRP) for Ad Hoc Networks," *Internet Draft draft-zone-routing-protocol-01.txt*, August 1998.
7. D. B. Johnson and D. A. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks," *Mobile Computing*, Chapter 5, pp. 153-181, Kluwer Academic Publishers, 1996.
8. C. E. Koksal, H. I. Kassab, and H. Balakrishnan, "An Analysis of Short-Term Fairness in Wireless Media Access Protocols," in *Proceedings of ACM SIGMETRICS*, Santa Clara, CA USA, June 2000.
9. H. Luo, S. Lu, and V. Bharghavan, "A New Model for Packet Scheduling in Multihop Wireless Networks," in *Proceedings of IEEE MOBICOM*, Boston, MA USA, August 2000.
10. J. P. Monks, V. Bharghavan, and W. W. Hwu, "Transmission Power Control for Multiple Access Wireless Packet Networks," in *Proceedings of IEEE LCN*, Tampa, FL USA, November 2000.
11. S. Murthy and J. J. Garcia-Luna-Aceves, "A Routing Protocol for Packet Radio Networks," in *Proceedings of ACM SIGCOMM*, Cannes, France, September 1997.
12. C. E. Perkins, E. M. Royer, and S. Das, "Ad Hoc On Demand Distance Vector (AODV) Routing," *Internet Draft draft-ietf-manet-aodv-04.txt*, October 1999.
13. R. Sivakumar, P. Sinha, and V. Bharghavan, "CEDAR: a Core-Extraction Distributed Ad hoc Routing algorithm," *IEEE Journal on Selected Areas in Communications (Special Issue on Ad-hoc Routing)*, vol. 17, no. 8, pp. 1454-1465, August 1999.