5.6.6 Proposal of Modified Version of TCP adapted to Large Delay Geo-Stationary Satellite Channels

Research Unit: CNIT - University of Genoa (DIST section)

Introduction

Matching the applications that use TCP with the advantages offered by satellites, it is natural to think of TCP/IP-based applications over satellite networks. The TCP/IP protocol family is not so suited for the satellite environment but, on the other hand, the diffusion of TCP/IP applications makes difficult to think of another protocol architecture, non-transparent to the user, dedicated to the satellite links.

This work is aimed at presenting the problems, the metrics and the solutions of TCP/IP-based applications over satellite networks. Characteristics of the channels, algorithms and control schemes are described. The reference technology is represented by Geostationary Orbit (GEO) satellites, because many tests have been performed over them and many real measures may be reported to show which is the actual effect of the decisions taken. Anyway, it will be also shown that the issue is not limited to GEO satellites. The issue, as well as the methodology and the metrics introduced, can also be applied in other environments, each of them characterized by a peculiar characteristic: the large delay per bandwidth product. Unfortunately, a popular transport layer protocol as TCP, and, in particular, its flow control, does not match the characteristics of a network where the product between the available bandwidth and the time that information requires to get to the destination and to have the confirmation of the arrival is large. Actually, the problem is hidden inside the TCP flow control. On the other hand, a large delay-bandwidth product characterizes not only GEO satellites but also many other environments as, for instance, broadband radio networks.

During the project our work was aimed at designing, implementing and experimenting protocols, topologies and applications that characterize the satellite network, properly designed, so to support Quality of Service (QoS) guaranteed data, audio and video services.

The future objectives of our protocol design are:

- Creation of an international network to store and broadcast data and information mainly dedicated to tele-learning.
- Creation of a multimedia library concerning tele-learning, which can be accessed through a satellite network via a Web interface.
- Creation and utilization of data, audio and video transmission tools, which allow the remote utilization of such data.

Metrics

On the basis of the applications mentioned, it is important to define some metrics to measure the performance. Most of the applications have a common root: they need a quick and reliable data transfer. As a consequence, the reference metrics are, essentially, two:

- The overall time to transfer data (e.g. file), measured in seconds [s].
- The amount of data transferred in each time unit (the throughput), measured in bytes per second [bytes/s].

Actually, the two metrics may be reduced to one if the final throughput is considered, being them depending on each other, but the distinction should be maintained because it helps understand the difference among the different schemes that will be presented in the following. For instance, even if a file transfer requires similar time to be concluded by using two different algorithms (and, as a consequence, the final throughput for the two alternatives is almost the same), it may have a different throughput in some instants of the transmission. For example, if one algorithm is more efficient than the other in the initial phase of the connection, it may be employed in applications that require a short file transfer, as tele-control.

The mentioned metrics are "objective" metrics of Quality of Service (QoS), and the results reported in the following sub-section concern objective metrics. Nevertheless, their correspondence with the user perception should not be neglected.

An Example

An important quantity is the "delay per bandwidth" product. As indicated in RFC 1323 [1], TCP performance does not depend upon the transfer rate itself, but rather upon the product of the transfer rate and the round-trip delay (RTT). The "bandwidthdelay product" measures the amount of data that would "fill the pipe", i.e., the amount of unacknowledged data that TCP must handle in order to keep the pipeline full. TCP performance problems arise when the bandwidth*delay product is large. In more detail, within a geostationary large delay per bandwidth product environment, the acknowledgement mechanism described takes a long time to recover errors. The propagation delay makes the acknowledgement arrival, provided by the TCP protocol, slow, and the throughput needs more time than in cable 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. In general, the problem is the delay of the network or, in more detail, the delay per bandwidth product of the network that has a devastating effect on the acknowledgement mechanism used by TCP.

Table 1 contains the exact values of the overall transmission time and of the throughput measured in the final phase of the connection (with fixed bandwidth and variable round trip time).

The implications on the different applications mentioned in the previous Section

RTT [ms]	Time [s]	Throughput [bytes/s]
100	3.1	215154
500	12.7	52278

Table 1: Overall transmission time and throughput in the final phase of the connection.

are simple to imagine. It is sufficient to think of a tele-learning system, where the student is waiting for the material (an image or a graph) at the screen, or to a remote control system aimed at activating an alarm or a robot. Moreover, within the satellite environment, the problems are different if a LEO (Low Earth Orbit), a MEO (Medium Earth Orbit) or a GEO satellite system is used [2].

Black Box Approach

In the following, two possible frameworks, where the different solutions to improve the performance of the transport layer over satellite channels may be classified, are proposed: the Black Box (BB) approach and the Complete Knowledge (CK) approach. The former implies that only the end terminals may be modified; the rest of the network is considered non-accessible (i.e. a black box). The latter allows tuning parameters and algorithm of the network components. Most of the state of the art has been based on the Black Box approach.

The classification proposed is not the only possible one and, probably, it is not exhaustive (i.e., not all the algorithms and methods in the literature can be classified within one of the two classes), but it is useful to understand the problems and to introduce the analysis proposed.

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 devices.

In the **Black Box**, the transport layer (TCP) is modified and tuned, by acting only on the end user terminals. The rest of the network is considered as a black box, where, even if the intermediate devices and the configurations are known, they cannot be modified. The behavior of TCP is enhanced by varying the value of the initial window (IW is measured in bytes; i.e., the notation IW=1 means IW=1 ·SMSS [bytes]) and of the buffer dimension. The latter is 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.

Complete Knowledge Approach

An alternative approach is supposing the **Complete Knowledge** of each network device (e.g. routers, modems 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 proprietary. The Complete Knowledge approach supposes the complete control of any network devices and allows a joint configuration of all the functional layers involved to get an optimized network resource management, aimed at improving the overall performance offered by the network. Within this framework it is feasible to propose a protocol architecture, designed for a heterogeneous network involving satellite portions, which conveys the benefits both from the Black Box and the Complete Knowledge approach. The advantages that derive from the Black Box approach are utilized to design a novel network architecture suited for satellite transport, where the transport layer is divided into two parts, one of them completely dedicated to the satellite portion of the network (Satellite Transport Protocol - STP), introducing, moreover, a relay entity layer to optimize the transmission. The following study is partially taken from [3-6].

The Satellite Transport Protocol (STP)

The transport layer will be properly designed to consider all the possible peculiarities of the application environment. Some guidelines (concerning the STL and its implementation through the Satellite Transport Protocol (STP)) may be also introduced. The modified version of the TCP proposed in the previous Section can be considered a former implementation of the transport layer and a basis for the design of STP. In details, the significant algorithm modifications introduced are:

- 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. A mathematical function along with the other parameters involved (e.g. the initial window “IW” and the receiver buffer “buf”) might help to 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. Also in this case another mathematical function might have the responsibility of this part.

The Real Test-Bed And Performance Comparison

The real test-bed is structured as follow: two remote hosts are connected through a satellite link by using IP routers. An average Round Trip Delay (RTT) of 511 ms has been measured. The TCP/IP protocol stack is used. The data link layer of the router uses HDLC encapsulation on the satellite side, where a serial interface is utilized, and Ethernet on the LAN side. 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 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. 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 transferring data of variable dimension (H [bytes], in the following) between the two remote sites. The application designed allows transferring a single file a time, which is a case often reported in the literature, both as a benchmark for working and as a configuration used in real environments. A type of file has been utilized, in the case shown, to perform the tests and to study the behaviour of the modified TCP, namely, a file of relevant dimension of about 2.8 Mbytes (2,800,100 bytes), indicated with H=2.8 Mbytes. The first analysis is dedicated to investigate the solution applied in the Black Box approach. We can observe that IW and buf govern the throughput in the initial phase. It is sufficient to observe figure 1, to see the better performance of the modified TCP, compared with classical TCP (IW=1, buf=64).

Finally, the comparison is aimed at giving a first idea of the further improvement of the STP (complete knowledge approach) with respect to the modified TCP configurations, already adapted to satellite channels in the Black Box approach.

Figure 2 shows the throughput versus time for the three configurations mentioned

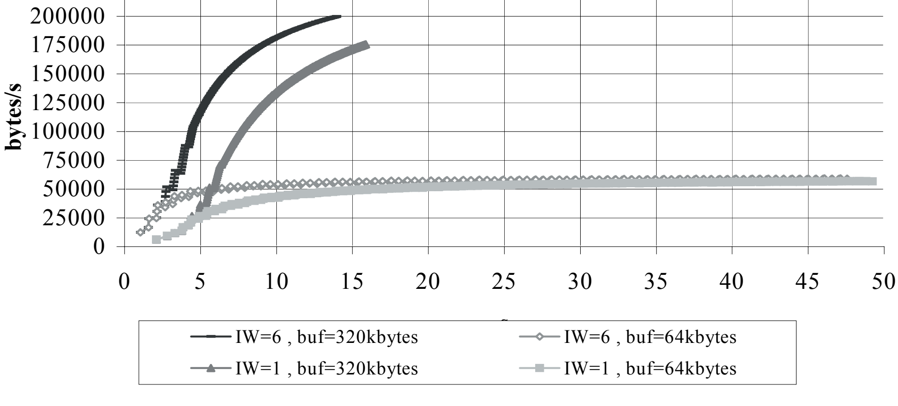

Fig. 1: Throughput (bytes/s) versus time for different values of the buffer length and of the initial congestion window, H=2.8 Mbytes.

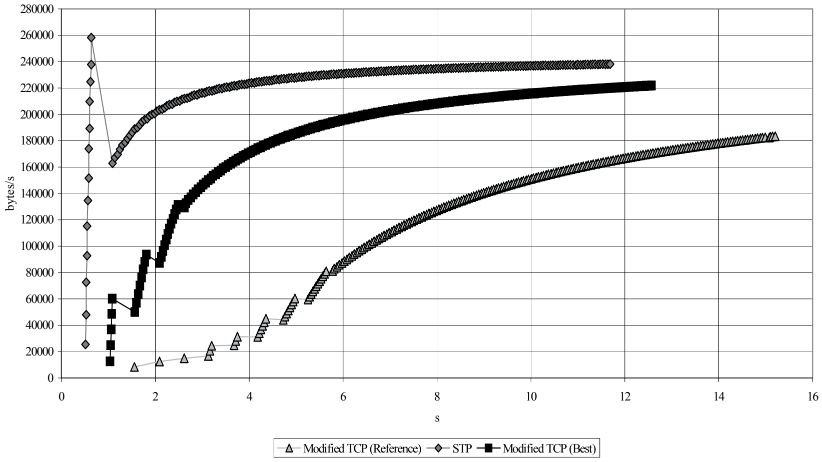

Fig. 2: Throughput vs time, H=2.8 Mbytes, mono-connection.

and a file transfer of 2.8 Mbytes. The overall transmission time is 15.2 s for the Reference configuration, 12.57 s for Modified TCP (called Best configuration) configuration and 11.69 s for STP. The gain of STL, computed in percentage as 100·(15.2-11.69)/ 15.2), is 23.1%, with respect to the Reference and 7% (100·(12.57-11.69)/12.57), with respect to the Best. STP has a shorter transmission time than the other configurations; thus the performance gain is actually the metric of "reduction" of the overall transmission time.

REFERENCES

[1] V. Jacobson, R. Braden, D. Borman, "TCP Extensions for High Performance", IETF, RFC 1323, May 1992.
[2] G. Maral, M. Bousquet, "Satellite Communications Systems - Systems, Techniques and Technology", III Ed,, John Wiley & Sons, Chichester, 1998.
[3] 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.
[4] M. Marchese "Performance Analysis of the TCP Behavior in a GEO Satellite Environment", Computer Communications, Special Issue on the Performance Evaluation of Telecommunication Systems: Models, Issues and Applications, Vol. 24, Issue 9, May 2001, pp. 877-888.
[5] M. Marchese, "Proposal of a Modified Version of TCP Adapted to Large Delay Satellite Channels", IEEE International Conference on Communications (ICC 2001), Helsinki, Finland, June 2001.
[6] M. Marchese, "A Parametric Approach to Improve the Performance of the Transport Layer in Satellite Communications", Globecom 2001, San Antonio, TX, Nov. 2001.