A Parametric Approach to Improve the Performance of the Transport Layer in Satellite Communications

Mario Marchese
CNIT - Italian National Consortium for Telecommunications, DIST - Univ. of Genoa Research Unit, University of Genoa, Italy
Via Opera Pia 13, 16145 Genova (Italy)
Phone: +39-010-3532985, Fax: +39-010-3532154

Abstract - **The paper begins suggesting a possible framework to classify the different approaches to modify the transport layer (namely, the TCP) over satellite channels. In one case, the satellite network is considered as a black box (Black Box - BB approach). Each modification is performed within the end terminals. In the other case, the satellite network is completely known and the transport layer parameter tuning and the consequent optimization of the overall performance is not independent from the choices followed within the satellite network (Complete Knowledge - CK approach). In both cases a complete parameterization of the algorithms at the transport layer is recommended to allow adaptability in different conditions. Concerning the CK approach, a first analysis proposed concerns the behavior of the queue length at IP layer in the intermediate routers. The IP buffer tuning represents a first step towards the implementation of a new transport protocol (STP - Satellite Transport Protocol) adapted to various satellite environments from LEO to GEO, whose interfaces towards the upper layers will keep the characteristics of the TCP interfaces but whose performance will be strongly improved. A possible generic architecture will be reported and some possible solutions to be implemented within the network architecture will be shown along with the results obtained by means of real measures in the field.**

I. INTRODUCTION

The problem of improving the performance of the transport layer (e.g. TCP) over satellite has been investigated in the literature for some years: reference [1] contains a first overview on the topic. More recently, reference [2] provides a summary about improved TCP versions as well as issues and challenges in satellite TCP. Reference [3] lists the main limitations of the TCP over satellite and proposes many possible methods to act. A recent tutorial on the topic is contained in [4]. A recent issue of International Journal of Satellite Communications is entirely dedicated to IP over satellite [5]. In more detail, reference [6] proposes a TCP splitting architecture for hybrid environments (see also reference [7]); reference [8] analyses the performance of web retrievals over satellite and reference [9] shows an extensive analysis of the TCP behavior by varying parameters as the buffer size and the initial congestion window. Reference [10] focuses also on the buffer management but in an ATM environment. Also International Standardization Groups as the Consultative Committee for Space Data Systems - CCSDS, which has already emitted a recommendation (reference [11]) and the European Telecommunications Standards Institute - ETSI [12], which is beginning its activity within the framework of the SES BSM group, are active on these issues.

The main observation, on which the work described in the following is based, is that the proposals of modifications at the transport layer in the satellite environment may be classified into two main frames: the Black Box (BB) approach and the Complete Knowledge (CK) approach. In the BB approach, only the end terminals (namely, the TCP/IP stack at the source and at the destination) may be modified because the satellite network, including intermediate network devices as routers, cannot be accessed. The CK approach allows tuning algorithms and parameters also inside the satellite network.

The paper presents a possible parameterization of the transport layer. The TCP is taken as a reference, the TCP interfaces towards the adjacent layers are kept but the algorithms are made parametrical in view of a future adaptation to various satellite environments as LEO (Low Earth Orbit) or GEO (Geostationary Orbit) channels. Then, a first action in the CK approach is proposed: an intervention on the queue length at IP layer in the intermediate routers. It is thought as a first step toward the implementation of future network architecture (Satellite Protocol Stack - SPS), where the satellite portion of the network is isolated and the protocols are adapted to the channel characteristics. In this context, some possible guidelines to design an efficient transport protocol for satellite environment (STP - Satellite Transport Protocol) to be implemented within the SPS architecture have been proposed and a preliminary performance analysis is presented.

The paper is structured as follows. The next section is dedicated to introduce the 'Black Box' and the 'Complete Knowledge' approaches. Section III summarizes the parameterization proposed for the transport layer. The analysis of the IP buffer is reported in Section IV. Section V contains the proposals concerning the satellite transport protocol. The investigation of the performance is shown in Section VI. The conclusions are reported in Section VII.

II. THE 'BLACK BOX' AND THE 'COMPLETE KNOWLEDGE' APPROACHES

The network configuration used in the tests and also the simplest network used in the literature for experiments is reported in Fig. 1. The box identified as APPLICATION PC may also represent a local area network (LAN). The system employs the satellite ITALSAT II (13° EST). It provides coverage in the single spot-beam on Ka band (20-30 GHz). The overall bandwidth is 36 MHz. Each satellite station can be assigned a full-duplex channel with a bit-rate ranging from 32 Kbits/s to 2 Mbits/s, this latter used in the experiments, and it is made up of the following components: satellite Modem, connected to the RF device; RF (Radio Frequency) Device; IP Router connected to the Satellite modem via RS449 Serial Interface and to the Application PC via Ethernet IEEE 802.310 BASE-T link; Application PC (Pentium III 500 MHz), source of the service under test.

The problem may be considered in two different ways from a network point of view. The first is considering the network as a black box ignoring each particular configuration of the used devices.

This approach has been used in previous works of the same author (e.g. [9] and [13],) and it has been modeled in Fig. 2. The TCP is modified and tuned by acting on the end user terminals. The rest of the network is considered as a black box.

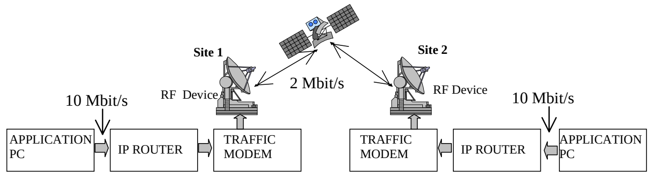

Fig. 1. Test-bed Network.

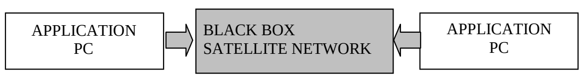

Fig. 2. Black box approach.

An alternative approach is supposing the complete knowledge of each network device (e.g. router, modem, and channel characteristics) and the possibility to modify the configurations to improve the performance of the overall satellite network (or of the satellite portion of the network). The approach is possible if the network is small and proprietary. An example of action to take is the intervention on the IP router (see Fig. 1).

III. TRANSPORT LAYER PARAMETRISATION

The parameters and the notation are substantially set following the standard in [14] and [15]. A C-like language is used for the description. The acronym cwnd stands for congestion window, smss for sender maximum segment size, and ssthresh for slow start threshold. FlightSize is the measure (in bytes) of the amount of data sent but not yet acknowledged, i.e., the segments still in flight. The real transmission window (TW) is set, in any case, to the minimum between cwnd and the minimum between the TCP buffer dimension at the source and the receiver's advertised window (rwnd), which is the half of the receiver TCP buffer length (TW=min{cwnd, min(source buff, rwnd)}). 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 the TCP to 64 Kbytes. The performance metrics considered are the throughput (i.e. the bytes received each second; the effective transmission capacity of the protocol) and the overall time required for the transmission.

TW=min{cwnd, min(source buff, rwnd)}	
Slow Start	cwnd=**IW**·smss; ssthr=**Th** ACK → cwnd=cwnd+**F** (# of received acks, cwnd)·smss
Congestion Avoidance	<cwnd> ACK → cwnd=cwnd+**G**(cwnd, •)
Fast Retransmit/ Recovery	ssthr=max{FlightSize/2, 2·smss}; cwnd=ssthr+3·smss; Delayed ACK → cwnd=cwnd+1·smss; cwnd=ssthr

Table 1. Modified TCP parameters.

Table 1 contains a parameterized version of the TCP, which allows maintaining the same interfaces of the TCP but should be regarded as a new Satellite Transport Protocol. The parameters IW and Th, along with the two functions F(·) and G(·) may be tuned following both the characteristics of the physical channel (delay, loss, bit error rate,…) and the network status (e.g. congestion). TCP commonly used sets: IW=1, Th=∞, F(·)=1 and G(·)=1. The function F(·) is aimed at regulating the size of the congestion window in the Slow Start phase. The characteristics of F(·) affects the increase of the window and, as a consequence, the transmission speed and the protocol performance. Possible proposals concerning the function F(·) for GEO links may be found in reference [13]. The function G(·) is aimed at regulating the behavior of the Congestion Avoidance algorithm. The modification of the congestion avoidance scheme has not provided outstanding results over GEO channels but it might be very useful in LEO or radio-mobile environments.

IV. A FIRST ANALYSIS IN THE CK APPROACH

The TCP connection may be roughly modeled as in Fig. 3, from the memory point of view. The TCP buffer serves the IP buffer at the channel speed (10 Mbits/s, in the test-network of Fig. 1), when there are some segments to serve. The server of the IP buffer, in the model, is the satellite modem, which works at 2 Mbits/s (when data are present). The formal notation (I(t), O(t), O'(t)), not fundamental in this paper, has been introduced to take into account the periods of inactivity. A certain amount of data may be kept in flight inside the pipe represented by the satellite channel (the capacity given by the (Round Trip Time - RTT) x (bandwidth available); e.g. if the GEO system in Fig. 1 is used: RTT (0.511 s) x Bandwidth (2 Mbits/s)). Actually, the memory defined includes the segments in flight plus the segments already at the destination but not yet acknowledged at the source. If there is at least one segment in the IP buffer and one segment at the destination, then O'(t)=O"(t). So, except for the initial phase, when O'(t) works and O"(t)=0, and the final phase of the connection, when the opposite situation happens, O"(t) should be equal to O'(t) for most part of the connection time.

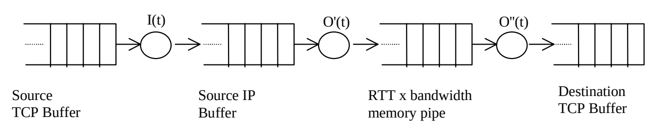

Fig. 3. Memorization model for a TCP connection.

In this context, the choice of the buffer dimensions, in particular the IP buffer length, which is fixed in the BB approach, and the dynamic of the segment arrivals (I(t)) is important. An excessive increase of the IP buffer length provides unpleasant effects, due to the Retransmission Timeout (RTO) of the TCP. A RTO time after sending a segment the TCP sets the Initial Window to 1·smss (IW·smss, in the parameterized case) and begins again the transmission by using the Slow Start algorithm. This situation is represented in Fig. 4, where one transfer of a 3 Mbytes file is shown by using three different configurations of the transport layer and the IP buffer. In more detail, the function F(·) is set to a constant value K=10; the variable IW=6; Th and G(·) are kept as in the TCP commonly used. Two configurations setting a TCP buffer of 320 Kbytes and, respectively, an IP buffer of 60 Kbytes (40 segments of 1500 bytes) and 144 Kbytes (96 segments), have been compared with a 10 Mbytes (TCP buffer) - 6.144 Mbytes (4096 segments, IP buffer) configuration. The configurations have been implemented in the test-bed and the results reported really measured in the field. A huge buffer, as in the (10 - 6.144 Mbytes) case, should guarantee a non-loss behavior. On the contrary, the configuration has a steep initial increase, due to the aggressive function F(·)=K=10. The throughput reaches 600 Kbytes/s after about 4 seconds but, after this phase, the performance dramatically decreases. The motivation is clear by analyzing the traces obtained. A short summary about them is reported in the following: after 4 seconds, the IP buffer contains about 1800 Kbytes (1200 segments). The transmission goes on but the last segments entering the IP buffer stay queued for a long time. After a time RTO, many segments are uselessly re-transmitted and the overall performance (the throughput and the overall transmission time) drastically decreases.

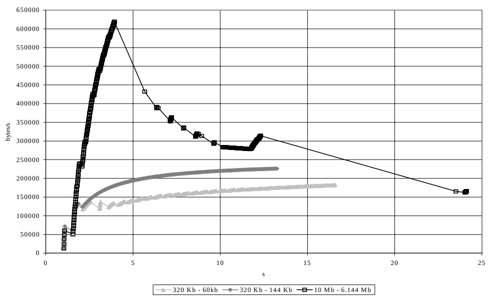

Fig. 4. Throughput vs time, 3 Mbytes, K=10, IW=6, IP-TCP buf.

It is important to remember that the TCP buffer length has been equally set both at the source and at the destination. It means that the window dimension effectively utilized is halved (see the TW setting in the previous section); i.e. it is 160 Kbytes in the 320 Kbytes case 5 Mbytes in the other case. The test reported in Fig. 4 has been obtained with only a connection in the network (mono-connection case). The negative effect may be enlarged by the presence of more than one connection. If the CK approach is adopted, it is possible to set the two values (the TCP and the IP buffer length) 'ideally'. If we suppose a RTT of 0.511 s, the maximum amount of data that can be stored including the segments in flight and the segments not yet acknowledged at the source, is the product between 2048 Kbits/s (the bandwidth available) and 0.511 s (see the model in Fig. 3). As a consequence, a TCP buffer of about 260 Kbytes/s assures a complete utilization of the channel (if a clear sky condition is supposed) and an empty IP buffer. A bigger TCP buffer does not improve the system performance but only makes necessary a longer queue at the IP layer. That is not a problem if the system avoids the RTO problem.

V. Satellite Transport Protocol (STP)

A proper network architecture called Satellite Protocol Stack (SPS) is reported in Fig. 5 and it is based on the work published in the literature about TCP splitting and spoofing [1, 2, 6, and 7]. The transport layer within the Satellite Protocol Stack (SPS) is called Satellite Transport Layer (STL) and implements a Satellite Transport Protocol (STP), suited for the specific environment and based on the CK approach and the parameterization of Section III. Anyway, a complete knowledge of the satellite network implies not only the use of the TCP and IP buffer management, as shown in the simple example of the previous section, but also the use of other information as the total number of active connections or the channel condition. From the protocol layering point of view, the key point is represented by the two Relay Entities, which are two gateways towards the satellite portion of the network. The SPS acts on the satellite links by using the necessary information because it has the knowledge and the control of all the parameters. The Relay Layer guarantees the communication between the satellite transport layer and the protocol used in the cable part (i.e. TCP).

Two possible alternatives may be chosen concerning the transport protocol: bypassing completely the concept of end-to-end service at the transport layer; preserving the end-to-end characteristic of the transport layer. In the first case, the connection at the transport layer is divided into two parts, dedicated, respectively, to the cable and the satellite part. The source receives the acknowledgement from the first Relay Entity, which opens other connections, with different parameters based on the current status of the satellite portion, and allocates the resources available. The Relay Entity on the other side of the satellite link operates similarly towards the destination. The transport layer of the cable portions is untouched. The end-to-end connection may be guaranteed only statistically. The second choice is aimed at preserving the end-to-end characteristic of the transport layer. In this case also the transport protocol in the terrestrial portion should be modified. A possibility may be dividing the transport layer into two sub-layers: the upper one, which guarantees the end-to-end characteristic, and the lower one, which is divided into two parts (as the overall transport layer in the first choice) and interfaces the STL. The terrestrial side of the lower transport layer may be also represented by the TCP. Fig. 6 shows the protocol architecture in this case. The transport layer is modified even if the interface with the adjacent layers may be the same as in the TCP.

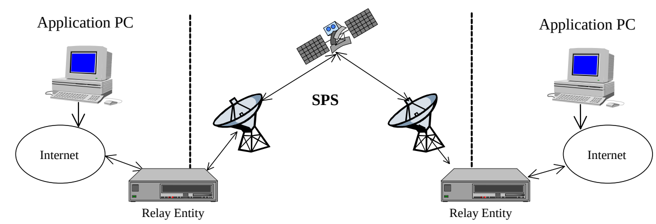

Fig. 5. SPS Architecture

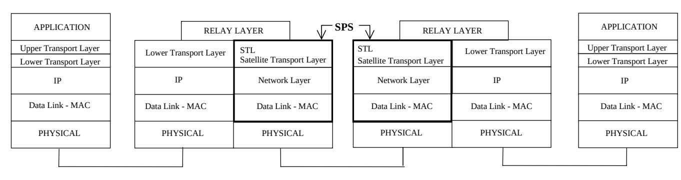

Fig. 6. End-to-end SPS Architecture.

The protocol stack is completely re-designed on the satellite side. The essential information concerning each layer (Transport and Network) of the terrestrial side may be compressed in the Relay Layer PDU (Protocol Data Unit). The Network Layer may use the structure of the IP layer (as in the case presented) but it may be properly designed together with the STL layer, so to optimize the performance of the overall transmission on the satellite side (see sections III and IV). Some more guidelines (concerning the STL and its implementation through the Satellite Transport Protocol (STP)) may be also introduced:

Slow Start Algorithm: the mechanism has no longer need of testing the network congestion at the Relay Entity, because the status is completely known. The algorithm has to rule the flow in accordance with the contemporary presence of other flows, whose characteristics are known. The function $F(\cdot)$, along with the other parameters involved (e.g. IW) should get to the aim. A proper tuning of the IP buffer is fundamental.

Congestion Avoidance Algorithm: the schemes currently used take into account only congestion conditions; a loss is attributed to a congestion event. Now, due to knowledge of the IP buffer status, a loss should be attributed mainly to transmission errors. The function $G(\cdot)$ has the responsibility of this part.

VI. RESULTS

The TCP configuration, adapted to the satellite GEO environment, identified as Modified TCP (Reference), applies an IW=2 and a TCP buffer of 320 Kbytes both at the source and at the destination. This choice is due to the fact that this configuration resulted as one of the most efficient and less dangerous, concerning the congestion risk, in the multi-connection case (see references [9] and [13]). It guaranteed a gain over 70% with respect to the TCP commonly used, IW=1, TCP buffer of 64 Kbytes. The new Complete Knowledge configuration, identified as STP, adapts the parameters to the different situations by choosing the best configurations, including the IP layer buffer tuning, time by time. The comparison is aimed at showing the further improvement of the STP with respect to a modified TCP configuration, already adapted to satellite channels in previous studies. Fig. 7 contains the throughput versus time for the two configurations mentioned and a file transfer of 3 Mbytes. The overall transmission time is 16.05 s for the Reference configuration and 12.54 s for STP. The gain, computed in percentage as 100·(16.05-12.54)/16.05), is 21.9%. STP has a shorter transmission time; thus the performance gain is actually the metric of 'reduction' of the overall transmission time, with respect to the reference configuration. Fig. 8 shows the behavior in the multi-connection case. The throughput in bytes/s is reported versus the number of connections in the network, each of them performing a file transfer of 3 Mbytes, for the Reference configuration and for STP. An improvement is noticeable up to 5 connections. After that the bandwidth available (2 Mbits/s) is not sufficient to match the requirements. The advantage is much more evident if a shorter transfer of 100 Kbytes is performed. Fig. 9 shows the same quantities as in Fig. 8, for a 100 Kbytes file transfer. The Modified TCP Reference configuration, although very convenient with respect to the TCP commonly used, may be strongly improved. The overall transmission time in the mono-connection case is 3.7 s, for the Reference case, and 0.6 s, for STP. It corresponds to a gain of 83.8%. The effect of such improvement in a remote control system (e.g. tele-robot, tele-control) may be simply guessed.

The last part of the results investigates the behavior of the new transport protocol when there are packet losses due to channel errors. The losses have been artificially introduced in the cases reported. The loss has been obtained by shutting down the modem for a fraction of second in the first phase of the connection. The IP router has been properly configured to avoid losses due to congestion. Fig. 10 reports the throughput versus time for a 3 Mbytes transfer in the mono-connection case. The packet loss is much more intense for STP, due to the aggressive behavior in the first phase of the connection, where the shut down happens, but it recovers thanks to the correct interpretation of the losses, which are not due to congestion, as estimated by the Reference configuration. Table 2 contains the gain in the same situation. The last two tables (Table 3 and Table 4) show the overall transmission time for the Modified TCP Reference configuration and the STP, respectively. The tables report the case with losses and with no losses. It is important to note that the Reference case is heavily affected by with presence of losses and this is due to the misinterpretation of the loss cause. The STP is robust and allows keeping a good performance also in the loss case: the difference among the loss and no loss case is only of 13.7 %.

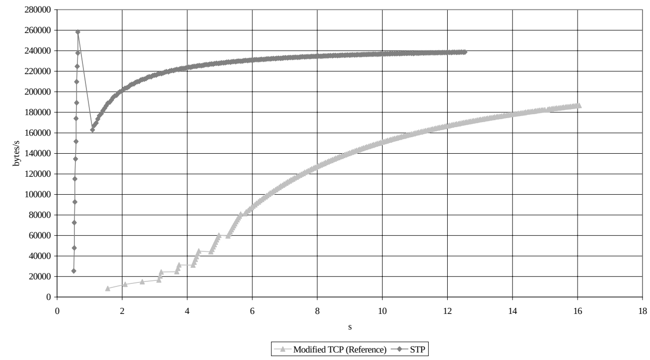

Fig. 7. Throughput versus time, 3 Mbytes, mono-connection.

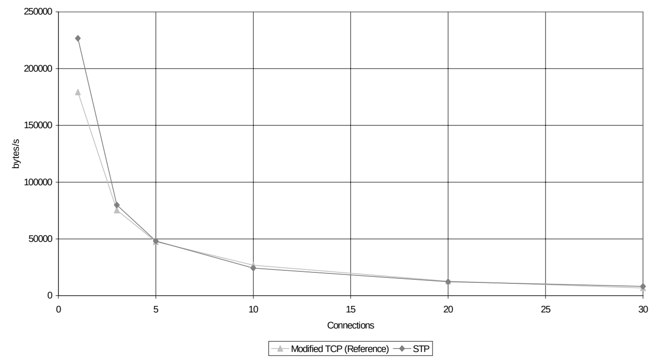

Fig. 8. Throughput versus number of connections, 3 Mbytes, multi-connection.

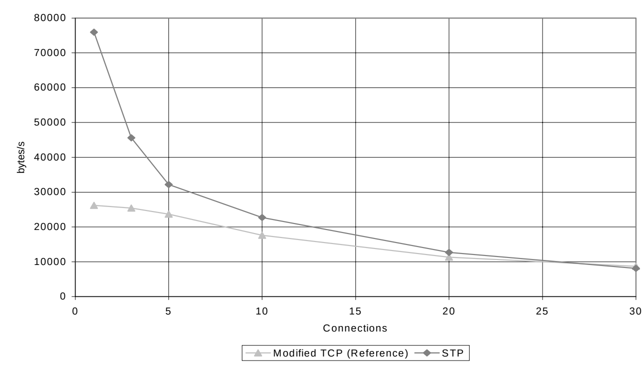

Fig. 9. Throughput versus number of connections, 100 Kbytes, multi-connection.

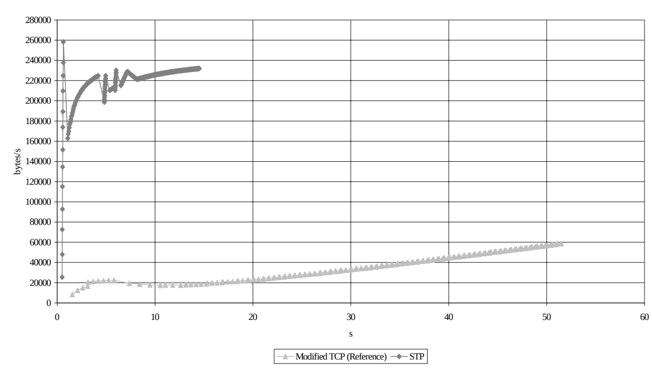

Fig. 10. Throughput versus time, 3 Mbytes, mono-connection, packet loss.

Transport Protocol	Overall Transmission Time [s]	Gain
Modified TCP (Reference)	51.5	-
STP	14.5	71.8 %

Table 2. Overall Transmission Time and Gain, 3 Mbytes file transfer, mono-connection, packet loss.

Modified TCP (Reference)	Overall Transmission Time [s]
Loss	23.6
No loss	13.2

Table 3. Overall Transmission Time, 3 Mbytes file transfer, mono-connection, comparison of loss and no loss, Reference.

STP	Overall Transmission Time [s]
Loss	14.5
No loss	12.5

Table 4. Overall Transmission Time, 3 Mbytes file transfer, mono-connection, comparison of loss and no loss, STL.

VII. CONCLUSIONS

The paper has presented some guidelines and a preliminary performance analysis of a new transport protocol (Satellite Transport Protocol - STP), studied to improve the performance at the transport layer in a satellite environment. The final aim would be to consider most of the possible situations both concerning the characteristics of the channels and of the networks (e.g. LEO, GEO, Radio Mobile).

The solution envisaged is a complete parameterization of the transport layer both to obtain a high degree of flexibility to match the different requirements and to bypass the limitations of the TCP algorithms (slow start and congestion avoidance). The result is an alternative protocol, which is heavily based on the TCP and which maintains the same interfaces towards the adjacent layers. The new protocol is completely transparent to the final user. As preliminary steps to the performance analysis of STP, the paper introduces: a classification of the possible approaches to improve the performance at the transport layer (the Black Box - BB and the Complete Knowledge - CK approach), a first way to operate at the IP layer and a parameterization of the transport protocol along with an architecture (Satellite Protocol Stack - SPS). The new protocol, implemented in the CK approach, is compared with an efficient modified version of the TCP taken from the literature. A GEO test-bed and real tests have been used to get the results.

STP has shown a good efficiency and a relevant improvement both in the mono and in the multi-connection case for all the tests performed. The improvement is outstanding for short file transfers. STP has also shown a meaningful robustness against the transmission errors.

REFERENCES

[1] C. Partridge, T. J. Shepard, "TCP/IP Performance over Satellite Links", IEEE Network, Vol. 11, n. 5, September/October 1997, pp. 44-49.

[2] N. Ghani, S. Dixit, "TCP/IP Enhancement for Satellite Networks", IEEE Comm. Mag., July 1999, pp. 64-72.

[3] M. Allman, S. Dawkinks, D. Glover, J. Griner, T. Henderson, J. Heidemann, S. Ostermann, K. Scott, J. Semke, J. Touch, D. Tran, "Ongoing TCP Research Related to Satellites", RFC2760, February 2000.

[4] C. Barakat, E. Altman, W. Dabbous, "On TCP Performance in a Heterogeneous Network: A Survey", IEEE Comm. Mag., Vol. 38, No. 1, January 2000, pp. 40-46.

[5] A. Ephremides, Guest Editor, International Journal of Satellite Communications, Special Issue on IP, Vol. 19, Issue 1, January/February 2001.

[6] V.G. Bharadwaj, J. S. Baras, N. P. Butts, "An Architecture for Internet Service via Broadband Satellite Networks", International Journal of Satellite Communications, Special Issue on IP, Vol. 19, Issue 1, January/February 2001, pp. 29-50.

[7] Y. Zhang, D. De Lucia, B. Ryu, Son K. Dao, "Satellite Communication in the Global Internet", http://www.wins.hrl.com/people/ygz/papers/inet97/index.html (last access August 22th, 2000).

[8] H. Kruse, M. Allman, J. Griner and D. Tran, "Experimentation and modelling of HTTP over satellite channels", International Journal of Satellite Communications, Special Issue on IP, Vol. 19, Issue 1, January/February 2001, pp. 51-69.

[9] M. Marchese, “TCP Modifications over Satellite Channels: Study and Performance Evaluation”, International Journal of Satellite Communications, Special Issue on IP, Vol. 19, Issue 1, January/February 2001, pp. 93-110.

[10] R. Goyal, R. Jain, M. Goyal, S. Fahmi, B. Vandalore, S. Kota, N. Butts and T. vonDeak, “Buffer Management and Rate Guarantees for TCP over satellite-ATM networks”, International Journal of Satellite Communications, Special Issue on IP, Vol. 19, Issue 1, January/February 2001, pp. 93-110.

[11] Consultative Committee for Space Data Systems (CCSDS), Space Communications Protocol Specification-Transport Protocol, CCSDS 714.0-B-1, Blue Book, May 1999.

[12] European Telecommunication Standards Institute - ETSI, http://www.etsi.org.

[13] M. Marchese "Performance Analysis of the TCP Behavior in a GEO Satellite Environment", Computer Communications Journal, Special Issue on the Performance Evaluation of Telecommunication Systems: Models, Issues and Applications, Vol. 24, Issue 9, May 2001, pp. 877-888.

[14] M. Allman, V. Paxson, W. S. Stevens, "TCP Congestion Control", IETF, RFC 2581, April 1999.

[15] S. Floyd, T. Henderson, "The NewReno Modification to TCP's Fast Recovery Algorithm", IETF, RFC 2582, April 1999.