Fairness and QoS in Multihop Wireless Networks

Jangeun Jun and Mihail L. Sichitiu[1]
Department of Electrical and Computer Engineering
North Carolina State University
Raleigh, NC 27695-7911
Email: {jjun, mlsichit}@ncsu.edu

Abstract— **In multihop wireless networks, fair allocation of bandwidth among different nodes is one of the critical problems that affects the serviceability of the entire system. Although there is significant research on the fairness issues in single-hop wireless networks, research on multihop fairness is rarely found in the literature. We study various queuing schemes for multihop wireless networks and examine the fairness and throughput performance of each scheme. Each scheme offers a different degree of fairness. While relatively simple queuing schemes require less hardware and processing budget, they inevitably lack good fairness and performance. In contrast, the scheme that provides fairness requires per-flow (i.e., network-layer flow) queuing. Furthermore, we show that in order to achieve the optimal bandwidth utilization, the medium access control (MAC) layer should be able to support different priorities. Without such a MAC-layer QoS scheme, in the worst case, the bandwidth utilization can be degraded by $O(N)$, where N is the number of end users. We theoretically investigate the pros and cons of different queuing schemes and verify the analytical results with detailed simulations.**

I. INTRODUCTION

Many wireless systems rely on multihop forwarding to reach destinations outside the direct range of the source. In the past decade, numerous ad hoc networking systems have been proposed and implemented for various applications [1]–[6]. Wireless sensor networks [7] emerged as a separate field in ad hoc networking (the traffic in these networks is primarily sensor data). Mesh networks [3] represent a promising alternative for broadband Internet access. Many other types of networks rely on multi-hop forwarding either to extend the range of an infrastructure-based wireless network (as in the case of bridged access points) or to avoid the need for infrastructure altogether (as in the case of pure ad-hoc networks).

Fairness is one of the most important properties of a computer network: when network resources are unable to satisfy demand, they should be divided fairly between the clients of the network. Most often, *min-max fairness* [8] is the desired fairness scheme. Under this scheme, the clients are split into two groups: the first group consists of the clients that cannot be completely satisfied by network resources; they all receive the same share of bandwidth. The second group is made up of the clients that need less bandwidth than their fair share; they receive exactly the amount of bandwidth that they ask for. Either of the two groups can be empty. A great deal of research has been done to ensure medium access control (MAC) layer fairness [9]–[11] and the result is that current wireless standards (e.g., IEEE 802.11 [9]) provide quite good MAC layer fairness. Unfortunately, as we will soon show, this does not ensure network layer fairness. An interesting practical mechanism providing proportional fairness is presented in [12]. With *proportional fairness*, clients that offer more load have a higher throughput, which may or may not be desirable depending on the particular applications of that network.

Another solution to the above problem can be provided by the transmission control protocol (TCP) [13] running over such a network. TCP's behavior over ad-hoc networks was recently the subject of considerable research (see [6], [14]–[20] and the references within). However not all traffic in multi-hop networks is TCP, and the fairness problem at the networking layer should be solved at the networking layer. After all, it is the job of the networking layer to provide fair resource allocations at the higher layers, and not vice versa. In this paper, several schemes for providing fairness and differentiated services at the network layer are investigated.

II. PROBLEM FORMULATION

A user node in a multi-hop network has to transmit both relayed and its own traffic. Therefore, besides the contention with other nodes for the same destination node, there is an inevitable contention between its own and relayed traffic. This contention does not occur in fixed wireless local loops or wireless LANs in infrastructure mode where user nodes are always at one-hop distance from the base station or the access point.

Consider the simple case depicted in Fig. 1(a) where two nodes (1 and 2) are having the same offered load G to be sent to the gateway (GW). Ideally, as the offered load at each of the nodes (G) increases, both nodes will receive the same share of the MAC layer throughput, B (see Fig. 1(b)). In practice, without a modified MAC or network layer, as the offered load increases, the node closest to the gateway (node 1 in Fig. 1(a)) will gradually but completely starve the node further away from the gateway (as shown in Fig. 1(c)). The results in Fig. 1(c) are obtained under the assumptions that the MAC layer is fair and that the traffic to be forwarded by node 1 (from node 2 to the gateway) is queued together (either in the forwarding engine or at the MAC layer) with the traffic originating at node 1.

As can be seen in Fig. 1(c), as the load at both nodes is increased, node 2 is gradually, but eventually completely

---

[1] This work was supported by the Center for Advanced Computing and Communication.

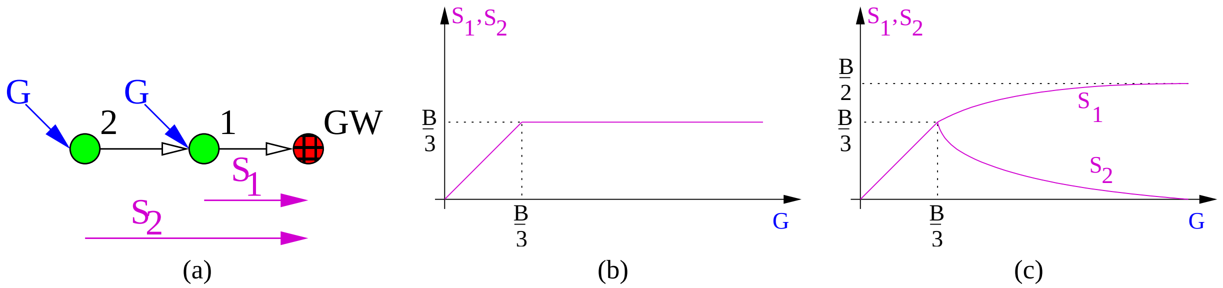

Fig. 1. (a) Fairness study of a two-node network forwarding packets to a gateway GW. The ideal (b) and real (c) throughputs of nodes 1 and 2 as a function of the offered load G.

starved by node 1. The overall system throughput (measured at the gateway) when the offered load is very high is $\frac{B}{2}$, where B is the throughput of the system when only node 1 forwards data to the gateway. So, the system operates at 50% efficiency *and* is unfair.

The unfair behavior observed in Fig. 1(b) is rooted in the fact that both the traffic originating at node 1, as well as the relayed traffic from node 2, are queued together at node 1. When the traffic load increases, the network cannot forward all data enqueued at node 1, and the queue starts to overflow. With a probability that increases with the offered load, the queue will be full when a new packet arrives from node 1, and it will be dropped immediately after it is received. The exact, expected throughput was determined theoretically, and it was verified using both OPNET and ns-2.

III. PROPOSED SOLUTIONS

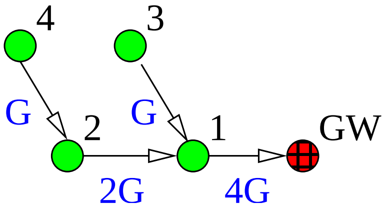

Fig. 2. A simple multihop wireless network with four user nodes and a gateway.

In this section, we propose various solutions to address the problem described in the previous section and explain the advantages and disadvantages of each scheme. For clarity, a simple multihop wireless network with four nodes as shown in Fig. 2 is used in the analysis. We assume that all the traffic flows are unidirectional toward the gateway. A more realistic case where traffic flows are bidirectional is discussed later. Each node in Fig. 2 can be modeled as a wireless router with network-layer queue(s) and MAC-layer queue(s) as shown in Fig. 3.

A. Isolate the originating traffic

Using the default queuing scheme in most routers, each node in Fig. 2 can be modeled as shown in Fig. 3(a) (i.e., only one network-layer queue accommodates both the originating and relayed traffic flows). Assuming that Fig. 3(a) shows the network and MAC queues of node 1 in Fig. 2, f_1 is the originating traffic flow and f_2, f_3 and f_4 are the relayed ones.

The basic MAC layer shown in Fig. 3(a)-(d) models a common IEEE 802.11-like MAC layer. In such a model, no QoS is supported, but a good MAC-layer fairness is provided.

With the basic queuing scheme in Fig. 3(a), it is clear that the traffic flow f_1 will receive more bandwidth and eventually starve others due to the problem described in the previous section.

Since this problem stems from the fact that both the relayed and the originating traffic share a common queue, the first solution that comes to ones mind is to use different queues for the relayed and for the originating traffic and to serve them in a round-robin fashion. This scheme will isolate the originating traffic which dominates the relayed traffic and protect the relayed traffic from being starved by the originating one.

In the network shown in Fig. 2, the maximum bandwidth share that each node can receive at the network layer (i.e., per-flow[1] or per-user bandwidth) is G. Since some nodes have to carry relayed traffic as well as their own traffic, the bandwidth that nodes 1, 2, 3 and 4 need to use at the MAC layer are $4G$, $2G$, G and G, respectively. The network-layer throughput ratio of 1:1:1:1 and corresponding MAC-layer throughput ratio of 4:2:1:1 can be maintained when the offered load at each node is less than the capacity of the network; however, as the offered load at each node is increased to exceed the network capacity, and eventually saturates the network (i.e., all the nodes equally share the MAC-layer bandwidth due to the MAC layer fairness), the resulting per-flow throughputs for nodes 1-4 converge to $\frac{B}{2N}$, $\frac{B}{8N}$, $\frac{B}{4N}$ and $\frac{B}{8N}$ respectively, where B is the theoretical maximum throughput (TMT) of the network [21], and $N = 4$ is the number of nodes.

Therefore, when the network saturates, the per-flow throughput ratio of nodes 1, 2, 3, and 4 is 4:1:2:1 and the per-node MAC throughput ratio is 1:1:1:1. Isolating the originating traffic by putting two fair queues at the network layer still shows significant unfairness of the per-flow throughputs, although the scheme is simple to implement and prevents the severe starvation of relayed traffic. Fairness is guaranteed with

---

[1] The term *flow* in this paper represents the distinguishable traffic flow entity at the network layer between end users and the gateway and not the flow in the transport layer or above. Likewise, the term *per-flow* is used as opposed to the term *per-node* which indicates the traffic at the MAC layer.

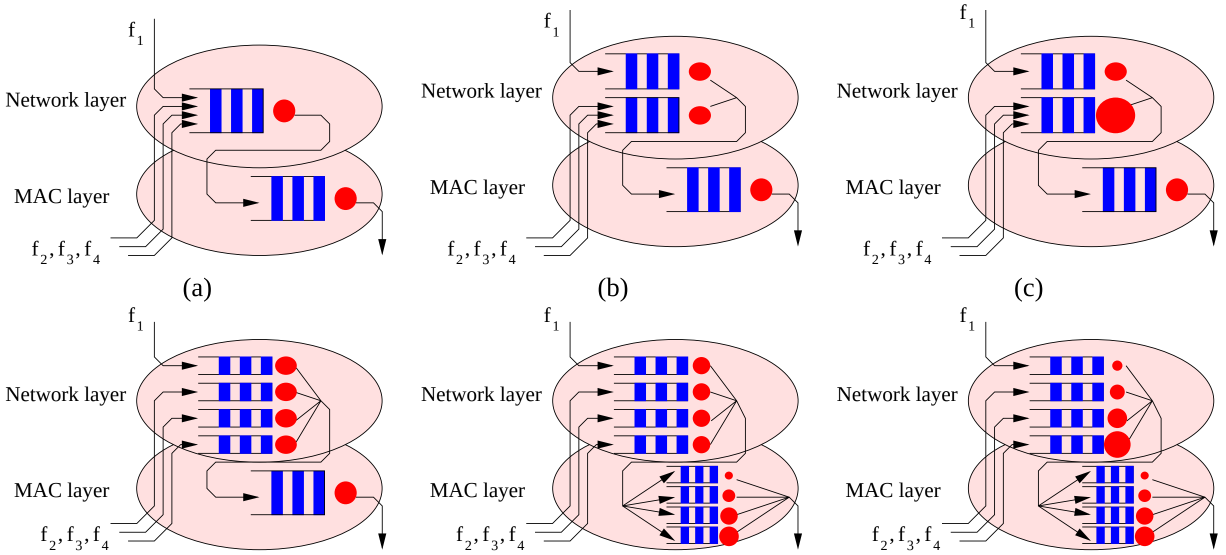

(a) (b) (c)

(d) (e) (f)

Fig. 3. Candidate queuing schemes for multihop wireless networks. (a) Single network-layer queue. (b) Two fair queues at the network layer. (c) Two weighted queues at the network layer. (d) Per-flow (per-user) fair queues at the network layer. (e) MAC-layer QoS is added to the scheme (d). (f) Weighted per-flow queues at the network layer with MAC-layer QoS support.

this scheme only when the length of a chain does not exceed two hops.

B. *Different weight on the relayed traffic*

To alleviate the unfairness shown in the previous section, one can assign different weights to each queue so that the queue for the relayed traffic will receive more bandwidth when needed. This scheme is modeled as shown in Fig. 3(c), where the larger disk indicates that a greater weight is given to the forwarding queue.

The weight of the forwarding queue can be fixed in all the nodes of the network, or different weights can be used, depending on the amount of relayed traffic at each node. The latter assumes that the amount of relayed traffic can somehow be determined in a distributed manner.

For the network shown in Fig. 2, the desirable weight ratios of the originating versus relayed traffic are 1:3 and 1:1 for nodes 1 and 2, respectively.

As in the case of the previous section, the network-layer throughput ratio of 1:1:1:1 and corresponding MAC-layer throughput ratio of 4:2:1:1 can be maintained when the offered load at each node is less than G. When the network saturates, the per-flow throughputs converge to $\frac{B}{4N}$, $\frac{3B}{16N}$, $\frac{3B}{8N}$ and $\frac{3B}{16N}$ respectively.

The resulting per-flow throughput ratio of nodes 1-4 is 4:6:3:3 and the MAC layer throughput ratio is 1:1:1:1, respectively. Giving more weights to the relayed traffic further alleviates the unfairness problem, and yet it does not achieve ideal fairness. The existing unfairness is the result of the asymmetric topology of the network, which is often the case in a real network. At node 1, incoming packets from nodes 2 and 3 are equally treated, although node 2 has to share its bandwidth with node 4.

C. *Per-flow queuing*

A more general approach is to use per-flow queuing. As shown in Fig. 3(d), packets of different flows are enqueued separately (based on their network-layer source address). We are considering the unidirectional flows towards the gateway for now.

When the network saturates, the throughput F_i of the i^{th} flow is:

$$F_i = \frac{B}{N} \times \frac{1}{4}. \tag{1}$$

There may exist networks lacking the resources to do per-flow queuing (e.g., in a large sensor network). Nevertheless, for many applications (especially in the access network where many multihop applications are encountered), per-flow queuing is a feasible strategy.

IV. MAC LAYER QOS AND BANDWIDTH EFFICIENCY

The per-flow fairness could be achieved by per-flow queuing at the network layer as shown in the previous section. But, when the offered load is high enough to saturate the network, a significant amount of bandwidth is wasted due to over-injected packets. The over-injection problem cannot be solved with a pure IEEE 802.11-like MAC layer exactly due to its fairness.

If the bandwidth loss due to over-injection is considerable, the fairness guarantee at the network layer become irrelevant. Fig. 4 shows a scenario where user throughput degrades as $O(N^2)$, where N is the number of user nodes. When the network saturates, all the nodes will have an equal share of bandwidth, $\frac{B}{N}$, at the MAC layer. At the network layer in node 1, this bandwidth is again fairly shared by N unique traffic flows. Thus, the worst-case per-flow throughput, F_{Worst} is

$$F_{Worst} = \frac{B}{N} \times \frac{1}{N}. \tag{2}$$

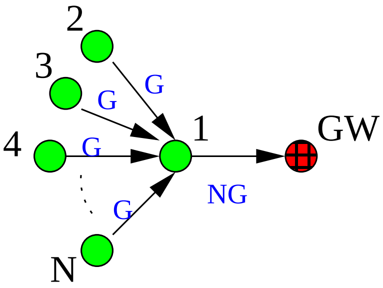

Fig. 4. A worst-case scenario for the bandwidth waste due to MAC fairness (i.e., lack of MAC layer QoS).

In order to avoid such a drastic performance degradation, a MAC layer QoS scheme is required. A MAC priority scheme similar to the one proposed in IEEE 802.11e [22] can be used to allocate a different bandwidth to each node. The MAC layer model in Fig. 3(e) shows the proposed scheme with MAC QoS enabled.

With QoS-enabled MAC layer, the throughput of i^{th} flow is calculated as

$$F_i = \frac{B}{2} \times \frac{1}{4}. \tag{3}$$

The factor N of the denominator in (1) is now replaced by a constant '2' because the MAC-layer bandwidth of $\frac{B}{2}$ is allocated to node 1 using the MAC-layer QoS. Compared to (1), the throughput is increased by 100

However, QoS schemes like IEEE 802.11e are designed for a single-hop network and may not work well in a multihop situation where bandwidth ratio between hidden nodes may not be maintained properly. A cross-layer scheme is also conceivable where the packet transmission is controlled at the MAC layer based on the information collected at the network layer. There are a number of different ways to do MAC QoS, but the details are beyond the scope of this paper.

V. DIFFERENTIATED SERVICE AND ASYMMETRIC TRAFFIC FLOWS

In this section, we examine how differentiated service can be provided by giving different weights to the network-layer queues. We also consider bidirectional traffic flows and explain how asymmetric bandwidth can be assigned.

A. *Differentiated service by weighted per-flow queuing*

Fairness is just a particular quality of service requirement. Fairness tries to ensure that all participants in the network receive the same share of the network bandwidth. Conceivably, for some networks some users should receive more bandwidth than others. For example, in a commercial wireless mesh network, some of the clients can have a "business class" subscription, and when the network is loaded, receive twice as much bandwidth as regular subscribers. Similarly, in a sensor network, some sensors can be more important than others (perhaps being closer to the observed phenomenon) and thus, should receive a larger share of the available bandwidth. The fairness schemes examined in the previous sections can be extended to provide differentiated service in multihop wireless networks.

The model shown in Fig. 3(f) depicts how differentiated service can be provided. The larger disks indicate the greater amount of per-flow bandwidth allocation.

In the network shown in Fig. 2, when the per-flow bandwidth ratio of nodes 1, 2, 3 and 4 is desired to be 1:2:3:4, the per-flow throughputs at saturation are $\frac{1}{23}B$, $\frac{2}{23}B$, $\frac{3}{23}B$ and $\frac{4}{23}B$ respectively.

B. *Bidirectional traffic flows and asymmetric bandwidth*

Although we used unidirectional traffic flows for the clarity of our analysis, traffic flows are bidirectional in a typical network. Packets originated from and destined for the same user node should be placed in separate queues. Giving different weights will enable asymmetric bandwidth allocation. Typical Internet users will benefit from larger centrifugal bandwidth (i.e., traffic flow away from the gateway), and server sites or sensor nodes will prefer larger centripetal bandwidth.

VI. SIMULATION

The analytical results presented above are verified through detailed simulations. In this section, we describe the architecture of the simulator and present the simulation results for different queuing schemes.

A. *Simulator Design*

We designed our simulator architecture in a similar way to that which is presented in [23]. The queue size and the queuing scheme is configured to model the different schemes under study. We assume IEEE 802.11 standard as the basic MAC layer, and the the TMT (i.e., B) is calculated based on the given MAC and physical (PHY) parameters. Specifically, we assume IEEE 802.11b, basic data rate of 11 Mbps and RTS/CTS scheme, and as a result, the TMT, B is 4.5 Mbps. To the best of our knowledge, there is no perfect MAC QoS scheme that works in a multihop wireless environment. Thus, we used a probabilistic approximation that models the ideal behavior of the MAC QoS. The network topology and the routing are manually configured to be the same as shown in Fig. 2. Exponential distribution is used for packet inter-arrival time, and the packet size is fixed to 1500 bytes.

B. *Simulation Results*

The simulation results for the queuing schemes in Fig. 3(a)-(f) are presented in Fig. 5(a)-(f), respectively. For all cases, the simulation result matches closely the analytically estimated values.

VII. CONCLUSION

In this paper we exposed a significant fairness problem existent practically in all wireless multihop networks. We proposed several network layer solutions to the fairness problem. The improvement in fairness is directly related to the resources investment. We showed that a network layer solution can restore fairness at the expense of bandwidth efficiency. This inefficiency can be dramatic in some topologies. It is shown

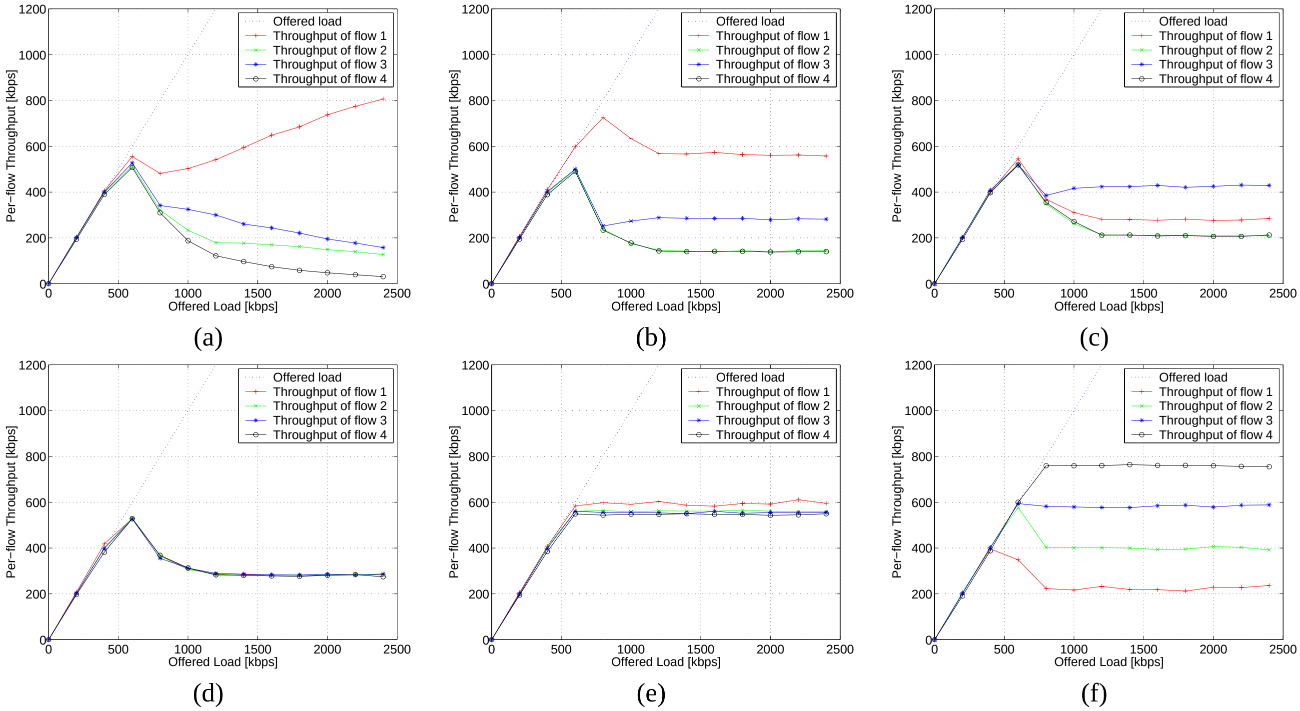

(a) (b) (c)

(d) (e) (f)

Fig. 5. Simulation results for different queuing schemes. (a) Single network-layer queue. (b) Two fair queues at the network layer. (c) Two weighted queues at the network layer. (d) Per-flow (per-user) fair queues at the network layer. (e) MAC-layer QoS is added to the scheme (d). (f) Weighted per-flow queues at the network layer with MAC-layer QoS support.

that a MAC layer providing priorities is able to restore network efficiency while maintaining network layer fairness. Finally, we generalized the fairness concept to enable differentiated services in multihop networks. The results are validated using simulations.

REFERENCES

[1] P. Gupta and P. R. Kumar, “The capacity of wireless networks,” *IEEE Trans. on Information Theory*, vol. 46, no. 2, Mar. 2000.

[2] S. Capkun, M. Hamdi, and J. P. Hubaux, “GPS-free positioning in mobile ad-hoc networks,” in *Proc. of the 34th Annual Hawaii International Conference on System Sciences*, Jan. 2001, pp. 3481–3490.

[3] B. Schrick and M. Riezenman, “Wireless broadband in a box,” *IEEE Spectrum Magazine*, no. 6, pp. 38–43, June 2002.

[4] M. W. Subbarao, “Dynamic power-conscious routing for MANETS: an initial approach,” in *IEEE Vehicular Technology Conference*, 1999, pp. 1232–1237.

[5] E. Royer and C. Toh, “A review of current routing protocols for ad-hoc mobile wireless networks,” *IEEE Personal Communications*, Apr. 1999. [Online]. Available: citeseer.nj.nec.com/royer99review.html

[6] M. Gerla, K. Tang, and R. Bagrodia, “TCP performance in wireless multihop networks,” in *Proc. of IEEE WMCSA 99*, New Orleans, LA, Feb. 1999. [Online]. Available: citeseer.nj.nec.com/gerla99tcp.html

[7] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, “A survey on sensor networks,” *IEEE Communication Magazine*, vol. 40, no. 8, pp. 102–116, Aug. 2002.

[8] L. Kalampoukas, A. Varma, and K. K. Ramakrishnan, “An efficient rate allocation algorithm for ATM networks providing min-max fairness,” UCSC-CRL-95-29 Computer Engineering Dept. University of California, Santa Cruz, Tech. Rep., June 1995.

[9] IEEE, “Wireless LAN medium access control (MAC) and physical layer (PHY) specification,” IEEE Std. 802.11, June 1999.

[10] V. Bharghavan, A. J. Demers, S. Shenker, and L. Zhang, “MACAW: A media access protocol for wireless LAN’s,” in *SIGCOMM*, 1994, pp. 212–225. [Online]. Available: citeseer.nj.nec.com/bharghavan94macaw.html

[11] C. L. Fullmer and J. J. Garcia-Luna-Aceves, “Floor acquisition multiple access (FAMA) for packet-radio networks,” in *SIGCOMM*, 1995, pp. 262–273. [Online]. Available: citeseer.nj.nec.com/fullmer95floor.html

[12] A. Woo and D. E. Culler, “A transmission control scheme for media access in sensor networks,” in *Mobile Computing and Networking*, 2001, pp. 221–235. [Online]. Available: citeseer.nj.nec.com/491433.html

[13] J. Postel, “Transmission control protocol,” RFC 793, Sept. 1981.

[14] S. Xu and T. Saadawi, “Does the IEEE 802.11 MAC protocol work well in multihop wireless ad hoc networks,” *IEEE Communications Magazine*, June 2001.

[15] S. Bae, K. Xu, S. Lee, and M. Gerla, “Measured analysis of TCP bahavior across multihop wireless and wired networks,” in *In Proc. IEEE Globecom 2002*, Taipei, Taiwan, R.O.C., Nov. 2002.

[16] K. Xu, S. Bae, M. Gerla, and S. Lee, “TCP behavior across multihop wireless and wired networks,” in *Proc. of WoWMoM 2002*, Atlanta, GA, USA, Sept. 2002.

[17] K. Xu and M. Gerla, “TCP over an IEEE 802.11 ad hoc network: Unfairness problems and solutions,” UCLA Computer Science Department, Tech. Rep. 020019, May 2002.

[18] K. Tang and M. Gerla, “Fair sharing of MAC under TCP in wireless ad hoc networks,” in *Proc. of IEEE MMT 99*, Venice, Italy, Oct. 1999.

[19] G. Holland and N. H. Vaidya, “Analysis of TCP performance over mobile ad hoc networks,” in *Proc. the ACM Intern. Conf. on Mobile Computing and Networking, (MobiCom)*, Seattle, WA, USA, Aug. 1999, pp. 207–218.

[20] P. Sinha, T. Nandagopal, N. Venkitaraman, R. Sivakumar, and V. Bharghavan, “WTCP: A reliable transport protocol for wireless wide-area networks,” *Mobile Computing and Networking*, pp. 231–241, 1999.

[21] J. Jun, P. Peddabachagari, and M. Sichitiu, “Theoretical maximum throughput of ieee 802.11 and its applications,” in *The 2nd IEEE International Symposium on Network Computing and Applications*, Apr. 2003.

[22] S. M. Sunghyun, “Ieee 802.11e wireless lan for quality of service.” [Online]. Available: citeseer.nj.nec.com/537394.html

[23] R. Luo, D. Belis, R. M. Edwards, and G. A. Manson, “A simulation design for link connection-oriented wireless mesh networks,” in *4th International Workshop on Mobile and Wireless Communications Network*, Sept. 2002, pp. 665–669.