

Detailed Analysis of TCP Fairness Issues in IEEE 802.11 Networks and their Solution using a Rate Control Approach

(Technical Report, University of Roma Tor Vergata)

Nicola Blefari-Melazzi¹, Andrea Detti¹, Ibrahim Habib², Alessandro Ordine¹, Stefano Salsano¹

¹ DIE, University of Rome "Tor Vergata"
{blefari, andrea.detti, alessandro.ordine, stefano.salsano}@uniroma2.it

² Department of Electrical Engineering, City University of New York
habib@ccny.cuny.edu

Abstract

In this paper, we study the problem of maintaining fairness for TCP connections in wireless local area networks (WLANs) based upon the IEEE 802.11 standard. Current implementations of 802.11 use the so-called Distributed Coordination Function (DCF) which provides similar medium access priority to all stations. Although this mode of operation ensures fair access to the medium at the MAC level, it does not provide any provisions for ensuring fairness among the TCP connections. TCP unfairness may result in significant degradation of performance leading to users perceiving unsatisfactory quality of service. We propose and analyze two solutions that are capable of enabling TCP fairness with minimal additional complexity. The proposed solutions are based on utilizing a rate-control mechanism in two modes: static or adaptive. They do not require modifying existing standards at the MAC or network layers. Hence, they are fully compatible with existing devices. Our performance analysis results prove the efficaciousness of our proposed solutions in achieving TCP fairness compared to existing approaches. We have, also, implemented the proposed solutions in an ad-hoc experimental test-bed, and performed measurements to demonstrate the validity of our approach and results.

This technical report is an extended version of a paper submitted for publication.

1 INTRODUCTION

1.1 Motivations and Problem Statement

Wireless LANs based on the IEEE 802.11 standard are shared media enabling connectivity in the so-called “hot-spots” (airports, hotel lounges, etc.), university campuses, intranet access in enterprises, as well as “in-home” for home internet access among others.

In all of the above-mentioned scenarios, WLAN coverage is based on “*Access Points*”, devices that provide the mobile stations with access to the wired network. These scenarios are often called “infra-structured” WLANs to distinguish them from the “ad-hoc” WLANs, where mobile stations talk to each other without using Access Points¹.

In the above scenarios, it is crucial to maintain fairness among the TCP connections competing for access to the shared media of the WLAN. By fairness among multiple TCP connections, we mean that any TCP engine would be capable of starting a connection with negligible delay, as well as achieving and maintaining a reasonable throughput. The latter, of course, depends on other competing TCP connections. Viewed this way, TCP fairness is, then, a mandatory prerequisite for enabling a satisfactory quality service for upper layer applications. However, we also have to specify that we are not requesting a “perfect” fairness, i.e., a perfectly balanced sharing of resources among all TCP connections (which can be seen as a “second order” objective). Rather, our main aim is to avoid the scenario of “*critical unfairness*” that is characterized by complete starvation of some TCP connections or, even, the inability of some TCP connections to start altogether.

This so called *critical unfairness* can arise in two cases: 1) interaction between upstream and downstream TCP connections, or 2) interaction between a set of upstream TCP connections. TCP connections are labeled as “downstream” or “upstream” (see Figure 1), depending upon the

¹ In the 802.11 terminology, an “infrastructured” WLAN is a BSS -Basic Service Set- (or an ESS -Extended Service Set- when multiple Access Points cooperate to provide access), while the “ad hoc” WLAN is an ISS (Independent Service Set).

direction of traffic flow. Downstream is used to describe traffic flowing from the wired network towards the mobile station (e.g., file downloads from a web server, video streaming, incoming e-mails), whereas Upstream is used to refer to traffic flowing from mobile stations to the wired network (e.g., e-mail posting, peer-to-peer file transmission, etc.).

In the following we introduce the two critical unfairness cases that will be thoroughly analyzed in the next section. In the first case, downstream TCP connections are penalized with respect to upstream ones. This is explained as follows: packets belonging to multiple downstream TCP connections are buffered inside the Access Point wireless interface. Note that the Access Point does not enjoy a privileged access to WLAN capacity, with respect to user terminals. Hence, a single station transmitting upstream packets will get the same priority as that of the Access Point which needs to transmit downstream packets heading towards many stations. Thus, downstream TCP connections suffer because of the arising congestion and corresponding packet losses happening in the download buffer at the Access Point. These losses in conjunction with TCP congestion control mechanism cause the starvation of downstream connections. This is defined as “critically” unfair.

The second case arises from the interaction of multiple TCP connections in the upstream direction. In this case, the Access Point wireless interface has to transmit TCP ACK packets traveling downstream towards stations in the WLAN. Also, in this case, we have a bottleneck because the Access Point can not access the medium with a priority higher than other stations. Hence, the Access Point buffer will be congested leading to severe loss of TCP ACK packets. Due to the cumulative nature of TCP ACKs, few connections will be able to “survive” and open their window, while the majority of connections will get starved. Note that this situation is not specific of our scenario; it can happen in whatever environment characterized by heavy losses of ACK packets. This case is also another example of “critical unfairness” which will be explained in more details in the next section.

It is worth-mentioning that the 802.11 standard also includes a different access control mechanism, called Point Coordination Function (PCF). An extension of the basic 802.11, namely the draft standard 802.11e, provides further mechanisms to control the allocation of the WLAN resources. Both the PCF and the 802.11e could be used to improve the fairness perceived at the application level. Unfortunately, the current status is that the large majority of existing WLAN cards and devices support neither the PCF functionality nor the 802.11e. Considering the lack of deployment of PCF and 802.11e, we focus on strategies to achieve TCP fairness by using the widely deployed DCF mechanism. Moreover, we focus our attention only on techniques that can be implemented within the Access Point (or in a nearby router), *without requiring changes in the 802.11 standard, nor any enhancement to mobile stations.*

1.2 Related Work and Basic Assumptions

The problem of WLAN fairness at MAC level has been analyzed in several papers, (e.g., [1], [2]); however, as we said earlier, MAC level fairness is not enough. Unfairness among TCP connections within a single WLAN was analyzed in [3]; in which the authors propose an elegant solution that addresses not only the so-called critical unfairness (i.e., it avoids connection starvations), but also a finer “per connection” fairness. However, the approach proposed in [3] has some drawbacks in terms of implementation complexity, dependence on connection Round Trip Times, and need to parse the TCP header (which cannot be accessed in case IPsec is used) (see Section 3). Unfairness among TCP and UDP flows in more complex topologies (i.e., with multiple WLANs) has been preliminary discussed in [4].

In this paper, we propose solutions aiming at avoiding critical unfairness (i.e., starvation) and at enforcing a fair sharing of radio bandwidth between the Access Point and the mobile stations. Our solution works on the aggregate TCP flows crossing the Access Point. Consequently, we do not provision a perfect “per connection” fairness. However, our solution is simple to implement, robust against variable Round Trip Times, and does not require parsing of TCP headers. Hence, it is also compatible with IPsec.

Our approach is based upon utilizing a rate-limiter, implemented via a Token Bucket Filter (TBF) [5]. It is characterized by two parameters: 1) the rate of generating tokens into the bucket (R), and 2) the capacity of the bucket (bucket depth B). The rate-limiter operates on the overall aggregate of uplink packets. The TBF generates tokens at rate R and puts them in the bucket. The rate limiter forwards arriving uplink packets only if there are tokens available in the bucket, otherwise uplink packets are dropped. Each arriving packet consumes a token. The TBF forces packets' loss when the uplink aggregate rate is higher than the TBF rate R and no token is left in the bucket. The TCP congestion control mechanisms, that are automatically enabled when losses are detected, reduce the transmission windows and consequently the number of transmitted packets. Thus, by setting the TBF rate R and bucket depth B , it is possible to suitably control the overall uplink rate.

We also assume that the parameter R can be either statically configured or it can be dynamically and adaptively varied as a function of an estimation of the attainable downstream throughput (this choice will be better motivated later on).

We will refer to these two alternatives as static rate control and dynamic rate control, respectively. We will show that these solutions are indeed very simple to implement and that the adaptive rate control is very effective in avoiding the starvation of TCP connections and resource wasting. In addition, our approach can provide the operator of the WLAN with a tool to controlling the sharing of WLAN resources between upstream and downstream applications.

As for the organization of the paper, in Section 2 we discuss the “basic” system model without rate control mechanisms and introduce, evaluate and comment some performance measures. In Section 3, we present our solutions based on rate control, to face the so called critical unfairness. The static approach is detailed in Section 4, together with the related performance evaluation; similarly, we present the adaptive rate control in Section 5, together with the related performance evaluation. In Section 6 we discuss conclusions and future work.

The document contains also some appendices concerning the following issues: i) the assessment of the impact of different TCP versions on our solution (in Appendix I); ii) the performance evaluation of our mechanisms for different values of the Round Trip Time, showing that our solution works independently from this parameter (in Appendix II); iii) a description of the EWMA algorithm used to estimate the rate of aggregate upstream and downstream connections (in Appendix III); iv) a description of our test-bed, which shows how our solution perfectly works in a real environment (in Appendix IV); v) the performance evaluation of our system when loaded with short-lived TCP sources and with a mix of short-lived and greedy sources, two traffic scenarios that extend the main one considered throughout this paper (in Appendix V); vi) the choice of the parameters of the token bucket filter (in Appendix VI).

2 SYSTEM MODEL AND PERFORMANCE MEASURES

We started our work by analyzing the performance of upstream TCP connections resulting from both mobile stations and fixed hosts, by means of simulations. Throughout this paper, we will not present the 95% confidence intervals of simulation results, in order to improve the neatness of the figures. However, such intervals are always less than 5%.

The simulation model is shown in Figure 1. A number of wireless stations are connected to an Access Point and exchange information with a host in the high-speed fixed network (this host being labeled as “*wired host*”). In particular, we consider N_{dn} wireless stations downloading information from a wired host and N_{up} wireless stations uploading information to a wired host.

As shown in Figure 1, in our simulation environment the wired host can be connected to the Access Point via a Fast Ethernet (100 Mb/s full duplex) LAN link, or via a generic duplex link with capacity C and one-way propagation delay D . The former represents the case in which the Access Point and the wired host are in the same Local Area Network (“local wired host”). The latter represents the case in which the wired host is remotely located somewhere in the Internet (“remote wired host”). We can set the Round Trip Time (RTT) between the wireless stations and the remote wired host to arbitrary values by choosing a proper value of D . In the same way, we

can emulate a bottleneck in the Internet connection toward the remote wired host by properly setting the capacity C . Simulations have been carried out by using the NS-2 simulator package (version 2.1b9a) [6]. Within this environment, we suitably defined the simulation scenario and wrote the necessary additional code; the most important simulation parameters that we adopted in this work are: IP packet size: 1500 bytes. Maximum TCP congestion window: 43 packets (64 kbytes). TCP version: Reno [8]. The latter choice stems from the fact that this is the version currently installed in commercial Microsoft Windows based PCs as well as in typical Linux distribution. However, we also tested TCP NewReno [9] and SACK [10], to verify that the phenomena under study are not specifically tied to the selected TCP version (see Appendix I). As regards the traffic loading the system, we assume two different traffic source models: i) greedy sources, that is TCP sources that have always information to transmit - this model is denoted in the literature also as “infinite file transfer” model; ii) short-lived TCP sources, modeling the download or the upload of a small amount of data. The short-lived traffic scenario is described and analyzed in the Appendix; the greedy sources traffic scenario is the main one, described and analyzed through the paper.

As for the downlink buffer, we will present simulations in which we vary its size to analyze the impact of this critical parameter (e.g., 50, 100, 300 packets). When it is not otherwise specified, the downlink buffer size is 100 packets, which, according to [3], is a typical value for commercial equipments.

For each connection, we evaluated the throughput. We denote by throughput the bit rate transported by the layer below IP, comprehensive of all upper layers overheads (including IP and TCP) and of the overhead resulting from TCP ACKs flowing in the reverse direction². Also, we denote by *upstream* the direction of an asymmetric TCP connection whose greater part of data flows from a mobile station to the wired network and by *uplink* the physical direction of packets

² This is a better indication on how the WLAN resources are being used, rather than the TCP goodput, which does not take into account the IP headers and the TCP ACKs.

going from a mobile station to the wired network. Similar definitions apply to *downstream* and *downlink*. This implies, for instance, that both uplink and downlink packets flow within an upstream TCP connection.

The figures of merit that we consider are: the total upstream throughput $R_{\text{up_tot}}$ (i.e., the sum of the throughputs of upstream TCP connections), the total downstream throughput $R_{\text{dn_tot}}$ (i.e., the sum of the throughputs of downstream TCP connections), and the total throughput R_{tot} (the sum of upstream and downstream throughputs).

Unfairness between upstream and downstream TCP connections occurs when $R_{\text{dn_tot}} \ll R_{\text{up_tot}}$, assuming that both upstream and downstream TCP connections are active and “greedy”.

To analyze unfairness among flows heading in the same direction, for instance in the upstream one, we evaluate the ratio between the standard deviation (σ_{up}) and the mean value ($R_{\text{up}}=R_{\text{up_tot}}/N$) of the throughput of upstream connections. If the ratio $\sigma_{\text{up}}/R_{\text{up}}$ is zero, then we have perfect fairness; otherwise unfairness increases as this ratio increases. The same applies for the downstream case, considering the downstream ratio $\sigma_{\text{dn}}/R_{\text{dn}}$ (with $R_{\text{dn}}=R_{\text{dn_tot}}/N$). In the following, this ratio will be called unfairness index.

2.1 Numerical results

We start our analysis with the case of “no rate-control”, that is without any special mechanism to improve the performance and we assume that wired hosts are locally connected to the Access Point (scenario “local wired host”). We also assume that $N_{\text{dn}}=N_{\text{up}}=N$ i.e., that the number of upstream connections is equal to the downstream ones. Simulation experiments lasted 600 seconds of simulated time; the throughput was measured during the last 200 seconds of each experiment, to avoid transient effects and operate in steady state conditions.

Figure 2 shows the upstream throughput, the downstream throughput and the total throughput (i.e., the sum of these two components) in the WLAN as a function of N ($N=N_{\text{dn}}=N_{\text{up}}$), when no rate control is implemented. For $N=1$, i.e., when there is only one upstream connection and one downstream connection, the overall bandwidth is fairly shared. The throughput of downstream

connections drastically decreases as N increases and is almost equal to zero for $N=4$. Thus, downstream connections do not succeed in perceiving their “right” bandwidth share, even with a moderate number of upstream connections. The total throughput slightly increases with N , as an indirect consequence of the increase in packets’ loss in the downlink buffer. When the losses are high, more TCP segments than TCP ACKs are transmitted in the WLAN. In fact, if there is no loss, there will be one TCP ACK transmitted for each TCP segment³. If a TCP segment is lost, no TCP ACK is transmitted. If a TCP ACK is lost, it means that a TCP segment has been transmitted while the corresponding TCP ACK is not transmitted. This means that the ratio between TCP ACKs and TCP segments decreases as the loss in the downlink buffer increases. Consequently, the total throughput will increase, since the shorter TCP ACKs (40 bytes) have a proportionally larger overhead as compared to TCP segments (1500 bytes).

We focus now on upstream connections. Figure 3 reports the throughput perceived by each upstream connection for $N=5$, 10 and 20 (always in the “no rate-control” case). It is evident that there is no fairness as the number of flows is greater than 5. In fact, some flows do not even succeed to starting transmission, while other flows seize all the WLAN resources. The bar charts show that for $N=5$ all flows are able to start. For $N=10$ and $N=20$ only 6 and 8 flows, respectively, actually use the WLAN resources. As anticipated, the ratio σ_u/R_u is a good gauge of unfairness and, as shown in Figure 4, it sharply increases for $N>5$.

2.2 Understanding the results

This section analyzes the results shown above. We state that the main cause of starvation, and unfairness, is the packet loss occurring in the downlink buffer. The causes of packet loss are first highlighted, then it is explained why such loss results in unfairness, and even starvation, of some connections.

³ If the “delayed ACK” mechanism is used, there will be one ACK for every two TCP segments, in the no loss case, but the effect of the loss is the same.

The packet loss in the downlink buffer may attain large values because of the DCF access mechanisms whose task is to, fairly, share the available capacity among all active entities, mobile stations, and Access Point alike. Since the Access Point does not enjoy a privileged access to WLAN capacity with respect to users' terminals and since it has to handle more traffic with respect to a single station; it is more likely that its downlink buffer becomes congested, with respect to the buffering resources of mobile stations.

We now investigate why this loss leads to TCP starvation of some connections. We start by looking at downstream connections. For such connections, a packet loss in the downlink buffer means a TCP segment loss; TCP segment losses trigger congestion control mechanisms which, in turn, cause a decrease of the TCP throughput. In addition, both at the beginning of a connection and after the occurrence of several segment losses, the TCP congestion window is small in order to prevent the use of fast retransmit mechanisms. Hence, most of the losses are recovered by means of the Retransmission Time Out (RTO) mechanism. Since the RTO doubles after each consecutive loss (and consecutive losses are likely in the above conditions), downstream connections experience long idle period, and even throughput starvation, as shown in Figure 2.

To support this hypothesis, we evaluated by simulations (see Figure 7, discussed below) the packet loss probability in the downlink buffer of the Access Point as a function of N . The main results are the following. When $N=1$, the downlink buffer at the Access Point is large enough to avoid loss; the downstream throughput is not starved. When N increases, the downlink buffer occupation increases as well, and even for small values of N the loss probability is great enough to starve all downstream connections (e.g., 20% loss for $N=3$).

Let us now turn our attention to upstream connections. In this case, a packet loss at the downlink buffer means the loss of a TCP ACK. For large values of such loss probability several consecutive ACKs of the same connection may be lost (e.g., in Figure 7, discussed below we show that the loss probability is in the order of 60 % when $N=10$). The impairments caused by

consecutive ACK losses worsen as the TCP congestion window decreases [7]. For instance, assuming ideal conditions, with ACK losses being the only cause of performance degradations, and assuming a congestion window equal to W packets, the sender will find itself in the Retransmission Time Out state, and thus reduce its throughput, only if W ACKs are lost. As a consequence, the greater W , the rarer are RTO events.

If we consider that the TCP congestion window increases when ACK segments are received, the probability of RTO events is maximized at the start of the connection. On the contrary, these events are always less likely to occur as the congestion window increases, and disappear once the window gets higher than a critical threshold (e.g., five packets). This chain of events is the cause of the behavior illustrated in Figure 3.

It is worth noting that upstream connections experience starvation for larger values of loss probabilities, as compared to downstream connections. For instance, in our scenario, the downstream starvation occurs when the loss probability is greater than 20% (and $N=3$), whereas upstream connections suffer this full service outage for loss probabilities greater than 50% (and $N=10$).

As a further corroboration of our interpretation, we present more simulation results obtained by varying the downlink buffer size. Figure 5 shows the total upstream throughput and the total downstream throughput as function of N for values of the buffer size ranging from 50 packets to a buffer size large enough to completely avoid loss phenomena, denoted by $B_{no\text{loss}}$. In our scenario, $B_{no\text{loss}}=2 \cdot N \cdot CW_max$, where CW_max is the TCP maximum congestion window size (expressed in packets of 1500 bytes, following NS-2 conventions).

Obviously, the loss probability decreases when the buffer size increases. Consequently, the number of stations representing the critical threshold beyond which such starvation occurs increases too. It is worth noting that for any buffer size, increasing the number of connections always leads to starvation conditions. For instance, for a buffer of 50 packets, starvation of

downstream connections occurs at N as small as two, whereas for a buffer of 300 packets, the critical threshold of N is seven (see Figure 5).

We observe that, with a downlink buffer of size $B_{no\text{loss}}$, all unfairness issues are automatically resolved. In fact, due to the lossless property, the congestion window of all TCP flows can reach its maximum value CW_{max} (this value being always expressed in packets of 1500 bytes). Therefore, once all TCP segments are in transit, the communication goes on into a “Stop&Wait” fashion (i.e., 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. Hence, it is easy to conclude that under these conditions all TCP connections get the same amount of bandwidth in the steady state. This is shown in Figure 6, where we plot the throughput of individual connections.

To verify our conjecture about the “Stop&Wait” behavior, we have complemented our throughput measurements (shown in Figure 5) by analyzing what happens at the MAC level under the same conditions. For instance, for $N=15$, with a downlink buffer size equal to $B_{no\text{loss}}$, we have found that the mean number of active stations at MAC level (i.e., stations that are transmitting or with backlogged traffic at the MAC layer), excluding the access point is only 0.76. Considering that the access point is always active, this means that communication is basically taking place one at a time (i.e., between the access point and one mobile station at a time). This implies that the WLAN is operating as if a polling scheme is in place and capacity is shared evenly among all stations.

On the other hand, when the buffer size is smaller than $B_{no\text{loss}}$, we have found that more than one station have packets stored in the MAC queue. These stations compete for medium capacity. There is no polling effect that can be observed and unfair access occurs. As a matter of fact, with a buffer size of 100 packets, we have found that the mean number of active stations at the MAC level is 3.16 in addition to the Access Point, thus proving our conjecture (note that with $N=15$, there will be 7 stations that are able to send their upstream connections, while all the other ones are starved).

This said, it would seem that, from the starvation problem point of view, the most sensible solution is to choose a size of the downlink buffer equal to $B_{no\text{loss}}$. This choice would also have the advantage of guaranteeing to all connections the same throughput, thus reaching a perfect fairness [3]. However, increasing the size of the downlink buffer has the disadvantage of increasing the queuing delay, with obvious consequences on the overall Round Trip Time experienced by TCP connections. For example, to support 8 upstream and 8 downstream connections without losses, we need a buffer size in the order of 600 packets. If we consider that: i) half of the buffer will contain TCP segments (1500 bytes); ii) the remaining half will be filled by TCP ACKs (40 bytes), iii) the downlink throughput is in the order of 2.5 Mb/s, then we can evaluate the queuing delay as being in the order of $300 \cdot 1500 \cdot 8 / 2.5e6 + 300 \cdot 40 \cdot 8 / 2.5e6 = 1.47$ s. Hence, TCP connections will experience a rather large Round Trip Time (RTT). In turn, increasing RTTs impair the throughput of short-lived TCP connections. This effect is not apparent in our figures since we are focusing on long-lived TCP ones. According to these considerations, the buffer size should be set considering the trade-off between maximizing throughput for long-lived TCP connection (large buffer) and minimizing RTT for short-lived TCP connection (short buffer). For the time being, we will follow our analysis by setting the downlink buffer size to a value of 100 packets.

The same settings of Figure 5 were used for Figure 7 to evaluate the downlink buffer loss rate as a function of N , and for different values of the buffer size. It is clear that increasing the buffer

size allows the handling, in a fair way, of a larger number of sources. However, as the number of sources increases, there is always a point beyond which the loss rate starts to increase, $N=6$ for $B=300$, $N=10$ for $B=500$ and $N=20$ for $B=1000$ (correspondingly, the total downlink throughput will start to decrease). Thus, the loss rate becomes greater than zero when the number of stations is greater than a threshold which increases with B and then it tends to an asymptotic value.

In the previous sub-section, we stated that the total throughput increases with N , and we anticipated that this would happen because as N increases the ratio between TCP ACK segments and overall TCP traffic decreases (see Figure 2). This phenomenon can be further explained with the ensuing considerations. For small values of the downlink buffer size, as N increases, the downstream connections tend to starve. The radio link capacity is used only by some upstream connections, and the AP downlink buffer contains mainly ACK segments. Additionally, as the downlink buffer loss probability increases, the ratio between the number of overall TCP segments and the number of ACKs exchanged over the air interface increases as well (as explained in our comments to Figure 2 in Section 2.1). The segments are significantly longer (1500 bytes) than the related ACKs (40 bytes), thus, the overhead at the MAC level decreases and the total throughput increases (from 4.5 Mbps to 5.3 Mbps for the range of N shown in Figure 2).

It is worth noting that the value of 4.5 Mbps is the maximum total throughput obtainable in any *lossless* scenario. In our scenario, TCP is more efficient for large values of the ACK loss probability. In fact, under these conditions a throughput of 5.3 Mbps is reached (see Figure 2). On the other hand, the advantage of enjoying an increased throughput has to be traded off with the above mentioned cons and in particular with the risk of heavy unfairness and connections starvation.

Before concluding the Section, we make two additional considerations:

- 1) Starvation phenomena are related to the TCP startup phase; this means that also other TCP versions will experience similar problems. This is shown to be actually true in the Appendix I for the cases of TCP SACK and TCP NewReno.
- 2) All of the above discussed unfairness phenomena happen when the overall system bottleneck is the WLAN radio interface. If the bottleneck is somewhere else in the fixed network, TCP connections will not be able to grab all the WLAN capacity and the WLAN will not introduce unfairness.

In the next Section, we propose our solution to the critical unfairness problem.

3 LIMITER BASED RATE CONTROL

As discussed in the previous Section, we could solve fairness impairments by setting the size of the downlink buffer of the Access Point to $B_{no\text{loss}}$. However, this approach has its disadvantages, some of which have been outlined above. More importantly, it is certainly not easy to influence manufacturers as to make them deploy Access Points with a minimum given buffer size, let alone the issue of all the devices already installed.

An alternative solution, proposed in [3] aims at controlling upstream and downstream rates so that no loss occurs in the downlink buffer. This is done by suitably modifying the window size advertised in TCP packets (such modification happening in the Access Point). In this case, fairness conditions are achieved since, as discussed in Section 2.2, each source operates in a “stop & wait” mode⁴.

However, we argue that, from an implementation point of view, this solution adds complexity since the device implementing this solution (for example the Access Point itself) must operate on a packet-by-packet basis and parse TCP headers, in order to modify the receiver advertised window. Moreover, it is also necessary to estimate, in real-time, the number of TCP flows

⁴ We tested the approach proposed in [3] by means of our simulator and we verified that the results in terms of throughput are equal to those shown in Figure 5, when the buffer size is equal to $B_{no\text{loss}}$.

crossing such device. This is by no means a trivial operation. The number of TCP connections can change very rapidly, as many TCP connections can be very short-lived. In addition, there can be open TCP connections that are long lived, but which have a minimal activity (for example a telnet connection). These should not be counted among the “greedy” connections competing for the WLAN access bandwidth.

Our proposal is to use a limiter-based Rate Control. This solution could be implemented within the Access Point or in a suitable router of the access network. Additionally, we require that the proposed mechanism operate transparently with actual WiFi mobile stations, based on the DCF defined in the 802.11 standard. Our approach purposely introduces packet losses that trigger the TCP congestion control mechanisms; fairness conditions are achieved by indirectly controlling the *aggregate* rate of TCP connections. The goal is to enforce a fairer sharing of WLAN capacity between the Access Point and the mobile stations. Our rate-limiter operates at the IP level, before the uplink buffer. Packets are dropped by the rate limiter with the aim of indirectly controlling the rate of uplink flows via the TCP congestion control. The rate limiter is a token bucket filter characterized by two parameters: 1) the rate of generating tokens into the bucket, R (expressed in Mb/s), and 2) the bucket size B_{bucket} (expressed in Mbit). The TBF generates tokens at rate R and puts them in the bucket. The rate limiter (and thus the AP) forwards arriving uplink packets only if there are tokens available in the bucket, otherwise uplink packets are dropped. Thus, the token bucket operates only as a dropper, i.e., it does not try to reshape non conforming packets and it does not need to queue packets. This makes its practical implementation very simple.

However, the interaction of the simple TBF described above with TCP can lead to annoying synchronization effects, as it is likely that the TBF drops packets in burst, causing several TCP connections to reduce their sending rate in an almost synchronized way. In order to avoid this effect we propose a modified version of the TBF, called Smoothed TBF (STBF). A Smoothed TBF introduces an additional loss, by randomly dropping packets, even when there would be

enough tokens in the bucket to forward packets, very much like a RED queue randomly drops packets before the queue gets full. In particular, let B_{bucket} be the bucket dimension, let H be the current size of the bucket, let $0 < Th < 1$ be the threshold level at which the STBF should start dropping packets. The STBF introduces a loss probability $P_{\text{drop}}(H)$ for an incoming packet as follows:

$$P_{\text{drop}}(H) = \begin{cases} 0 & \text{if } H > Th \cdot B_{\text{bucket}} \\ \frac{Th \cdot B_{\text{bucket}} - H}{Th \cdot B_{\text{bucket}}} & \text{if } H < Th \cdot B_{\text{bucket}} \end{cases}$$

We verified by means of simulations that the STBF avoids pseudo-periodic loss pattern that are instead observed when using the simple TBF, and that may lead to synchronization among TCP connections.

We also assume that the parameter R can be either statically configured or it can be dynamically and adaptively varied. We will refer to these two alternatives as static rate control and dynamic rate control and we will analyze them in the following Section 4 and 5, respectively.

4 STATIC RATE CONTROL

The Figure 8 shows the static rate-limiter as it would be implemented within the access point. However, the rate-limiter could also be implemented in an external device (e.g., a router) connected to the access point. The latter solution is especially suitable for a short term scenario or for a test-bed. For example, a Linux based router connected to an “off-the-shelf” access point could constitute an interesting test-bed, useful to make practical experiments. Such solution is depicted in Figure 9.

We chose the parameters of the Token Bucket filter, i.e., the rate R and the bucket size B_{bucket} , by means of a simulation study having the aim of identifying such parameters so as to avoid starvation and to provide a fair access possibility to all applications, while maximizing the

overall throughput. This simulation study is presented in Appendix VI; the final choice is $R=2.3$ Mb/s, $B_{\text{bucket}}=500$ packets of 1500 bytes and $T_h=0.9$. We stress that we need to run simulations to choose the parameters of our mechanism only in this static case, which is introduced only as a study case; in the adaptive case we will introduce a procedure to evaluate in real time the parameters of our mechanism, avoiding the need of simulations.

4.1 Numerical results

In order to evaluate the performance of the proposed rate-control mechanism in the static case, we use the same scenario adopted in Section 2, and measure throughput and fairness by varying the number of station N , with $N_{\text{dn}}=N_{\text{up}}=N$. We first address the “local wired host” scenario. We compare the performances of 3 solutions: i) no rate control, ii) the lossless rate control solution proposed in [3], iii) our proposed static rate limiter. This comparison is reported in Figure 10, which shows the total throughput as a function of N . The lossless rate control of [3] and our static rate limiter attain almost identical performances: the total throughput is almost constant as a function of the number of sources.

The total upstream and the total downstream throughput are reported in Figure 11. The lossless rate control of [3] achieves a perfect fairness, as the upstream and downstream curves cannot be distinguished. Our rate limiter solution, however, is slightly less effective in terms of fairness when there are few sources ($N=2$ to 5), since in this case the upstream connections receive a smaller capacity.

We will come back to this effect later on in this Section. Here we just observe that we have obtained the important result of avoiding the so-called critical unfairness, observed in the “no rate control” case, without the disadvantages of the lossless rate control in terms of complexity.

xxx

To continue our analysis, we evaluate the ratio between the standard deviation (σ_{up}) and the mean value ($R_{\text{up}}=R_{\text{up_tot}}/N$) of the throughput of upstream connections (the so-called unfairness index). We recall that if the ratio $\sigma_{\text{up}}/R_{\text{up}}$ is equal to zero, then we have perfect fairness;

otherwise the greater the ratio, the greater the unfairness. This ratio is shown in Figure 12 as a function of N ; the lowermost curve reports this ratio evaluated for the case of our proposed solution, the other curve is relevant to the same scenario but without rate control. The ratio relevant to the lossless rate control proposed in [3] is not plotted since it is always equal to zero for any N (perfect fairness). Our solution shows small unfairness among upstream flows (in this case the unfairness increases with N), which can be considered as tolerable, given the relatively small price paid to obtain it (in terms of added system complexity); in any case, our solution performs better than the “no rate-control” case.

So far, we analyzed throughput performance when the number of upstream connections is equal to that of downstream ones ($N_{dn}=N_{up}$). We will now examine what happens when N_{dn} is different than N_{up} , in the following four cases:

- 1) N_{up} equal to 3, N_{dn} varying between 1 and 15
- 2) N_{up} equal to 10, N_{dn} varying between 1 and 15
- 3) N_{dn} equal to 3, N_{up} varying between 1 and 15
- 4) N_{dn} equal to 10, N_{up} varying between 1 and 15

Figure 13 compares the performance of our proposed rate control solution with that of the “no-rate control”. We observe that our static rate-limiter is always capable of avoiding the critical unfairness, which is instead evident in the case of no-rate control. In addition, the static rate limiter is capable of enforcing a reasonably good sharing of resources among the upstream and downstream stations close to the ideal target of equal sharing.

As already discussed, our solution does not reach the finer “per connection” fairness. Consider for example cases (1) and (2); as the number of downstream connections increases from 1 to 15, the downstream throughput increases only slightly. The throughput of each downstream connection is roughly inversely proportional to N_{dn} , whereas the throughput of each upstream connection decreases only slightly with N_{dn} .

However, the “per connection” fairness was not among our goals. We are satisfied as long as the critical unfairness is avoided and resources are not wasted by using a low complexity solution.

xxx

To conclude the analysis of the rate limiter in the static case, we have performed a simulation study related to the impact of the variable round trip times (RTTs) that TCP connections may experience. With this analysis we have also verified that the proposed rate limiter approach, differently from the lossless rate control proposed in [3], is not affected by limitations related to the RTT. In this simulation study (reported in Appendix II) we consider the “remote wired host” scenario (see the right part of Figure 1). TCP connections are terminated on the “remote wired host” and we evaluate the throughput by varying the RTT between the Access Point and the remote wired host itself. This simulation study shows that the performance of the lossless rate control of [3] start to worsen for RTTs greater than 300 ms, whereas the performance of our proposed rate limiter does not depend on RTT at all.

The results presented so far show that the static rate control mechanism enforced by a Token Bucket Filter is effective in avoiding the critical unfairness. In the next Section we highlight the limitations of this static approach and propose an adaptive mechanism.

5 ADAPTIVE RATE CONTROL

The static mechanism described in the previous Section has two major problems: 1) if downstream connections are not present at all or they are not “greedy”, the capacity of upstream connections is un-necessarily limited to the rate R of the Token Bucket Filter, leading to a waste of resources; 2) our static mechanism assumes that the overall capacity C of the WLAN is known and used as an input parameter to the mechanism, since we need to set $R \approx C/2$; however, in general, this capacity is not known since different stations can attach to the Access Point with different physical rates (from 1Mb/s to 11Mb/s in 802.11b). To solve these problems we proceed as follows.

Let us denote by C the WLAN capacity and by R the rate of the token bucket filter. If the downstream connections are limited to a rate $R_{\text{down}} < C - R$, then the static rate limiter causes a waste of capacity in the order of $C - R - R_{\text{down}}$. The idea of the adaptive rate control is to increase the rate of the token bucket filter in these conditions up to $R' = C - R_{\text{down}}$ so that no capacity is wasted. When the downstream connections become again greedy, the rate of the token bucket filter is suitably reduced. The proposed mechanism adapts the rate R via discrete steps of amount R_{step} (Mb/s). This adjustment is performed periodically, with an interval of T_p (ms). The choice whether to increase or decrease the token bucket rate is based on a control rule, which takes into account the estimation of uplink and downlink traffic and the information about the packet losses at the AP downlink queue. The uplink and downlink traffic can be estimated outside the AP, and the packet loss information can be extracted from the AP using for example SNMP. Therefore, it is possible to implement the adaptive rate limiter solution in an external device.

Our proposed adaptive mechanism works as follow: each T_p milliseconds we estimate the “instantaneous” throughput (here we use the same definition of throughput given in Section 2, i.e., the bit rate transported by the layer below IP, comprehensive of all upper layers overheads, including IP and TCP) crossing the AP in uplink (R_{up}) and in downlink (R_{down}) - actually, we use a throughput averaged over a short period, in the order of hundreds of milliseconds. For such estimation, we use an Exponentially Weighted Moving Average (EWMA) algorithm (reported in Appendix III), which is very simple to implement⁵. We denote by C_{max} the total estimated throughput (i.e., $C_{\text{max}} = R_{\text{up}} + R_{\text{down}}$). At the same time, we monitor the number of losses at the AP downlink buffer (in a real life environment this can be done by SNMP). If no packet losses are observed in the last interval of duration T_p , this means that the downlink buffer is not congested. On the contrary, if there is at least one packet lost this means that the downlink buffer is full. In the former case, we can give more room to the upstream connections (increasing the

⁵ As a comparison, we note that [3] requires to estimate the *number* of TCP connections, which as said above is more difficult to implement.

rate of the Token Bucket Filter), in the latter case we reduce the rate of the Token Bucket Filter to avoid the risk that upstream connections will increase too much their rate, ultimately leading to starvation of TCP connections. The adaptive algorithm, which runs periodically at the time instants $T=kT_p$, can be expressed as:

```

Rup: estimated throughput from the WLAN interface to the AP at time k·Tp;
Rdown: the estimated throughput from the wired interface to the AP at time k·Tp;
NL : number of packets lost at the downlink queue in the time [(k-1)Tp, k Tp )
- at time k·Tp the TBF rate R is changed according to:
  if NL = 0
  then
    first_loss = true
    R = min (R+Rstep, C_max_theor);
  else
    target_rate = (Rdown + Rup)/2
    if first_loss
    then
      first_loss = false
      R = max( min (R-Rstep , Rup-Rstep), target_rate)
      Tokens = min(Tokens, Bbucket *Th)
    else
      R = max (R-Rstep, target_rate);

```

In case of loss, the bucket rate is decreased down to the target rate, set as one half of the estimated current capacity. When there is the first loss event after a sequence of time intervals without loss in the downlink buffer, the rate is set to the current evaluated rate in the upstream direction R_{up} , minus R_{step} and the number of tokens in the Token Bucket Filter is set to the threshold value where it can start dropping packets.

The parameter R_{step} controls the maximum speed of rate increase and decrease (equal to R_{step}/T_p). Too small values of R_{step} may make difficult the startup of new downstream connection (that need a certain amount of free capacity) and may reduce the efficiency when the bandwidth needs to be increased after a sudden reduction of the capacity required by downlink connection. On the contrary, too large values of R_{step} may give rise to significant throughput oscillations due to interactions with the underlying TCP congestion control mechanisms.

In the next sub-section, we evaluate numerically the effectiveness of our solution. We have empirically chosen a value of R_{step} equal to 200kb/s, since such choice provides good performance, in our case study. We also point out that this choice is a matter of a trade-off

between convergence time and granularity of the mechanism and that our analysis has shown that it is not a critical one, in the sense that the system is loosely sensitive to this parameter.

5.1 Numerical Results

To demonstrate the effectiveness of the adaptive rate limiter, we resort to a time-based analysis, by performing a simulation experiment in which the number of active sources varies with time. The time-schedule of the number of active upstream and downstream connections is reported in Table 1. For the same connection activity pattern, we have simulated the system by using three different approaches: i) no rate control; ii) our static rate control with $R=2.3$ Mbit/s and $B_{\text{bucket}}=500$ packets of 1500 bytes and $T_h=0.9$; iii) our adaptive rate control algorithm with the parameters reported in Table 2.

Figure 14 reports the temporal evolution of the upstream and downstream throughput, without rate-control. We observe that when there are active upstream connections (i.e., during the intervals $0 \div 400$ and $500 \div 550$ seconds), all downstream connection are starved. In addition, we have analyzed upstream starvation phenomena and registered the occurrence of such phenomena when more than six upstream connections are active (the related numerical results are not reported here for space limitations).

Figure 15 reports the temporal evolution of the throughput obtained by using the static rate limiter. We note that critical starvation of downstream connection has been avoided. When both upstream and downstream connections are present, their total throughputs are comparable, as expected by the choice of $R=2.3$ Mbps. Nevertheless, this figure shows the necessity of an adaptive mechanism in order to avoid a waste of resources. In fact, during the 100-150 seconds time interval, when there are no downstream connections, the upstream connections are not able of obtaining more than 2.3 Mbps, thus wasting half of the radio capacity.

Finally, Figure 16 reports the temporal evolution of the throughput obtained by using the adaptive rate limiter. The proposed mechanism is effective in granting all the capacity to the upstream connections during the 100-150 seconds time interval. Moreover, the sudden

throughput decrease and increase, occurring after variations of the number of connections, prove that the reaction of the adaptive control is fast enough and that the resource waste is very limited.

6 CONCLUSIONS AND FURTHER WORK

In this paper, we addressed fairness issues in a wireless access network based on the IEEE 802.11b standard, operating in DCF mode at 11 Mbps. We proposed a solution based on a “rate limiter”, operating on the uplink traffic. The rate of the rate limiter can be set statically or dynamically in response to network traffic conditions. Since the rate limiter enforces a limitation on the rate of upstream TCP connections, the remaining WLAN capacity remains available to downstream connections. When the rate is statically set, the system may waste resources, when TCP downstream connections are not greedy, i.e., when they do not use all available capacity.

Our proposed rate limiter mechanism avoids critical starvation in all considered scenarios and is independent of RTTs. Simulation results show the effectiveness of the proposed adaptive control in a dynamic scenario where the number of upstream and downstream connections varies with time. Simulation results are confirmed and validated *in a real ad-hoc developed test-bed* (the test bed and some selected measurements are reported in Appendix IV).

Coming to possible extensions of this work, a straightforward one is the support of a mix of TCP and UDP traffic (supporting real time services). We are addressing this issue by taking into account the capacity consumed by UDP flows in the adaptive rate limiter setting, either on a pre-reservation basis or on the basis of a dynamic estimation. The solution will also consider a separate queuing of UDP and TCP packets in the access point.

Finally, we note that throughout all this work we assumed an ideal behavior of the IEEE 802.11b radio link. In a more realistic environment, several factors (per-station link rate adaptation, physical layer impairments and transmission errors, MAC layer impairments, such as hidden terminals, etc.) contribute to a time-variant reduction of the WLAN capacity.

7 APPENDIX

In this Appendix we analyze the fairness performance of our system in two scenarios: i) system loaded with short-lived sources; ii) system loaded with a mix of short-lived and greedy sources. We recall that, in our scenario, unfairness phenomena are due to losses occurring in the downlink buffer. Greedy up-stream connections produce a heavy loading of the downlink buffer, for long periods of time. As a consequence, downstream connections perceive high segment losses, severely limiting their performance with respect to the upstream connections, which mainly experience ack losses. In these conditions, critical unfairness shows up.

On the other side, short-lived upstream connections generate a lighter load of the downlink buffer since their TCP congestion windows do not have the time to reach large values. As a consequence, competing downstream connections succeed in better accessing buffer resources and enjoying better fairness performance than in the previous case. In these conditions, critical unfairness does not occur (at least in our scenario; it is clear that we are talking of a continuous process and that by increasing the source duration, sooner or later we will reach critical unfairness conditions; the definition of short-lived sources is rather a qualitative one).

7.1 *Short-lived TCP sources*

Short-lived sources model the download or the upload of a small amount of data (in our setting 100 kbytes). We considered a dynamic scenario in which we activate 120 upstream connections and 120 downstream connections in the time interval lasting from $t_0=20$ s to $t_1=50$ seconds; connections are activated one after the other, every 30/240 s, alternating an upstream one and a downstream one. Connections remain active until they transfer all their data. The system is loaded so that the WLAN capacity is saturated and the system works in conditions similar to the greedy sources scenario.

Figure 23 shows the aggregate throughput of upstream and downstream connections and the overall throughput, without rate control. We can see that: i) the overall throughput reaches the system capacity; ii) data transmission ends at time $t_2=78$ seconds showing that all connections

succeed in completing their transmission; iii) the system reaches a good level of fairness between upstream and downstream connections, with a slight prevalence of the former ones (curves cross in the last part of the simulations only because upstream connections are ending their transmission before than the downstream ones and thus leave space for the latter ones); iv) critical unfairness never occurs.

Figure 24 shows the same setting but with our adaptive, dynamic rate control in action: our mechanism has a limited effect since the system is already operating in good fairness conditions; the mechanism slightly improves fairness between upstream and downstream but is not able to reach a perfect sharing. The reason is that the greater traffic burstiness of short-lived connections yields an over estimation of the parameter C_{\max} (defined in Section 5) and thus limits the upstream traffic to a greater value (i.e., 2.5 Mbit/s) than the optimal one (i.e., 2.3 Mbit/s).

7.2 *Mix of short-lived and greedy sources*

This traffic scenario comprises the same short-lived connections of the previous case, activated in the same way, plus three greedy upstream connections and three greedy downstream connections activated at time $t_0=0$.

Figure 25 shows the aggregate throughput of both short-lived and greedy upstream and downstream connections plus the overall throughput, without rate control. We can see that: i) greedy upstream connections cause the critical starvation of greedy downstream connections; ii) short-lived connections are heavily penalized: their throughput is much less than that of the greedy upstream connections; we also stress that we are speaking of 240 short-lived connections and thus the throughput perceived by the single connection can be very small or even zero.

Figure 26 shows the same setting but with our rate control in action. The system should be observed when all sources are active and have reached a stable state, i.e., in the time interval lasting from $t_0=30$ s to $t_1=40$ seconds. The effect is: i) greedy downstream connections perceive the same throughput of the greedy upstream ones; ii) short-lived traffic gets now a greater share of system resources; iii) overall, capacity is fairly shared among the four classes; we point out

that although the curves show that the overall short-lived throughput is greater than the greedy one, we must remember that we have 240 short-lived connections against 6 greedy connections. If we take into account this, we find that the capacity is indeed fairly shared among the four classes.

We can conclude by saying that our mechanism is effective in all conditions, even if, when the system is loaded with short-lived traffic only, it is not really necessary.

8 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", IEEE INFOCOM 2003, March 30-April 3, 2003, San Francisco, USA.
- [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] S. Shenker, J. Wroclawski, "Network Element Service Specification Template", RFC 2216, September 1997
- [6] The Network Simulator - ns-2, <http://www.isi.edu/nsnam/ns/>
- [7] H. Balakrishnan, V. N. Padmanabhan, and R. H. Katz, "The Effects of Asymmetry on TCP Performance", ACM Mobile Networks and Applications (MONET) Journal, Vol. 4, No. 3, 1999
- [8] W. Stevens, "TCP Congestion Control", RFC 2001
- [9] RFC 2582: "The NewReno Modification to TCP's Fast Recovery Algorithm".
- [10] M. Mathis et al. "TCP Selective Acknowledgement Options", RFC 201
- [11] Linux Advanced Routing & Traffic Control - LARTC, <http://lartc.org/>

9 APPENDIX I – ANALYSIS WITH DIFFERENT VERSIONS OF TCP

In order to assess the impact of different TCP version, we have repeated the simulation analysis described in section 2.A by using TCP NewReno and SACK. The results confirm that both the upstream/downstream critical unfairness and the critical unfairness among upstream connections are not dependent on the TCP version.

For shortness we report the simulation results only for the critical unfairness between upstream connections. Figure 17 reports the throughput of each upstream station, while varying both the number of stations (N) and the TCP version.

We can see that the occurrence of the starvation phenomena is not related with the TCP version. We note that the number of upstream connections that are able to start (i.e., with a throughput greater than zero) is similar for each TCP version taken into account. For instance, let us focus our attention on the N=10 case; the graph shows that only seven upstream connections among ten achieve throughput values greater than zero; that is, three connections are starving.

This independence of the TCP version is quite to be expected, being the starvation a problem tied with the TCP startup phase, during which the different TCP versions behave in a similar way.

10 APPENDIX II – VALIDATION OF THE STATIC RATE LIMITER FOR DIFFERENT RTT

If the server is connected to the WLAN via an high speed Local Area Network, the round trip time is not an issue. In this case the lossless rate control is an optimal solution, and it represents the upper bound of the performances. Here, we want to analyze what happens when connections experience an higher Round Trip Time. This is mentioned as an open issue in [3]. As we will show, in this case our proposed rate limiter solution may work better then the lossless rate control.

We consider now the “remote scenario” (see the right part of Figure 1) and terminate the TCP connections on a remote host in the fixed network, considering different values of the RTT of the link between the Access Point and the wired host..

Figure 18 shows the total, upstream and downstream throughput as a function of the number of stations, for an RTT value of 200 and 400 ms, comparing the two solutions: “lossless” rate control proposed in [3], and our static rate limiter. For an RTT of 200 ms the results are still comparable with the scenario where the wired host is locally connected to the Access Point. The only difference appears when the number of stations is high ($N=30$): the lossless rate control is not able to reach the maximum throughput.

For an RTT of 400 ms, the performance of the lossless rate control is definitely worse, as it never reach the maximum WLAN throughput. This is due to the imposed limitation on the TCP Congestion Windows of the TCP connections. The well know throughput limit of the window based protocols ($R \leq W/RTT$), where W is the window size, comes into play and reduces the achievable total throughput. On the other hand, the rate limiter solution is not affected by the increase in the Round Trip Time, both in terms of the total throughput and in terms of upstream/downstream fairness. Figure 19 provides another insight on the RTT problem, as it reports the total throughput versus RTT for the “no rate control”, lossless rate control, and rate limiter solutions for a fixed number of connections ($N=15$). The throughput in case of lossless rate control decreases starting from $RTT \cong 250$ ms.

11 APPENDIX III – RATE ESTIMATION ALGORITHM

When a packet arrives at a rate estimator, an EWMA (exponentially weighted moving average) algorithm is used to estimate the rate R as follows:

T : arrival time of the new packet
 L : size of the packet (expressed in bits)
 R_{old} : previous value of the EWMA estimate
 T_{old} : previous time of update of the EWMA estimate

$$R = L / \min(T - T_{old}, \tau) \cdot \alpha + R_{old} \cdot (1 - \alpha)$$

$$T_{old} = T$$
 where $\tau = 600$ ms, $\alpha = \min((T - T_{old}) / \tau, 1)$

where τ represent a time constant which can be set in the order of magnitude of the ratio buffer size/channel capacity, for example considering 100 packets of 1500 bytes at 2 Mb/s leads to 600 ms.

When the rate estimate is used at an instant T , it can be updated considering the elapsed time since its last update:

T : current time when the rate estimation is used

$$R_{updated} = R \cdot (1 - \alpha)$$
 where $\alpha = \min((T - T_{old}) / \tau, 1)$

12 APPENDIX IV – TEST-BED AND MEASUREMENT CAMPAIGN

In order to validate our simulation results we performed a measurements campaign using an ad hoc developed real test-bed, shown in Figure 20.

All the stations are associated to the Access Point. The AP is connected to a Linux PC (where our mechanisms are implemented) via a Fast Ethernet link (100 Mbps); the Linux PC is linked via a Fast Ethernet link (100 Mbps) to the server. A suitable application running on each mobile station and on the servers generates traffic and estimates the throughput. We first consider the case in which no fairness enforcer mechanism is in place and in a “static source” scenario, with $N=5$. The results are shown in Figure 21: critical unfairness phenomena are evident and confirm the simulations results shown in Figure 2.

The static Token Bucket Filter has been implemented in the Linux PC by using the “*tb*f queue discipline”, provided in a SW tool called “Traffic Control” [11].

Figure 22 shows the results obtained by using the static rate limiter with a TBF rate $R=2.7$ Mbit/s. We chose this value instead of the value $R=2.3$ Mbit/s used in the simulations because current TCP implementations (and thus the TCP used in the test-bed) use the “delayed ACK” mechanism, which yields an increase of the total WLAN throughput. As a consequence, the fairness bound and the R parameter have to be increased to 2.7 Mbit/s. Figure 22 confirms that the rate limiter avoids critical unfairness and provides a satisfactory level of fairness.

13 APPENDIX V – ANALYSIS WITH SHORT-LIVED SOURCES

In this Appendix we analyze the fairness performance of our system in two scenarios: i) system loaded with short-lived sources; ii) system loaded with a mix of short-lived and greedy sources. We recall that, in our scenario, unfairness phenomena are due to losses occurring in the downlink buffer. Greedy up-stream connections produce a heavy loading of the downlink buffer, for long periods of time. As a consequence, downstream connections perceive high segment

losses, severely limiting their performance with respect to the upstream connections, which mainly experience ack losses. In these conditions, critical unfairness shows up.

On the other side, short-lived upstream connections generate a lighter load of the downlink buffer since their TCP congestion windows do not have the time to reach large values. As a consequence, competing downstream connections succeed in better accessing buffer resources and enjoying better fairness performance than in the previous case. In these conditions, critical unfairness does not occur (at least in our scenario; it is clear that we are talking of a continuous process and that by increasing the source duration, sooner or later we will reach critical unfairness conditions; the definition of short-lived sources is rather a qualitative one).

13.1 Short-lived TCP sources

Short-lived sources model the download or the upload of a small amount of data (in our setting 100 kbytes). We considered a dynamic scenario in which we activate 120 upstream connections and 120 downstream connections in the time interval lasting from $t_0=20$ s to $t_1=50$ seconds; connections are activated one after the other, every $30/240$ s, alternating an upstream one and a downstream one. Connections remain active until they transfer all their data. The system is loaded so that the WLAN capacity is saturated and the system works in conditions similar to the greedy sources scenario.

Figure 23 shows the aggregate throughput of upstream and downstream connections and the overall throughput, without rate control. We can see that: i) the overall throughput reaches the system capacity; ii) data transmission ends at time $t_2=78$ seconds showing that all connections succeed in completing their transmission; iii) the system reaches a good level of fairness between upstream and downstream connections, with a slight prevalence of the former ones (curves cross in the last part of the simulations only because upstream connections are ending their transmission before than the downstream ones and thus leave space for the latter ones); iv) critical unfairness never occurs.

Figure 24 shows the same setting but with our adaptive, dynamic rate control in action: our mechanism has a limited effect since the system is already operating in good fairness conditions; the mechanism slightly improves fairness between upstream and downstream but is not able to reach a perfect sharing. The reason is that the greater traffic burstiness of short-lived connections yields an over estimation of the parameter C_{\max} (defined in Section 5) and thus limits the upstream traffic to a greater value (i.e., 2.5 Mbit/s) than the optimal one (i.e., 2.3 Mbit/s).

13.2 *Mix of short-lived and greedy sources*

This traffic scenario comprises the same short-lived connections of the previous case, activated in the same way, plus three greedy upstream connections and three greedy downstream connections activated at time $t_0=0$.

Figure 25 shows the aggregate throughput of both short-lived and greedy upstream and downstream connections plus the overall throughput, without rate control. We can see that: i) greedy upstream connections cause the critical starvation of greedy downstream connections; ii) short-lived connections are heavily penalized: their throughput is much less than that of the greedy upstream connections; we also stress that we are speaking of 240 short-lived connections and thus the throughput perceived by the single connection can be very small or even zero.

Figure 26 shows the same setting but with our rate control in action. The system should be observed when all sources are active and have reached a stable state, i.e., in the time interval lasting from $t_0=30$ s to $t_1=40$ seconds. The effect is: i) greedy downstream connections perceive the same throughput of the greedy upstream ones; ii) short-lived traffic gets now a greater share of system resources; iii) overall, capacity is fairly shared among the four classes; we point out that although the curves show that the overall short-lived throughput is greater than the greedy one, we must remember that we have 240 short-lived connections against 6 greedy connections. If we take into account this, we find that the capacity is indeed fairly shared among the four classes. We can conclude by saying that our mechanism is effective in all conditions, even if, when the system is loaded with short-lived traffic only, it is not really necessary.

14 APPENDIX VI – CHOICE OF THE PARAMETERS OF THE TOKEN BUCKET FILTER

In the following we present a subset of results of a simulation study having the aim of selecting the parameters of the Token Bucket filter, i.e., the rate R and the bucket size B_{bucket} . Ideally, we want to choose such parameters so as to avoid starvation and to provide a fair access possibility to all applications, while maximizing the overall throughput.

We started by studying the effect of the parameter R on the performance and we found out that there is a range of R where the throughput is maximized and fairness is achieved. In particular, we measured the total upstream and downstream throughput by varying the rate R from 1.3 Mb/s to 3.3 Mb/s, for several values of the number of sources $N_{\text{dn}}=N_{\text{up}}=N$. Such throughputs are plotted in Figure 27 as a function of R for $N=3$ and $N=15$, as an example (other values of N give rise to the same behavior).

These results show that capacity is never wasted by varying the rate R , as the total throughput is constant. The token bucket rate limiter enforces an upper bound to the upstream throughput. The upstream throughput is closer to such bound when the number of sources is large. For instance, for $R=2.3$ Mb/s the upstream throughput is 2.25 Mb/s for $N=15$ and 2 Mb/s for $N=3$.

This behavior can be explained by considering that the TCP connections are operating in congestion avoidance, and their throughput fluctuates in time. When the rate limiter introduces packet losses, the TCP upstream connections suddenly reduce their throughput. Afterwards, they slowly increase it again. These throughput values fluctuate around (but mostly below) a value equal to R/N . For small values of N , the size of oscillations is higher than for high values of N and the gap between the upstream throughput and the token bucket rate is correspondingly larger, since the token bucket rate limiter drops comparatively more packets.

Thanks to this analysis, we can choose a value for the rate R , with the obvious aim of providing good fairness performance. As a case study, in this paper we set $R=2.3$ Mb/s.

We performed a similar analysis to select the bucket size B_{bucket} . We evaluated by means of simulations the upstream throughput as a function of B_{bucket} for several values of $N_{\text{dn}}=N_{\text{up}}=N$ and

for several values of the rate R . In Figure 28, we report a small selection of these simulation results. The upstream throughput is plotted for $N=3$ and $N=15$, and for two values of R : 2 Mb/s (the two lowermost curves) and 3 Mb/s (the two uppermost curves).

We observe that the upstream throughput gets closer to the bound of the token bucket rate as the bucket size increases. However, the derivative of such curve decreases with the bucket size. As a case study, in this paper we set $B_{\text{bucket}}=200$ packets (of 1500 bytes).

Summing up, in the following we will set the parameters of the Token Bucket filter to $R=2.3$ Mb/s and $B_{\text{bucket}}=2.4$ Mb (corresponding to 200 packets of 1500 bytes).

FIGURES AND TABLES




Figure 1 - Reference simulation scenario




Figure 2 - Upstream, downstream and average total throughput ("local wired host" scenario) without rate control mechanisms




Figure 3 - Throughput of upstream connections (“local wired host scenario”)

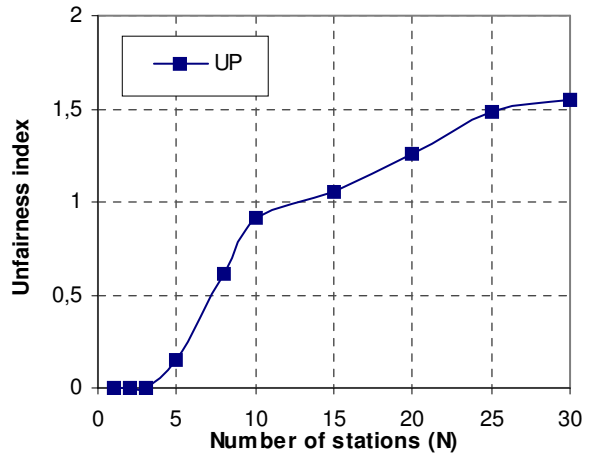


Figure 4 - Ratio σ_{up}/R_{up} for upstream connections (“local wired host scenario”)

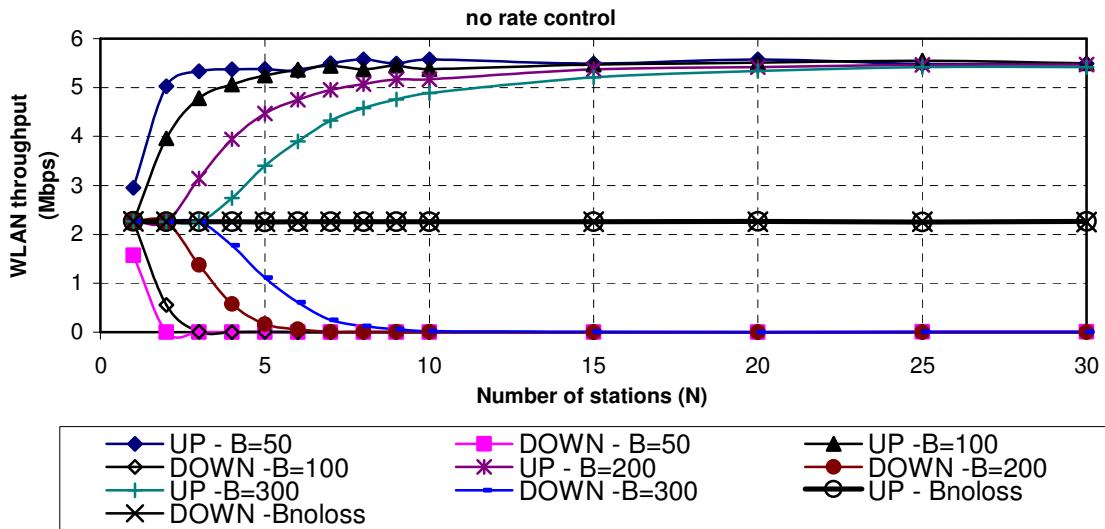


Figure 5 - Total upstream and total downstream throughput

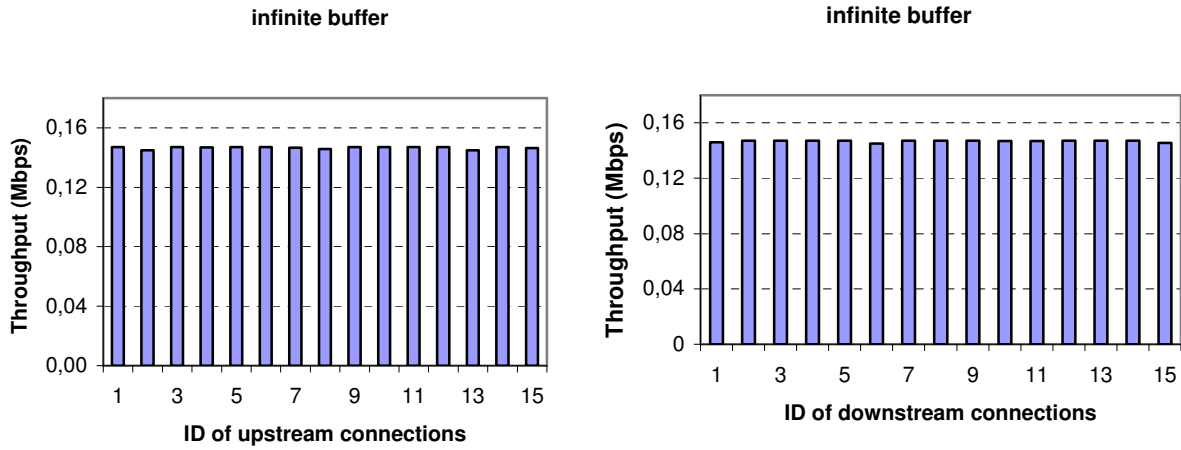


Figure 6 - Fairness achieved in lossless conditions

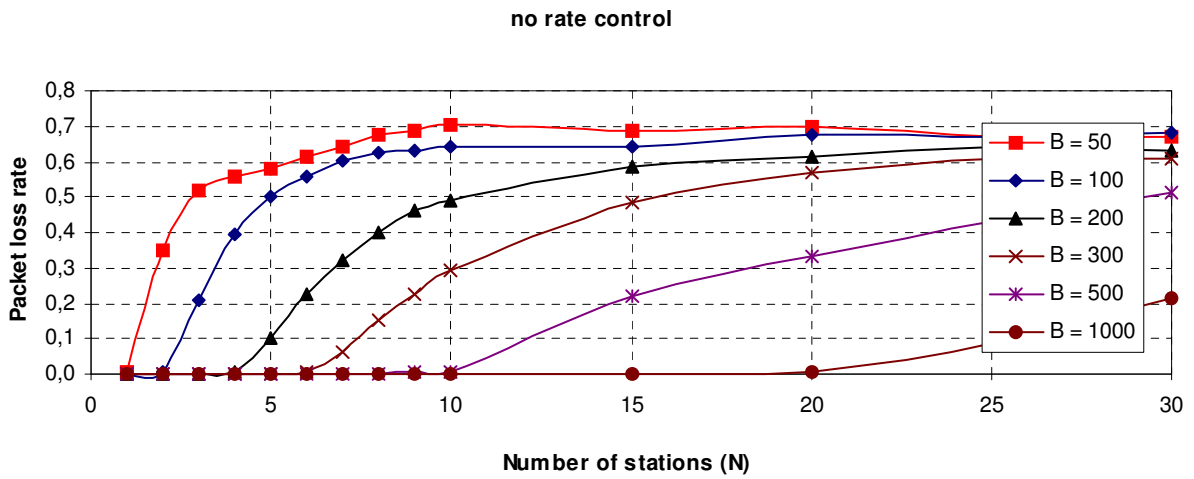


Figure 7 - Packet loss rate in the downlink Access Point buffer ("local wired host" scenario)




Figure 8: Rate control solution based on a rate-limiter

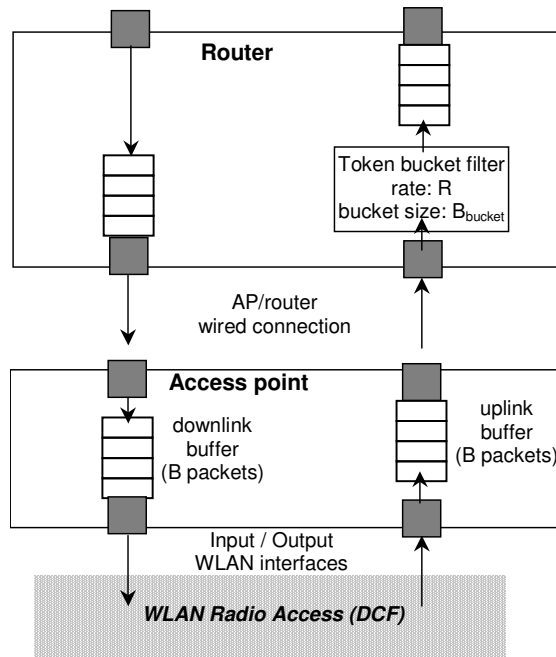


Figure 9: Rate-limiter placed on an external device




Figure 10 - Average total IP level throughput

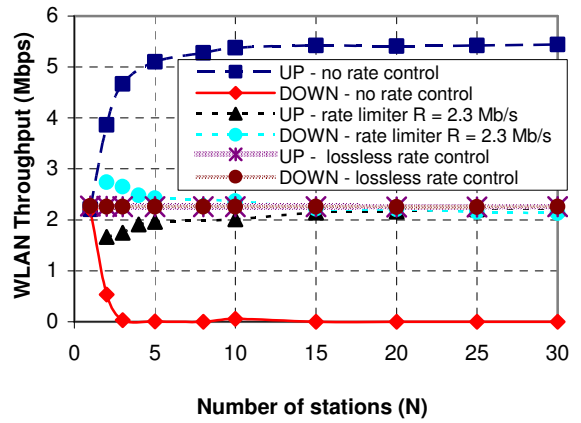


Figure 11 - Upstream and downstream throughput

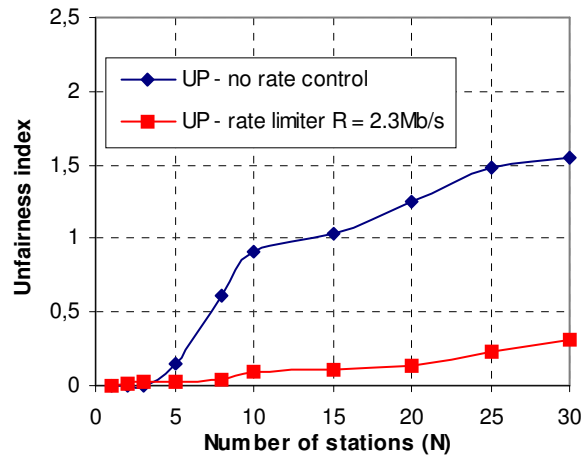


Figure 12: Upstream fairness

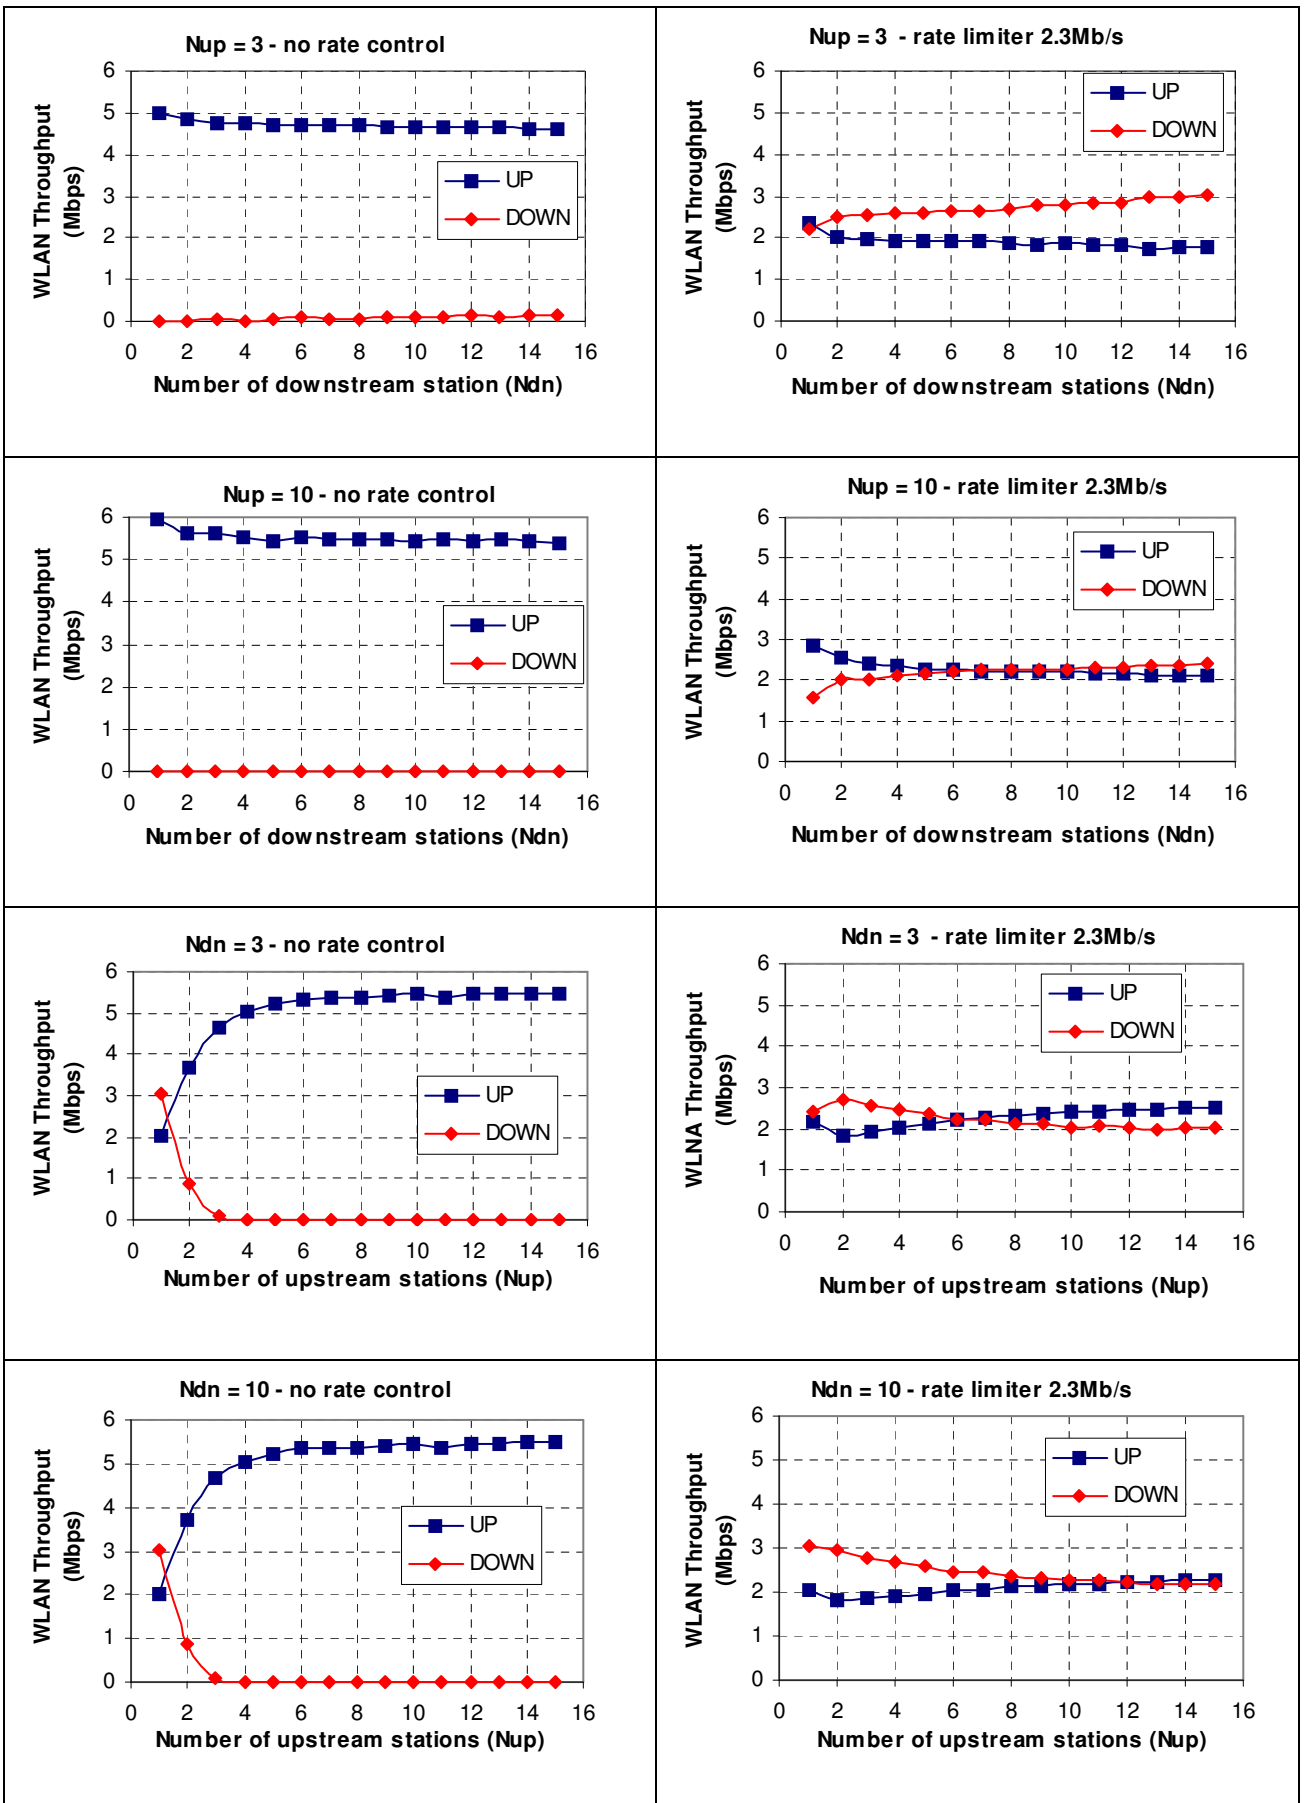


Figure 13 : Comparison between fixed rate-control solution and the “no rate control” case

| Time (sec) | 0-50 | 50-100 | 100-150 | 150-200 | 200-250 | 250-300 | 300-350 | 350-400 | 400-450 | 450-500 | 500-550 |
|-----------------------|------|--------|---------|---------|---------|---------|---------|---------|---------|---------|---------|
| No. active downstream | 3 | 3 | 0 | 3 | 6 | 6 | 10 | 10 | 10 | 10 | 10 |
| No. active upstream | 3 | 10 | 10 | 10 | 10 | 6 | 6 | 6 | 0 | 0 | 3 |

Table 1 - Time-schedule of the simulation experiments

| Parameter | Value |
|---|----------|
| AP downlink buffer B (packets) | 100 |
| TBF bucket size B_{bucket} (bytes) | 500*1500 |
| T_h | 0.9 |
| R_{step} (kbps) | 200 |
| T_p (ms) | 300 ms |

Table 2 - Adaptive rate limiter parameters

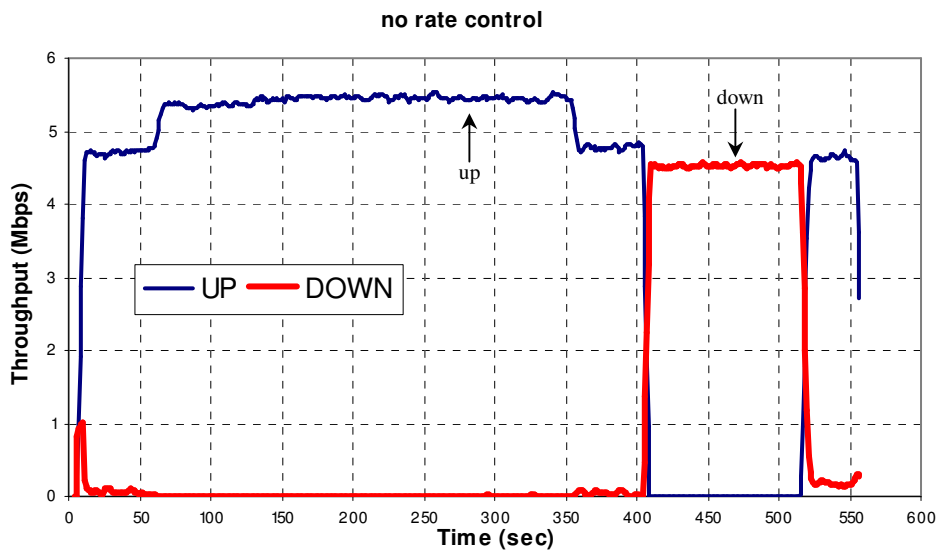


Figure 14 - Time evolution of upstream and downstream throughput without rate control




Figure 15 - Time evolution of upstream and downstream throughput with static rate control

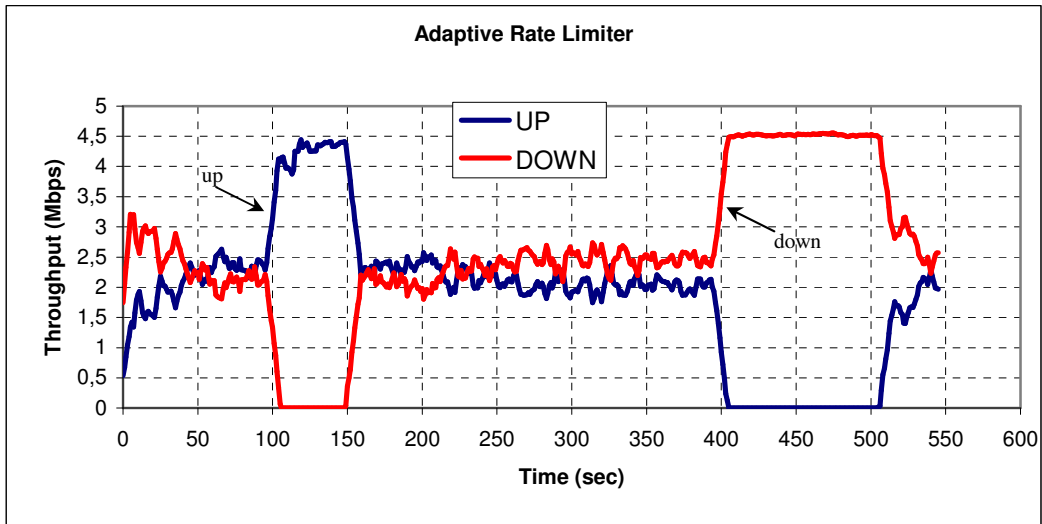


Figure 16 - Time evolution of upstream and downstream throughput with adaptive rate control

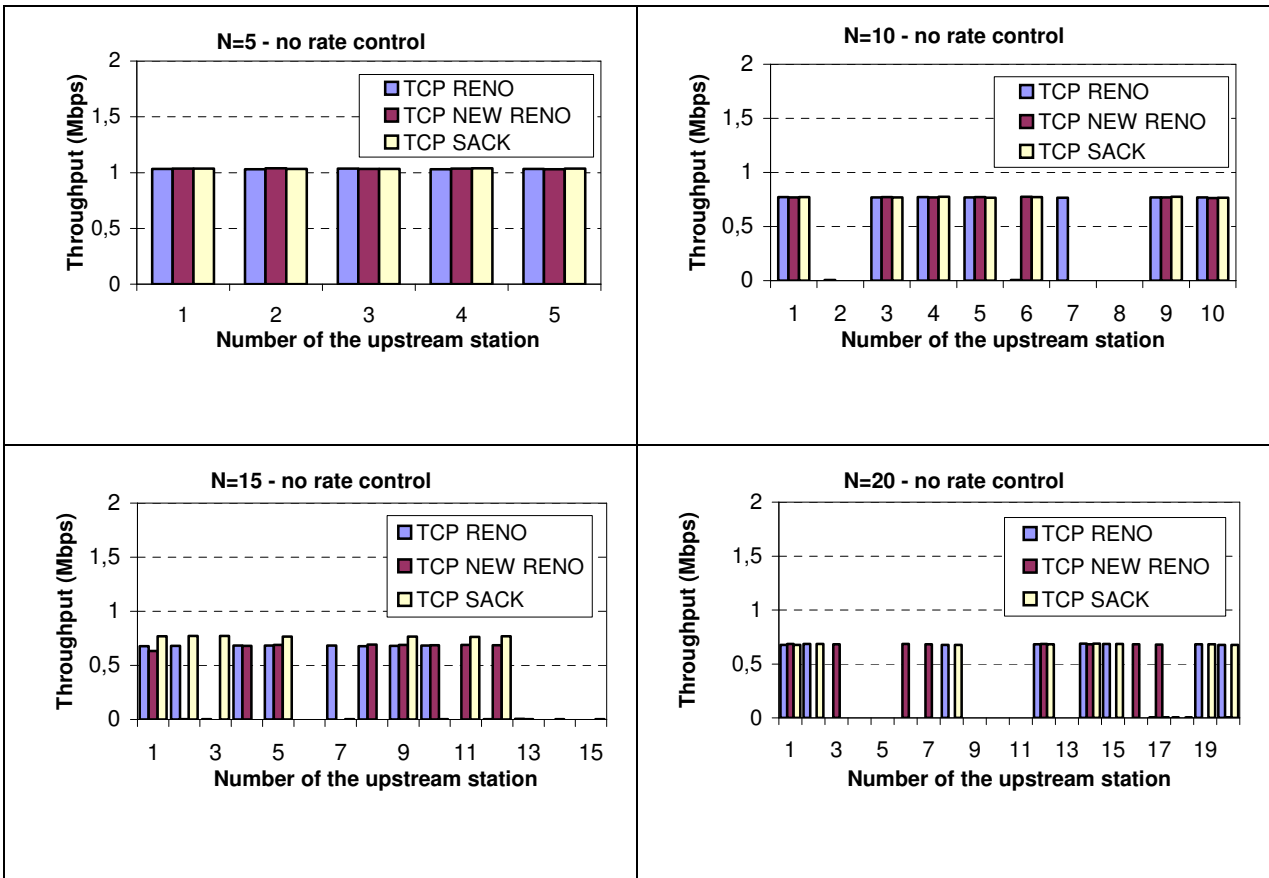


Figure 17: Upstream TCP connection throughput in the “no rate control” scenario

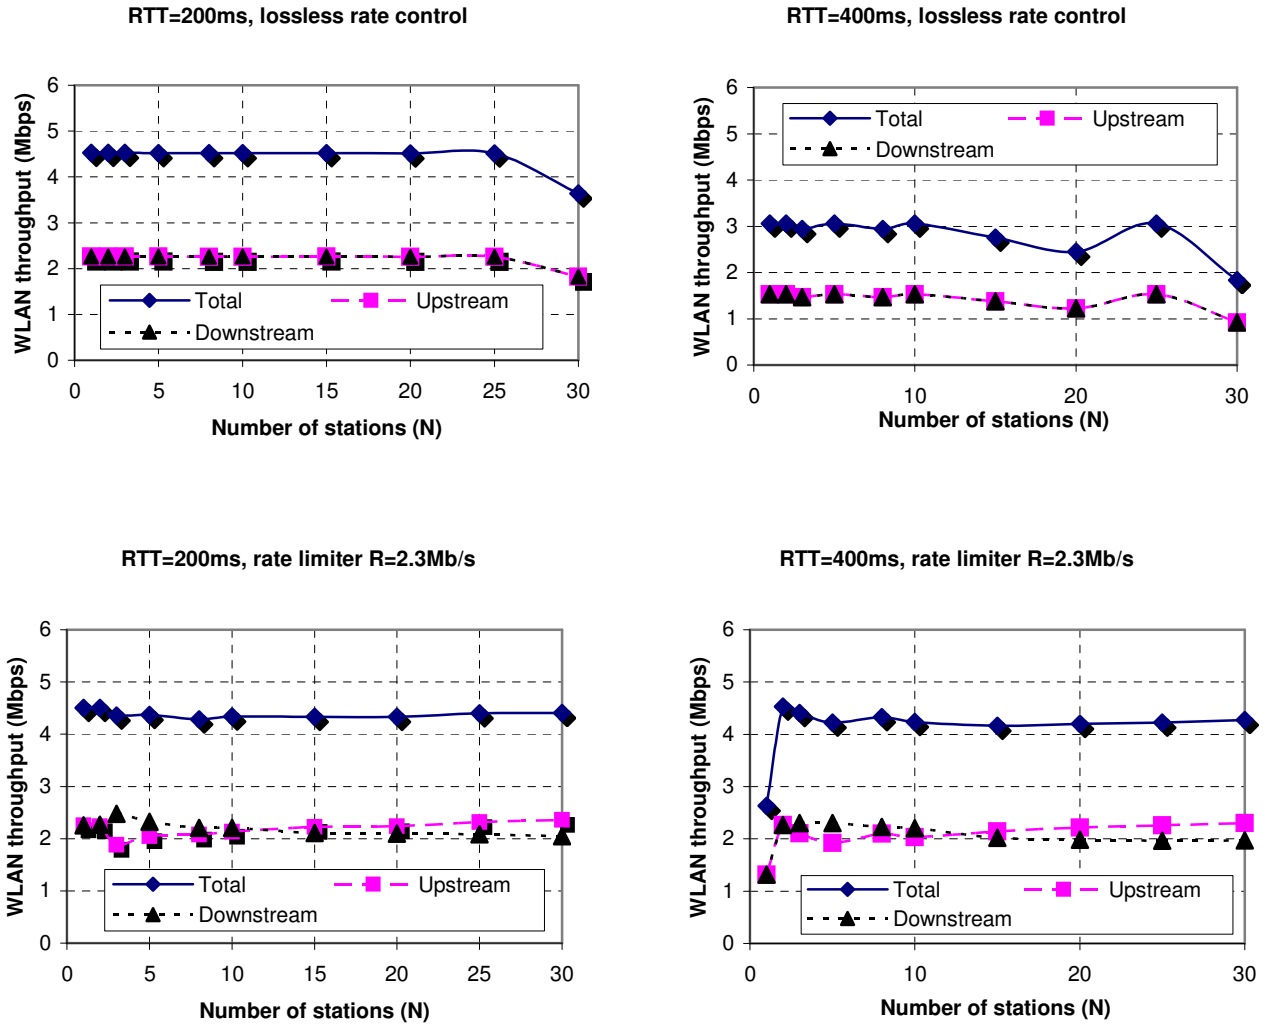


Figure 18: Throughput versus N in the “remote wired host” scenario (RTT=200, 400); comparison between lossless rate control and static rate limiter.

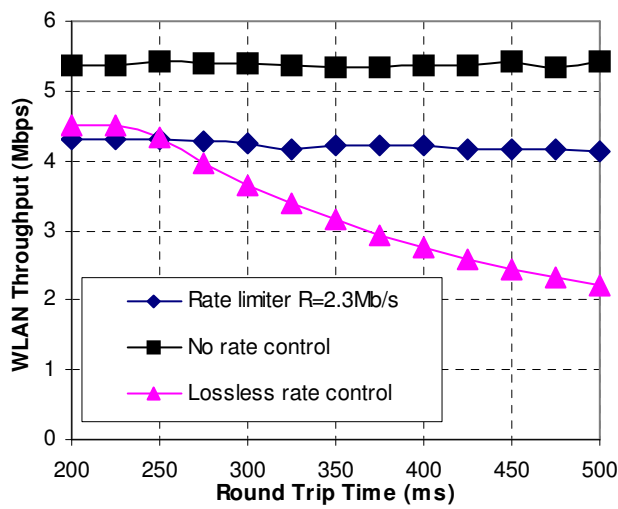


Figure 19: Total throughput vs. RTT, $N=15$ (“remote wired host” scenario)




Figure 20: Test-bed configuration




Figure 21: Throughput measured without any rate control mechanism (N=5)




Figure 22: Throughput measured with Static TBF R=2.7kbps (N=5)




Figure 23 – Aggregate throughput of upstream and downstream connections (short-lived sources scenario)




Figure 24 - Aggregate throughput of upstream and downstream connections (short-lived sources scenario), with adaptive rate control




Figure 25 - Aggregate throughput of upstream and downstream connections (mix of short-lived and greedy sources)




Figure 26 - Aggregate throughput of upstream and downstream connections (mix of short-lived and greedy sources), with adaptive rate control

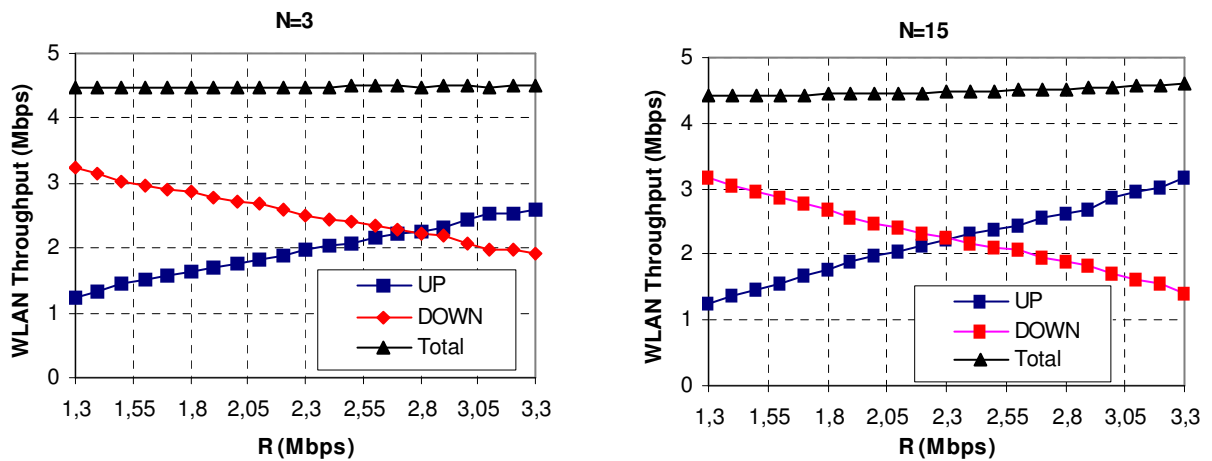


Figure 27 - Total, upstream and downstream throughput versus rate limiter R

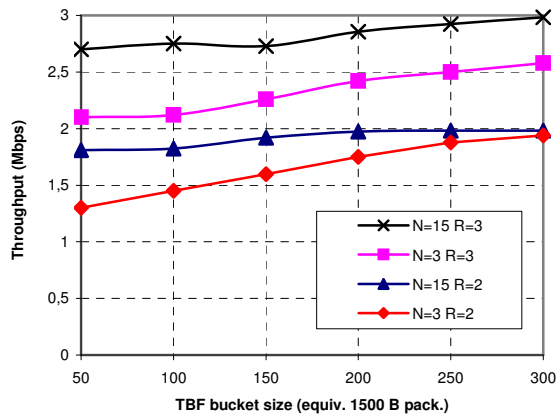


Figure 28 - Upstream throughput as a function of the bucket size