10

Some simulation results about TCP connections in ATM networks

M.Ajmone Marsan, A.Bianco, R.Lo Cigno, M.Munafò
Dipartimento di Elettronica, Politecnico di Torino
Corso Duca degli Abruzzi 24, 10129 Torino – Italy
Tel. +39 11 5644000 – Fax. +39 11 5644099
email: {ajmone,bianco,locigno,munafo}@polito.it

Abstract

We discuss simulation results concerning the performance of the TCP protocol when running over high-speed ATM networks. Two network topologies are considered: a simple network topology, comprising just two ATM switches and supporting 3 TCP connections, and a candidate Italian ATM network topology comprising ten ATM switches and supporting 6 TCP connections. In all simulation scenarios the TCP traffic is mixed with some background traffic whose level is taken as a variable parameter. Both the background traffic and the TCP traffic are either unshaped, or shaped according to the GCRA algorithm.

The effect of the background traffic on the TCP protocol performance is discussed, varying the buffering capacity within nodes as well as the peak bit rate that each TCP connection is allowed to use. Numerical results clearly show that shaping the TCP traffic according to fixed parameters significantly improves both the goodput and the efficiency of the TCP connections with respect to the case in which no traffic shaping is implemented. Moreover, the performances achievable with an *adaptive* shaping of the TCP traffic (using a simplified version of the ABR ATM transfer capability) can be observed to be extremely satisfactory.

Keywords

ATM, simulation, TCP, traffic control, traffic shaping, ABR

1 INTRODUCTION

The evolution of the ATM standards and products towards the LAN market clearly indicates that the first ATM networks will be mainly used to transport data traffic for business applications. Even in the long run, however, data traffic is expected to remain a relevant part of the load in ATM networks. It is thus very important that the high-level protocols used for the implementation of data applications be carefully investigated with respect to their adaptability to the ATM environment.

TCP (Transport Control Protocol) is today the de facto standard transport protocol for data applications in the LAN, MAN and WAN areas. Many experts believe that TCP

D. D. Kouvatsos (ed.), *ATM Networks*

for a long time to come will remain the most frequently used transport protocol in the ATM environment, even if it has been recognized that TCP is not specifically tailored to high bandwidth-delay product networks.

Some studies of the behaviour and performance of TCP when used in ATM networks already appeared in the literature (Romanow 1994, Meempat 1994, Bianco 1994, Ajmone 1995[2], Perloff 1995). Our work concentrates on the effect that the heterogeneous traffic present in the network, that we call background traffic, may have on the TCP performance. The importance of the presence of background traffic goes beyond the reduction of the bandwidth available to TCP, since background traffic interferes with the TCP behavior by altering the probability of cell losses within node buffers. Moreover, we also investigate the influence of "traffic shaping" on the TCP performance. Shaping the TCP traffic at the network ingress may be a reasonable approach to allow the network to control the TCP source rate, without requiring a substantial rewriting of the TCP code itself. Note however that a negotiation phase between the user and the network is necessary in order to agree on a given peak cell transmission rate; this rate will limit the throughput of the TCP connection, even during periods of low network load, when the throughput achievable by the TCP connection could be higher. A possible solution to this drawback is the use of shaping devices that can adapt the peak cell transmission rate of a TCP source according to feedback signals conveyed by the network. Such a solution was foreseen by the ATM Forum within the ABR (Available Bit Rate) ATM transfer capability (ATM Forum 1995). We investigate the viability of this solution by studying the effectiveness of a simplified version of ABR.

2 PERFORMANCE RESULTS

The results presented in this paper are obtained via simulation with CLASS, an ATM network simulator recently developed at Politecnico di Torino (Ajmone 1995[1]). To obtain a model for the TCP protocol, we adapted the officially distributed C code of the BSD 4.3-reno release (Jacobson 1990), without considering the delayed and selective ACK options (for details see Ajmone 1995[2]). The simulation software was validated by comparison with measurements performed on an experimental ATM LAN; furthermore, an approximate analytical model is being developed for a simple network configuration.

In all the simulation scenarios that we considered, TCP connections are supposed to perform a long file transfer from a TCP transmitter to a TCP receiver: the TCP transmitter sends only data segments, and the TCP receiver returns only ACK segments. TCP sources operate in sustained overload: segments are always ready at the transmitter when an ACK is received. The size of the buffers at the TCP transmitters is set to a large value that avoids any loss at the source during the fragmentation process of a TCP segment into ATM cells. The TCP receivers are assumed to be fast enough and to have enough buffer space so as to avoid losses. The maximum window size is set to a value that allows a single TCP transmitter to obtain the full available bandwidth on the link. It is supposed that the TCP protocol always transmits segments of 9140 bytes (9180 bytes including IP and TCP overhead), the suggested maximum segment size for TCP over ATM; TCP segments are divided in cells by the AAL5 sublayer (requiring the addition of 8 overhead bytes).

The background traffic messages are generated according to a Poisson process, with a truncated geometric message length distribution with mean equal to 20 cells and maximum length 200 cells; the background traffic is segmented according to the AAL3/4 sublayer.

The burstiness of both the TCP connections and the background traffic can be controlled with a shaping device that operates according to an adaptation of the GCRA (Generic Cell Rate Algorithm) recommended by ITU-T for traffic policing in ATM networks (ITU-T 1992).

A GCRA shaper is based on the control of the cell interdeparture time by delaying cells that are scheduled for transmission too early. The basic parameters of the GCRA shaper are the *bandwidth allocation factor* β, which is the amount of bandwidth allocated to the connection relative to its mean bandwidth, and the *cell delay variation tolerance* τ which is the amount of time that a cell is allowed to "accelerate" with respect to its expected arrival time. When the background traffic is shaped, we assign to each connection $\beta = 1.2$ and $\tau = 0$ in the case of the simple 2-node network, while in the Italian network the bandwidth allocation factor is $\beta = 1.5$.

Numerical results are presented as curves referring to two performance indices:

- the useful throughput, called *goodput*, at the TCP receivers, obtained considering the received data, but discarding all the faulty and the retransmitted segments;
- the *efficiency* of the TCP connections, i.e., the ratio between the goodput and the total offered load of TCP connections.

Curves are plotted as functions of the background traffic load, expressed in Mbit/s of user data; the background load on the link can be computed multiplying the abscissa values by 53/44. The TCP goodput is instead expressed in Mbit/s of user data for uniformity with what is generally done in literature, considering the whole 9180 byte segments. Thus, in order to obtain the link utilization, the TCP goodput must be divided by the efficiency and multiplied by a factor 53/48 (AAL5 is used) and added to the background load. The background traffic is formed by 9 different connections (in addition, an identical background traffic flows on the backward link).

Simulations were either run until the receiver throughput reached a 98% precision with 95% confidence, or stopped after about one minute of simulated time. However, with the exception of the case when ABR-like services are simulated, the 2-node network with 100 Mbit/s background traffic was so overloaded that one of the TCP connections was forced to close by the backoff mechanism, a symptom that the network is not working properly. For this reason the results of the simulation runs with a background traffic load equal to 100 Mbit/s must be interpreted very carefully.

2.1 The two-node network

We first consider a very simple network, whose topology is sketched in Fig. 1, and comprises only two ATM switches. The data rate on each channel is 150 Mbit/s, and channel L_0, linking the two ATM switches, is the system bottleneck. Three TCP connections share the network resources with a variable amount of background traffic.

The performance of this very simple ATM network was studied in detail in (Ajmone 1995[2]) as a function of three variables: the TCP connection length (the network span), the background traffic load and characteristics, and the TCP traffic shaping parameters.

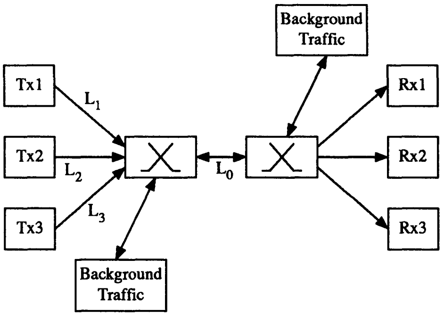

Figure 1 The simulated two-node ATM network

We only report here some figures that help the reader understand the novel results and also serve as a reference when considering more complex scenarios.

The results presented in Fig. 2, are obtained when all the links L_0, L_1, L_2, L_3 in Fig. 1 have length 500 km, while the links from the second ATM switch to the TCP receivers are assumed to have negligible length. The TCP connections length is thus 1000 km. The results are presented as a function of the background traffic load, for two values of the node buffer size in front of the congested link L_0: 1000 and 5000 cells, with shaped background traffic. As expected, when no shaping is performed on the TCP traffic (dotted lines with square markers), the TCP goodput steadily decreases with increasing background traffic; on the contrary, an increase in the node buffer size results in an increase of the TCP goodput, even if this increase is not very significant, as can be observed comparing the lines with the square markers in the left-hand side figures.

The results when the traffic on TCP connections is shaped are presented on the same charts with the plus and diamond markers for the cases of 50 and 15 Mbit/s shaping, respectively, assuming a cell delay variation tolerance $\tau = 0$. These shaping values correspond to 1/3 and 1/10 of the bottleneck link capacity. The results are rather interesting, showing that smoothing the burstiness of the traffic offered to the network allows TCP connections to better exploit the available resources. In particular, when a 50 Mbit/s shaping is enforced on TCP connections and no background traffic is present, the TCP connections completely saturate the link capacity, since they grab 149.4 out of 150 available Mbit/s, while without shaping the goodput does not exceed 83 Mbit/s, for 5000-cell node buffers. In any case, the goodput achieved with a 50 Mbit/s shaping is always greater than the unshaped goodput, regardless of the node buffer size and the background load.

The situation is slightly different when we analyze the curves with 15 Mbit/s shaping. In this case the TCP goodput is limited by the shaping function, not by the window mechanism, and it remains constant until the background load is increased to 75 Mbit/s. In this case, for high background traffic load, the goodput is greater than the one obtained in the cases without shaping and with 50 Mbit/s shaping. It is interesting to notice that in the case of 100 Mbit/s background traffic load, when the network is clearly overloaded, the performance of the TCP connections is basically the same for the three cases that were considered.

More insight can be achieved by looking at the efficiency of the TCP protocol (the

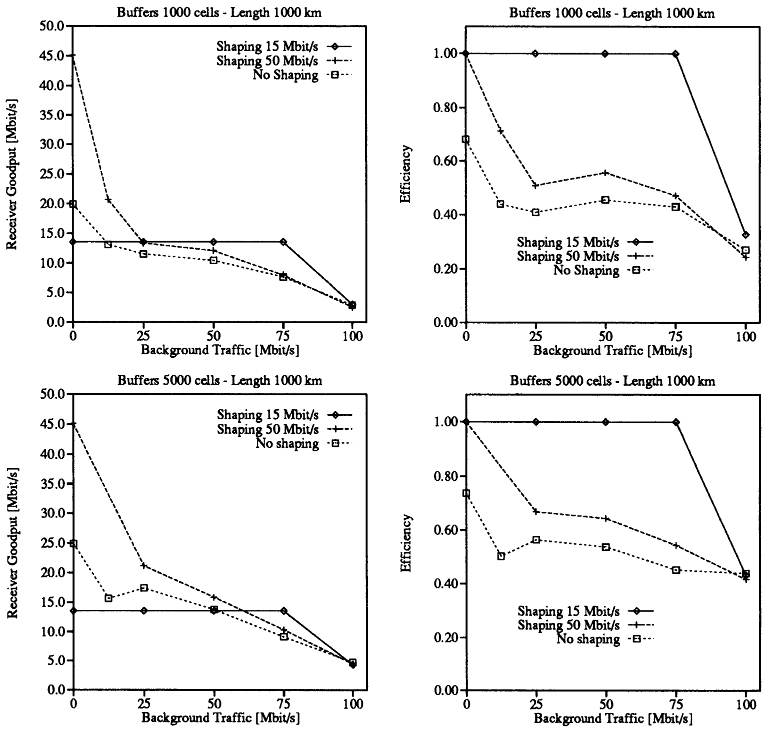

Figure 2 Average goodput and efficiency of the TCP connections for the simulated two-node network with 1000 km connections, as a function of the node buffer size and the background load; the background traffic is shaped

charts on the right-hand side of Fig. 2). These curves clearly show that as soon as the total traffic offered to the network exceeds the bottleneck link capacity, the efficiency of the TCP protocol becomes very poor, dropping to 0.5 or even less. Moreover, the more bursty is the traffic, the poorer is the efficiency. This result is also confirmed by simulation runs without shaping of the background traffic, where the TCP performance (not reported in the graphs) is even poorer. Indeed, the only acceptable, even amazingly good, situation is the one with 15 Mbit/s shaping, whose efficiency remains equal to 1 (no segment loss was recorded) up to a background traffic load equal to 75 Mbit/s; when the background load reaches 100 Mbit/s, the network is, as already stated, overloaded in all cases. It is interesting to notice the fact that with node buffer size equal to 1000 cells, the efficiency of TCP without shaping and with 50 Mbit/s shaping seems to increase slightly for background traffic load 50 and 75 Mbit/s after dropping to about 0.5 for background traffic load 50 Mbit/s. This phenomenon might be due to statistical fluctuations, but we believe that it is more probably due to phasing phenomena like those examined in (Romanov 1994, Bianco 1994).

The second set of results that we discuss considers TCP connections with different

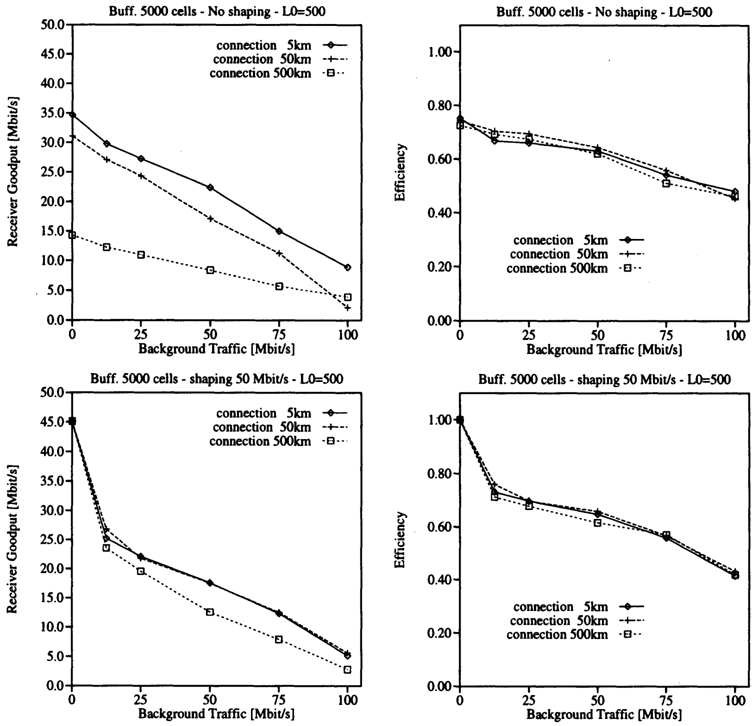

Figure 3 Goodput and efficiency of the TCP connections for the two-node network with 1000, 550, and 505 km connections, as a function of the shaped background traffic load, with buffer size equal to 5000 cells

lengths. This situation may be very common in reality, and it deserves investigation, since the TCP control mechanism is known to be biased against longer connections. In this scenario the goodput of each connection is separately taken into account and plotted. With respect to the simulation scenario presented in Fig. 1, the bottleneck link length L_0 is set to 500 km, while the lengths of the links L_1, L_2 and L_3 are set respectively to 5, 50 and 500 km, resulting in connections whose lengths are 505, 550 and 1000 km. Fig. 3 reports the results for buffer size equal to 5000 cells, when the TCP connections are either unshaped or shaped at 50 Mbit/s. The shaping at 15 Mbit/s is not reported for the sake of brevity since all of the connections obtain exactly the same goodput.

When the TCP connections are unshaped, it can be noticed that the goodput obtained by the connections is inversely proportional to the connection length, as expected, since the TCP throughput is roughly inversely proportional to the round trip delay. The unfair behavior is clear, and it can be remedied by adopting a 50 Mbit/s shaping. In this case the goodput difference between the connections of length 505 and 550 km is negligible, while the connection with length 1000 km still gets a lower bandwidth, but with low background traffic loads this difference becomes less significant.

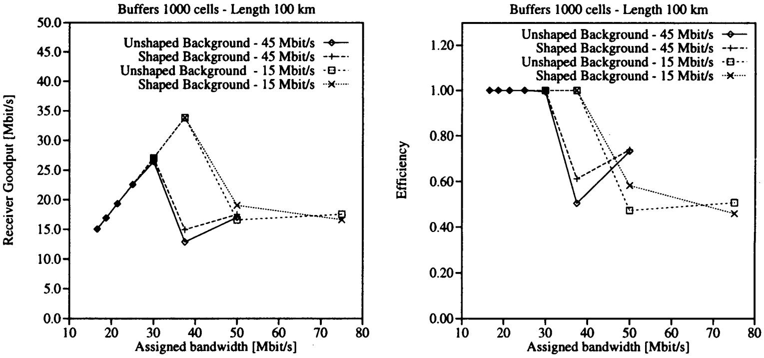

Figure 4 Goodput and efficiency of the TCP connections for the two-node network with three 100 km connections, as a function of the assigned bandwidth, with 1000-cell buffers; the background traffic load is set to either 15 or 45 Mbit/s

It is interesting to notice that the efficiency of the connection is independent from the connection length, even if no shaping is performed on the connections. This means that the losses due to buffer overflow are roughly proportional to the bandwidth grabbed by the connection.

Let us now consider what happens if we draw the results as a function of the bandwidth assigned to each TCP connection. In this case, with reference to Fig. 1, we set all the links lengths to 50 km, thus simulating 100 km connections. The size of the buffer in front of the congested link L_0 is set to 1000 cells. Fig. 4 reports the results obtained in this scenario. Each graph contains two pairs of curves: the first pair is obtained with a background traffic level equal to 45 Mbit/s, while the second one is obtained with background traffic level equal to 15 Mbit/s. Curves within each pair refer either to the case of shaped background traffic or to the case of unshaped background traffic.

The difference between the curves obtained by shaping the background traffic and those obtained by letting the background traffic remain unshaped is negligible because the buffer is big enough to accomodate the background traffic bursts. All curves show the following behaviour: if the assigned bandwidth is such that the link is not overloaded (assigned bandwidth up to 30 Mbit/s with 45 Mbit/s background, and up to 37.5 Mbit/s with 15 Mbit/s background), then the TCP connections efficiency sticks to one and hence the average goodput increases linearly. As soon as the link becomes overloaded, the TCP efficiency drops to about 0.5 and the goodput decreases.

The results presented in Fig. 4 show that, at least statically, it is possible to identify a shaping rate that optimizes the throughput obtained by TCP connections as a function of the background load; this same value also allows the maximum exploitation of the network resources without QoS reduction. This consideration suggests the exploitation of *adaptive* shaping algorithms, like those specified in the ABR ATM transfer capability, (ATM Forum 1995) for the transport of TCP connections.

In order to investigate the performance of the TCP protocol on an ABR-like transfer capability, a simplified version of ABR was implemented in the CLASS simulator. This implementation follows the ATM Forum guidelines (ATM Forum 1995) but considers only the key aspects of the algorithms, neglecting all the details that are believed not to play a key role in determining the system performance. For this reason, in order not to create confusion, we shall refer to this scenario as *adaptive shaping*, rather than ABR.

The main features of adaptive shaping are as follows.

- The interaction between the end users and the network is managed through special RM (resource management) cells that are transmitted in band. RM cells convey only a ternary feedback to sources: either Increase Rate (IR), or Keep Rate (KR), or Decrease Rate (DR). This ternary feedback is intended to guide the behavior of the source shaper.
- TCP sources shape their traffic, and introduce in their cell flow one RM cell every $N_{\rm RM}$ information cells; the feedback of the RM cell is always initialized to IR by the source.
- TCP receivers route RM cells back toward their corresponding sources, without changing the feedback carried in the RM cells.
- ATM switches monitor the traffic on the forward link, trying to identify any congestion situation. However, switches can modify the feedback in RM cells only when these reach the switch in their backward trip, while returning to the source (this is done in order to reduce the distance and hence the delay between the control point and the source). The feedback in an RM cell can be modified only from IR to either KR or DR, or from KR to DR. This is done in order to avoid the danger that nodes closer to the source set the feedback to more optimistic values than nodes farther away, that may be experiencing congestion.
- ATM switches determine their congestion state by monitoring the occupancy of the buffer associated with the link on which forward RM cells are routed. Congestion is determined using two thresholds T_l and T_h. If the buffer is filled below T_l then the switch does not modify the feedback carried in RM cells, that thus keeps its current value (IR if not previously reduced by other nodes); if the buffer is filled above T_h then the node sets the feedback in RM cells to DR; if the buffer is filled between the two thresholds then the feedback value is set to KR (unless it was already set to DR, in which case it remains DR).
- Source shaping devices always follow the indication that is contained within an RM cell; the time needed to adapt the rate is negligible. The transmission rate R_T can only be set to a value that divides the capacity of the link, i.e., $R_T = \frac{C}{N_S}$ where C is the capacity of the link and N_S is an integer. When a source shaping device receives an RM cell with a DR feedback, the value of N_S is increased by 1; instead, when an RM cell carries an IR feedback, the value of N_S is decreased by 1. This introduces a quantization effect that in some cases, especially when the bandwidth requirements of the connections are high, may affect the performance of the system, introducing oscillations in the transmission speeds.
- Switches are able to enforce fairness in the partition of the bandwidth among connections.

Numerical results were obtained with $N_{\rm RM} = 32$ (one RM cell every 32 information cells), and by setting the two thresholds T_l and T_h to 1%, and 50%, respectively, of the switch buffer size.

Fig. 5 reports the results obtained with adaptive shaping in the two-node network sce-

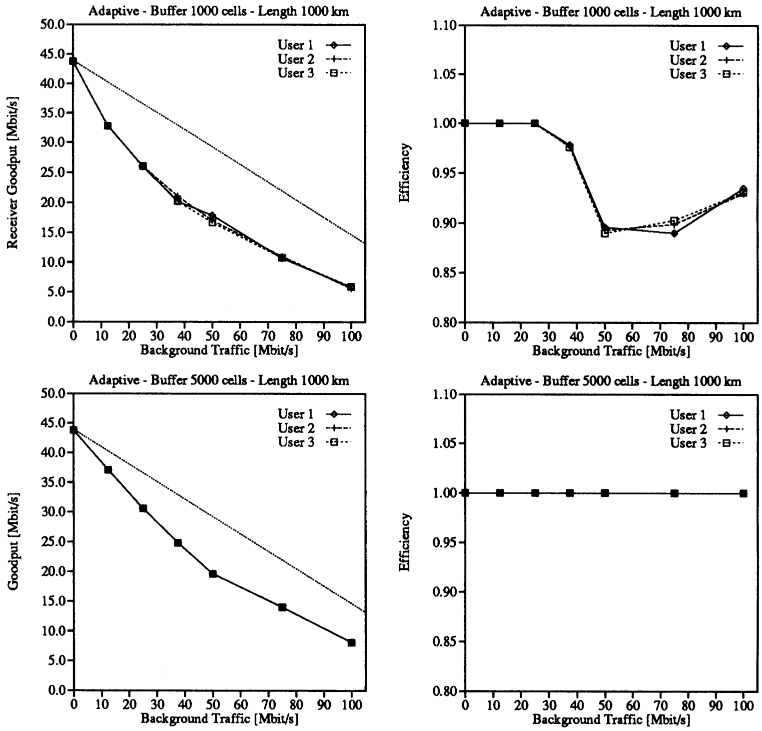

Figure 5 Goodput and efficiency of the TCP connections for the two-node network with adaptive shaping, with 1000 km connections, as a function of the background load, with buffer size equal either to 1000 or to 5000 cells

nario with 1000 km connections. These results are to be compared with those reported in Fig. 2; however, here throughput and efficiency are plotted separately for each TCP connection. The dotted straight line represents the available free bandwidth for each connection (obtained by subtracting the background and RM traffic loads from the available data rate, dividing the result by the number of TCP connections, and multiplying by 48/53 to account for the ATM cell overhead). The maximum allowed transmission rate is 50 Mbit/s for each TCP connection. Observe that the scale of the efficiency plots is greatly magnified with respect to the one in Fig. 2. The performance improvements that can be obtained with adaptive shaping are quite remarkable, and the great increase in efficiency must be noted in particular. A further important consideration is that the performance is now much better with 5000 cell buffers than with 1000 cell buffers: the reason for this difference is that these buffers must be large enough to absorb the transient phase between the congestion detection and the transmitter adaptation, whose duration is proportional to the network span. The efficiency in the case of 1000 cell buffers is affected by the already mentioned granularity in the shaper rates. Indeed, at medium loads the transmission rate keeps oscillating between roughly 15 Mbit/s and 50 Mbit/s; a trans-

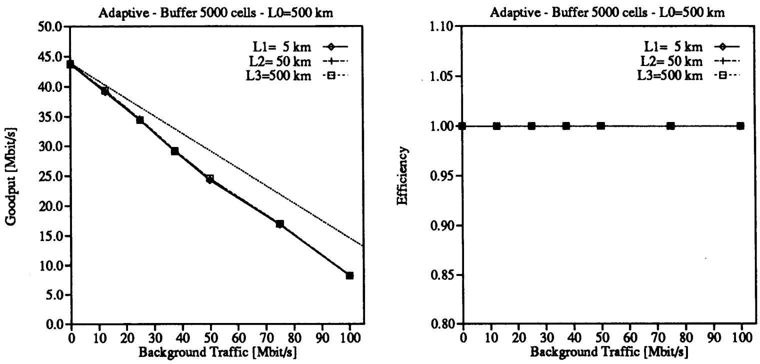

Figure 6 Goodput and efficiency of the TCP connections for the two-node network with adaptive shaping, with 1000, 550, and 505 km connections, as a function of the shaped background traffic load, with buffer size equal to 5000 cells

mission rate of 50 Mbit/s is clearly too much for the network load, but since the allowed rates above 25 Mbit/s are only 30, 37.5 and 50 Mbit/s, the transmission rate is increased too fast, thus leading to buffer overflows. This phenomenon is attenuated at high loads because the transmission rates are lower and the rate granularity becomes negligible.

Fig. 6 reports the results obtained with adaptive shaping when the TCP connections have different lengths. This figure can be compared with Fig. 3. In this case the advantage of a mechanism that allows the transmission rate to be controlled, and the fairness among connections to be enforced, leads to a striking performance improvement: all TCP connections have efficiency one, all of them receive the same amount of bandwidth which corresponds to a very high fraction of the available bandwidth. It is quite interesting to notice that the overall performance in this case is better than the one obtainable when all TCP connections are 1000 km long. Since shorter connections are easier to control, this means that all connections benefit from the presence of short connections: a behavior which is exactly the opposite of the one observed in the case of TCP connections without adaptive shaping.

2.2 The Italian network

The candidate Italian network topology comprises ten ATM switches, located in the major Italian cities, and is shown in Fig. 7; the buffering capacity at all nodes is set either to 100 or to 1000 cells per port, and the user transmission buffer sizes are set to quite large values in order to avoid losses at the source. Six TCP connections can be identified: Mi-Ro, To-Fi, Ve-To, Ro-Ba, Ba-Pa and Pa-Ba. The total amount of background traffic in the network is equal to 0.8, 1.0 and 1.2 Gbit/s, and the traffic distribution is highly asymmetric, the network having essentially two hot spots in Rome and Milan. The complete workload distribution for the background traffic is reported in Table 1. Table 2 presents the background load, measured as a percentage of the link capacity, of all the links crossed by the six TCP connections.

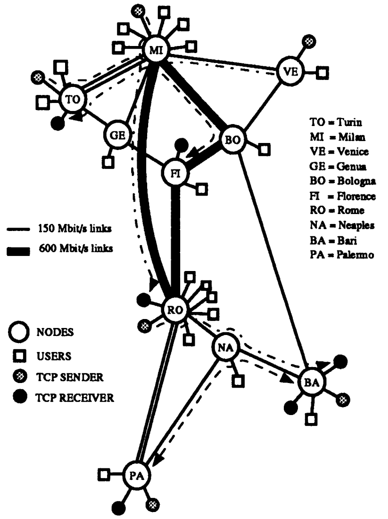

Figure 7 Topology of the Italian network

Node	MI	TO	GE	VE	BO	FI	RO	NA	BA	PA
MI	0	52	20	2	18	3	206	17	2	17
TO	52	0	3	0	3	5	32	3	0	3
GE	20	3	0	0	2	2	12	1	0	1
VE	2	0	0	0	0	0	1	0	0	0
BO	18	3	2	0	0	2	11	1	0	1
FI	35	5	2	0	2	0	22	2	0	2
RO	206	32	12	1	11	22	0	10	1	10
NA	17	3	1	0	1	2	10	0	0	1
BA	2	0	0	0	0	0	1	0	0	0
PA	17	3	1	0	1	2	10	1	0	0
Total	369	101	41	3	38	70	305	35	3	35

Table 1 Traffic matrix used in the simulation of the Italian topology; the traffic is generated by the node in the column and goes to the node in the row; the relations are expressed in thousandths of the global generated traffic

TCP Conn	0.8 Gbit/s			1.0 Gbit/s			1.2 Gbit/s		
	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd
Mi-Ro	0.41			0.51			0.61		
To-Fi	0.294	0.114	0.11	0.367	0.143	0.13	0.44	0.165	0.15
Ve-To	0.072	0.30		0.090	0.377		0.098	0.43	
Ro-Ba	0.25	0.06		0.32	0.074		0.38	0.088	
Ba-Pa	0.06	0.028		0.074	0.035		0.088	0.042	
Pa-Ba	0.028	0.06		0.035	0.074		0.042	0.088	

Table 2 Measured load of the background traffic, as a percentage of the link capacity, on the links crossed by the TCP connections

All the TCP connections have a window size that allows a transmission speed of up to 150 Mbit/s if no shaping is applied by the source. The Mi-Ro TCP connection is carried on a link with much available capacity, since either 60%, or 50%, or 40% of a 600 Mbit/s channel is available for it, respectively, in the three cases of background load. The Ba-Pa and Pa-Ba connections are running on very lightly loaded links, while the three TCP connections To-Fi, Ve-To, and Ro-Ba run over 150 Mbit/s channels whose loads could significantly influence the TCP performances. The Ro-Ba and Pa-Ba connections interfere with one another on the Na-Ba link.

We present results for either shaped or unshaped background traffic. Simulations were run considering five possible scenarios for all the TCP connections: unshaped TCP connections, TCP connections shaped at a link speed equal to either 25Mbit/s, or 37.5Mbit/s, or 50Mbit/s, which means that the TCP goodput can be at most either 22.5, or 33.8, or 45.1 Mbit/s, and finally TCP connections with adaptive shaping. All these scenarios are simulated considering node buffer lengths of either 100 or 1000 cells.

Figures 8 and 9 present the goodput and efficiency for the six TCP connections, as a function of the global background network load; curves refer to TCP connections shaped either at 25 Mbit/s, or at 37.5 Mbit/s, or at 50 Mbit/s, or unshaped, with background traffic either shaped or unshaped; the node buffers are 100 cells long. Figures 10 and 11 reports the results for the same scenarios, but with node buffer size 1000 cells.

The first consideration concerns the buffering capacity within ATM switches. When the buffers are quite small (100 cells) the burstiness of the background traffic has a great impact on the performance of the TCP connections. Instead, when the buffering capacity is large enough to absorb the bursts of cells generated by the background traffic sources (1000 cells buffers), the influence on the TCP connections of the burstiness of the sources becomes negligible.

Let's now consider each one of the six connections separately, since all of them exhibit peculiar behaviours that are worth discussing.

The Mi-Ro connection has a lot of spare bandwidth to exploit; thus it operates with

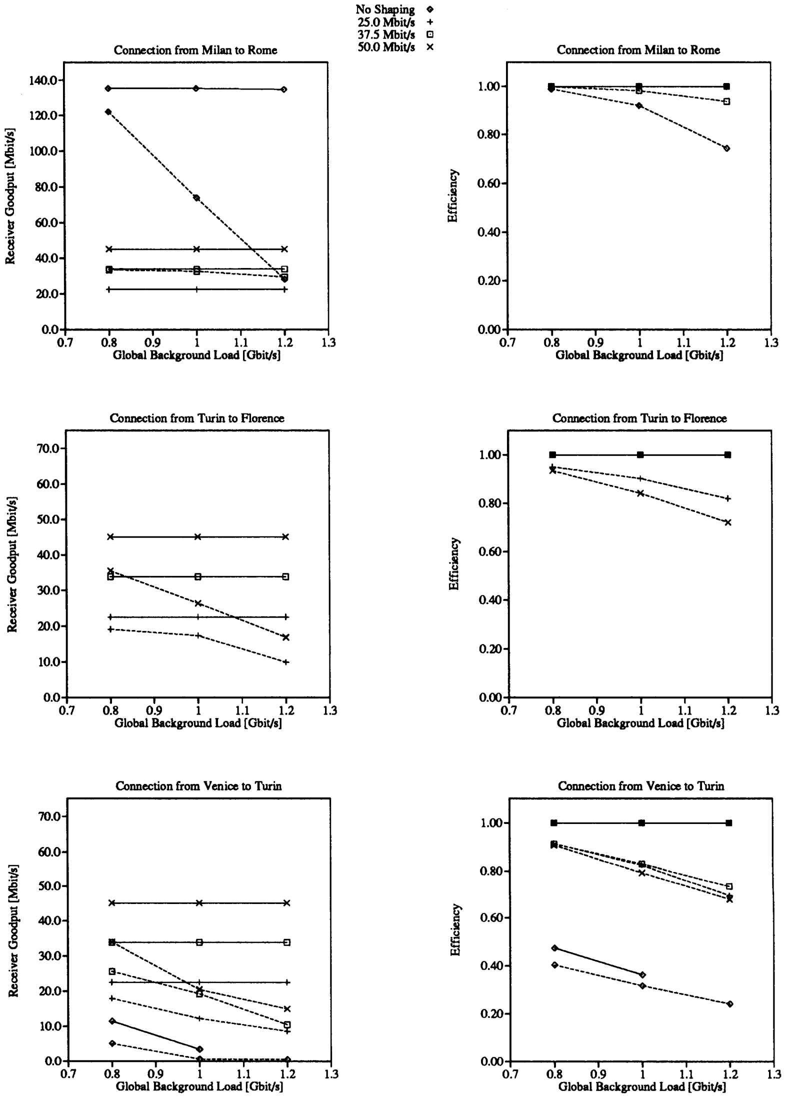

Figure 8 Average goodput and efficiency of the TCP connections for the Italian network as a function of the background load, when the node buffers sizes are 100 cells; solid lines refer to shaped background traffic, whereas dashed lines refer to unshaped background traffic

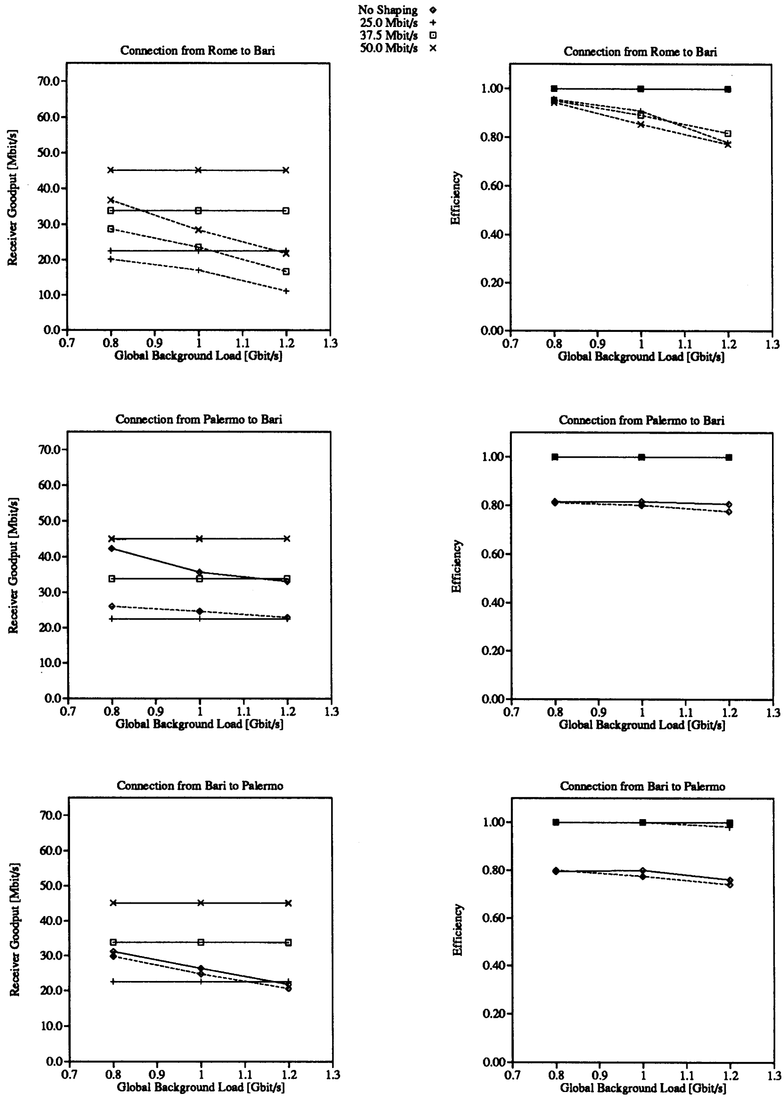

Figure 9 Average goodput and efficiency of the TCP connections for the Italian network as a function of the background load, when the node buffers sizes are 100 cells; solid lines refer to shaped background traffic, whereas dashed lines refer to unshaped background traffic

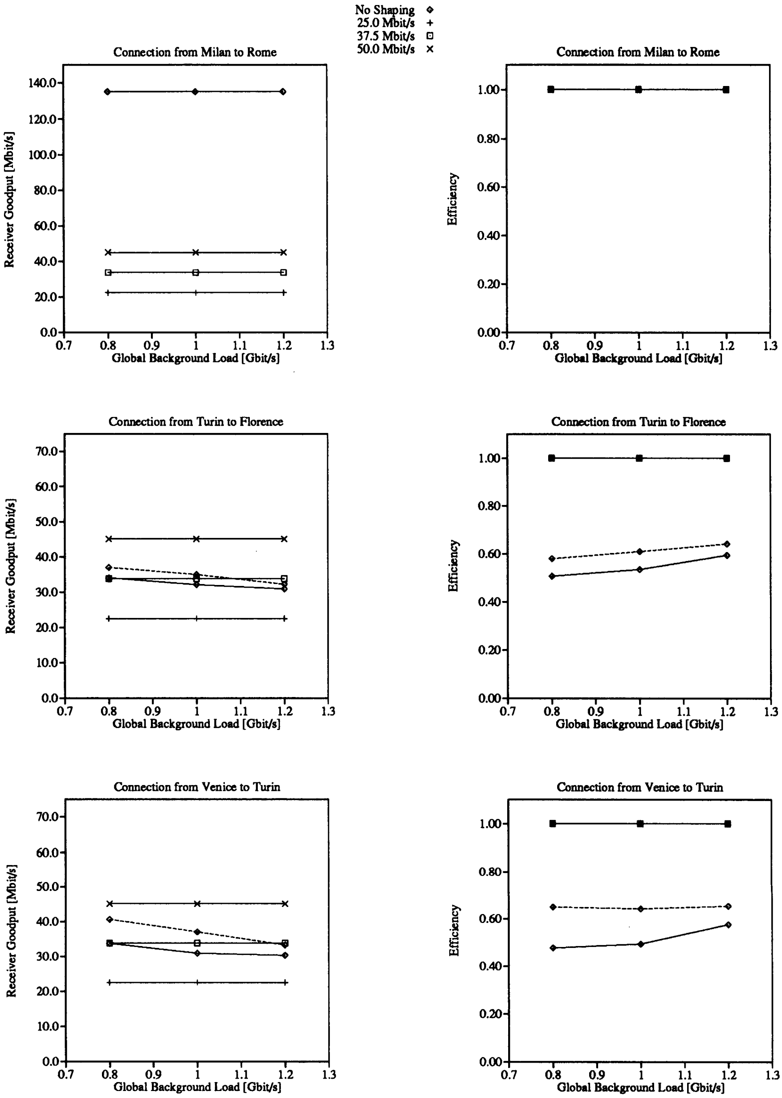

Figure 10 Average goodput and efficiency of the TCP connections for the Italian network as a function of the background load, when the node buffers sizes are 1000 cells; solid lines refer to shaped background traffic, whereas dashed lines refer to unshaped background traffic

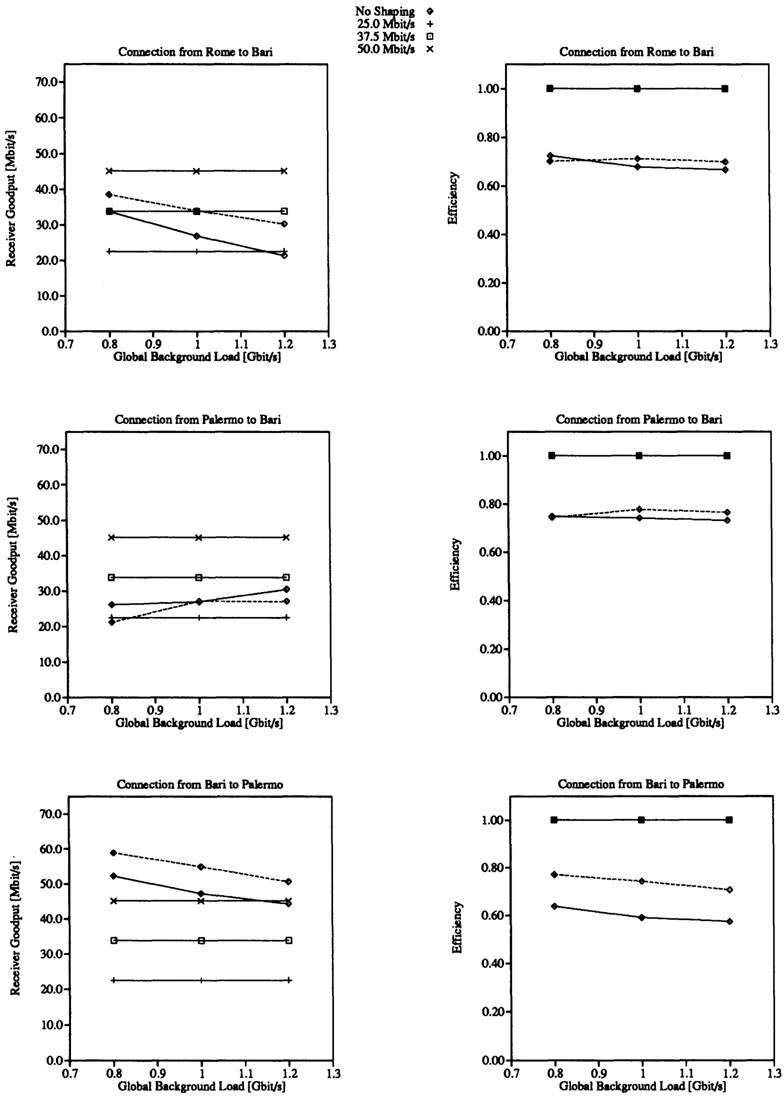

Figure 11 Average goodput and efficiency of the TCP connections for the Italian network as a function of the node buffer size and the background load, when the node buffers sizes are 1000 cells; solid lines refer to shaped background traffic, whereas dashed lines refer to unshaped background traffic

efficiency one, grabbing all the resources it can, in all case except one: if neither the TCP traffic nor the background traffic are shaped and the node buffers are small (100 cells), losses occur in the node buffer, so that both the TCP goodput and efficiency significantly decrease with the increase in the background load.

The To-Fi and Ve-To connections have a similar amount of available resources to exploit, and they behave similarly. Both of them completely use their assigned bandwidth when both their traffic and the background traffic are shaped, or when the node buffers are large enough to absorb the background traffic burstiness. On the other hand, both connections suffer significantly when their traffic is not shaped, obtaining very poor efficiency and a remarkable reduction in goodput. It is important to observe that the missing points in the graphs, like for instance those referring to the To-Fi connection without shaping of the background and TCP traffics with 100 cells buffers, correspond to simulations where the TCP connections were closed due to the TCP backoff mechanism. This behavior clearly shows that without some kind of rate control the TCP flow control mechanism is not able to work properly in high speed networks. The last effect to be observed is that when the buffers size is 1000 cells and the TCP traffic is not shaped, TCP achieves better performance if also the background traffic is not shaped; this is due to the fact that when the background traffic is not shaped, cells are lost in bursts, thus concentrating the losses on a smaller number of TCP segments.

Let's now come to the connections that interact in Neaples: Ro-Ba and Pa-Ba. Also in this case some points are missing due to the closure of TCP connections, and the same general considerations presented above apply here too. Moreover, it can be observed that the interaction between the two TCP connections on a lightly loaded link does not seem to jeopardize performance.

Finally, consider the Ba-Pa connection. This connection runs alone on a lightly loaded set of links, and its performance is consequently quite good. The most interesting aspect to be noted in this case is what happens to the TCP connections when no shaping is used, and the buffer size is increased from 100 to 1000 cells. The goodput of the connection increases significantly with the larger buffers, but the efficiency remains very poor, below 0.8, and in fact it is even reduced when the buffer size increases. This phenomenon once again confirms that TCP by itself is not suited to high speed networks, since it wastes a great amount of resources.

Figure 12 refers to the case of adaptive shaping, with the characteristics described in the previous section, but with no enforcement of fairness among connections. The circular markers refer to the case of 100 cell buffers, while the square markers refer to the case of 1000 cell buffers; black markers refer to unshaped background traffic, white markers to shaped background traffic. The buffer thresholds are set to $T_l = 5$ and $T_h = 50$ in the case of buffer size 100, and to $T_l = 10$ and $T_h = 500$ in the case of buffer size 1000. The maximum transmission rate is 75 Mbit/s for all connections, except Mi-Ro that is allowed to transmit up to 150 Mbit/s. First of all, it must be noted that when buffers are 1000 cells long all connections attain efficiency 1, a result which is quite a success in itself; moreover, the throughput of the connections is higher than that obtained without shaping or with fixed shaping. Also in the case with 100 cell buffers the benefits of an adaptive shaping policy are quite evident: both the efficiency and the throughput are higher than without shaping. However, in this case buffer sizes are not large enough to allow a smooth control of the sources, i.e., they cannot accommodate all the cells that are transmitted at high speed by the source before it receives the RM cell with the DR feedback, hence the

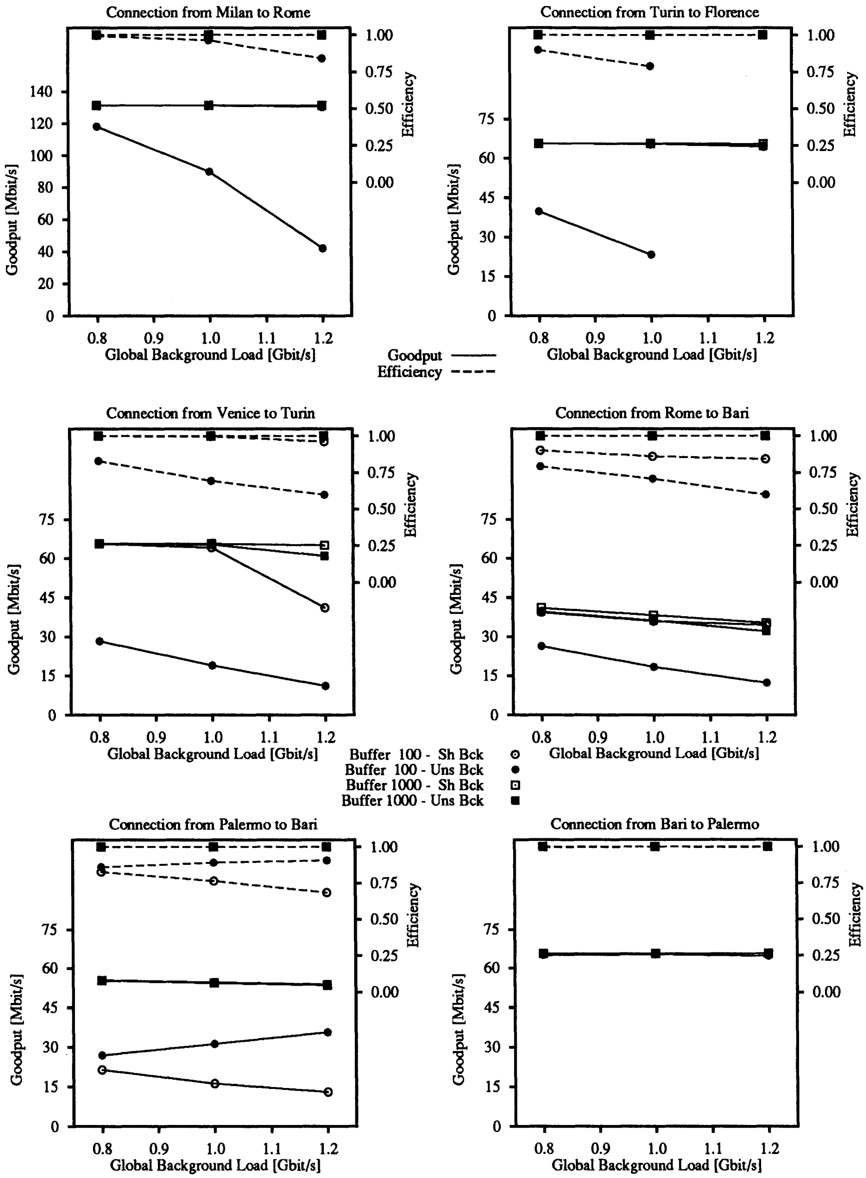

Figure 12 Average goodput and efficiency of the TCP connections in the case of adaptive shaping for the Italian network as a function of the node buffer size and the background load

efficiency of the TCP connections is not always one. Notice that in these conditions the To-Fi connection is still forced to close when the background load is 1.2 Gbit/s. One more aspect worth noticing is the behavior of the Pa-Ba connection with buffer size 100 cells, without shaping of the background traffic. In this case the TCP connection throughput increases when the background increases. This is due to the fact that for this connection the true bottleneck is the Na-Ba link, where most of the traffic is due to TCP, while the competing TCP connection also suffers from the quite heavily loaded Ro-Na link, hence when the background traffic increases the Ro-Ba connection is forced to reduce its rate and the Pa-Ba connection can exploit the bandwidth on the Na-Ba link freed by the Ro-Ba connection. This shows that an adaptive shaping is indeed able to exploit dynamically the available bandwidth. Moreover, it seems that, if the number of connections competing for the bandwidth is small, an adaptive shaping scheme may work well even if fairness is not enforced by the network.

3 CONCLUSIONS

The performance of the TCP protocol when running over ATM networks was studied through simulation, in two network scenarios, considering the TCP connections goodput and efficiency as significant performance parameters.

The variable parameters of the study are the background traffic load, the buffering capacity in the ATM switches, the traffic shaping parameters of both TCP and background traffic and the length of the TCP connections.

Numerical results clearly show that shaping the traffic on the TCP connections greatly improves the TCP performance, and also indicate that ABR-like solutions may be quite advantageous.

An important advantage of the shaping approach for TCP connections is that shaping techniques can be applied with no modification of the TCP protocol itself.

4 REFERENCES

Ajmone Marsan M., Bianco A., Do T.V., Jereb L., Lo Cigno R., Munafò M. (1995^1) ATM Simulation with CLASS. *Performance Evaluation Journal*, **24**, 137-159.

Ajmone Marsan M., Bianco A., Lo Cigno R., Munafò M. (1995^2) Shaping TCP Traffic in ATM Networks. *IEEE ICT'95*, Bali, Indonesia.

ATM Forum/95-0013R8 (1995), ATM Forum Traffic Management Specification 4.0.

Bianco A. (1994) Performance of the TCP Protocol over ATM Networks. *IEEE ICCCN'94*, San Francisco, CA USA.

ITU-T Recommendation I.371 (1992) Traffic Control and Congestion Control in B-ISDN. Geneve, Switzerland.

Jacobson V. (1990) Berkeley TCP evolution from 4.3-tahoe to 4.3-reno. *Eighteenth IETF*, Vancouver, BC Canada.

Meempat G. (1994) Interactions of Reliable Stream Data Transport Protocols with Rate Contol Algorithms. *IEEE ICCCN'94*, San Francisco, CA USA.

Perloff M., Reiss K. (1995) Improvements to TCP Performance in High-Speed ATM Networks. *Communications of the ACM*, **38-2**, 90-100.

Romanow A. and Floyd S. (1994) Dynamics of TCP Traffic over ATM Networks. *ACM SIGCOMM'94*, London, UK.

5 ACKNOWLEDGMENTS

This work was supported in part by a research contract between Politecnico di Torino and CSELT, in part by the EC through the Copernicus project 1463 ATMIN, and in part by the Italian Ministry for University and Research.

6 BIOGRAPHIES

Marco Ajmone Marsan is a Full Professor at the Electronics Department of Politecnico di Torino, in Italy. He holds a Dr. Ing. degree in Electronic Engineering from Politecnico di Torino, and a Master of Science from the University of California, Los Angeles. Since November 1975 to October 1987 he was at the Electronics Department of Politecnico di Torino, first as a Researcher, then as an Associate Professor. Since November 1987 to October 1990 he was a Full Professor at the Computer Science Department of the University of Milan, in Italy. His current interests are in the fields of performance evaluation of data communication and computer systems, communication networks and queueing theory.

Andrea Bianco is a Research Assistant at the Electronics Department of Politecnico di Torino. He holds a Dr.Ing. degree in Electronics Engineering and a Ph.D. in Telecommunications Engineering both from Politecnico di Torino. Since May 1991 to December 1994 he has been with the Electronics Department of Politecnico di Torino as a Ph.D. student. In 1993 he spent one year visiting HP Labs in Palo Alto, CA, working on performance analysis of the TCP protocol over ATM networks. Since November 1995 he is with the Dipartimento di Elettronica of the Politecnico di Torino. His current research interests are in the field of access protocols for all-optical networks, performance analysis of ATM networks, and formal description techniques.

Renato Lo Cigno is a Research Engineer in the Electronics Department of Politecnico di Torino. He received a Dr. Ing. degree in Electronics Engeneering from Politecnico di Torino in 1988. Since then he has been with the telecommunication research group of the Electronics Department of Politecnico di Torino, first under various research grants and contracts, then as a staff member. His research interests are in communication networks simulation and performance analysis.

Maurizio Munafò is a Research Engineer in the Electronics Department of Politecnico di Torino. He holds a Dr.Ing. degree in Electronics Engineering and a Ph.D. in Telecommunications Engineering, both from Politecnico di Torino. Since November 1991 he has been with the Electronics Department of Politecnico di Torino, where he has been involved in the development of an ATM networks simulator. His research interests are in simulation and performance analysis of communication systems.