TCP Fairness Issues in IEEE 802.11 Based Access Networks

Andrea Detti[1], Emiliano Graziosi[2], Valerio Minichiello[3]
Stefano Salsano[3], and Vincenzo Sangregorio[2],

[1] CNIT
andrea.detti@uniroma2.it
[2] Radiolabs
{graziosi, sangregorio}@radiolabs.it
[3] DIE, University of Rome "Tor Vergata"
stefano.salsano@uniroma2.it

Abstract. This paper deals with application level fairness issues in IEEE 802.11 Wireless LANs. The current implementations of 802.11 use the so called Distributed Coordination Function (DCF) that gives the same medium access priority to all the stations: realizing fairness at the MAC level. Nevertheless, fairness at MAC level may not correspond to fairness at application level, especially in presence of TCP controlled applications operating in an infrastructured network topology, such as an hot-spot. The goal of this paper is twofold: firstly, we investigate the fairness issues experienced by TCP based applications in an infrastructured WLAN; then we propose a rate-control mechanism, which seems appealing for its implementation practicability.

1 Introduction

In the last few years the IEEE 802.11 – Wireless LAN standard is experiencing a tremendous growth thanks to cheaper and easier-to-install components. Unlicensed spectrum and the interoperability granted by adherence to standards and to certifications (such as Wi-Fi) also contributed to this growth.

The so-called "hot-spot" coverage seems to be the short term killer-scenario, which provides access to nomadic users in both public (hotels, airport, railway station, etc.) and private (offices, conferences, etc.) environments. This type of coverage is based on the *Access Point* devices that provide the mobile stations with access to the wired network. The scenario is often called "infrastructured" WLAN to distinguish from the "ad-hoc" WLANs, where the mobile stations talk each other without using the Access Points.

Within this scenario it is important to guarantee a fair access to the WLAN bandwidth to the applications that are running on the mobile stations.[1] This type of fair-

[1] In the 802.11 terminology, an "infrastructured" WLAN is a BSS (Basic Service Set) or an ESS (Extended Service Set) if multiple Access Points are cooperating to provide access, while the "ad hoc" WLAN is an ISS (Independent Service Set).

ness, that we will name “application level fairness”, is a critical issue in several 802.11 scenarios. In order to understand why, let us consider the fairness mechanisms in the 802.11 standard.

In the basic 802.11 standard the fairness is assured at MAC level by means of two alternative techniques: the Distributed Coordination Function (DCF), and the Point Coordination Function (PCF). The former is completely distributed and aims to give the same access priority to all the wireless devices (i.e., mobile stations and Access Point). As we will discuss hereafter, the DCF is not able to provide application level fairness as it just provides MAC level fairness. The PCF grants more control to the Access Point, which can decide “who and when” can get the radio resource. In the PCF, the Access Point performs the scheduling of transmission resources by means of a “polling” technique which gives the permissions to transmit to the mobile stations. An extension of the basic 802.11, namely the draft standard 802.11e, provides further mechanisms to control the allocation of transmission resources.

Using their scheduling mechanisms, both the PCF and the 802.11e can be used to reach the application level fairness. Unfortunately, the current status is that the large majority of the existing WLAN cards and devices does not support the PCF functionality, nor the 802.11e; hence, we investigate on how it is possible to guarantee the application level fairness for DCF access. Moreover, we focus our attention only on those techniques that may be implemented within the Access Point and that do not require any enhancement (protocol, software…) on the mobile station.

The paper is organized as follow: in section 2 we describe the application fairness problem in an infrastructured 802.11 WLAN with DCF access control; in section 3 we deeply analyze this problem and provide a classification of two solving approaches; in section 4 we explain our solution and compare its effectiveness with respect to another recently proposed technique; finally, in section 5 we outline conclusions and future works.

2 Problem Statement

The problem of WLAN fairness at MAC level has been analyzed in several papers, for example see [1], [2]. The unfairness at higher levels is discussed in [3], [4]. In [3] the focus is on unfairness between the “uplink” and “downlink” TCP flows in a single WLAN. [4] discusses unfairness for both TCP and UDP flows in more complex topologies (i.e. with multiple WLANs).

In this paper we address an “infrastructured” scenario similar to the one considered in [3], where a wireless LAN provides access to a wired network for a set of mobile stations. The mobile stations exchange data with “fixed” hosts in the wired network using TCP flows. In particular there will be “upstream” and “downstream” applications running on the mobile stations. With the term downstream application we mean an application for which the main amount of data flow is from the wired network towards the mobile station (e.g., file download from a web server, video streaming, receiving e-mail); on the contrary, upstream application are those ones for which the main amount of data is directed from the mobile station towards the wired network (e.g., e-mail posting, peer-to-peer file transmission, etc.).

The most relevant application level unfairness is that TCP flows in "downlink" direction are penalized with respect to TCP flows in the uplink direction originated from multiple stations. In fact all the TCP flows in the downlink direction must be transmitted through the Access Point wireless interface to the stations in the WLAN. In this case the transmission capacity from the AP to all the receiving stations becomes the bottleneck. The problem is that the access to WLAN bandwidth is equally distributed between all the transmitting entities, stations and Access Point: the Access Point has no privileged access to WLAN bandwidth in the basic 802.11 configuration (DCF). There is fairness at the MAC level, but this does not correspond to fairness at the service level: the data transmitted by the AP belong to several connections, while the data transmitted by a station may belong to a single connection.

Another critical "application level unfairness" concerns the unfairness among flows in the same direction (uplink or downlink). When the MAC level is operating in congestion situation, the TCP is not able to provide a fair share among competing flows: some flows may grab most of the resources, while other flows could not be able to transmit for a long time.

Note that the uplink/downlink unfairness is the main issue analyzed in [3], which provides some analytical models to describe some particular cases and proposes a solution to reach a fair share of bandwidth (see section 4 hereafter). To the best of our knowledge, the issue of unfairness between flows in the same direction is a novel contribution.

In order to better understand the importance of the unfairness issue, let us analyze some simulation results. The reference scenario is reported in Fig. 1. It consists of $2 \times N$ mobile stations and an Access Point connected to a wired network. The Access Point has an uplink and a downlink buffer to store the packets received from the radio interface and from the wired interface respectively. On each of the first *N* stations a single "downstream" application is running; on each of the others *N* stations, a single "upstream" application is running as well. We assume that the applications has always something to transmit, i.e. we are considering a set of "infinite file transfers". We also assume that the wired network connecting the Access Point to the fixed hosts has a large capacity, so that the bottleneck is in the WLAN access (we considered 802.11b at 11 Mb/s). In particular in the simulation we assumed a single fixed hosts connected with a Fast Ethernet (100 Mb/s full duplex) to the Access Point. All the TCP receivers and sources in the wired network are placed on this single host.

Using the NS-2 simulator (version 2.1b9a) [5], we measured the throughput of the different TCP connections. We used simulation runs of 800 simulated seconds and we evaluated the throughput over the last 200 seconds for each connection.

The first important merit figure for the application level fairness is the ratio between the average throughput for uplink connections (denoted by R_u) and for downlink connections (denoted by R_d). If the ratio R_u/R_d is equal to 1, we have a perfect uplink/downlink fairness, the greater the ratio, the greater the unfairness. To analyze the fairness in the same direction, we consider the ratio between the standard deviation and the mean of the throughput values (considering one direction at a time).

If the ratio (σ_u/R_u for uplink and σ_d/R_d for downlink) is 0 we have perfect fairness, the greater the ratio, the greater the unfairness.

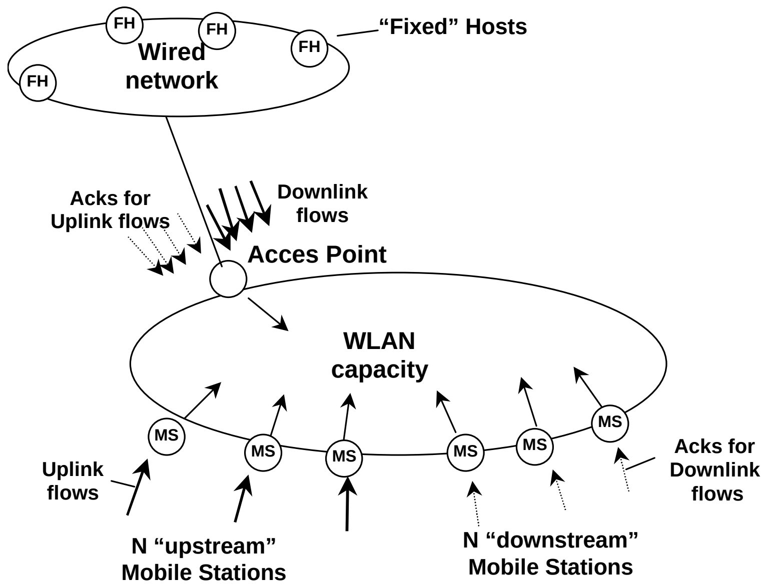

Fig. 1. Reference scenario

Fig. 2 reports the simulation results for 30 TCP connections (i.e. N=15) that exchange data on the radio interfaces when only DCF is employed; 15 connections are in upstream and 15 in downstream. We note that the most part of flows does not start at all both in up and downstream direction. The average throughput of downstream TCP flows is very low (R_d=0.088 kb/s) with a relatively high standard deviation (σ_d=0.2 kb/s). On the other side, some upstream flows reach a good throughput value (e.g., 700 kb/s), and their average throughput is significantly greater (R_d=360 kb/s) than the downstream one; nevertheless, also in upstream a large standard deviation (σ_u= 350 kb/s) occurs. As we can see, the downstream flows are highly penalized with respect to the upstream ones and, moreover, both in up and downstream there is a great unfairness also among the flows in the same direction. In particular, the ratio R_u/R_d is as high as 4111, σ_u/R_u=0.97 and σ_d/R_d =2.11. Additional results for different scenarios are reported in [6].

It is important to remark that the above discussed unfairness issues take only place when the system bottleneck is in the Wireless LAN radio interface. To understand why, consider for example a set of remote fixed hosts located somewhere in the Internet. If the connections from the Access Point to those hosts are limited in bandwidth, the TCP connections will not be able to increase their throughput and to grab all the WLAN capacity. Therefore the WLAN will not introduce any unfairness in this case.

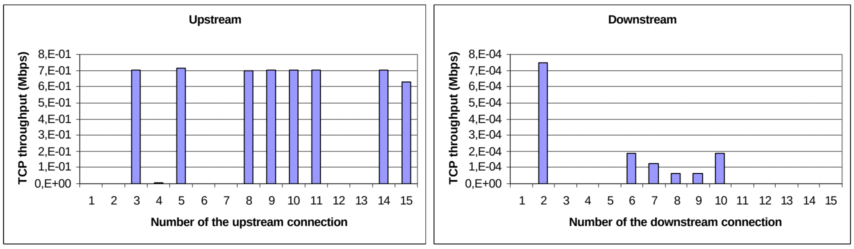

Fig. 2. TCP throughput for 30 TCP connections (15 upstream and 15 downstream); IP packet size is 1500 bytes; uplink and downlink buffer size is 100 packets, which is a typical value for commercial equipment

3 Problem Analysis

The main unfairness reason is that the access to radio interface is equally distributed between all the transmitting entities, i.e. stations and Access Point. In other words, the Access Point has no privileged access to bandwidth while it has to manage more communication processes with respect to the single station. Hence, the MAC layer downlink buffer of the Access Point is more stressed with respect to the ones in mobile stations and this results in a high packets loss probability.

Let us consider the uplink/downlink unfairness. For the downstream flows, a packet loss within the downlink buffer means a TCP segment loss resulting in a lowering of the TCP throughput due to the congestion control mechanism. For the upstream flows a packet loss within the downlink buffer means the loss of a TCP ACK. Due to the cumulative nature of TCP ACKs, these loss events do not afflict the TCP throughput of a flow whose Congestion Window is large enough.

Let us consider the unfairness between flows in the uplink direction. The high loss rate can prevent some flows from starting, because of the TCP timeout mechanisms. Therefore there will be some flows that will reach a high throughput and that are “insensitive” to the loss of TCP ACKs and some flows that are not able to transmit at all.

So, the losses on the Access Point downlink buffer appears to be the reason of unfairness. To prove this let us assume to have a downlink buffer of size *B* able to store all the possible segments coming from the downstream connections and all the ACKs coming from the upstream connection, in formula :

$$B \geq 2 \cdot N \cdot CW \qquad (1)$$

where CW is the TCP maximum congestion window size. In these conditions, due to the lossless environment, the congestion window of all TCP flows reaches its maximum value CW. Therefore, once all the TCP segments are in fly, the communication goes on in a “Stop&Wait” fashion: a generic TCP source can transmit only one segment and then must wait for the corresponding TCP ACK. The downlink TCP packets and the TCP ACKs for the uplink flows get stored in the downlink buffer in the Access Point. When the Access Point sends a downlink TCP packet, the corresponding “downstream” mobile station is enabled to send the TCP ACK; when the Access Point sends a TCP ACK for an uplink flow, the corresponding “upstream” mobile station is enabled to send a new TCP segment. Hence, the Access Point gets half of the overall bandwidth for the downlink transmission, while the stations equally share the remaining bandwidth for their uplink transmission. It is easy to conclude that all the TCP flows, in the steady state, get the same amount of bandwidth. This conclusion is confirmed by the simulation results reported in Fig. 3.

Let us notice that the TCP congestion control mechanisms, in the lossless case, leads the DCF MAC to operate in a polling fashion (as it would be done by the PCF), so assuring the application level fairness.

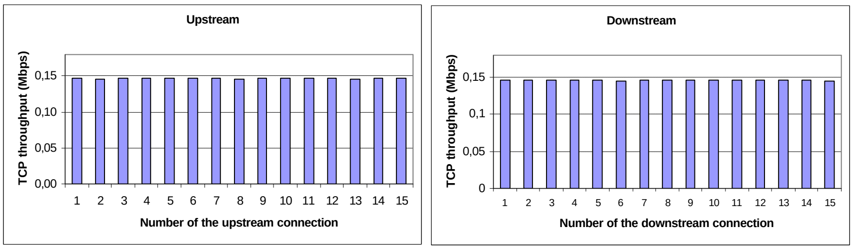

Fig. 3. TCP throughput for 30 TCP connections (15 upstream and 15 downstream); IP packet size is 1500 bytes; uplink and downlink buffers with size $B \geq 2 \cdot N \cdot CW$

4 Proposed Solutions

Obviously, the solution of having a buffer size like in eq. (1) is only theoretical. In order to achieve application level fairness we envisage the use of a rate control mechanism. This could be implemented within the Access Point or even in a router of the access network. As a requirement, we want that this mechanism is able to operate independently of the mobile stations, relying on the standardized DCF of 802.11 MAC.

In this context two approaches may be distinguished: lossless rate control and lossy rate control. The lossless approach aims to control the upstream and downstream rates so that no loss occurs in the downlink buffer. With no loss, the application level fairness can be achieved as discussed in section 3. On the other hand, the lossy approach aims to reach the fairness by controlling the flow rates but without caring about the packet loss.

In the following we discuss these two approaches and compare them. The technical solution for the lossless rate control is the one proposed in [3]; while for the lossy mechanism we present our novel solution.

4.1 Lossless Rate Control Solution

In [3] the authors realize a lossless rate control constraining the value of TCP receiver advertised windows (i.e., CW), contained within the TCP ACKs, to be at most $B/(2 \times N)$. Therefore, the maximum amount of TCP segments or ACKs in fly for each TCP connection is equal to this value. As the downlink buffer is loaded by N downstream connections carrying segments and N upstream connections carrying ACKs, this buffer never fills up and the application level fairness is achieved.

From an implementation point of view this solution is quite complex. The device that implements this solution (for example the Access Point itself) should: i) operate on a packet-by-packet basis and examine the TCP headers in order to modify the receiver advertised window; ii) know the number of crossing TCP flows; iii) know the size of the Access Point buffer.

4.2 Lossy Rate Control Solution

Our solution equally divides the TCP/IP overall radio capacity R between the Access Point and the terminal stations by means of an IP level rate-limiter, placed at the output of the uplink buffer (Fig. 4).

With TCP/IP overall radio capacity R, we mean the maximum achievable throughput at the IP level when the radio interface is loaded by TCP connections and no loss occurs. This capacity accounts both for TCP segments and TCP ACKs. We have evaluated this capacity by means of simulation in the "infinite buffer size" scenario described in section 2. The resulting value for R is about equal to 4.6 Mbps [(2)]. We limit to $R/2$ only the rate of uplink transmissions, since the rate of the downlink transmission will be the remaining $R/2$ due to the TCP congestion control mechanisms.

On the basis of the classification defined in the previous paragraph, this mechanism belongs to the lossy group because it does not assure the lossless behaviour neither in the downlink nor in the uplink buffers. Instead, we drop packets on the uplink buffer with the goal to indirectly control the rate of uplink flows via the TCP congestion control.

From an implementation point of view, this solution is much simpler with respect to the lossless solution. It only requires the knowledge of the TCP/IP capacity R and does not modify any TCP/IP header, so maintaining the end-to-end behavior of the transport protocol. The knowledge of the number of TCP flows, that is used in the lossless rate control is not strictly needed now. Actually, this knowledge can be useful if the number of active uplink flows is different from the number of active downstream flows. The rate of the uplink limiter could be dynamically modified in order to follow the ratio between the number of uplink and downlink flows. Anyway this is left for further study.

---

[2] We notice that this value is quite independent from the number of flows $2 \cdot N$. This behaviour has been verified with simulation runs not reported in the paper.

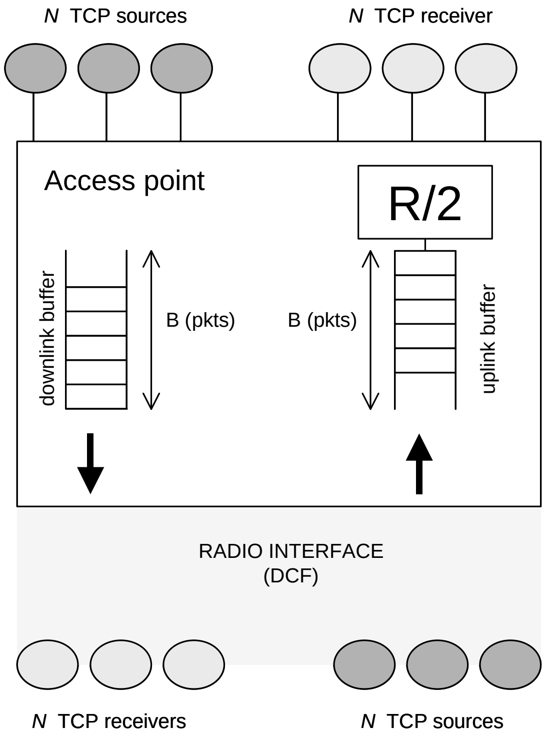

Fig. 4. "Lossy" rate control solution based on a rate-limiter

4.3 Numerical Result

In this section we report the comparison between the different solution. The simulation results are obtained in the reference scenario with N=15.

Fig. 5, Fig. 6 report the TCP throughput experienced on each downstream and upstream connections for both the lossy solution and the lossless one with B=100 packets; in addition Table 1 summarize the average results.

We notice that the lossless rate control solution presents homologous performance as the case of infinite buffer. The lossy technique is able to achieve uplink/downlink fairness (R_u/R_d=1.007), while it is slightly unfair for flows in the same direction. Moreover, from Table 1 we notice for the lossy solution a slight decrease in the overall achieved throughput.

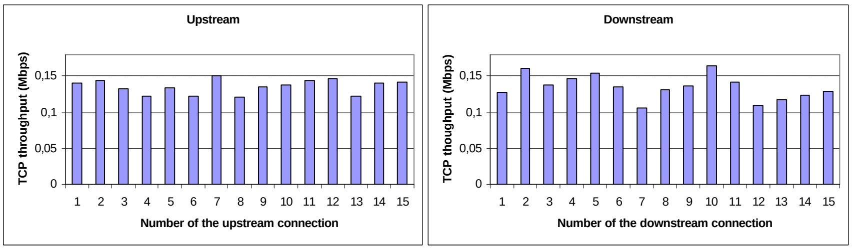

Fig. 5. Proposed Lossy rate control. TCP throughput for 30 TCP connections (15 upstream and 15 downstream); IP packet size is 1500 bytes; uplink and downlink buffer size is 100 packets; uplink rate limiter with *R*=4.7 Mbps

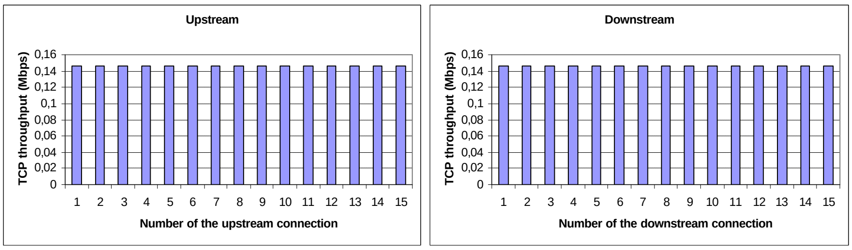

Fig. 6. Lossless rate control [3]. TCP throughput for 30 TCP connections (15 upstream and 15 downstream) IP packet size is 1500 bytes; uplink and downlink buffer size is 100 packets.

Table 1. Summary of simulation results.

Solution Type	**Average upstream throughput R_u (kb/s)**	**Uplink unfairness s_u/R_u**	**Average downstream throughput R_d (kb/s)**	**Downlink unfairness s_d/R_d**	**Uplink/ downlink unfairness R_u/R_d**
None (sec. 2)	370	0.973	0.09	2.111	4111
Infinite Buffer (sec. 3)	146.6	0.005	146.7	0.005	0.999
Lossless rate control (sec. 4.1)	146.6	0.000	146.6	0.001	1.000
Lossy rate control (sec.4.2)	136	0.071	135	0.128	1.007

This worsening is due to the lossy behavior of our solution, and it is the price to be paid for the much simpler implementation. So, it is expected that with the increase in Access Point buffer size this inefficiency will vanish. This is confirmed by the plot reported in Fig. 7, which reports the average uplink and downlink TCP throughputs versus the buffer size, for the two rate control techniques and for the case with no rate control. The uplink throughput for the case with no rate control is the highest line in the order of 0.35 Mb/s. The downlink throughput for this case is not distinguishable from the x axis, due to the huge uplink/downlink unfairness. The uplink and downlink throughput for the lossless rate control are coincident at 0.146 Mb/s. The proposed lossy technique does not perform equally well when the buffer size takes small values. With the increases of the buffer size the performance become very close to the ideal performance of the lossless rate control.

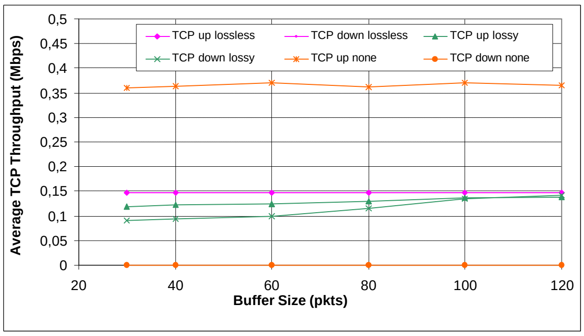

Fig. 7. Average TCP throughput varying the Access Point buffer size in the case of lossless rate control, lossy rate control, and no rate control mechanism

5 Conclusion and Future Work

In this paper we considered the issue of application level fairness in a wireless access network based on IEEE 802.11b standard operating in DCF mode at 11 Mbps. We considered the interaction between the TCP congestion control mechanisms and the WLAN. We first examined an ideal solution (no loss on the Access Point buffer) which guarantees fairness between TCP flows. Then we described and compared, by means of simulations, two different fairness enforcing techniques, named lossless rate control and lossy rate control.

The lossless mechanism, already proposed in the literature [3], gives slightly better results in terms of fairness and throughput at the expense of a higher complexity introduced in the Access Point. In fact, the number of active TCP flows must be know at any time and the TCP header of each TCP ACK must be read and modified in order to limit the advertised receiver window. On the other hand, our lossy solution works without changing any header and does not require any additional configuration knowledge.

It is worth to notice that both techniques are useful in situations for which the radio interface is congested and represents the system bottleneck.

Future work will concern several open issues:

- an analytical estimation of the TCP/IP overall radio capacity *R*
- the impact of the rate and of the buffer size assigned to the uplink rate limiter on the fairness performance
- the transitory regime of TCP connections
- the behavior of the considered mechanisms in absence of radio congestion

- the analysis of more complex scenarios: asymmetric number of uplink and downlink connections, dynamic scenarios where the number of connections varies with time, introduction of some TCP connections towards remote hosts (these connection have a limited bandwidth and a higher Round Trip Time)

References

[1] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan, "Achieving MAC layer Fairness in Wireless Packet Networks", ACM Mobicom 2000, Boston, USA, Aug. 2000

[2] N.H. Vaidya, P. Bahl, and S. Gupta, "Distributed Fair Scheduling in a Wireless LAN", MobiCom 2000, Boston, USA, Aug. 2000

[3] S. Pilosof, R. Ramjee, Y. Shavitt, P. Sinha: Understanding TCP fairness over Wireless LAN, INFOCOM 2003

[4] G. Bianchi, N. Blefari-Melazzi, E. Graziosi, S. Salsano, V. Sangregorio, "Internet access on fast trains: 802.11-based on-board wireless distribution network alternatives", IST Mobile & Wireless Communications Summit 2003, 15-18 June 2002, Aveiro, Portugal

[5] The Network Simulator - ns-2, http://www.isi.edu/nsnam/ns/

[6] Project IST 2001-39097 FIFTH, Deliverable D8, "Mobility Management and Routing Specification Document, I1", section 4.2 "Train Network Architecture and Fairness/QoS Issue"