Transport Protocol and Resource Management for Satellite Networks: Framework of a Project

Davide Adami °, Mario Marchese *, Giacomo Morabito ▽, Michele Rossi ', Luca Veltri ◊
CNIT - Italian National Consortium for Telecommunications
° University of Pisa Research Unit, Via Diotisalvi, 2, 56126, Pisa (Italy)
* University of Genoa Research Unit, Via Opera Pia 13, 16145 Genova (Italy)
▽ University of Catania Research Unit, Viale A. Doria, 6, 95125, Catania (Italy)
' University of Ferrara Research Unit, Via Saragat, 1, 44100, Ferrara (Italy)
◊ University of Rome Research Unit, CORITEL, Via Anagnina, 203, 00040, Roma
e-mail: davide.adami@cnit.it, mario.marchese@cnit.it, giacomo.morabito@cnit.it, michele.rossi@cnit.it, luca.veltri@cnit.it

Abstract

The paper describes the protocol definition within the framework of the Project "Transport Protocol and Resource Management for Mobile Satellite Networks", funded by the European Space Agency ESA) and carried out by an Italian group composed of Marconi Mobile, CNIT and Etnoteam. In more detail, the Project is aimed at designing, implementing and testing a protocol stack adapted to the specific characteristics of a satellite communications system.

The protocol stack is based on the TCP/IP suite adapted to the channel characteristics. The objective is the optimization of both the transport protocol performance for a satellite network environment and the efficient utilization of network resources. This has been achieved without re-designing the protocol interfaces, so that they will keep the same characteristics of the interfaces currently used. These characteristics should get the target of maximizing the system performance and, in the same time, allow the utilization of standard applications so reaching a high degree of portability.

The protocol architecture proposed in the paper is mainly targeted to GEO (Geostationary Orbit) satellite communication systems, where the high round-trip delay heavily affects the system performance, but other environments, as radio-mobile and LEO (Low Earth Orbit) satellite systems, characterized by fading and high bit-error rate, are not excluded. The protocol stack has a high degree of flexibility to allow an efficient adaptation to these characteristics.

I. Introduction

Satellite networks offer clear advantages compared to terrestrial networks but the characteristics of the satellite environments are very different and introduce new and complex problems. On the other hand, to assure the simple and efficient use of widespread applications, it is necessary to maintain the interfaces commonly used and, in the same time, guarantee a high level of service quality. The problem is not trivial because it means that the TCP/IP suite has to be kept as a reference and that Quality of Service (QoS) algorithms have to be applied, although both the transport layer and the QoS schemes are heavily affected by the peculiarities of the satellite links. Reference [1] contains a discussion of the current work and progress being done in transport layer (namely TCP, defined in reference [2]) research related to satellites and outlines mechanisms that may help the TCP to use efficiently the bandwidth available over the satellite channel. Improvements to TCP and proposals of modifications to the transport scheme are widely treated in the literature and a new structure is suggested in the paper. In the same time, it is necessary to establish a mechanism, strictly connected with the transport structure, to reserve a specific bandwidth. Being a real network composed of terrestrial and wireless portions, the solution should not be limited to the satellite links, where bandwidth allocation algorithms should take into account concepts as fading and outage, but it should provide an overall end-to-end solution, strictly integrated with the transport.

The paper describes the general framework of the Project, the aim and the scope, and presents the transport system structure. A solution based on the use of a specific transport protocol dedicated to the satellite portion is introduced. The end-to-end characteristic of the transport layer is completely preserved. The TCP interfaces towards the adjacent layers are maintained. From the transport point of view, the key tools are represented by the Relay Entities, whose role is to separate the terrestrial portions from the satellite parts of the network. The Relay Entities are composed of two different protocol stacks, oriented, respectively, to the non-satellite (e.g. cable) and to the satellite portion, and linked by a relay layer. The transport stack is integrated with the QoS architecture both within the terminals and within the Relay Entities. A standard QoS architecture for IP networks is taken as reference and it is extended to match requirements specific of satellite links.

The paper is structured as follows. Section II contains the aim and the scope of the paper. Section III introduces the transport layered architecture; the functional requirements and the protocol characteristics of each layer are listed and explained. The QoS architecture and the resource management specification are contained in Section IV. Section V contains some preliminary results. Section V presents the conclusions.

II. AIM AND SCOPE

The paper contains the definition of the layered protocol architecture and of each protocol that compose a new transport architecture, aimed at improving the performance of the communications over a satellite environment and at guaranteeing a specific level of Quality of Service. The work covers the Protocol Definition of a Project called "Transport Protocol and Resource Management for Mobile Satellite Networks", funded by the European Space Agency (ESA) and carried out by Marconi Mobile, as project leader, CNIT and Etnoteam. The Project is aimed at designing, implementing and testing a protocol stack adapted to the specific characteristics of a satellite communications system. The protocol stack is based on the TCP/IP philosophy adapted to the channel characteristics. The objective is the optimization of both the transport protocol performance for GEO satellite networks and the efficient utilization of network resources. This will be achieved without re-designing the protocol interfaces, so that they will keep the same characteristics of the interfaces currently used. These characteristics should get the target of maximizing the system performance and, in the same time, allow the utilization of standard applications so reaching a high degree of portability. The protocol stack will have a high degree of flexibility to allow an efficient adaptation to different channel characteristics. The reference architecture is reported in Fig. 1. The network is composed of terrestrial portions, represented by the Internet in the figure, and of a satellite portion (a backbone, in this case). The latter is isolated from the rest of the network by using Relay Entities. Only two of them are shown in the figure but, actually, one Relay Entity is required whenever a satellite link is accessed. The transport layer of this new Satellite Protocol Stack (SPS) is called Satellite Transport Layer (STL) and it implements a Satellite Transport Protocol (STP), suited for the specific environment.

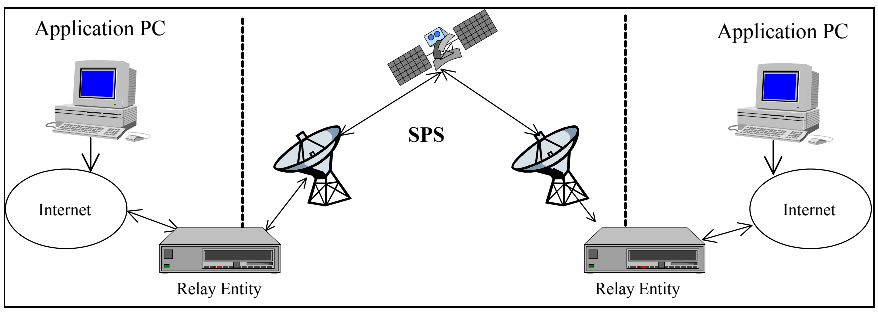

Fig. 1. Scope.

The architecture proposed may be a valid alternative both in case the satellite portion represents a backbone network (Fig. 1) and both in case it represents an access network where the Relay Entity is a simple tool, directly attached to the Application PC. It may be also a hardware card inside the Application PC. In this case, it may be a simple plug-in module, as a network or a video card to be inserted inside the PC. An user accessing the Internet (or other terrestrial network) through a satellite network represents a further case of Access Network, where the new architecture may be applied. Being an Access Network case, the plug-in module, mentioned above, may be applied.

III. SATELLITE ADAPTED PROTOCOL ARCHITECTURE (SAPA)

A. Global Layered Architecture

The two Relay Entities are gateways towards the satellite portion of the network. The Satellite Protocol Stack (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). The design is aimed at preserving the end-to-end characteristic of the transport layer. The transport layer is divided into two sub-layers: the upper one (Upper Transport Layer - UTL), which guarantees the end-to-end characteristic, and the lower sub-layer (Lower Transport Layer - LTL), which is divided into two parts and interfaces the STL. The terrestrial side of the lower transport layer may be also represented by the TCP. Fig. 2 shows the layered protocol architecture. The transport layer is modified even if the interface with the adjacent layers may be the same as in the TCP. The Upper Transport layer will include the TCP and the UDP implementation to allow a full compatibility with a different architecture. The transport layer at the destination will be properly identified during the set-up.

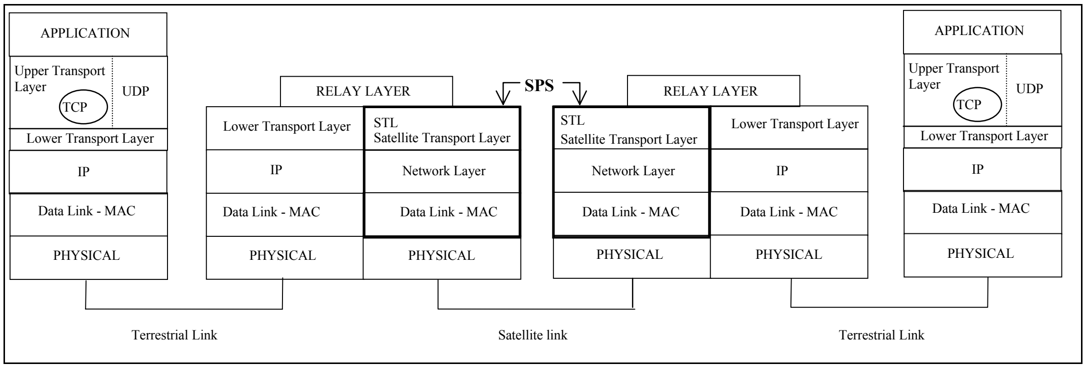

Fig. 2. End-to-end SPS Architecture

B. Upper Transport Layer

The main capabilities of the Upper Transport Layer are indicated in the following

- Upper Transport Layer shall guarantee the end-to-end semantic of the transport layer
- Upper Transport Layer shall work with TCP standard interfaces towards the application and toward the underlying layer (i.e. LTL)
- The completeness of the overall transfer shall be guaranteed
- Upper Transport Layer shall include TCP (and UDP) to assure a correct internetworking with a different architecture (e.g. a full TCP/IP architecture) at the destination
- Upper Transport Layer must "understand" if there is an Upper Transport layer at the destination or if there is TCP or UDP
- Upper Transport Layer must address and manage the failure problem when the communication fails at the Lower Transport Layer or a link is dropped.
- Upper Transport Layer shall be capable of establishing a route to the destination.
- Upper Transport Layer could extend its capabilities via options field within the header.(e.g. the suspend / resume transfer options)

In this perspective the double transport UTL / LTL layer, which is batch ACK- based in its upper part, allows:

- To have a simple and structured functional layering (e.g. each layer has a specific and clear role)
- To avoid complex signaling mechanisms within the network (including relay entities and end terminals)
- To shift the complexity towards the network border
- To manage persistent sources where the connections are not explicitly dropped (e.g. the FIN packet is not transmitted).
- To keep information about the connection status; even if no batch is re-transmitted at the moment, information about the connection status can be very useful for future extensions and for the storage of data (a special Resource Management API will be provided).
- To obtain the same results of the ACK filtering (presented below) by ruling the batch dimension; the possibility to have this type of regulation is a degree of freedom that can be studied in future extensions.

For what concerns the end – to – end characteristic, an alternative approach is represented by Ack Filtering – based mechanism. On the other hand, it increases the complexity of the network, compels to use a signaling mechanism, mixes the functional duties of each single layer, completely loses the information about the connection status and does not offer any possibility of future extension. As a consequence, in the following we refer exclusively to the double transport layer UTL / LTL approach because it has been the choice for the design of the protocol. It is important to remember that the protocol is originated from the TCP and inherits the philosophy of that protocol. In particular, it acts as the TCP for what concerns the UTL connection opening and the positive acknowledgement mechanism. UTL differs from TCP specification in the loss recovery management, which is not necessary because it is performed by LTL.

C. Lower Transport Layer

The Lower Transport Protocol Header is the same as standard TCP because LTL acts like TCP. The distinction between TCP and LTL is made using the Protocol Type field located in the IP level header. The LTL acts transparently in presence of TCP and UDP traffic. The same state machine as the TCP is used.

D. *Relay Entity*

An idea concerning the Relay Entity is reported in Fig. 3 (where the architecture of a Relay Entity is shown). 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), i.e. a specific unit of information created in the Relay Layer. The Data Link layer (the Medium Access Control sub-layer, in this case) offers to the upper layer a Bandwidth Reservation service, a sort of Bandwidth Pipe available to the Network Layer, which can itself reserve resources for the Transport Layer. The Resource Management provides the tools to perform the reservation. The Network Layer may use the structure of the IP layer but it may be also properly designed together with the STL layer, so to avoid the possibility of the event 'congestion' and to optimize the performance of the overall transmission on the satellite side. The Network Layer may reserve resources by using either the Integrated Services [3] or the Differentiated Services [4] approach (see the Resource Management to make a choice), considering the two possibilities offered in the IP world. Anyway, the aim is to create a bandwidth pipe (Relay Entity - to - Relay Entity, in the satellite portion, concerning the relay), so to offer a dedicated channel to a single connection or to a group of connections at the transport layer. If no bandwidth reservation is provided, the pipe shown in Fig. 3 may be simply represented by the transfer capacity of the physical interface. In this case, all the connections of the STL share the same portion of bandwidth and the STL design must take it into account.

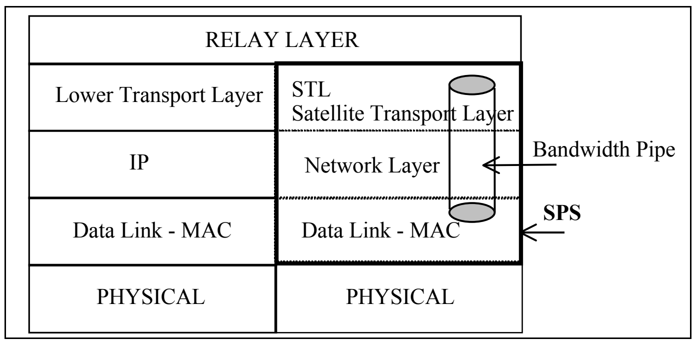

Fig. 3. Relay Entity.

IV. QoS Architecture

A. *QoS Specification*

This section describes the mechanisms for the support of the Quality of Service (QoS) within terrestrial-satellite scenarios.

The considered QoS scenario is based on a standard QoS architecture for general IP networks; however, such architecture is extended in order to cover also some specific satellite requirements. Two QoS profiles are considered:

- A per flow-based treatment, that implies a fine granularity of QoS provisioning, but leading to some scalability limits,
- A per class of traffic treatment, in which each class gathers all flows with similar QoS requirements; this approach can achieve more overall scalability, with the disadvantage of some leaks for the QoS strictness.

The first solution appears more suitable for satellite access links, and it can better fit for a scenario where the Relay Entity is an access entity. The second is more suitable in a core network scenario where the satellite links replace backbone links.

All the QoS reservation and support mechanisms must be considered as unidirectional; if some QoS has to be provided for both directions, two explicit reservations/procedures should be considered. The general QoS architecture is depicted in Fig 4.

The following QoS entities are defined:

- Policy Decision Point (PDP). A server entity that has a database with the topology and the network resources mapped into some service-class states inside of the domain that it controls. It takes admission decisions relying on the information that characterizes the ingress-to-egress flows, provided by an Edge Device/Router (ER) that acts as QoS client (i.e. the requesting entity). The PDP must implement also a QoS signaling server used to perform QoS requests.
- Local PDP (LPDP). An entity resident generally on the edge devices (the ERs) that relies on an initial resource configuration made by a PDP (e.g. the total capacity for the various classes from "its" ingress point (the ER) to other possible end-points for the same domain). It can take local admission decisions for the assigned resources. In case of it can't handle locally the request, it makes an explicit request to the PDP for outsourcing band from ingress to an egress point. It also implements the client (for communication with the PDP) and server (to communicate with the users who wants to request QoS for a certain bandwidth link from an ingress point to an egress one inside the specified domain) for the QoS signaling.
- Policy Enforcement Point (PEnP). It is located on the edge devices and it is configured by the PDP or LPDP to policy the traffic at the ingress of the domain.

- QoS differentiating (and marking) mechanisms resident in the edge devices (the ERs) to map the ingress traffic into the classes of service inside the domain (generally located at the PEnP points).
- The signaling protocol used for interaction between the different entities (e.g. COPS, SNMP, H.323 or SIP).
- Edge Router (ER). A boarder entity that include IP router, PEnP, LPDP...
- Core Router (CR). An IP router (as light as possible) inside the QoS domain (in the backbone).
- Terminal (T) The end-user point with the applications that can make QoS requests using a QoS API, using a local signaling protocol client ("inside" the terminal) in order to perform QoS requests.
- Relay Entity (RE). The *satellite specific entity* that may implement, for specific scenarios, a simple PEnP and LPDP.

In figure Fig. 5 and Fig. 6, the general protocol architecture for the Terminals and for the Relay Entities are shown. Fig. 6 is equivalent to Fig. 3 but it includes also QoS management.

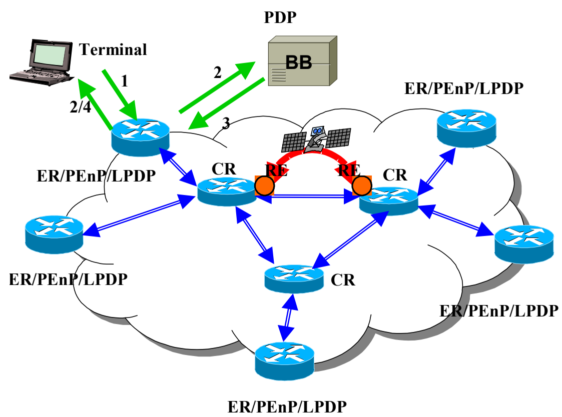

Fig 4. QoS enabled network architecture.

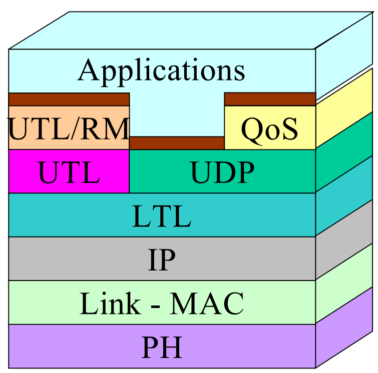

Fig. 5. Terminal's protocol architecture (including QoS and RM).

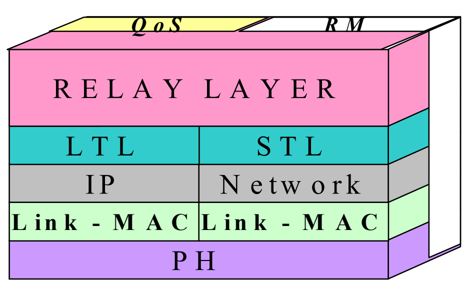

Fig 6. Relay Entity protocol architecture (including QoS and RM).

B. Resource Management Specification

The Resource Management can be performed with or without resource reservation (i.e. providing or not QoS guarantees as result of a certain QoS request).

In case of a certain form of QoS is requested, the Resource Management can use the same QoS information (and hence the same signaling mechanisms). In this case a "pipe" (strict or not) can be created according to the reservation.

In case of no QoS is explicitly requested, the RM can be performed according to different roles. Managing a resource means:

- Optimizing the use of a given amount of some resources (for example, bandwidth) that is logically shared between users.
- Deciding when to request/release part of the resource based on different information and logic of treatment/optimizing, implemented.

Different approaches to RM can be based on the different scenarios and actors:

- QoS signaling based

  In case of existing explicit signaling protocol for reserving resources with a certain QoS, it can be used this information to perform RM. The assigned resource utilization can be optimized inserting Best Effort (BE) traffic when it is possible. In this case the resource management mechanisms can act request/release of resources as a component of the QoS provisioning.

- Measures based

  Based on the measurement of the traffic class queues and adapting the bandwidth pipes on the class-flow load. Note that this RM mechanism can also be considered for QoS aware flows, when no per-flow hard RM is implemented. In this approach the information of the resource utilization results from queues' measurement. Even in this case the main idea is to insert BE (Best Effort) traffic into an under-utilized bandwidth pipe. The success of this mechanism strictly depends on the effective rapidity of changing the available bandwidth (BW) and it ought to be stabilized by thresholds to eliminate the effects of close loops (feedbacks).

- Specific information based

  A dedicated mechanism relying on information added to the streams by RM/satellite aware applications;

  The applications may give (or not) some supplementary information that can be used to predict the traffic generated. In this way it is possible to optimize the bandwidth assigned to a specific application flow, and prevent resource congestion (requesting more BW when unsupported traffic is expected).

Although all the three approaches could be used, the third one (*RM based on a specific information bound with the data flow, by means of ad-hoc protocol/mechanism*) seems to be the best for the Project. It is important to note that such solution does not exclude the possibility to use (in addition) the previous ones.

A RM mechanism is defined that exalts the possibilities offered by using the satellite architecture but in the same time is compatible with other approaches.

Why: The application can give (more) information about the data traffic in such a way that some RM entities on the data path can use it to optimize the resource utilization.

What: The application may give various information such as: the length of the total ready-data (write(buff,len)), time lengths, data rates (for example in the form of data (Byte) for a time interval (seconds)), other...

The use: To give information for the "near-future" traffic so the RM can adapt the pipes associated to the applications flow…

Where to put-it: The solution should be implemented within the terminals (Fig. 5) and within the Relay Entities (Fig. 6).

V. RESULTS

In this Section we report some performance results obtained by means of simulation, relative to the presented architecture. We refer to a single FTP-like connection between a sender wired application and a receiver wired application in a scenario similar to the one depicted in Fig. 2. In order to prove the effectiveness of the proposed network architecture (SAPA), in our simulation study we have considered two different scenarios. In the first one, a traditional TCP connection is considered, here relay nodes simply act as repeaters that forward the traffic over the satellite channel. Moreover, the TCP protocol experiences a large round trip delay given by the sum of the round trip time characterizing the two terrestrial networks and the one due to the satellite network. In the second case the whole SAPA architecture is considered. In the next, some preliminary results are reported for an error-free case. We label as B_{wired} the bandwidth characterizing the terrestrial links, with the term B_{sat}, instead, we refer to the bandwidth available on the satellite network. W_{max} is the maximum window size (expressed in bytes) for what concern the TCP and the LTL protocols. In Fig. 7, we report the trace of transferred LTL packets (each LTL packet has a dimension of a full segment, MSS=1024 bytes in the results reported here) by considering a satellite bandwidth equal to 128Kbps, whereas in Fig. 8, we report the same graph relative to the case where B_{sat}=2Mbps. In both cases, the presented architecture is able to speed-up End-to-End performance. This is due to two main reasons. First of all the LTL Slow-Start phase is improved by the presence of relay nodes that are able to decrease the round trip time experienced by each LTL protocol. Moreover, the data transfer in not limited by the dimension of W_{max} as happens in the standard TCP case. With the SAPA architecture, in fact, the satellite channel pipe is always full filled by transmitted packets thanks to the presence of the STP protocol. The first aspect is dominant where the satellite bandwidth is at a low value (see Fig. 7), while the second is the most important as the satellite bandwidth-delay product increases (see Fig. 8).

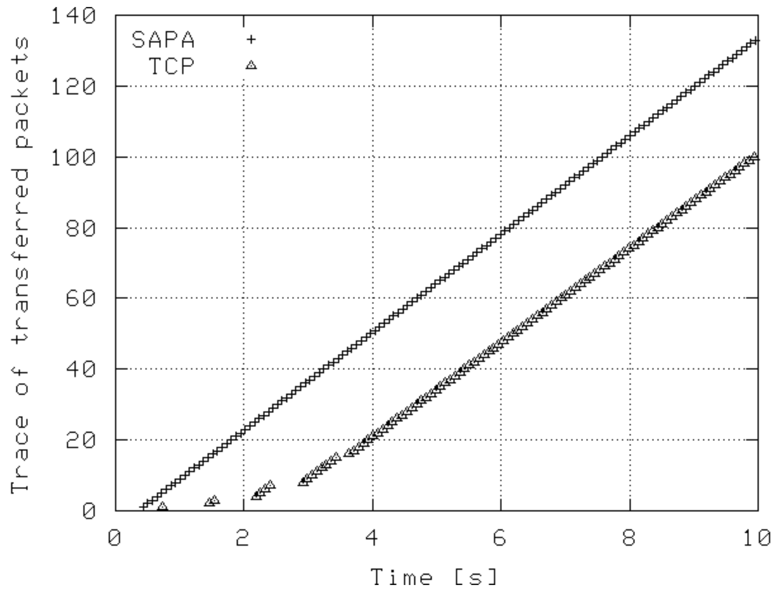

Fig 7.Trace of packets transferred, comparison between SAPA and a standard TCP connection. B_{sat} =128Kbps, B_{wired} =2Mbps, RTT (satellite)=0.5 s, RTT (terrestrial) = 100 ms, TCP/LTL W_{max} =32 Kbytes.

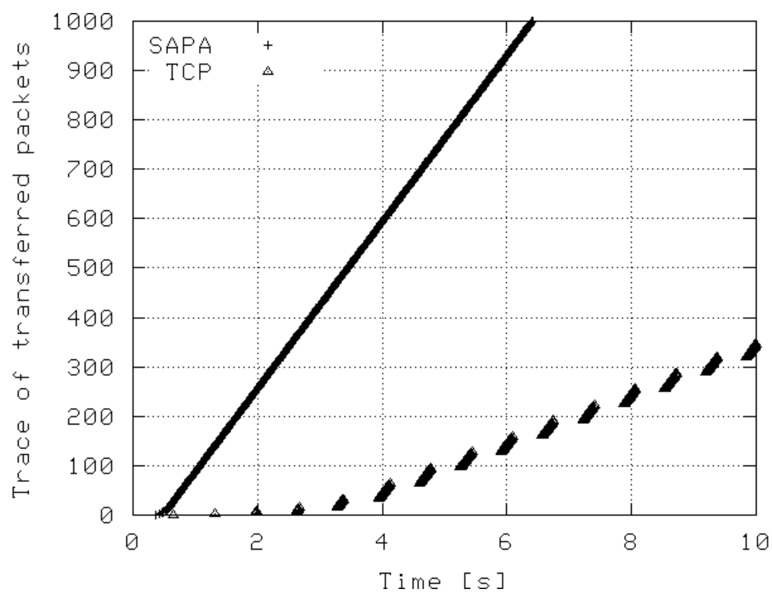

Fig 8. Trace of packets transferred, comparison between SAPA and a standard TCP connection. B_{sat}=2Mbps, B_{wired}=2Mbps, RTT (satellite)=0.5 s, RTT (terrestrial) = 100 ms, TCP/LTL W_{max}=32 Kbytes.

VI. CONCLUSIONS

A new Satellite Adapted Protocol Architecture has been proposed in the paper. The reference network includes satellite portions, which have been isolated from the rest of the network by using Relay Entities. The Relay Entities contain a protocol stack properly designed to improve the performance on the satellite portion. A Transport Layer, called STL - Satellite Transport Layer, a Network

Layer and a Data Link Layer, compose it. STL implements the Satellite Transport Protocol (STP), which adapts its parameters and algorithms depending on the performance offered by the lower layers. For example, Network and the Data Link layers may implement bandwidth reservation schemes, which, together with the STP, may improve the performance of the overall system. The overall architecture proposed works both if the satellite portion represents a backbone network and if it is an access network. The architecture includes a Resource Management that can be performed with or without resource reservation. In case of a certain form of QoS is requested, the Resource Management can use the same QoS information (and hence the same signaling mechanisms). In this case a "pipe" (strict or not) can be created according to the reservation. In case of no QoS is explicitly requested, the RM can be performed according to different roles. Managing a resource means: optimizing the use of a given amount of some resources (for example, bandwidth) that is logically shared between users and deciding when to request/release part of the resource based on different information and logic of treatment/optimizing, implemented. The interaction between the transport architecture and the QoS/RM architecture is presented.

The results reported show the efficiency of the architecture presented.

ACKNOWLEDGEMENTS

This work has been supported by the European Space Agency (ESA) under the contract ESA 14956/00/NL/ND "Transport Protocol and Resource Management for Mobile Satellite Networks".

REFERENCES

[1] 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", IETF, RFC 2760, February 2000.

[2] Information Sciences Institute, University of Southern California, "Transmission Control Protocol - Darpa Internet Program - Protocol Specification", IETF, RFC 793, September 1981.

[3] Integrated Services Working Group, http://www.ietf.org/html.charters/intserv-charter.html.

[4] Differentiated Services Working Group, http://www.ietf.org/html.charters/diffserv-charter.html.