INTERNATIONAL JOURNAL OF SATELLITE COMMUNICATIONS
Int. J. Satell. Commun. 2001; **19**:93–110

TCP modifications over satellite channels: study and performance evaluation

Mario Marchese*,†

Consorzio Nazionale Interunversitario per le Telecomunicazioni (CNIT), University of Genoa Research Unit, Department of Communications, Computer and System Science (DIST), University of Genoa, Via Opera Pia 13, 16145 Genova, Italy

SUMMARY

The paper presents a study and an analysis of the performance offered by TCP over a GEO (geostationary orbit) satellite link. The characteristics of satellite channels (notably the large round trip delay) heavily influence the TCP flow control, which is essentially based on acknowledgements. The effect is a delay in the acknowledgement reception and, consequently, in the delivery of messages, with respect to cabled networks. The drawbacks of using TCP in a satellite environment may be mitigated by a proper tuning of some TCP parameters. The behaviour of the protocol, as a consequence of variations in the buffer length of both transmitter and receiver and in the initial congestion window, is investigated in the paper and a proper configuration that drastically improves performance (measured by the throughput in bytes/s and by the overall transmission time) is proposed. Two test environments have been used to evaluate the proposed modifications: a real testbed, composed of two remote hosts connected through a satellite channel, and a satellite network emulator, composed of three PCs (two of them representing two hosts, the third one reproducing the behaviour of the satellite link). In both environments, a ftp-like application designed for the aim has represented the reference application; three different file sizes have been used and the different effect of the tuning, depending on the transfer length, has been evidenced; the analysis includes both the single application case and the multiple application case, where several connections a time share the satellite link.

KEY WORDS: satellite communications; network protocols; TCP/IP

1. INTRODUCTION

Owing to the inherent broadcast capability of satellites, which can connect remote sites when there is no terrestrial infrastructure, offering, at the same time, relatively high-speed channels, there is a growing interest in providing interconnection and multimedia services by using satellite links (see Reference [1], for an overview about the significant role of satellites in the next future).

* Correspondence to: Mario Marchese, CNIT—Italian National Consortium for Telecommunications, DIST—University of Genova, Via Opera Pia 13, 16145 Genova, Italy.
† E-mail: dama@dist.unige.it

Contract/grant sponsor: Italian space Agency; contract/grant number: ASI-ARS-99-205

Several projects have been dedicated to this issue; we mention here just a few. Most of the research reported in the paper has been developed in the framework of Reference [2], which is a project, called ASI-CNIT in the following, aimed at analysing the problems related to a satellite or terrestrial/satellite interconnection, from the point of view of both transmission and network protocols. The project is funded by the Italian Space Agency (ASI), and managed by the Italian Consortium for Telecommunications (CNIT), a research centre composed by 18 Italian universities. It is divided into two integrated lines: an experimental activity of multimedia services over a terrestrial-satellite network and a study activity for the system evolution concerning network protocols, integration of satellite and terrestrial networks, both wired and wireless, medium access techniques, resource allocation, development of terminal equipment and user interfaces. Adami *et al.* [2] uses the ITALSAT satellite, and it works at 2 Mbits/s with an antenna of 1.8 m in the Ka-band (20–30 GHz). The project is strictly related to other European activities planned by the European Space Agency (ESA), as, for instance, DICE (Direct Inter-establishment Communications in Europe) [3], proposed as a multi-point videoconference system via satellite. The DICE system can connect different sites and guarantee image and voice reception without any switching operation. It works at 384 kbits/s and uses the Eutelsat satellite with a VSAT antenna of 2.4 m. Another ESA activity is the CODE network (Co-operative Data Exchange), which is a VSAT system dedicated to LAN inter-connection with a user bit rate of 64 kbits/s. CODE has a mixed topology, with a downlink channel bit rate of 2 Mbits/s. Some experiments have been carried out on the Ka-Band (20–30 GHz) with the Olympus and the DFS Kopernikus satellites, in cooperation with Deutsche Bundespost (see References [3] for an overview and [4] for more details). The European program about the development of 'Switched VSAT' [5] is also worth mentioning. In this system, ATM will be the reference technology (see Reference [6] for a list of related projects) for the provision of different services, ranging from the interconnection of LANs to videoconference services. Concerning the next future, the Advanced Research in Telecommunications Systems element 3 (ARTES-3) program (see References [7] for an overview of the ESA programs including ARTES-3 and [8], for more details about the outline, the ongoing projects and the documentation of the ARTES-3 initiative) should be mentioned, where a research concerning multimedia satellite systems and novel pioneering systems, strictly related to the research activity of the ASI-CNIT project, will be carried out.

It is very important to highlight the research activities on satellite organized by NASA (summarized in Reference [9]); the research topics include: MPEG-2 over ATM over satellite, TCP over ATM QoS and, from the protocol point of view, TCP over satellite. Concerning this last, of special interest is the NASA's ACTS experiments program [9] and, in particular, the experiments identified as 118*x* [10], whose objectives, in part, were 'to promote the development of interoperable, high-performance, TCP implementations across multiple computing/ operating platforms' [10], for the efficient delivery of advanced services over satellite. The Advanced Communications Technology Satellite (ACTS) has been used to perform the tests. The mentioned activity is, at least concerning the internetworking, very similar to the topics addressed by [2] and by this paper. Other interesting projects, where the object is the development of new technology for high-speed satellite systems and of new services, are investigated (see Reference [11], for an example) by the Communications Research Laboratory (CRL) and by NASDA [12].

Some projects, or parts of them, as in References [2,8,10] case, are dedicated to satellite interconnection using the TCP/IP suite, whose adaptation to a satellite environment is the subject of this paper. TCP/IP protocols are used in the Internet and are the base for any Internet application (as file transfer, www, and the like) and for any further internetworking within the future Global Internet using satellite links. The characteristics of a satellite network are very different from wired or wireless networks some protocols have been originally designed for. As a consequence, the performance of protocols may be very different. For instance, the delay of a remote login or file retrieval in a satellite environment may be unacceptable for the user. So, when constructing a network containing a satellite link, or, as in this paper, when using a single satellite link to connect remote sites, there are many considerations to take into account and many points where to act [13]. Among others, the data link protocol, where retransmissions could heavily affect the performance of the protocols above it and of the whole network; the application protocol, whose implementation, as shown in Reference [14], is really important to improve the performance; the configuration of the interconnection tools (e.g. routers), as regards the scheduling algorithm implemented and the buffer dimensions. Another important issue in this environment is the performance of the transport protocol (see Reference [15] for one among the first works touching this issue). If the transport protocol does not offer high performance the network throughput may become really low and, as a consequence, the quality perceived by the users may be poor.

The object of the paper is the investigation of the transport protocol and the proposal of some modifications to improve performance. TCP [16] has been chosen as transport protocol due to the application environment described above.

The problem is currently investigated [17–23]. Along with the already-mentioned projects, an IETF Working Group is completely dedicated to it [24]. A review of the problems due to the utilization of TCP in a satellite environment is presented in Reference [13], while an overview of the research performed in the field may be found in References [18,25]. The proceedings of Reference [26] include a very interesting session about 'TCP enhancement' where References [27–29] list the main limitations of the TCP over satellite and propose some methods for performance improvement.

This paper, which contains an extension of the activity presented in Reference [30] and new results, studies TCP performance by tuning two TCP parameters, namely the buffer dimension and the initial congestion window. The analysis itself allows to suggest some solutions and to take decisions. The behaviour has been investigated by using an ftp-like application suited to transfer files of various dimensions. Both the single and multiple connection case are evaluated. The experimental environment is based on a real testbed, where two hosts are connected through a Geostationary Orbit (GEO) [31] satellite link by means of routers, and on a satellite network emulator.

The paper is structured as follows. Section 2 contains the description of the main problems related to TCP in a satellite environment. Section 3 describes the test network and the application used to get the results. The proposals about the modified TCP and the performance analysis are contained in Section 4. Section 5 presents the conclusions and some possible future activities.

2. TCP/IP IN A SATELLITE ENVIRONMENT

As stated in the Introduction, the use of the TCP/IP stack in satellite environments is really important: many applications are based on this protocol family and any connection with the Internet of the future would imply the use of TCP/IP. The network used heavily affects the behaviour of the transport protocol; even in a satellite environment, the problems are different if a LEO (Low Earth Orbit), a MEO (Medium Earth Orbit) or a GEO satellite system is used [31].

Issues related to each environment are listed in Reference [13]. A geostationary system, where 'there is an inherent delay in the delivery of a message over a satellite link due to the finite speed of light and the altitude of communications satellites' [13] is approached in this work and the round trip time (RTT) is above 500 ms. TCP has not been designed for these network characteristics: the high delay to receive acknowledgements decreases performance and makes the quality perceived by the users really poor. On the other hand, there are also positive aspects: the RTT is approximately constant, the connectivity is guaranteed in geostationary systems (not in LEO, for instance), the transmission errors measured are very low, at least in the testbed used. Moreover, if the application field has a limited extension (i.e., a single satellite link or a small number of hops), as in the testbed considered in this paper, the characteristics of the links and of the devices traversed are well known, and the congestion aspects may be treated differently than in a large cabled network, where the number of devices in the path is potentially large and difficult to control. So, even if the performance of TCP is poor, improvements can be obtained by properly tuning some parameters and modifying algorithms. The decision of taking into account a single link may seem a limitation. It is surely true that the presence of a long heterogeneous path may heavily influence the protocol behaviour, but it is also true that the satellite part might receive a different treatment and attention with respect to the cabled parts of the networks; methodologies as TCP splitting [18,23,32,33] and TCP spoofing [18,33] bypass the concept of end-to-end service by either dividing the TCP connection into segments or introducing intermediate gateways, with the aim of isolating the satellite link.

The problem of enhancing TCP performance in a satellite environment is approached in two ways in the literature. Either the application level is modified to better adapt to the network characteristics [14], or the transport protocol itself is revised, which is the approach followed in this paper.

A short summary of the TCP characteristics related to the paper's topic may be of help before introducing the testing environment and the protocol modifications presented.

A NewReno TCP is used under the 2.2.1 version of the Linux kernel. The parameters are substantially set following the standard in References [34,35]. The notation used herein has been introduced in Reference [34].

The transmission begins with the slow start phase, where the congestion window (cwnd) is set to 1 segment ($1 \cdot$ smss, where smss, measured in bytes, stands for sender maximum segment size, which is bottlenecked by the Ethernet link in the testbed, as shown in the next section), and the slow start threshold (ssthresh) to a very high value (infinite). With each received acknowledgement (ack), cwnd is increased by $1 \cdot$ smss. If the value of cwnd is less than ssthresh, the system uses slow start. Otherwise, the congestion avoidance phase is entered, where cwnd is incremented by $1 \cdot$ smss at each RTT. More precisely, cwnd is increased by $1 \cdot$ smss after receiving a number 'cwnd' of acknowledgements. If there is a loss, a packet is considered lost after three acks that carry the same acknowledgement number (duplicated acks), the system enters the fast retransmit/fast recovery algorithm and performs a retransmission of the missing segment, without waiting for the retransmission timer to expire. In the TCP version used by the 2.2.1 Linux kernel, ssthresh was set to cwnd/2. The tests have been performed by setting ssthresh to the maximum between FlightSize/2 and $2 \cdot$ smss, as indicated in Reference [34], where FlightSize is the measure (in bytes) of the amount of data sent but not yet acknowledged, i.e., the packets still in flight. The cwnd is set to (ssthresh $+ 3 \cdot$ smss). When the error is recovered, i.e., when the lost packets have been successfully retransmitted, the value of cwnd is set to ssthresh. The real transmission window (TW) is set, in any case, to the minimum between cwnd and the receiver's advertised window (rwnd), if the buffer

Table I. TCP parameters.

TW = min{cwnd, buf}	
Slow start [cwnd < ssthresh]	cwnd = 1 · smss ssthresh = ∞ ACK → cwnd = cwnd + 1 · smss
Congestion avoidance [cwnd ⩾ ssthresh]	⟨cwnd⟩ACK → cwnd = cwnd + 1
Fast retransmit/recovery	ssthresh = max{FlightSize/2, 2 · smss} cwnd = ssthresh + 3 · smss Delayed ACK → cwnd = cwnd + 1 · smss cwnd = ssthresh

space of the transmitter does not represent a bottleneck. In this last case, the transmission buffer governs the transmission speed. The receiver window rwnd has been measured to be 32 kbytes at the beginning of the transmission. The receiver buffer space is automatically set by TCP to 64 kbytes. The selective acknowledgment (SACK) mechanism is utilized [35].

In a more schematic way the procedures listed above may be summarized in Table I; a C-like language is used for the description.

As stated above, in a geostationary environment, the mechanism described takes a long time to recover errors. The propagation delay makes the acknowledgement arrival slow and cwnd needs more time than in cabled networks to grow. If, for example, just one segment was sent, it takes at least one RTT to be confirmed. The throughput is very low, even in the slow start phase. This heavily affects performance of the applications based on TCP, as should be clear from the measurements reported in the next sections.

The paper proposes a study and a tuning of the receiver/transmitter buffer space and of the initial congestion window. The TCP window scale option [36] is used to allow enlarged buffer values be effective. The other algorithms are untouched. The performance metrics considered are the throughput and the overall time required for transmission.

3. TEST NETWORK

The results presented in the following have been obtained by using two different tools: a real testbed, described in Section 3.1 and a satellite network emulator, which is used to simulate the satellite network, described in Section 3.2. The application used to perform the tests is described in Section 3.3.

3.1. Real testbed

The real testbed is shown in Figure 1: two remote hosts are connected through a satellite link by using IP routers. The TCP/IP protocol stack is used. The data link level of the router uses HDLC encapsulation on the satellite side, where a serial interface is utilized, and Ethernet on the LAN side. The main characteristics of the radio frequency (RF) devices used for the tests are contained in Table II; more details are described in Reference [37], while detailed measurements of the

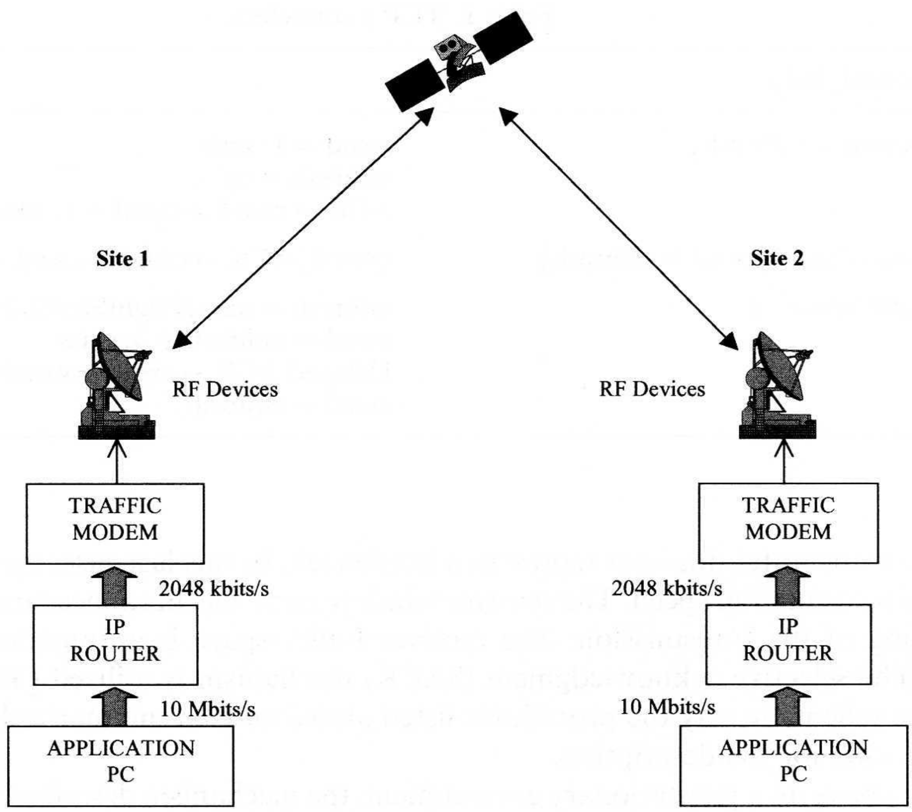

Figure 1. Testbed.

satellite channel characteristics are reported in Reference [2]. A raw bit error rate—BER (i.e. BER with no channel coding) approximately of 10^{-2} has been measured; the utilization of a sequential channel coding with code rate $\frac{1}{2}$, to correct transmission errors, has allowed to reach a BER of about 10^{-8}. As a consequence, the data link protocol 'sees' a reliable channel. The system offers the possibility of selecting the transmission bit rate over the satellite link and a bit rate of 2048 kbits/s has been used for the tests.

3.2. *Network emulator*

The satellite environment described in the previous section has been also emulated by using three PCs connected as in Figure 2: two of them (PC1 and PC3) represent the remote hosts; PC2 emulates the satellite channel and imposes a delay corresponding to the real measures on the satellite link (about 250 ms each hop) and a bandwidth of 2048 kbits/s. The application program NistNet is installed in PC 2. NistNet allows to impose the mentioned delay and the bandwidth bottleneck to adapt to the real situation. The emulator provides accurate results and may be successfully used to substitute the real testbed when it is not available or when the tests would be too long and expensive to be performed in the real environment. It has actually been used to tune parameters and to complete the analysis (see Sections 3.3 and 4 below).

3.3. *Test application*

The application used to get the results is a simple ftp-like one, i.e., a file transfer application located just above the TCP. It allows to transfer data of variable dimension (H (bytes) in the

Table II. RF device characteristics (SSPA stands for solid-state power amplifier).

Antenna diameter	1.8 m
TX gain	52 dB
RX gain	48.5 dB
SSPA power	37 dBm (5 W)
Polarization	V
TX centre frequency	29.75 GHz
RX centre frequency	19.95 GHz

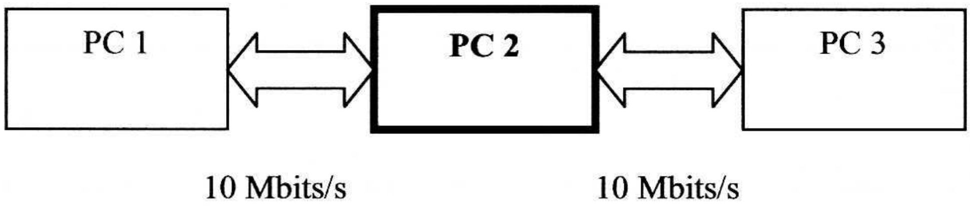

Figure 2. Emulation testbed.

following) between the two remote sites. The application designed allows to transfer a single file at a time, which is the case mostly reported in the literature, both as a benchmark for working and as a configuration used in real environments [14]. Also in this paper, most part of the analysis is performed by transferring one variable dimension file. Three types of files have been utilized to perform the tests and to study the behaviour of the modified TCP: a big file of 11 Mbytes ($H = 11$ Mbytes), a file of relevant dimension ($H = 2.8$ Mbytes) and a small file ($H = 100$ kbytes). The choice among them is indicated in the text in Section 4. Tests with files of 2.8 Mbytes and 100 kbytes have been performed through the real testbed, while the emulator has been used for the tests with $H = 11$ Mbytes. The values obtained with the emulator should be considered from a qualitative point of view; they cannot be quantitatively compared with the measured values since, even if the tuning of the emulator has been accurate, it cannot completely match the real testbed. The multiple connections case, reported in Section 4 to show the effect of a loaded network on the modified TCP, is obtained by activating consecutively a fixed number (N, in the following) of connections. Due to the high speed of the activation, the effect is loading the testbed with N connections at the same time, even for transfers of short files. The multiple connections case is investigated by using only the real testbed.

It should be highlighted that most of the Internet applications, as browsing, are mainly based on file transfer. A quick and secure file transfer guarantees a high quality of service to the users. Many network applications, e.g. distance-learning, use massive file transfer. The work has taken this application as a reference because it is thought as fundamental for most applications of interest.

4. RESULTS

The analysis is dedicated to investigate the behaviour of TCP by varying the value of the initial window (IW, measured in bytes, as indicated in Section 2; i.e. the notation IW $= 1$ means

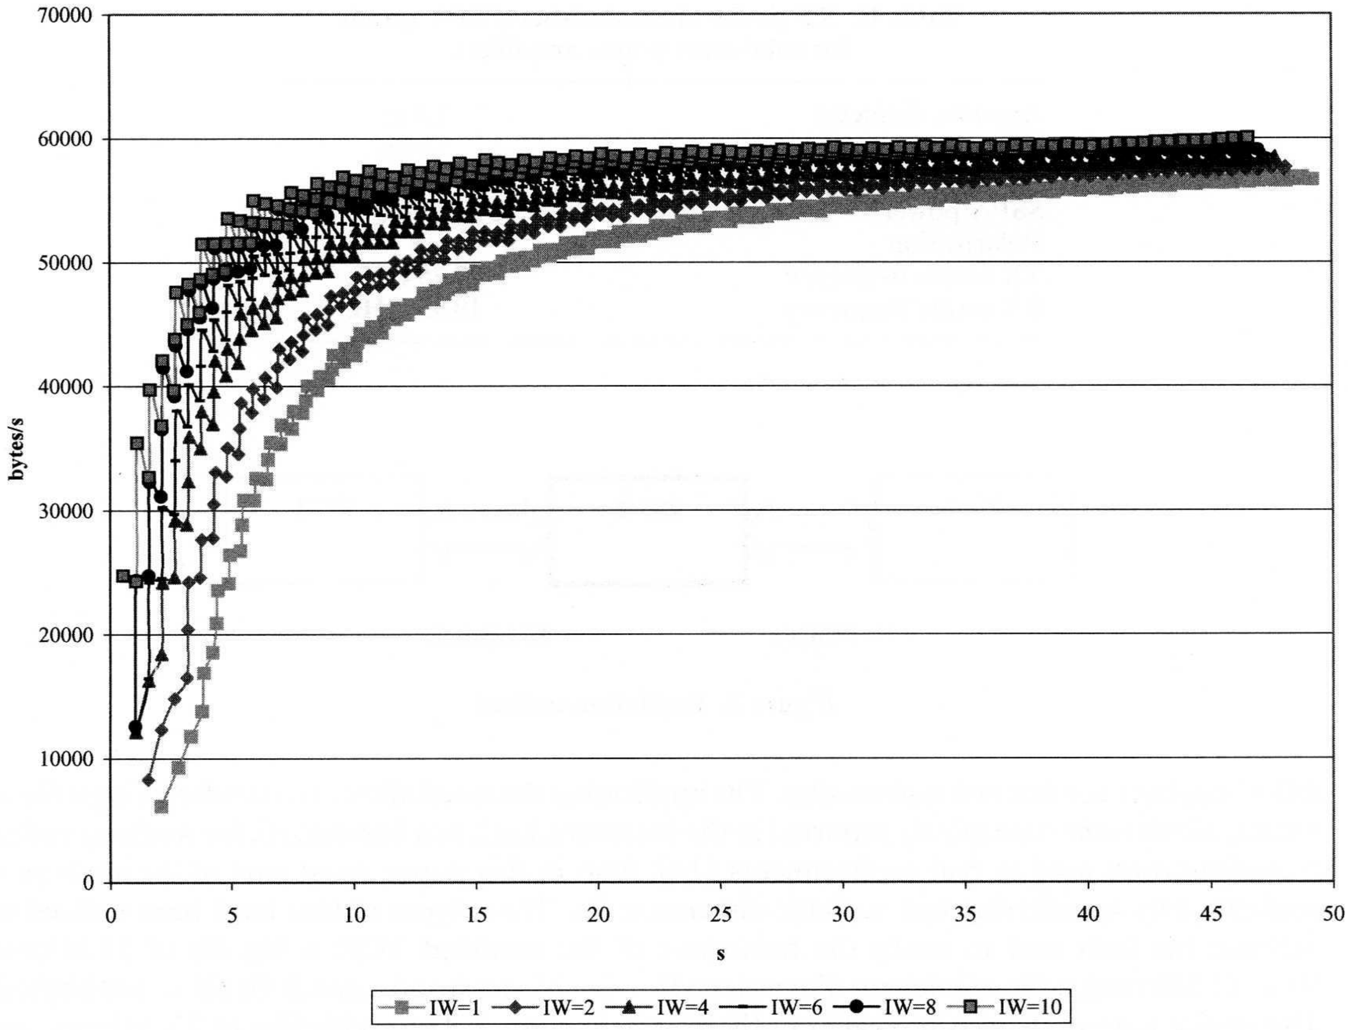

Figure 3. Throughput (bytes/s) versus time for different values of the initial congestion window (IW), buf = 64 kbytes, $H = 2.8$ Mbytes.

IW $= 1 \cdot$ smss) and of the buffer dimension, intended as the memory availability in bytes, for source and destination, which is kept equal; i.e. the buffer has the same length both for the source and the destination. It is identified with the variable 'buf' in the following.

Concerning the initial congestion window, the issue has been treated in the literature. Simulation studies, though not for the specific satellite environment [38], show the positive effect of increased IW for a single connection. Poduri and Nichols [38] clarify the strict dependence of the performance on the application environment and suggests that 'larger initial windows should not dramatically increase the burstiness of TCP traffic in the Internet today'. IW is set to 1 in classical TCP, as stated in Section 2. The performance improvement (i.e. the reduction of the time required for the whole transmission and the higher throughput) provided by varying the value of IW is shown in Figure 3 and Table III for a buffer (buf) of 64 kbytes, which is the value set by TCP if no modification is imposed. A file transfer of 2.8 Mbytes is performed in this case.

The improvement obtained by increasing the initial window may be noted, both for throughput and for the time needed to complete the transmission. The overall transmission time is evidenced in Table III, along with the gain in percentage with respect to TCP taken as a reference

Table III. Overall transmission time for different values of the initial congestion window (IW), buf = 64 kbytes, H = 2.8 Mbytes.

IW	Transmission time (s)	Per cent gain
1	49.21	–
2	48.13	2.2
4	47.69	3.1
6	47.55	3.4
8	46.99	4.5
10	46.57	5.4

(buf = 64 kbytes, IW = 1). The gain is computed as follows: if T_{REF} is the reference transmission time ($T_{\mathrm{REF}} = 49.21$ s, in Table III) and T is a generic transmission time, the percentage gain is defined as

$$\text{Per cent gain} = \begin{cases} \dfrac{T_{\mathrm{REF}} - T}{T_{\mathrm{REF}}} \times 100, & \text{if } T < T_{\mathrm{REF}} \\ 0 & \text{otherwise} \end{cases}$$

If $T_{\mathrm{REF}} \leqslant T$, there is no gain. For example, referring to Table II, if $T = 46.57$ s, per cent gain $= (49.21 - 46.57)/49.21 = 5.4$ per cent.

The advantage of using an increased IW is more evident for shorter transfers. The transmission time necessary to perform a 100 kbytes transfer is reported in Table IV, for different values of IW, along with the gain in percentage obtained with respect to the reference (buf = 64 kbytes, IW = 1).

The improvement obtained is noticeable. The best results for shorter files are due to the initial window behaviour in the TCP: IW, as will be remarked after the analysis concerning the buffer dimension, is mainly responsible of the behaviour in the first part of the transmission. As a consequence, short file transfers receive a real advantage from an increased IW: it should be observed, in fact, that even if the difference in time is small, it might be really important for specific applications (e.g. remote controls). In both cases, Figure 3/Tables III and IV, there is no congestion and TCP is more aggressive, depending on the initial window; if cwnd, as presented in Section 2, starts from a higher value, more data can be sent without waiting a too long time. The tuning of the initial congestion window allows to mitigate the problem introduced by the channel delay.

As regards the second parameter, Figure 4 shows the behaviour of the throughput (bytes/s) versus time for different values of the buffer length for IW = 1.

The classical TCP, currently used in cabled networks, is labeled as buf = 64 kbytes. It is clear that, in this situation, the increase in speed is very slow, the TCP is drastically blocked by the satellite delay, the transmission window cannot increase its length because the buffer dimension represents an actual bottleneck for the system. The time required to transfer data is shown in Table V, along with the percentage gain with respect to the reference (IW = 1, buf = 64 kbytes). If the buffer dimension is higher, the number of packets in flight is increased, as well as the system

Table IV. Overall transmission time for different values of the initial congestion window (IW), buf = 64 kbytes, H = 100 kbytes.

IW	Transmission time (s)	Per cent gain
1	4.42	–
2	3.76	14.9
4	3.21	27.4
6	2.70	38.9
8	2.66	39.8
10	2.62	40.7

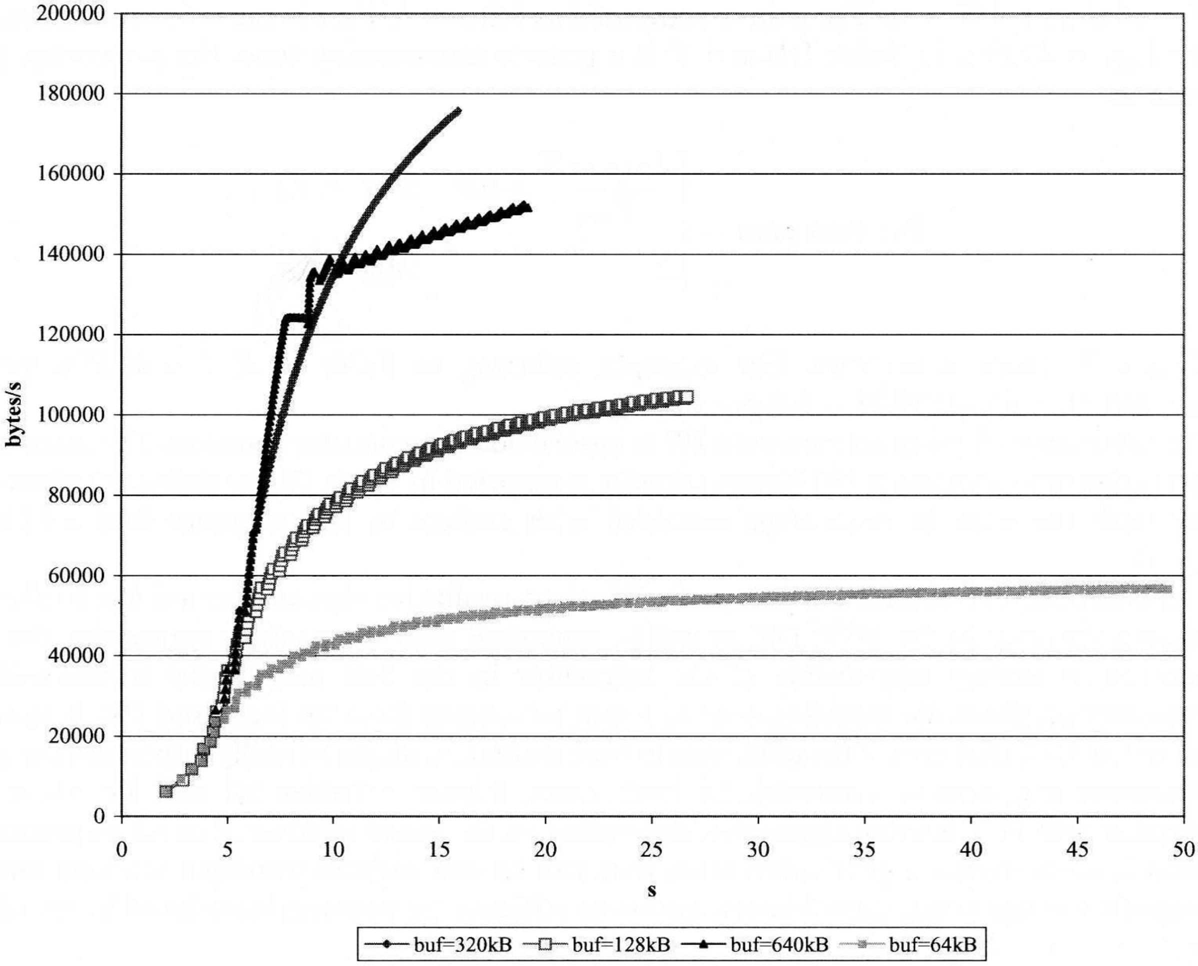

Figure 4. Throughput (bytes/s) versus time for different values of the buffer length, IW = 1, H = 2.8 Mbytes.

performance. The improvement is outstanding, as shown by the transfer time reduced to about 15 s in the best case, corresponding to an improvement of about 68 per cent.

The congestion window represents the network bottleneck until it reaches the buffer length, which happens after few seconds, then the buffer rules the system. This is true if there is no

Table V. Overall transmission time for different values of the buffer, IW = 1, $H = 2.8$ Mbytes.

buf (kbytes)	Transmission time (s)	Per cent gain
64	49.21	–
128	26.67	45.80
320	15.87	67.75
640	26.67	61.03

congestion. Due to the presence of just one active connection, congestion can be generated only by a too aggressive behaviour of the TCP; in this case, the router between the host and the modem saturates. This behaviour can be evaluated by observing the line labeled as 640 kbytes. The congestion window increases exponentially for about 7 s and, before the buffer length represents the bottleneck, the system saturates and the fast retransmit strategy is applied. At this point, the increase is linear, due to the congestion avoidance scheme.

The buffer length is very important for the performance of the system; it rules the congestion window by imposing a bottleneck to its increment. A short buffer drastically limits performance, but an excessively long buffer makes the system congested. When the system is congested, the throughput is strongly reduced even if the efficiency is high at the beginning of the connection. The congestion issue deserves a particular attention because, even if the qualitative behaviour does not change, the specific measures heavily depend on the network devices' configuration. In more detail some TCP segments might be lost 'due to the inability of the router to handle small bursts' [39]. The router ability of managing bursts is heavily affected by the buffer dimensions in the router, by the buffer management strategy implemented, by the presence, or not, of a bias against bursty traffic and, as a consequence, by the queue discipline used. No parameter tuning has been realized in the routers used to perform the tests; default configurations have been maintained. Even if the design of a commercial tool where the strategies suggested in the paper were implemented would require a precise optimization and the knowledge of every configuration detail, the generality of the research is not affected by the particular configuration used. The observation is due to the possibility of obtaining different quantitative measures if the same testbed is utilized with a different router configuration. Similar observations may be reported concerning other network devices (e.g., modems).

The effect of congestion, already envisaged in Figure 4, is evidenced in Figure 5, where the throughput versus time for different values of the initial congestion window (IW) is shown for a value of buf of 640 kbytes and $H = 2.8$ Mbytes.

In this case there is the combined action of increasing both the initial window and the buffer length: the network saturates after few seconds. This behaviour results from Figure 6(a) and 6(b), too, where the throughput is shown versus time for different values of buf ($H = 2.8$ Mbytes). IW is set to 6, as this appears to be the best suited value from the previous graphs. The results reported in Figure 6(a) and 6(b) should appear in the same graph, the scales, in fact, have been maintained; they have been divided to allow a better comprehension.

Figure 6(a) and 6(b) allow the detection of the 'ideal' combination of IW and buf values for the configurations taken into account in the paper. Values of buffer over 384 kbytes saturate the intermediate router; a value of 320 kbytes seems to guarantee a very high gain with respect to the classical solution of 64 kbytes and, at the same time, a certain margin against saturation.

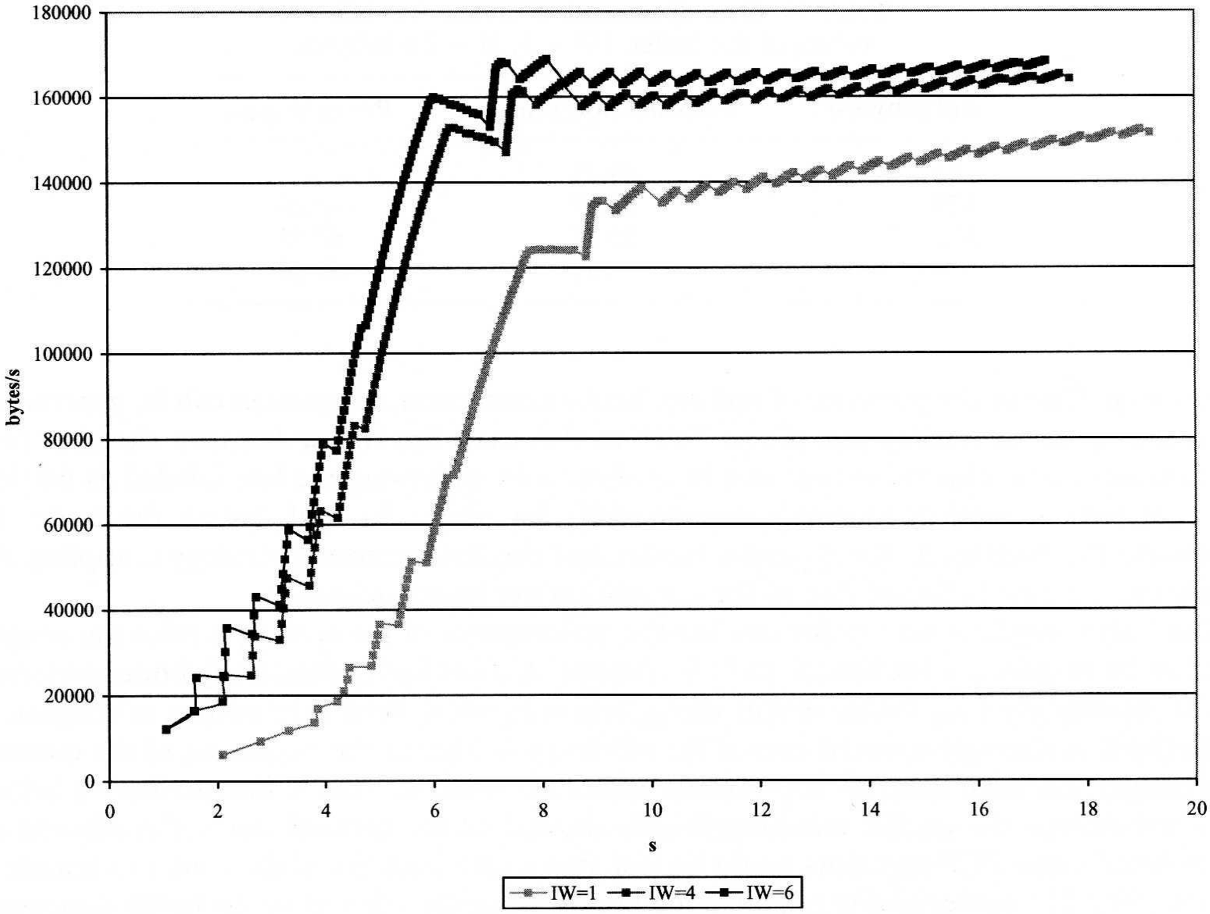

Figure 5. Throughput (bytes/s) versus time for different values of the initial congestion window (IW), buf = 640 kbytes, $H = 2.8$ Mbytes.

The percentage gain of the overall transfer time is reported in Table VI versus the buffer length, up to the value of 384 kbytes, which is the maximum length to avoid saturation, with respect to the (IW = 6, buf = 64 kbytes) configuration. The same quantities, up to 320 kbytes, with the same reference, are reported in Tables VII and VIII for $H = 100$ kbytes and 11 Mbytes, respectively. As is evident from the tables, the gain is light for short files, since the buffer does not represent a bottleneck for the system and IW rules the efficiency, while it is quite meaningful for larger transfers.

Table IX allows to evidence the effect of the parameter tuning performed in the paper. The table contains the combination of the two parameters analyzed (IW and buf), the time required for the overall transmission and the gain in percentage obtained with respect to the basic configuration (IW = 1, buf = 64 kbytes). The measures obtained with the real testbed and $H = 2.8$ Mbytes have been chosen to summarize the results of the paper. The gain in the overall transmission time (up to 71.63 per cent) is mainly due to the buffer length, which may represent a real bottleneck for the system. It has to be remembered that the effective transmission window is the minimum between cwnd and the receiver's advertized window (rwnd), which is strictly dependent on the receiver buffer length. A large buffer guarantees that the bottleneck of the system (concerning the packets in flight) is not so severe: if the buffer is short the congestion window quickly reaches the buffer length. On the other hand, the throughput in the initial phase is governed by IW. It is sufficient to

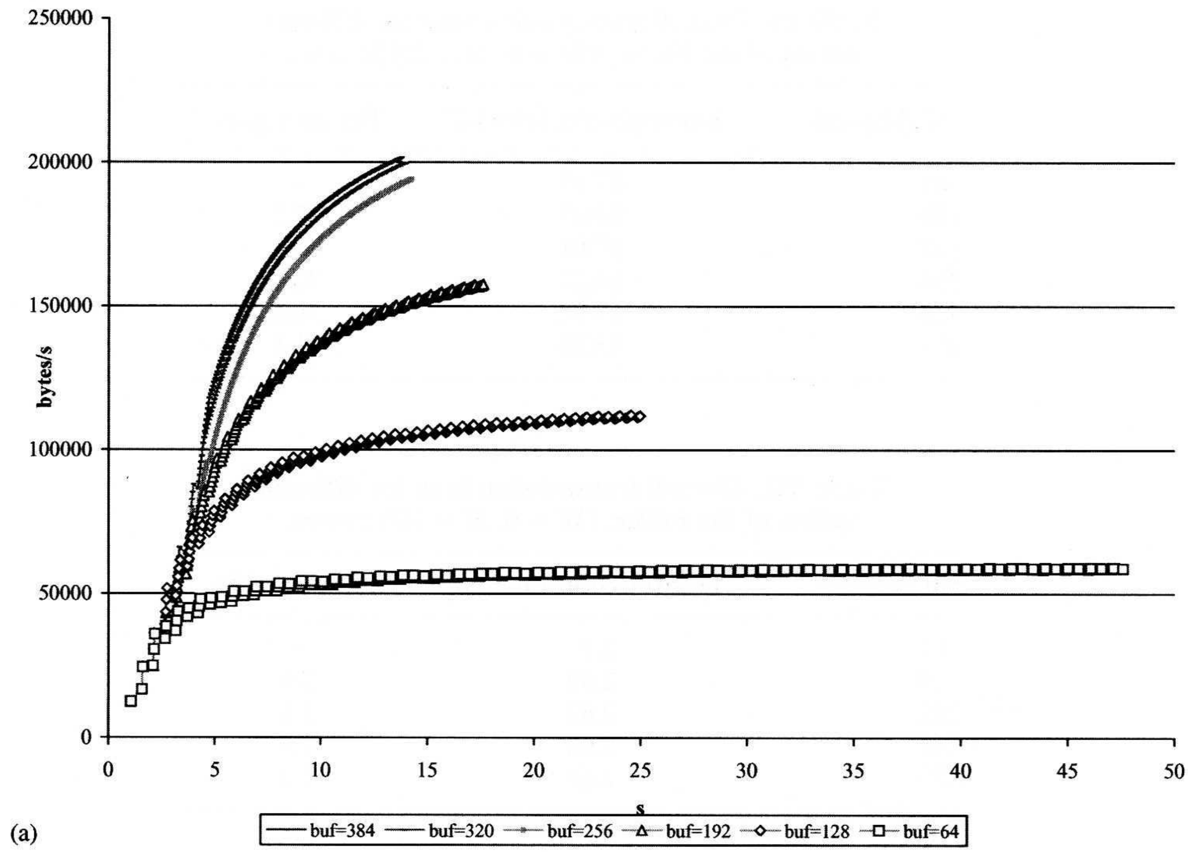

(a)

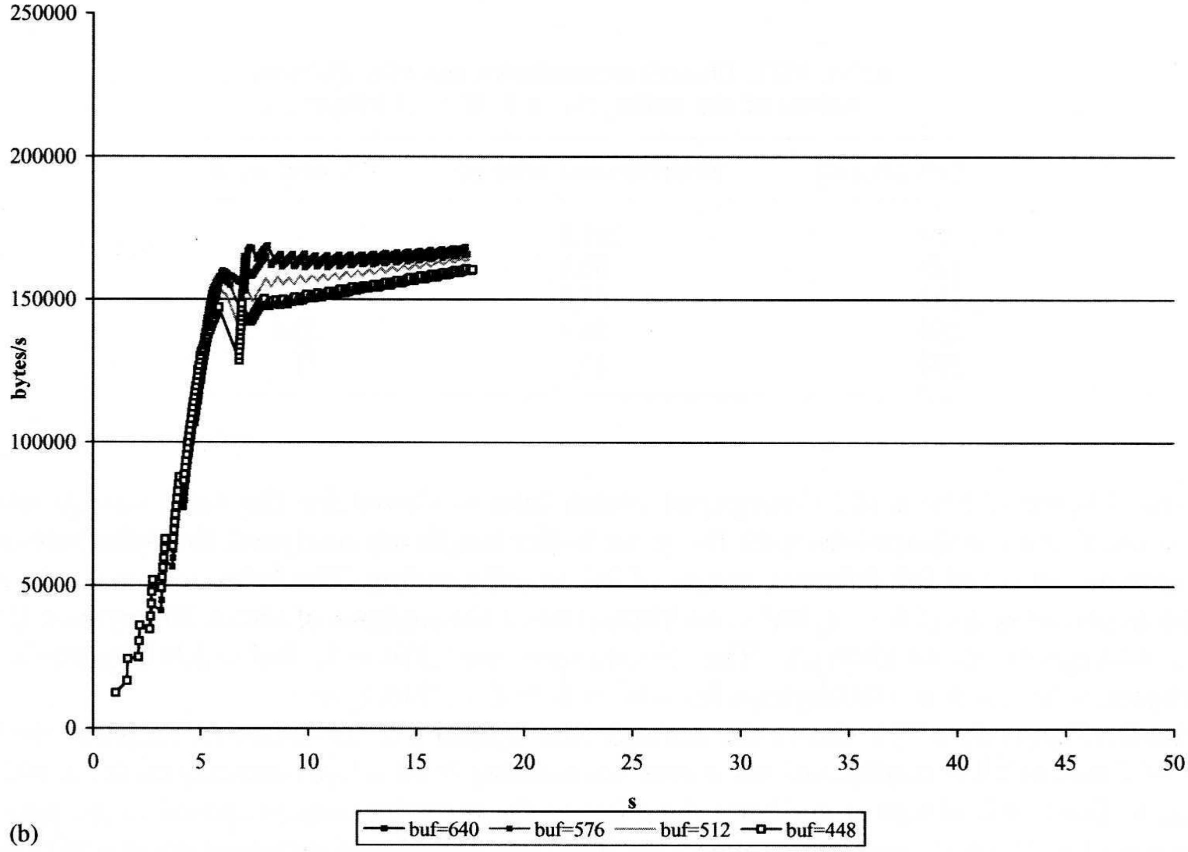

(b)

Figure 6. Throughput (bytes/s) vs time for different values of the buffer length, IW = 6, $H = 2.8$ Mbytes.

Table VI. Overall transmission time for different values of the buffer, IW = 6, H = 2.8 Mbytes.

buf (kbytes)	Transmission time (s)	Per cent gain
64	47.55	–
128	24.97	47.5
192	17.61	63.0
256	14.24	70.1
320	13.96	70.6
384	13.76	71.1

Table VII. Overall transmission time for different values of the buffer, IW = 6, H = 100 kbytes.

buf (kbytes)	Transmission time (s)	Per cent gain
64	2.7	–
128	2.69	0.4
192	2.67	1.1
256	2.66	1.5
320	2.65	1.9

Table VIII. Overall transmission time for different values of the buffer, IW = 6, H = 11 Mbytes.

buf (kbytes)	Transmission time (s)	Per cent gain
64	191.1	–
128	95.4	50.1
192	64.8	66.1
256	56.1	70.6
320	55.5	71

observe Figure 7, where the throughput versus time is shown for the same configurations of Table IX. If the configurations with the same buffer length are analyzed, the difference between the increase in speed for different values of IW is outstanding. The behaviour after 5 s may be taken as an example: (IW = 1, buf = 64 kbytes) has a throughput of about 26 kbytes/s, (IW = 6, buf = 64 kbytes) of 48 kbytes/s. The throughput for (IW = 1, buf = 320 kbytes) is about 35 kbytes, whereas it is 118 kbytes/s for (IW = 6, buf = 320 kbytes).

Finally, Figure 8, which shows the average throughput per connection, compares the behaviour of different TCP configurations versus the number N of active connections for a 100 kbytes transfer. The test is aimed at verifying the effect of the modifications proposed in the presence of a network loaded with multiple connections. The case considered is typical when a WWW site is shared by several users, who try to download remote data. Four configurations are taken into

Table IX. Comparison of TCP configurations by varying the initial congestion window and the buffer length, $H = 2.8$ Mbytes.

IW, buf (kbytes)	Transmission time (s)	Per cent gain
1, 64	49.21	–
6, 64	47.55	3.37
1, 320	15.87	67.75
6, 320	13.96	71.63

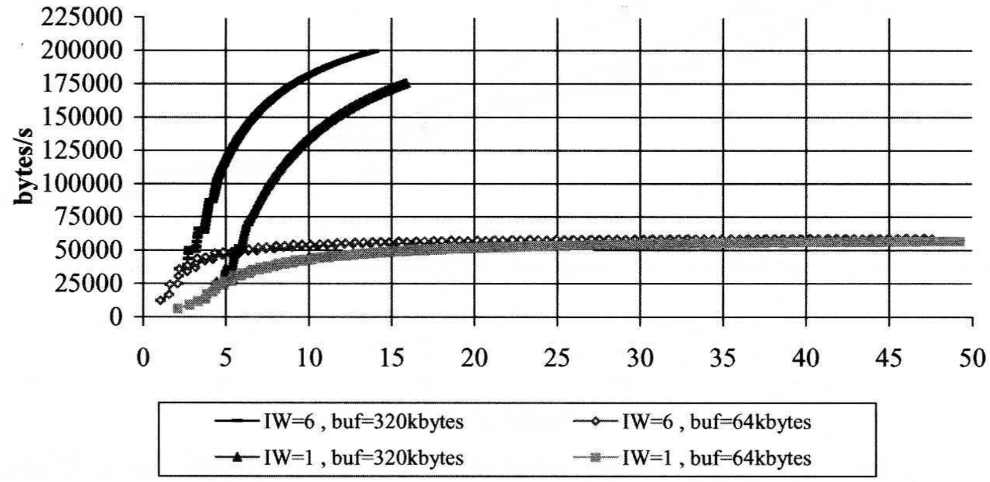

Figure 7. Throughput (bytes/s) vs time for different values of the buffer length and of the initial congestion window, $H = 2.8$ Mbytes.

account: the reference configuration (IW = 1, buf = 64 kbytes); (IW = 1, buf = 320 kbytes), where only the buffer is varied; and two configurations where both buf and IW are increased: (IW = 2, buf = 320 kbytes) and (IW = 6, buf = 320 kbytes).

It is important to note that the performance improvement due to modified values of IW and buf is not lost by the higher traffic load. The gain in throughput is evident up to 15 active connections; for larger values of N, the traffic load due to the number of connections in progress makes the TCP insensitive to the modifications introduced.

5. CONCLUSIONS

The paper has presented an analysis of TCP behaviour by tuning parameters as the initial congestion window and the buffer length. A ftp-like application, i.e. a file transfer application located just above the TCP, has been used to test the system.

The TCP version that is commonly used in cabled networks, taken as a reference in the paper, results inefficient. The delay imposed by the satellite link heavily affects the acknowledgement mechanism and drastically reduces the overall throughput. The time required for the whole transmission is very long and the quality of service perceived by the users is poor.

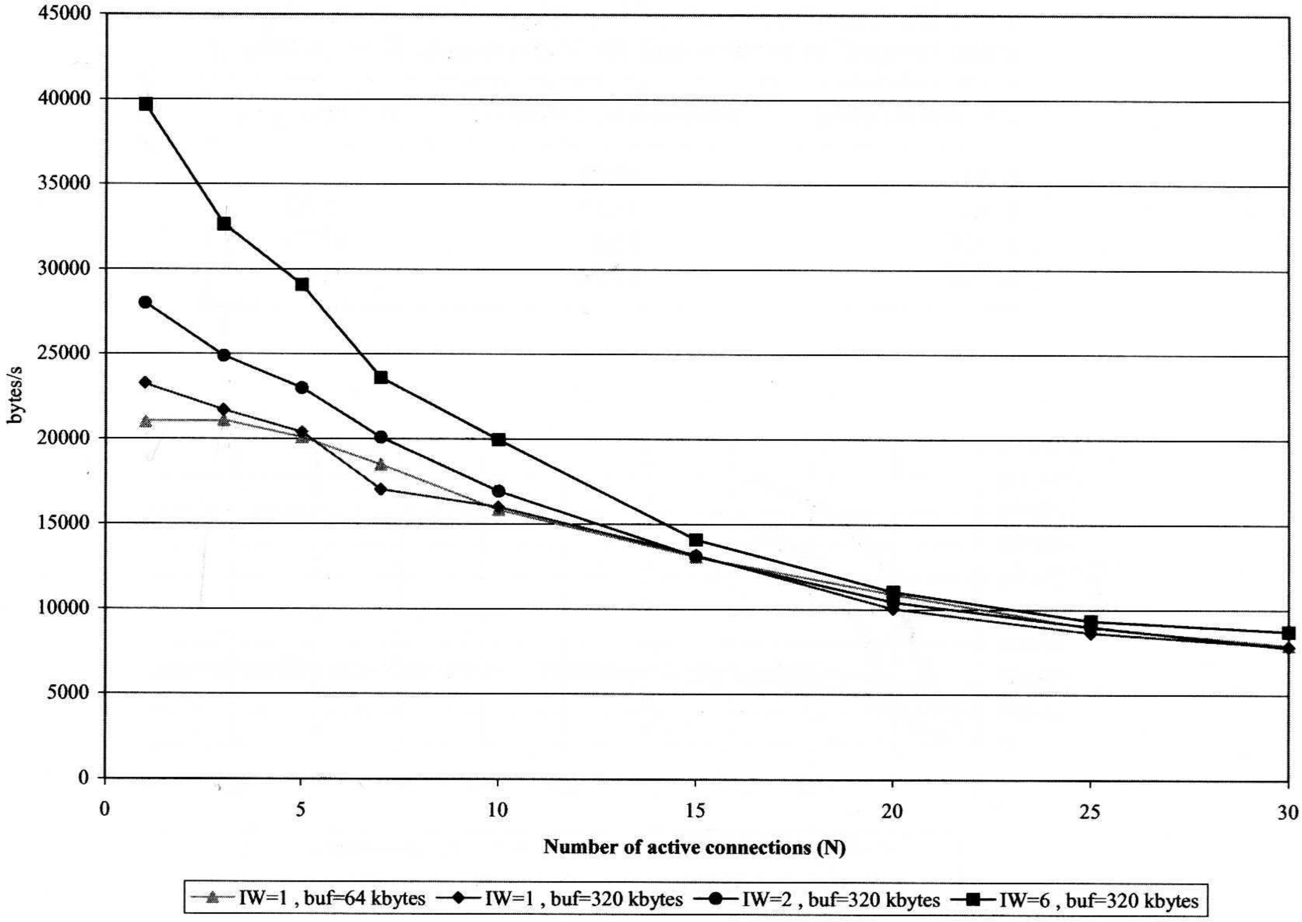

Figure 8. Throughput (bytes/s) versus the number of active connections (N), H = 100 kbytes.

The improvement due to an increased dimension of the transmitter/receiver buffer and of the initial congestion window is outstanding. A gain above 71 per cent is reached with respect to the reference configuration. A proper tuning of the two quantities is strongly necessary to get better performance and to avoid congestion. A long buffer dimension reduces the overall transmission time, while the initial congestion window rules the throughput in the first phase of the connection. Moreover, a large buffer, along with an extended initial window, makes TCP more aggressive: on one hand, this improves the network throughput but, on the other hand, if the tuning is not precise, the risk of saturation is high. An 'ideal' tuning may be found. Both the single and multiple connection cases have been taken into account in the paper. The behavioural analysis of slow start and congestion avoidance schemes, as well as of fast retransmit/recovery algorithms will be the object of future research.

ACKNOWLEDGEMENTS

This work has been supported by the Italian Space Agency (ASI) under the CNIT contract ASI-ARS-99-205 'Integration of Multimedia Services on Heterogeneous Satellite Networks'.

REFERENCES

1. Kota S, Jain R, Goyal R. Broadband satellite network performance, Guest Editorial. *IEEE Communications Magazine* 1999; **37**(7):94–95.
2. Adami D, Bozzardi A, Dardari D, Marchese M, Ronga LS. Integration of multimedia services on satellite heterogeneous networks. *CNIT Internal Report* (in Italian).
3. The DICE/CODE Project Page. In http://www-inw.tu-graz.ac.at/dice/welcome.htm (last access 27 October 1999).
4. Fairhurst G, Spracklen CT, Samaraweera N, Ord K. A report on the performance of the IP networking software used by CODE. *Contract Report*, January 1993.
5. ESA Switched VSAT (XVSAT). In http://www.erg.abdn.ac.uk/users/silas/systems/xvsat.html (last access 27 October 1999).
6. Summary Information about ATM Projects. In http://www.erg.abdn.ac.uk/users/silas/systems/index. html (last access 27 October 1999).
7. European Space Agency. Shaping the future of Europe in space: which programmes, which needs? In http://www.esa.int/Info/99/info06.html (last access 27 October 1999).
8. ARTES 3. The ESA Multimedia Initiative, Home Page. In http://www.estec.esa.nl/artes3 (last access 27 October 1999).
9. Ivancic WD, Brooks D, Frantz B, Hoder D, Shell D, Beering D. NASA's broadband satellite networking research. *IEEE Communications Magazine* 1999; **37**(7):40–47.
10. Brooks DE, Buffinton C, Beering DR, William AW, Ivancic D, Zernic M, Hoder DJ. ACTS 118x Final Report, High speed TCP interoperability testing. *NASA/TM*—1999-209272, July 1999.
11. Kadowaki N, Yoshimura N, Ogawa Y, Komiya N, Araki T. Ka-band gigabit communications technology satellite: development status and connection control scheme. *Proceedings of Fifth Ka-band Utilization Conference*, Taormina, Italy, October 1999; 19–25.
12. National Space Development Agency of Japan (NASDA), Home Page. In http://www.nasda.go.jp/ index_e.html (last access 27 October 1999).
13. Allman M, Glover D, Sanchez L. Enhancing TCP over satellite channels using standard mechanism. *IETF*, *RFC 2488*, January 1999.
14. Allman M. Improving TCP performance over satellite channels. *Master Thesis*, Ohio University, June 1997.
15. Pink S. TCP/IP on gigabit networks. In *Measurement, Implementation and Optimization of Internet Protocols. Ph.D. Dissertation*, Royal Institute of Technology, Sweden, October 1993.
16. Information Sciences Institute, University of Southern California, Transmission control protocol—darpa internet program—protocol specification. *IETF*, *RFC 793*, September 1981.
17. Henderson TR, Katz RH. Transport protocols for internet-compatible satellite networks. *IEEE Journal on Selected Areas in Communications* (*JSAC*) 1999; **17**(2):326–344.
18. Partridge C, Shepard TJ. TCP/IP performance over satellite links. *IEEE Network* 1997; **11**(5):44–49.
19. Lakshman TV, Madhow U. The performance of TCP/IP for networks with high bandwidth-delay products and random loss. *IEEE/ACM Transactions on Networking* 1997; **5**(3):336–350.
20. Allman M. On the generation and use of TCP acknowledgments. *ACM Computer Communication Review* 1998; **28**(5).
21. Durst R, Miller GJ, Travis EJ. TCP extensions for space communications. *Proceedings of MOBICOM'96*, New York, USA, November 1996; 15–26.
22. Department of Engineering, Aberdeen University. Monitoring and optimization of CODE protocols. Estec Contract 9980/92/NL/LC, March 1994.
23. Ghani N, Dixit S. TCP/IP enhancement for satellite networks. *IEEE Communications Magazine* 1999; **37**(7):64–72.
24. IETF Working Group tcpsat. In http://www.ietf.org/html.charts/tcp-charter.html (last access 27 October 1999).
25. Allman M, Glover D, Griner J, Heidemann J, Ostermann S, Scott K, Semke J, Touch J, Tran D. Ongoing TCP research related to satellites. IETF, Internet Draft, draft-ietf-tcpsat-res-issues-05.txt, November 1998.
26. Technical Workshop. Internet protocols over satellite, ESTEC, Noordwijk, January 1999, In http://www.estec.esa.nl/artes3/workshop/index.html (last access 27 October 1999).
27. Samaraweera N. Transparent enhancements of TCP for internet service provision over DVB-S networks. Presentation 1 and Presentation 2, In http://www.estec.esa.nl/artes3/workshop/nihal.html (last access 27 October 1999).
28. Hurley P. Performance and optimisation of TCP over satellite links. In http://www.estec.esa.nl/artes3/workshop/paul.html (last access 27 October 1999).
29. van Oosten G. TCP throughput over a satellite link simulator. In http://www.estec.esa.nl/artes3/workshop/gertjan.html (last access 27 October 1999).
30. Marchese M. Modifications to TCP's parameters to enhance performance over satellite channels. *Proceedings of Fifth Ka-band Utilization Conference*, Taormina, Italy, October 1999; 343–350.
31. Maral G, Bousquet M. *Satellite Communications Systems—Systems, Techniques and Technology* (3rd edn). Wiley: Chichester, 1998.

32. Bharadwaj VG, Baras JS, Butts NP. An architecture for internet service via broadband satellite networks. *Center for Satellite and Hybrid Communication Networks Technical Report, CSHCN T.R: 99–12*, In http://www.isr.umd.edu/TechReports/CSHCN/ (last access 27 October 1999).
33. Zhang Y, De Lucia D, Ryu B, Dao SK. Satellite communication in the global internet. In http://www.wins.hrl.com/people/ygz/papers/inet97/index.html (last access 27 October 1999).
34. Allman M, Paxson V, Stevens WS. TCP congestion control. *IETF, RFC 2581*, April 1999.
35. Floyd S, Henderson T. The NewReno modification to TCP's fast recovery algorithm. *IETF, RFC 2582*, April 1999.
36. Jacobson V, Braden R, Borman D. TCP extensions for high performance. *IETF, RFC 1323*, May 1992.
37. Adami D, Conti A, Dardari D, Marchese M, Ronga LS. CNIT-ASI project integration of multimedia services on heterogeneous satellite networks: an overview. *Proceedings of Fifth Ka Band Utilization Conference*, Taormina, Italy, October 1999; 219–226.
38. Poduri K, Nichols K. Increasing TCP'S initial windows. *IETF, RFC 2415*, September 1998.
39. Allman M, Floyd S, Partridge C. Increasing TCP's initial window. *IETF, RFC 2414*, September 1998.

AUTHOR'S BIOGRAPHY

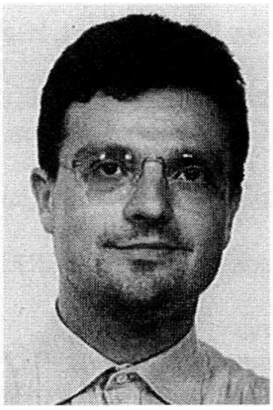

Mario Marchese was born in Genoa, Italy in 1967. He got his 'Laurea' degree cum laude at the University of Genoa, Italy in 1992 and his PhD (Italian 'Dottorato di Ricerca') degree in 'Telecommunications' at the University of Genoa in 1997. In 1998 he got a post-doctoral position at the University of Genoa, Italy. From January 1999 he has been a research scientist at the Italian National Consortium of Telecommunications (CNIT), University of Genoa Research Unit, where he is responsible of the Work Package TCP/IP in the ASI-CNIT project 'Integration of multimedia services in etherogeneous satellite networks'. His research interests include ATM networks, best-effort multimedia networks, TCP/IP protocols, satellite networks, transport protocols for satellite links.