

ECN Marking for Congestion Control in Multihop Wireless Networks

Vasilios A. Siris and Despina Triantafyllidou*

Institute of Computer Science (ICS),
Foundation for Research and Technology – Hellas (FORTH),
P.O. Box 1385, GR 711 10 Heraklion, Crete, Greece
`dtriant@ics.forth.gr`

Abstract. In this paper we propose an approach to increase TCP's fairness in multihop wireless networks, using ECN as a congestion signalling mechanism. The novel idea we introduce is that the marking probability at a wireless node is a function of the aggregate utilization in the node's neighborhood, which is determined by the sum of the receiving rates of all nodes within its collision domain. A node's received rate can be communicated to neighboring nodes by piggy-backing it on control packets, such as CTS and RTS messages, or data packets. Simulation results demonstrate that our approach can improve TCP's fairness in a multihop wireless network compared to drop tail queueing, while achieving the same aggregate throughput. Moreover, the proposed approach yields smaller average packet delay and delay jitter compared to drop tail queueing.

1 Introduction

The efficient control and management of wireless network resources has been an area of active research. Enhancing TCP performance over a wireless ad hoc network is a challenge, due to the constraints posed by the limited capacity of the wireless spectrum and the medium access contention among the wireless nodes. On the other hand, the number of users accessing the Internet through IEEE 802.11 wireless links is growing dramatically. The above motivate the need for efficient and fair congestion control over such networks.

Multihop environments pose three specific issues related to TCP's performance. First, changing the position of nodes causes path interruptions, resulting in repeated TCP timeouts. Efficient retransmission techniques to overcome this problem have been proposed. The second problem has to do with the fact that while there exists a TCP window size for which a highest throughput is achieved, TCP tends to operate with a higher window size, leading to packet losses and reduced throughput [4]. The third problem is the unfairness introduced by such networks [6].

* Authors are also with the Dept. of Computer Science, University of Crete.

In this paper we focus on the unfairness problem, and propose an Explicit Congestion Notification (ECN) marking approach that correctly accounts for the level of congestion within different collision domains of a multihop network. The approach combines two key ideas: first, it uses ECN as an end-to-end signalling mechanism for conveying congestion information across a multihop wireless network; second, marking at a wireless node is performed using a modification of the load-based marking (LBM) algorithm proposed in [12]. This modification suggests that the marking probability is a function of the aggregate utilization within the node's collision domain.

Although the application of ECN to wireless networks is not new, e.g. see [8, 9], its application as a common signalling mechanism for conveying congestion information in wired and wireless networks, in a way that takes into account the particular characteristics of the underlying wireless technology was first proposed in [11], for the case of 3G networks based on Wideband CDMA. However, IEEE 802.11 networks differ from 3G WCDMA based cellular networks, hence the marking procedure for each should be different. In the case of a single hop wireless network, as discussed in [11], the marking probability for the wireless channel, which is a single shared resource, can be taken to be a function of the level of congestion, expressed by the aggregate utilization measured at the access point, which includes both the uplink and the downlink transmission throughput.

In the case of a multihop wireless network, unlike a single hop network, there is no access point acting as a gateway for all the packets travelling in the wireless network. Moreover, due to attenuation in the wireless medium, contention is location-dependent, hence the total utilization can no longer accurately reflect the level of congestion in the wireless network. For this reason we consider the utilization within a collision domain, which is given by the sum of the receiving rates of all nodes within the collision domain; a node's receiving rate can be communicated to all its neighbors within its collision domain by piggy-backing this information on control packets, such as CTS and RTS messages, or data packets. Such a scheme has the goal to broadcast location-dependent congestion information within a single collision domain, which is analogous to the use of the RTS/CTS mechanism for informing nodes of the presence of other neighboring nodes with which their transmission can collide. Indeed, as the experimental results demonstrate, our approach can address the unfairness that is observed in hidden and exposed terminal topologies. Each node, based on the sum of the received rates of all neighboring nodes within its collision domain, computes a marking probability using a load-based marking (LBM) algorithm [12]. Hence, our proposal combines this local (collision domain specific) broadcasting of congestion information with ECN marking, which provides the end-to-end communication and accumulation of congestion information.

The rest of the paper is structured as follows. In Section 2 we present our approach for load-based marking over multihop wireless networks. In Section 3 we present and discuss simulation results demonstrating that the proposed approach can improve TCP's fairness over a multihop wireless network. In Section 4 we

present a brief overview of related work, identifying where it differs from the work presented in this paper. Finally, in Section 5 we conclude the paper, identifying related ongoing and future work.

2 Congestion Signaling in a Multihop Wireless Network

In this section we describe our approach for measuring the load and signaling congestion in a multihop wireless network. The approach is motivated by the observation that in such networks the level of congestion is no longer an individual characteristic of a single resource, but a shared feature of all the nodes within the same collision domain. Each node, depending on the packets it sends, receives or senses within its range, can obtain a different view of the level of congestion in its surrounding area. In fact, it itself also contributes to congestion. The above location-dependent information must be taken into consideration, in order to effectively and fairly control the usage of the wireless channel by all nodes. Given the congestion information from all its neighbors, a node determines the probability with which it will mark packets at its output queue. Then, based on the aggregate congestion along the whole path its packets traverse, which is signaled using ECN marks, the node responds by adjusting its transmission window according to the TCP congestion control protocol.

According to the approach proposed in this paper, the first key idea to achieve the above goal is to implement ECN as the end-to-end signalling mechanism for communicating congestion information within the multihop wireless network. Secondly, marking is performed at each node using a load-based marking (LBM) algorithm [12], modified to consider the appropriate measure of the level of congestion in a multihop wireless network.

2.1 ECN as the End-to-End Signalling Mechanism

Explicit Congestion Notification (ECN) has been approved as an IETF proposed standard [10]. With ECN, congestion of a network link is explicitly signaled by having routers set the CE (Congestion Experienced) bit located in the IP header, rather than implicitly signaled through lost packets as is the case with TCP's current operation. ECN can thus provide an early warning of incipient congestion, before packets start to be dropped, thus avoiding their retransmission. Hence, ECN can, to a large extent, avoid packet drops and the corresponding overhead of retransmitting lost packets.

ECN is appropriate for wireless networks, since in wireless networks non-congestion related losses due to wireless channel corruption can occur with a non-negligible probability. However, ECN alone cannot solve the problem of TCP decreasing its throughput in the case of non-congestion related losses. To address this either TCP's reaction to losses must be modified, such as by identifying and differentiating between congestion and non-congestion related losses, or link-layer mechanisms, such as forward error correction and retransmission over the wireless link, should hide losses due to corruption from TCP.

Our proposal for using ECN goes one step further from addressing the issue of congestion and non-congestion related losses, which we assume are handled by IEEE 802.11 MAC's link-layer retransmission mechanism. Our approach proposes to use ECN to convey congestion information across a multihop wireless network. For wired networks, marking is performed at the output link of routers. On the other hand, for a multihop wireless network we propose to implement marking at the output queue of every wireless node, considering the aggregate utilization within the nodes collision domain. In the following section we describe how this aggregate utilization can be measured in a distributed fashion.

2.2 Load-Based Marking (LBM) in a Multihop Environment

In a multihop network, there is no pre-defined link between any two nodes. Instead, nodes share the wireless channel and compete for it in a distributed manner using the MAC protocol. Thus, congestion cannot be tracked to a specific node, but to the entire area around it. Hence, the aggregate utilization of this area should be taken as an indication of the level of congestion of the wireless resource. Moreover, since there is no single shared buffer that is used for the packets flowing in the same area, a RED (Random Early Detection)-like marking algorithm, where the packet marking probability is a function of an average queue length, cannot be applied.

Apparently, for each node we need to define this area over which the aggregate utilization should be estimated. Interference with a node may be caused both by one-hop neighbors and by two-hop neighbors. However, the set of nodes that interfere with each other also depends on the traffic destination. For example, consider the multihop network shown in Figure 1, where the transmission range for each node extends up to its immediate neighbor. Assume the case where node 2 transmits to node 1, and node 3 transmits to node 4. Even though 2 and 3 are immediate neighbors, they can both simultaneously perform their transmissions, since each is two hops away from the other's destination (this is the so-called exposed terminal scenario). Similarly, which two hop neighbors interfere also depends on the destination. Hence, if both nodes 2 and 4 transmit to 3, then they are in the same collision domain (this is the so-called hidden terminal scenario). On the other hand, if 2 transmits to 1 and 4 transmits to some other node on its right (not shown in the figure), then nodes 2 and 4 are not in the same collision domain.

Next we discuss how a node can compute the aggregate utilization within its collision domain. First note that summing the transmission rates of neighboring




Fig. 1. All nodes apply ECN marking based on the congestion within their collision domain




Fig. 2. LBM has three parameters: time interval t_{avg} over which aggregate utilization is measured, minimum utilization threshold ρ_0 , and slope parameter α . Marking probability is a piecewise linear function of the average utilization

nodes is not correct, since in the scenario mentioned above where node 2 transmits to node 1, and node 3 transmits to node 4 (exposed terminal), because 2 and 3 do not interfere, node 2 should not consider node 3’s transmission rate. Moreover, considering the transmission rate wouldn’t account for contention due to two hop neighbors, for example in the case both nodes 2 and 4 transmit to node 3. Summing the utilization measured by neighboring nodes is for the same reason incorrect; indeed, such an approach has another subtle problem since it can result in considering some transmissions twice (once through the transmission rate and once through the receiving rate). Another option is to sum the received rates of all neighbors. In the case where the transmission range is the same for all nodes, and the transmission range is the same with the interference range, then this approach would indeed consider the received rates of only those neighbors that a node’s transmission can reach, hence can interfere with other nodes transmitting to the same destination.

Based on the previous discussion, we propose that the marking probability is a function of the aggregate utilization over some time interval t_{avg} , measured by dividing the sum of the received throughput of all nodes within the same collision domain, with the wireless channel capacity. Because we are using TCP, t_{avg} should be at least a few RTTs, to allow the system to obtain an accurate estimate of the average rate. Also, it should be such that the system can track traffic changes, e.g. of number of users, therefore not too large.

The marking probability can have a piecewise linear dependence on the aggregate utilization, as illustrated in Figure 2. The marking probability is zero when the average utilization is less than ρ_0 . For utilization values ρ larger than ρ_0 , the marking probability is given by $\min\{\alpha(\rho - \rho_0), 1\}$. Hence, the extended load-based marking algorithm for multihop environments has the same three LBM parameters: the time interval t_{avg} over which the aggregate utilization is measured, the minimum utilization threshold ρ_0 , and the slope parameter α .

3 Simulations, Results and Discussion

In this section we present and discuss the simulation results comparing the proposed marking approach with drop tail queueing. The results show that the

proposed approach can achieve higher fairness compared to drop tail queueing, while achieving the same utilization; as utilization we consider the ratio of the throughput -including the packet header overhead- and the wireless capacity, which is 11 Mbps. Furthermore, we show that the proposed approach results in smaller packet delay and delay jitter.

3.1 Experiment Scenarios

Experiments were conducted using the NS-2 simulator, with the topologies shown in Figure 3. The transmission range equals interference range, both being 250 m. The scenario in Figure 3(a) implements the hidden terminal case. Traffic flows from node 1 to node 2, and from node 4 to node 3. Marking is performed at the senders, i.e. nodes 1 and 4, and the marking probability at both of these nodes considers the sum of the receiving rate at nodes 2 and 3. The scenario in Figure 3(b) implements the exposed terminal case, with traffic flowing from node 2 to node 1, and from node 3 to node 4. In this case, node 2 calculates a marking probability based on the receiving rate of node 1, whereas node 3 calculates a marking probability based on the receiving rate of node 4. The scenario in Figure 3(c) is a multihop scenario. Traffic flows from node 3 to node 1 through node 2, and from node 4 to node 5. Each of the three sending nodes, i.e. nodes 3, 2, and 4 performs marking based on the received rate of its one-hop receiver.

Ftp flows transfer files whose sizes follow a pareto distribution of average 500 KBytes and 5 MBytes, and shape parameter 1.5. The throughput for each flow is given by the ratio of the total received traffic -data and overhead- and the duration of the file transfer, taken as the interval between the time the first SYN packet is sent, and the time the last ACK is received. Each flow's start-time was randomly selected from the interval $[0, 5]$ seconds. We used TCP Reno. In all the experiments the measurement interval parameter of the LBM scheme is 500 ms. On the other hand, we consider different values for the minimum utilization threshold and the slope parameter.

The IEEE 802.11 MAC layer performs retransmissions of corrupted packets. To model the bursty (time correlated) errors in the channel, we consider a multi-state error model similar to that presented in [2], consisting of a three-state

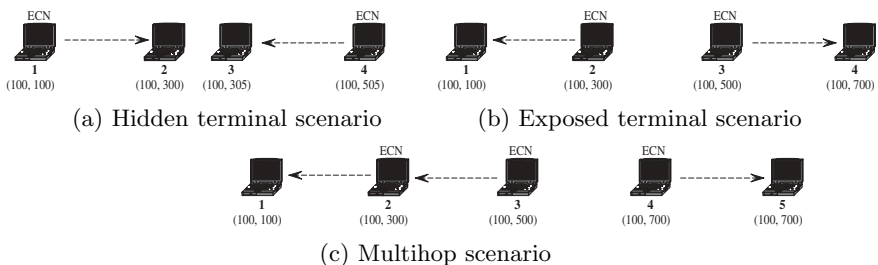


Fig. 3. Simulation scenarios

Table 1. State transition matrix

Current state	Next state		
	G	B	VB
G	0.95	0.02	0.03
B	0.10	0.20	0.70
VB	0.20	0.03	0.77

discrete-time Markov chain. The time slot is equal to the slot of the 802.11b MAC sublayer, i.e. $20 \mu s$. In the good state (G), transmission is error free. In state B and VB errors occur uniformly with probability p , and $5.5 \times p$, respectively. The probability of state transition is shown in Table 1. In our experiments we consider the case of error free transmission, and the case where errors occur according to the above model with an average error probability 1%, achieved for $p = 0.01$.

3.2 Fairness and Throughput

The graphs we present next show the average fairness, and the corresponding 95% confidence interval from 15 independent runs of the same experiment. As a measure of fairness we consider the widely used metric given in [3]:

$$\text{Fairness Index} = \frac{\left(\sum_{i=1}^N x_i\right)^2}{N \sum_{i=1}^N x_i^2},$$

where x_i is the rate of flow i and N the number of flows. The fairness index takes values in the interval $(0,1]$, with a higher value indicating higher fairness.

Figures 4(a) and 4(b), for the hidden terminal scenario of Figure 3(a), show that the proposed LBM scheme achieves higher fairness compared to drop tail (DT) queueing, for both short and long ftp flows, and for a range of values of the LBM minimum utilization threshold ρ_0 . Figure 5(a), which is for the exposed terminal scenario of Figure 3(b), also shows that fairness improves. Observe that with short ftp flows, the fairness improvement appears to be smaller for smaller values of the threshold parameter; this can be attributed to the fact that very small thresholds correspond to very early, and thus, inefficient marking.

Comparison of Figure 4(a) with 4(b) shows that the improvements are larger for longer ftp flows. This is also the case in the exposed terminal scenario. Results also show that for long ftp flows the improvement is larger in the exposed terminal scenario; on the other hand, the improvement is similar in both scenarios, in the case of small ftp flows. To illustrate what the different values of fairness imply, note that a fairness index equal to 0.8 corresponds to the case where the two flows achieve a throughput of 2.09 Mbps and 0.75 Mbps. On the other hand, a fairness index equal to 0.97, corresponds to the case where the two flows achieve a throughput of 1.40 Mbps and 1.87 Mbps. For an error free channel, fairness is also improved with the proposed approach, compared to DT.

In all the above experiments, the aggregate throughput achieved by both the LBM scheme and drop tail queueing is the same, as shown in Figure 5(b); observe

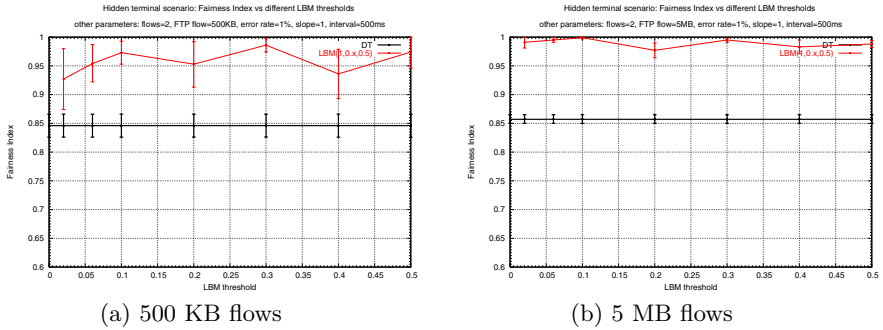


Fig. 4. Fairness for the hidden terminal scenario, for different LBM minimum utilization thresholds ρ_0 and average error probability 1%

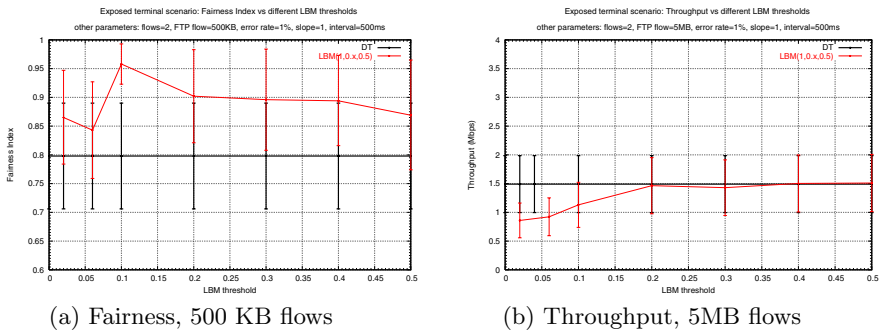


Fig. 5. Throughput and fairness for the exposed terminal scenario, and different LBM minimum utilization thresholds ρ_0 , and average error probability 1%

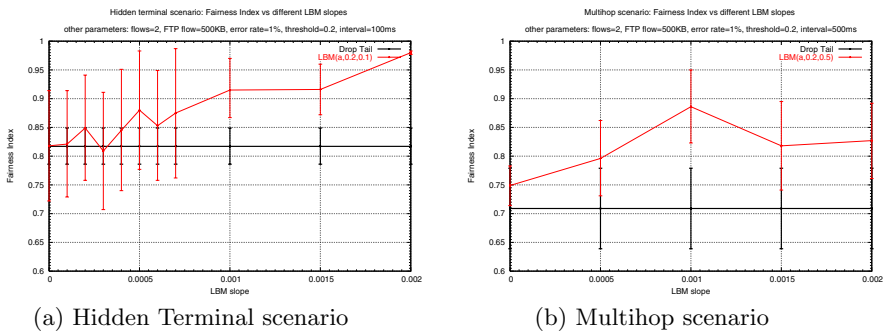


Fig. 6. Fairness for different LBM slope parameters, and for 500 KB flows and average error probability 1%

that the throughput decreases for small values of the utilization threshold, which is also the case with the fairness index, and which corresponds to a rather early marking. Hence, improved fairness is not achieved at the expense of decreased throughput, as is the case with the NRED scheme proposed in [6]. Conclusions are identical for the hidden terminal scenario. In Figures 6(a) and 6(b), LBM shows higher fairness compared to DT, for different slope parameters. Observe that the improvement decreases for smaller values of the slope parameter. This can be explained by the fact that a smaller slope yields a less aggressive marking.

3.3 Packet Delay and Delay Jitter

Table 2 shows the average and standard deviation of the packet delay over the wireless link. The packet delay is given from the time the packet is enqueued in the output queue of the TCP sender, until the corresponding ACK is received. Thus, it also includes queuing delay. Observe that LBM marking scheme achieves a smaller average delay and delay jitter compared to drop tail queuing, as indicated by the smaller values of the standard deviation.

Table 2. Average and standard deviation of packet delay over the wireless link. File size=1 MByte, loss prob.=1%, LBM: $\alpha = 1$, $\rho_0 = 0.1$, $t_{avg} = 500$ ms

Terminal scenario Flow		DT		LBM	
		avg. delay(ms)	std.dev.	avg. delay(ms)	std.dev.
Hidden	1	97.6	77.7	13.77	13.1
Hidden	2	60.9	57.7	9.24	26.1
Exposed	1	41.3	26.5	22	19.8
Exposed	2	49.6	65.9	13.9	37.9

4 Related Work

In this section we present a brief overview of the related work. Schemes considering ECN marking fall within the end-to-end approaches for improving TCP's performance over wireless networks, and have been proposed for wireless networks in [8, 9, 5]. The new contribution of this paper is to propose an approach for computing the ECN marking probability that takes into account the location-dependent nature of contention in multihop wireless networks.

Other approaches consider using a RED-like mechanism [9] or some other level of congestion, such as a congestion price [1]; however, as we argue in this paper, a shared buffer does not exist in a wireless network, hence such marking algorithms do not reflect the underlying resource sharing model. The work in [6] explores the relationship between TCP unfairness and early network congestion in wireless multihop networks, suggesting a scheme called Neighborhood-RED (NRED) for marking, based on the aggregate (incoming and outgoing) queue size of a node's neighborhood. The underlying idea of this approach is that

improvement of spatial reuse can lead to better fairness. However, in this work fairness is achieved at the expense of decreased throughput; this is not the case with the approach proposed in this paper as our simulation results demonstrate.

The work in [4] considers a packet drop scheme (which the authors note can also be adapted to a mark probability scheme), where the drop probability is a piecewise linear function of the number of MAC layer retransmissions, which can signal network overload, when they exceed some minimum number. Such a scheme, called Link-RED, was shown to increase the performance in the case of multi-hop wireless networks. Finally, the work in [7] considers the general framework for congestion control schemes using a utility-based modelling approach. The proposed marking scheme is a concave function of the traffic arriving at a link, when this rate is larger than some minimum capacity value. Our approach differs in that the marking probability is a linear function of the aggregate utilization in a collision domain.

5 Conclusions

The main contribution of this paper is to use the aggregate receiving rate within a node's collision domain to compute an ECN marking probability. Simulations show that the proposed approach, operating with TCP congestion control, can achieve higher fairness compared to drop tail queuing, while achieving the same aggregate throughput. The approach, also, yields smaller packet delay and delay jitter.

We are currently performing experiments with more complex topologies, to investigate the influence of the LBM parameters on the performance. Such results can guide their proper selection, and provide directions on how they can be dynamically adjusted. Having focused on fairness, an interesting research topic is how marking can also be used to improve the aggregate throughput in a multihop wireless network.

References

1. S. Athuraliya, S. H. Low, V. H. Li, and Q. Yin. *REM: Active Queue Management*. IEEE Network, pages 48-53, May/June 2001.
2. M. Bottigliengo, C. Casetti, C.F. Chiasserini, and M. Meo. *Short-term Fairness for TCP Flows in 802.11b WLANs*. Proc. of IEEE INFOCOM'04, 2004.
3. D. Chiu and R. Jain. *Analysis of the Increase/Decrease Algorithms for Congestion Avoidance in Computer Networks*. Computer Networks and ISDN Systems, 17:1-14, June 1989.
4. Z. Fu, P. Zerfos, H. Luo, L. Zhang, and M. Gerla. *The Impact of Multihop Wireless Channel on TCP Throughput and Loss*. Proc. of IEEE INFOCOM'03, 2003.
5. H. Inamura, G. Montenegro, R. Ludwig, A. Gurtov, and F. Khavizov. *TCP over 2.5G and 3G wireless networks*. RFC 3481, February 2003.
6. L. Qi Y. Shu K. Xu, M. Gerla. *Enhancing TCP Fairness in Ad Hoc Wireless Networks Using Neighborhood RED*. Proc. ACM MOBICOM'03, 2003.

7. S. Kunniyur and R. Srikant. *End-to-End Congestion control Schemes: Utility Functions, Random Losses and ECN Marks*. IEEE/ACM Trans. on Networking, 11(5):689-702, October 2003.
8. G. Montenegro, S. Dawkins, M. Kojo, V. Marget, and N. Vaidya. *Long Thin Networks*. RFC 2757, January 2000.
9. F. Peng, S.D. Cheng, and J. Ma. *A Proposal to Apply ECN to Wireless and Mobile Networks*. Proc. INET'00, 2000.
10. K.K. Ramakrishnan, S. Floyd, and D. Black. *The Addition of Explicit Congestion Notification (ECN) to IP*. RFC 3168, September 2001.
11. V.A. Siris. *Resource Control for Elastic Traffic in CDMA Networks*. Proc. of ACM MOBICOM'02, 2002.
12. V.A. Siris and D. Triantafyllidou. *Seamless Congestion Control over Wired and Wireless IEEE 802.11 Networks*. Proc. of Networking 2004, 2004.