

Seamless Congestion Control over Wired and Wireless IEEE 802.11 Networks

Vasilios A. Siris* and Despina Triantafyllidou

Technical Report 327 - October 2003

Institute of Computer Science (ICS)

Foundation for Research and Technology – Hellas (FORTH)

P.O. Box 1385, GR 711 10 Heraklion, Crete, Greece

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 different traffic and load conditions. Simulation results demonstrate that our approach achieves higher fairness compared to drop-tail queueing, while achieving the same utilization, and can effectively control the average packet delay over the wireless link.

Keywords: congestion control, Explicit Congestion Notification (ECN), load and delay monitoring, fairness

1 Introduction

An increasing number of users is accessing the Internet and enterprise intranets through wireless links, and IEEE 802.11 based wireless LANs (WLANs) in particular. This number is expected to grow dramatically with the proliferation of

*Corresponding author: V. A. Siris, vsiris@ics.forth.gr, tel./fax: +30 2810 391726/391601. The authors are also with the Dept. of Computer Science, University of Crete

wireless hotspots and enterprise WLANs. Moreover, emerging multimedia services over wireless networks will have different bandwidth and delay requirements, and, compared to wired networks, there is a limited ability to increase the capacity of wireless networks, since the capacity is determined by the available wireless spectrum. All the above motivate the need for efficient and fair congestion control over heterogeneous networks that include both wired and wireless links.

In this paper we propose an 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 [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 [14], for the case of 3G networks based on Wideband CDMA. 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. In particular, according to our approach the marking probability for the wireless link is a function of the aggregate traffic load in both the uplink and the downlink. Moreover, the same marking algorithm is used for both the uplink and the downlink, since both directions share the same resource (wireless spectrum); it is for the same reason that we consider the aggregate traffic for computing the marking probability. The application of a load-based marking algorithm for wired networks, and its interaction with various end-system congestion control algorithms was investigated in [15].

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 [13]. 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 char-

acteristics and the resource sharing model of IEEE 802.11 WLANs, hence differs from the marking procedure for wired links; furthermore, as indicated above, different wireless link technologies have different characteristics hence can require different marking procedures. Interestingly, our approach maintains TCP's end-to-end operation and semantics, hence adheres to the end-to-end argument stated in [12]: congestion control is performed at the end-systems since it is there where aggregate information on the level of congestion for the whole end-to-end path exists. The particular characteristics of the underlying link technology is hidden by 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 that demonstrate that the proposed approach increases fairness and can be used to effectively control the average packet delay over the wireless link. 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 Seamless congestion control

In this section we describe our approach for seamless congestion control over wired networks and WLANs based on the IEEE 802.11 protocol. 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 the following subsections we discuss each of the above three key ideas.

2.1 ECN as a common signalling mechanism

Explicit Congestion Notification (ECN) has been approved as an IETF proposed standard [11]. 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 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. In particular, as we discuss in the next subsection, the access point is responsible for packet marking in both the uplink and the downlink directions, based on measurements of the aggregate utilization that takes into account the traffic flowing in both directions. Hence, the marking procedure for the wireless link takes into account its particular characteristics and resource usage model, and hence differs from the marking algorithm used for a wired link.

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 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




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

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, 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 load-based marking (LBM) algorithm has three parameters: the time interval t_{avg} over which the aggregate utilization is measured, the minimum utilization ρ_0 , and the slope parameter α .

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, for the downlink of WCDMA networks a RED-like mechanism, where the packet marking probability is a function of the average queue size at the base station, can be applied. On the other hand, a RED-like mechanism cannot be applied to the uplink since, similar to IEEE 802.11 WLANs, there is no single buffer shared by packets originating from different mobile hosts.

2.3 LBM adaptation

Next we discuss how the LBM algorithm can dynamically adapt to varying traffic and load conditions, arguing that the most appropriate parameter to achieve




Figure 2: With load-based marking (LBM), the marking probability is a piecewise linear function of the average utilization. The LBM algorithm has three parameters: the time interval t_{avg} over which the aggregate utilization is measured, the minimum utilization ρ_0 , and the slope parameter α .

effective adaptation is the minimum utilization parameter ρ_0 .

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 [6]: 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; indeed, the difference

$d_{max} - d_{min}$ determines the allowed variation of the average delay in different traffic and load situations.

3 Simulation results and discussion

In this section we present and discuss 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 (11 Mbps). Furthermore, we show that the proposed approach can achieve smaller average delay over the wireless link, and can be used to effectively control this average delay by dynamically adapting to different traffic and load conditions.

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 50 KBytes and 500 KBytes. The throughput for each flow is given by the ratio of the total transmitted traffic (data and overhead) and the duration of the file transfer, where the latter is the interval between the time the first SYN packet is sent by the sender, and the time the last acknowledgement is received. The start time of each ftp flow was randomly selected from the interval $[0, 0.5]$ seconds (for 50 KBytes average file size) and the interval $[0, 5]$ seconds (for 500 KBytes average file size). Finally, in the experiments we use TCP Reno.

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 the fairness index given by [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 is the total number of flows. The fairness index takes values in the interval $(0,1]$, with a higher value indicating higher fairness.

Figures 3(a) and (b) shows the fairness and throughput respectively, for drop tail (DT) queueing and load based marking (LBM) with slope parameter $\alpha = 1$, minimum threshold $\rho_0 = 0.15$ and 0.2 , and averaging interval $t_{avg} = 500$ ms; we have found that it is sufficient to set this interval to be a few times the round trip time. Figure 3(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 3(b) shows that 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 experiment 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 [10]. Figure 4(a) and (b) show the fairness and throughput respectively, for DT and LBM, with different parameters. The conclusions from these graphs are identical to the above.

In the case of LBM with $\rho_0 = 0.15$, the utilization for 5 and 10 flows is smaller. Latter in this section we present and discuss results showing how the minimum threshold ρ_0 affects the average packet delay.

Figures 5(a) and (b), which are for $RTT = 20$ ms and 100 ms respectively, support the above conclusion regarding the improved fairness of LBM. Figure 6(a) 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. Figure 6(b) shows that, for different wireless loss probabilities, the throughput achieved by LBM and DT is the same; indeed, it decreases as the loss probability increases.

The results in Figures 7(a) and (b) are for average file size 50 KBytes. Figure 7(a) shows that, as above, LBM achieves higher fairness than DT. Figure 7(b) shows that LBM can achieve the same throughput as DT, by appropriately setting the minimum threshold parameter ρ_0 . Moreover, observe that, as expected, a higher value of ρ_0 can lead to lower throughput.




Figure 3: Fairness and throughput for different number of FTP flows. LBM parameters: $\alpha = 1$, $\rho_0 = 0.15, 0.2$, $t_{avg} = 500$ ms. Average file size = 500 KBytes, RTT = 50 ms, wireless loss probability = 0.01.




Figure 4: Fairness and throughput for different number of FTP flows. LBM parameters: $\alpha = 1, 0.7, 1.3$, $\rho_0 = 0.2$, $t_{avg} = 500$ ms, 100 ms. Average file size = 500 KBytes, RTT = 50 ms, wireless loss probability = 0.01.




Figure 5: Fairness index for different RTT values. LBM parameters: $\alpha = 1$, $\rho_0 = 0.15, 0.2$, $t_{avg} = 500$ ms. Average file size = 500 KBytes, wireless loss probability = 0.01.






Figure 6: Fairness index and throughput for different wireless loss probabilities. LBM parameters: $\alpha = 1$, $\rho_0 = 0.15, 0.2$, $t_{avg} = 500$ ms. Average file size = 500 KBytes, RTT = 20 ms, 20 FTP flows.



(a) Fairness index



(b) Throughput

Figure 7: Fairness and throughput for different number of FTP flows. LBM parameters: $\alpha = 1$, $\rho_0 = 0.15, 0.2$, $t_{avg} = 500$ ms. Average file size = 50 KBytes, RTT = 20 ms, wireless loss probability = 0.01.

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.
3	23.5	8.1	13.1	3.4
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

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 this 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
3	13.1	3.4	2.2	0.20
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 Related work

In this section we present a brief overview of representative work on improving the performance of TCP over wireless links, and more specifically on the application of ECN in wireless networks, identifying how it differs from the approach presented in this paper. Approaches for improving the performance of TCP over wireless links fall into the following three categories [2]: end-to-end approaches, link-layer approaches, which include TCP-unaware and TCP-aware approaches, and split connection approaches.

End-to-end approaches include proposals that can distinguish between losses due to packet corruption, in which case an Explicit Loss Notification signal can be sent, and losses due to congestion. ECN also falls within this category, and has been proposed for wireless networks in [8, 9, 5]. However, ECN alone without any help from the link layer cannot adequately address the issue of corrupted packets, since senders will still decrease their congestion window in response to packets lost due to corruption. Moreover, by reducing TCP's window only when ECN signals are received poses danger in the case of high congestion, in which case ECN packets will also be lost [8, 9].

Our proposal goes one step further from addressing the issue of congestion and non-congestion related losses, and considers using ECN to convey congestion information of both the wired and wireless links. Furthermore, due to the resource sharing model of IEEE 802.11 WLANs, we propose an approach for marking packets based on the aggregate wireless link utilization and a procedure for adapting the marking algorithm to varying traffic and load conditions. Other approaches consider using a RED-like mechanism [9] or some other level of congestion (a congestion price) [1]; however, as we argue in this paper, a shared

buffer does not exist in WLANs hence such marking algorithms do not reflect the underlying resource sharing model.

The work in [4] also consider a drop probability scheme (which the authors note can also be adjusted to a mark probability scheme), where the drop probability is a linear increasing function of the number of MAC layer retransmissions, when these 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 of congestion control schemes using a utility-based modelling approach. The authors propose a marking scheme which is a concave function of the traffic arriving at a link, when this rate is larger than some minimum capacity value. Moreover, this minimum capacity parameter is adjusted independently at each link, such that the incoming traffic rate is equal to some percentage of the total capacity. Our approach differs in that the marking probability is a linear function of the aggregate traffic flowing in both directions of the wireless link.

Link-layer approaches include automatic repeat request (ARQ) schemes that retransmit corrupted packets only over the wireless link (i.e. between the wireless host and the access point) and forward error correction schemes. Different link layer schemes might be appropriate for different transport layer protocols (TCP and UDP) and application types [16]. Interestingly [2], the performance of TCP with link-layer retransmissions can potentially be worse than in the absence of link-layer retransmissions, when the timeout of TCP and link-layer retransmissions is close, or when duplicate acknowledgements are transmitted by TCP receivers, even though the link layer eventually retransmits the lost packet; the latter is due to out-of-order delivery of packets over the wireless link. To avoid such negative interactions, TCP-aware approaches have been proposed, such as the SNOOP protocol, which combines the link-layer retransmission of corrupted packets and the suppression of duplicate acknowledgements.

In this paper we assumed that link-layer retransmission, which are supported by the IEEE 802.11 MAC layer suppress packet loss due to corruption from the TCP layer. However, our proposal can be combined with more advanced link-layer proposals that support retransmission of corrupted packets and in-order packet delivery.

5 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 (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. Experimental results demonstrate that our approach can achieve higher fairness compared to drop tail queueing, while achieving the same throughput. Moreover, it can achieve a smaller and controlled packet delay over the wireless link. Further experimental work is investigating the transient behaviour of the proposed algorithm, in a dynamic environment where flows arrive and leave; the behaviour in this case will depend both on the slope parameter of the LBM algorithm, and the minimum utilization step size of the LBM adaptation procedure.

Ongoing work seeks to investigate 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. Finally, a related research area we are currently investigating is service differentiation over IEEE 802.11 wireless LANs.

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] H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz. A comparison of mechanisms for improving TCP performance over wireless links. *IEEE/ACM Trans. on Networking*, 5(6):756–769, December 1997.

- [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, S. Lu, L. Zhang, and M. Gerla. The impact of multihop wireless channel on TCP throughput and loss. In *Proc. of IEEE INFOCOM'03*, 2003.
- [5] H. Inamura, G. Montenegro, R. Ludwig, A. Gurtov, and F. Khafizov. TCP over 2.5G and 3G wireless networks. RFC 3481, February 2003.
- [6] F. P. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [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. Magret, 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. In *Proc. INET'00*, 2000.
- [10] K. Pentikousis and H. Badr. An evaluation of TCP with Explicit Congestion Notification. To appear in *Annals of Telecommunications.*, 2003.
- [11] K. K. Ramakrishnan, S. Floyd, and D. Black. The Addition of Explicit Congestion Notification (ECN) to IP. RFC 3168, September 2001.
- [12] J. Saltzer, D. Reed, and D. Clark. End-to-end arguments in system design. *ACM Trans. Comp. Sys.*, 2(4):277–288, November 1984.
- [13] S. Shakkottai, T. S. Rappaport, and P. C. Karlsson. Cross-layer design for wireless networks. *IEEE Commun. Mag.*, pages 74–80, October 2003.
- [14] V. A. Siris. Resource control for elastic traffic in CDMA networks. In *Proc. of ACM MOBICOM'02*, 2002.

- [15] V. A. Siris, C. Courcoubetis, and G. Margetis. Service differentiation and performance of weighted window-based congestion control and packet marking algorithms in ECN networks. *Computer Communications Journal*, (26):314–326, 2003.
- [16] G. Xylomenos and G. C. Polyzos. Link layer support for quality of service on wireless Internet links. *IEEE Personal Commun. Mag.*, 6(5):52–60, 1999.