

Seamless Congestion Control over Wired and Wireless IEEE 802.11 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
vsiris@ics.forth.gr

Abstract. We present a new approach for seamless congestion control over heterogeneous networks containing wired and wireless IEEE 802.11 links. The approach uses ECN (Explicit Congestion Notification) as a common signalling mechanism for conveying congestion information from both wired and wireless links. Two additional novel aspects of the approach are that ECN marking for a wireless link, due to the way resources are shared, is performed for both the uplink and the downlink based on measurements of the aggregate traffic in both directions, and the marking mechanism dynamically adapts to varying traffic and load conditions. Simulation results demonstrate that our approach achieves higher fairness compared to drop-tail queueing, and can effectively control the average packet delay over the wireless link.

1 Introduction

The number of users accessing the Internet and enterprise intranets through wireless links, and IEEE 802.11 wireless LANs (WLANs) in particular, is expected to grow dramatically with the proliferation of wireless hotspots and enterprise WLANs. Hence, there is a need for efficient and fair congestion control over heterogeneous networks that include both wired and wireless links.

In this paper we propose a new approach that enables TCP, in conjunction with Explicit Congestion Control (ECN), to operate seamlessly over heterogeneous networks. The approach combines three key ideas: First, it uses ECN as the common end-to-end signalling mechanism for conveying congestion information from both wired and wireless links; second, marking for the wireless link is performed using a load-based marking (LBM) algorithm, where the marking probability is a function of the aggregate utilization; third, the load-based marking algorithm dynamically adapts to varying traffic and load conditions in order to achieve an average packet delay over the wireless link within a target range.

Although the application of ECN to wireless networks is not new, e.g. see [1, 2], its application as a common signalling mechanism for conveying congestion information in wired and wireless networks, in a way that takes into account the

* The authors are also with the Department of Computer Science, University of Crete

particular characteristics of the underlying wireless technology was first proposed in [3], for the case of 3G networks based on Wideband CDMA. However, as we discuss in this paper, IEEE 802.11 WLANs differ from 3G WCDMA based cellular networks, hence the marking procedure for each should be different.

It is interesting to position our approach with respect to the cross-layer paradigm, which has emerged as an effective way for designing efficient network protocols over wireless link technologies [4]. Cross-layer design departs from the strict layer separation, which has been the traditional approach for network protocol design. Our approach follows the cross-layer design paradigm in the sense that the proposed ECN marking procedure takes into account the particular characteristics and the resource sharing model of IEEE 802.11 WLANs, hence differs from the marking procedure for wired links. Interestingly, our approach maintains TCP's end-to-end operation and semantics, hence adheres to the end-to-end argument stated in [5]: congestion control is performed at the end-systems, which is where aggregate information on the level of congestion for the whole end-to-end path exists. The particular characteristics of the underlying link technology are hidden from the TCP layer through appropriate design of the ECN marking procedure.

The rest of the paper is structured as follows. In Section 2 we present our approach for seamless end-to-end congestion control, in Section 3 we present and discuss simulation results, and finally in Section 4 we conclude the paper.

2 Seamless Congestion Control

2.1 ECN as a Common Signalling Mechanism

Explicit Congestion Notification (ECN) has been approved as an IETF proposed standard [6]. 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 due to congestion and the corresponding overhead of retransmitting lost packets.

Our approach for using ECN goes one step further, and proposes to use ECN to convey congestion information from both the wired and the wireless links. For wired networks, marking is performed at the output link of routers, whereas for a wireless link, marking is performed at the access point, Figure 1. Indeed, the marking procedure for wired and wireless links need not be, and as we argue in this paper, should not be the same.

2.2 Load-Based Marking (LBM)

For WLANs based on IEEE 802.11, both the uplink and the downlink share the same resource (wireless spectrum). Hence, the aggregate throughput in both

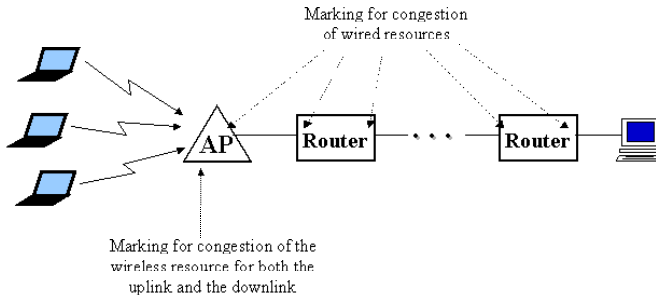


Fig. 1. Routers are responsible for marking based on the congestion level at each of their wired links, whereas the access point (AP) is responsible for marking based on the congestion level of the wireless resource (spectrum).

directions should be taken as an indication of the utilization, and hence 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 both directions, a RED (Random Early Detection)-like marking algorithm, where the packet marking probability is a function of an average queue length, cannot be applied.

Based on the above discussion, we propose that the probability of marking a packet, flowing either in the uplink or the downlink direction the wireless link, is a function of the aggregate utilization over some time interval t_{avg} , taking into account traffic flowing in both directions. The marking probability can have a piecewise linear dependence on the aggregate utilization: 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 $\min\{\alpha(\rho - \rho_0), 1\}$.

It is interesting to note that in the case of 3G cellular networks based on Wideband CDMA, unlike the case of WLANs based on IEEE 802.11, the uplink and downlink directions use different frequency bands. Furthermore, in the downlink of WCDMA networks there is a shared buffer located at the base station; this is not the case in the uplink direction, where each mobile has its own local buffer. Hence, similar to 802.11, a RED-like mechanism cannot be applied in the uplink direction of WCDMA networks.

2.3 LBM Adaptation

Recall that the LBM algorithm has three parameters: the time interval t_{avg} over which the average utilization is measured, the minimum utilization ρ_0 , and the slope parameter α . The time interval t_{avg} determines how quickly the algorithm adjusts the marking probability to changes of the aggregate utilization, and the timescale over which congestion is detected. Typically, t_{avg} will be set to some number of round trip times, in order to obtain stable measurements of the load. The slope parameter α affects the reactivity and the stability of the marking algorithm [7]: A higher slope would yield a more reactive algorithm since a small

change of the utilization would give a large change of the marking probability. Finally, for a fixed slope parameter α , the minimum utilization ρ_0 determines the marking probability for a given aggregate utilization, hence the utilization achieved in the steady state.

Based on the above discussion, we propose the following procedure for adjusting the minimum utilization parameter ρ_0 : The average delay for transmitting a packet over the wireless link is measured in a non-intrusive manner, i.e. the delay is measured for actual data packets transmitted over the wireless link. The parameter ρ_0 is adaptively adjusted so that the average delay is within a target interval $[d_{min}, d_{max}]$. In particular, ρ_0 is increased when the average delay is less than d_{min} , and is decreased when the average delay is greater than d_{max} . The magnitude of the change in each increase or decrease step is determined by the minimum utilization step size $\Delta\rho_0$. The values of d_{min} and d_{max} are related to the target packet delay requirements over the wireless link.

3 Simulation Results and Discussion

In this section we present and discuss simulation results comparing the proposed marking approach with drop tail queuing. Our experiments were conducted using the ns-2 simulator. The topology simulated is the one shown in Figure 1. Traffic flows from the fixed host to the wireless hosts, i.e. from right to left. In the experiments the IEEE 802.11 MAC layer performs retransmission of corrupted packets; losses due to corruption are assumed to be independent (non-bursty). We consider ftp flows that transfer files whose sizes follow a pareto distribution with average 500 KBytes.

3.1 Fairness and Throughput

The graphs we present next show the average and 95% confidence interval, from 10 independent runs of the same experiment. As a measure of fairness we consider




Fig. 2. Fairness for different number of FTP flows and different loss probabilities. LBM parameters: $\alpha = 1$, $\rho_0 = 0.15, 0.2$, $t_{avg} = 500$ ms. Average file size = 500 KBytes.

Table 1. Average and standard deviation of packet delay (in ms) over the wireless link. File size = 500 KBytes, RTT = 50 ms, loss prob. = 0.01, LBM parameters: $\alpha = 1, \rho_0 = 0.2, t_{avg} = 500$ ms

# of flows	DT		LBM	
	avg	std.dev.	avg	std.dev.
5	39.1	16.4	15.9	4.3
10	65.9	17.3	32.9	14.3
15	92.9	19.1	57.4	8.7
20	101.1	10.7	71.9	15.0

the fairness index given by Fairness Index = $\frac{(\sum_{i=1}^N x_i)^2}{N \sum_{i=1}^N x_i^2}$, where x_i is the rate of flow i and N is the total number of flows. The fairness index takes values in the interval (0,1], with a higher value indicating higher fairness.

Figure 2(a) shows that LBM achieves better fairness compared to DT; moreover, the difference between the fairness achieved by LBM and DT is larger for a larger number of flows. Figure 2(b) shows the fairness for different packet loss probabilities over the wireless link. Observe that the difference between the fairness achieved with LBM and DT is larger for smaller loss probabilities. The utilization achieved by both DT and LBM, with $\rho_0 = 0.2$, is identical. The fact that the use of ECN does not result in higher utilization compared to DT should not be that surprising, since experiments for wired networks also show that, for an appropriately dimensioned network, TCP with ECN does not achieve higher throughput compared to TCP with drop tail queueing [8].

3.2 Packet Delay and LBM Adaptation

Table 1 shows the average delay and the standard deviation of the delay over the wireless link. Observe that LBM achieves a smaller average delay and delay jitter, as indicated by the smaller values of the standard deviation.

Next we investigate the dynamic adaptation feature of the proposed approach. Assume that the goal is to maintain the average delay within the interval $[d_{min}, d_{max}] = [10, 15]$ milliseconds. To achieve this, the value of ρ_0 is increased when the average delay is smaller than d_{min} , or is decreased when the average delay is larger than d_{max} . Moreover, we assume that the adjustment of ρ_0 is performed in steps of $\Delta\rho_0 = 0.01$; this step size affects how fast the algorithm adapts to changes in the network load. The results appear in Table 2, and show that by adjusting the minimum threshold parameter ρ_0 we can effectively control the average delay, such that it remains inside the target interval.

The results in Table 2 do not illustrate the dynamic behaviour of the LBM adaptation procedure, e.g. how fast the algorithm adapts to sudden changes in the traffic and load conditions, that result in the average delay obtaining values outside the target interval. The transient behaviour will depend on the step parameter $\Delta\rho_0$ and the difference $d_{max} - d_{min}$; results for such dynamic behaviour will be included in an extended version of this paper.

Table 2. Minimum utilization threshold ρ_0 for a different number of flows, when the target delay interval is [10, 15] ms.

# of flows	avg delay	std.dev.	throughput (Mbps)	ρ_0
5	13.7	3.4	1.7	0.19
10	13.8	2.5	2.2	0.13
15	14.5	2.4	2.3	0.07
20	14.0	2.4	2.2	0.05

4 Conclusions

We presented an approach for seamless congestion control over heterogeneous networks containing wired and wireless IEEE 802.11 links. The approach combines three key features: First, it uses ECN as the common end-to-end signalling mechanism for conveying congestion information from both wired and wireless links; second, marking for the wireless link is performed using a load-based marking algorithm, where the marking probability is a function of the aggregate utilization; third, the load-based marking algorithm dynamically adapts to varying traffic and load conditions.

Further experimental work is investigating the transient behaviour of the proposed algorithm, in a dynamic environment where flows arrive and leave, different shapes of the marking probability curve (convex and concave, rather than piecewise-linear that we consider in this paper), and different measures of congestion of the wireless medium, such as the delay to access the wireless medium and the throughput measured in packets per time unit. Other interesting areas are the application of the proposed approach to multi-hop wireless networks and its combination with more advanced link-layer retransmission mechanisms.

References

1. Montenegro, G., Dawkins, S., Kojo, M., Magret, V., Vaidya, N.: Long thin networks. RFC 2757 (2000)
2. Peng, F., Cheng, S.D., Ma, J.: A proposal to apply ECN to wireless and mobile networks. In: Proc. INET'00. (2000)
3. Siris, V.A.: Resource control for elastic traffic in CDMA networks. In: Proc. of ACM MOBICOM'02. (2002)
4. Shakkottai, S., Rappaport, T.S., Karlsson, P.C.: Cross-layer design for wireless networks. IEEE Commun. Mag. (2003) 74–80
5. Saltzer, J., Reed, D., Clark, D.: End-to-end arguments in system design. ACM Trans. Comp. Sys. **2** (1984) 277–288
6. Ramakrishnan, K.K., Floyd, S., Black, D.: The Addition of Explicit Congestion Notification (ECN) to IP. RFC 3168 (2001)
7. Kelly, F.P., Maulloo, A., Tan, D.: Rate control in communication networks: shadow prices, proportional fairness and stability. Journal of the Operational Research Society **49** (1998) 237–252
8. Pentikousis, K., Badr, H.: An evaluation of TCP with Explicit Congestion Notification. To appear in *Annals of Telecommunications*. (2003)