Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Computer Networks 49 (2005) 797–815
Computer Networks
www.elsevier.com/locate/comnet

Validation of a QoS architecture for DVB–RCS satellite networks via the SATIP6 demonstration platform ☆

Antonio Pietrabissa [a,*], Tiziano Inzerilli [a], Olivier Alphand [b], Pascal Berthou [b], Thierry Gayraud [b], Michel Mazzella [c], Eddy Fromentin [d], Fabrice Lucas [d]

[a] University of Rome ''La Sapienza'', Dipartimento di Informatica e Sistemistica, via Eudossiana 18, 00184 Rome, Italy
[b] LAAS/CNRS, 7 Avenue du Colonel Roche, 31077 Toulouse cedex 4, France
[c] ALCATEL SPACE, 5 rue Noël Pons, 92737 Nanterre, France
[d] SILICOMP-AQL, Rue de la Chataigneraie, CS51766, 35517 Cesson Sevigne, France

Received 5 July 2004; received in revised form 23 November 2004; accepted 6 January 2005
Available online 26 March 2005

Responsible Editor: E. Altman

Abstract

Full integration of satellite technology in future terrestrial infrastructures requires support for high-quality broadband bi-directional communications. Research efforts in the field of satellite communications are currently oriented in the study of QoS-aware solutions for DVB–S and DVB–RCS which allow seamless deployment in the Internet. In this paper, the QoS architecture of the SATIP6 project, sponsored within the 5th EU Research Programme Framework, is presented, along with the implemented demonstrator and the obtained results. The QoS architecture is organized into two main modules, the Traffic Control and Access Control modules, whose aims are (i) to provide for differentiated service of conveyed IP flows and (ii) to achieve efficient utilization of uplink bandwidth, respectively. Experimental results obtained through the developed demonstration platform are reported and discussed to assess the effectiveness of the designed solution in terms of both service differentiation and efficient utilization of satellite resources.

Keywords: DVB–RCS; QoS; BoD; Access control; Traffic control

☆ This work was supported by the SATIP6 project (contract IST-2001-34344), financed by the European Commission within the 5th Research Programme Framework.
* Corresponding author. Tel.: +393358135654.
E-mail addresses: pietrabissa@dis.uniroma1.it (A. Pietrabissa), inzerilli@dis.uniroma1.it (T. Inzerilli), alphand@laas.fr (O. Alphand), berthou@laas.fr (P. Berthou), gayraud@laas.fr (T. Gayraud), michel.mazzella@space.alcatel.fr (M. Mazzella), eddy.fromentin@aql.fr (E. Fromentin), fabrice.lucas@aql.fr (F. Lucas).

1389-1286/$ - see front matter
doi:10.1016/j.comnet.2005.01.018

Nomenclature

ACK	Acknowledgement
AF	Assured Forwarding
ATM	Asynchronous Transfer Mode
BE	Best Effort
BER	Bit Error Rate
BoD	Bandwidth-on-Demand
CAC	Connection Admission Control
CoS	Class of Service
DAMA	Demand Assignment Multiple Access
DLB	Dual Leaky Bucket
DVB	Digital Video Broadcasting
DVB–RCS	DVB–Return Channel via Satellite
DVB–S	DVB–Satellite
EDF	Earliest Deadline First
EF	Expedited Forwarding
FCA	Free Capacity Assignment
FTP	File Transfer Protocol
GEO	Geostationary Earth Orbit
GPS	Generalized Processor Sharing
GW	Gateway
IETF	Internet Engineering Task Force
IP	Internet Protocol
LAN	Local Area Network
MAC	Medium Access Control
MF-TDMA	Multi Frequency-Time Division Multiple Access
MPEG	Moving Picture Experts Group
NCC	Network Control Centre
NRT	Non-Real Time
NTP	Network Time Protocol
QoS	Quality of Service
RC-EDF	Rate Controlled-EDF
RT	Real Time
SOHO	Small Office/Home Office
ST	Satellite Terminal
TCP	Transmission Control Protocol
TBTP	Terminal burst time plan
TDM	Time-division multiplexing
UDP	User Datagram Protocol
WFI	Weighted Fairness Index
WS	Work Station

1. Introduction

Driven by the growth of the World Wide Web and the Internet, satellite revenues in future are increasingly likely to come from the delivery of Internet Protocol (IP)-based applications and services, either to complement terrestrial broadband services, or to offer added-value services in some niche markets. Satellite access, rather than long-distance transport, is seen as a particularly convenient element of the overall telecommunication infrastructure as it provides ubiquitous and broadband access to anyone deploying a satellite terminal, both for single residential users and in Small Office/Home Office (SOHO) or corporate networks.

Satellite systems presently being defined and developed are based on different technologies and are tailored to support specific multimedia services. There is a large interest in the study of Geostationary Earth Orbit (GEO) system solutions which aim at providing a transparent integration of satellites in Internet networks. In particular, the satellite technology which is recently attracting most interest is DVB–RCS (Digital Video Broadcasting–Return Channel via Satellite), which complements DVB–S (Digital Video Broadcasting–Satellite) allowing interactivity and bi-directional communications without the need of a return channel via a terrestrial network.

Optimized transport of IP applications over a DVB–S/DVB–RCS satellite segment is the focus of the SATIP6 project, sponsored by the European Commission. Major efforts in the definition of the SATIP6 architecture were dedicated to the design of the Quality of Service (QoS) provisioning support. In the literature, providing QoS to GEO satellite networks has been subject to extensive researches (see, for instance, [1,2]). However, as observed in [3], most of the references in the literature analyze the performances of applications (video, voice, file transfer, etc.) without considering realistic resource management procedures, but only rather simplistic approximations, which cannot catch all of their effects on delay, jitter and other constraints. For instance, in [4,5], the access delay caused by the *access control* protocols is approximated by a fixed delay, which does not vary with respect to the statistical characteristics of the traffic and to the network load conditions. In this respect, the main contribution of this work is that it presents final experimental results obtained through a validation campaign with the SATIP6 demonstration platform, which supports an actual QoS provision framework, testing real applications. The work reported here has been recently presented before the ETSI's TC-SES Broadband Satellite Multimedia (BSM) Group [6].

The current QoS framework which has been implemented in the testbed is organized into two modules: (i) *traffic control* module and (ii) the *access control* module, aimed at providing IP flows with differentiated service and at achieving efficient exploitation of uplink bandwidth, respectively. In this paper, a proposal for the integration of the two approaches is presented and relevant experimental results are discussed.

The paper is organized as follows: Section 2 presents the SATIP6 scenario, along with some related work; Section 3 presents the *traffic* and *access control* modules (Sections 3.1 and 3.2, respectively), and the proposed interaction between them (Section 3.3); Section 4 presents the demonstration platform and the obtained experimental results; finally, in Section 5 the conclusions are drawn.

2. Satellite access scenario

The SATIP6 scenario, shown in Fig. 1, consists of a geostationary satellite network with onboard switching capabilities, Ka-band MF-TDMA (Multi Frequency-Time Division Multiple Access) uplinks and Ku-band TDM (Time Division Multiplexing) downlinks. Satellite Terminals (STs) provide single PC or Local Area Networks (LAN) with the access to the network, while Gateways (GWs) allows the connection with Internet core networks. The satellite network resources are managed by a Network Control Centre (NCC). The uplink access from each ST is managed through DVB–RCS interface using the ATM (Asynchronous Transfer Mode)-option (i.e., IP packets are segmented into ATM cells before being transmitted on the uplink), whereas transmissions from GWs are implemented through DVB–S interfaces. STs and GWs are boundary devices between the satellite and terrestrial links and play an important role in access to satellite resources and hence in QoS provision. In the SATIP6 architecture both

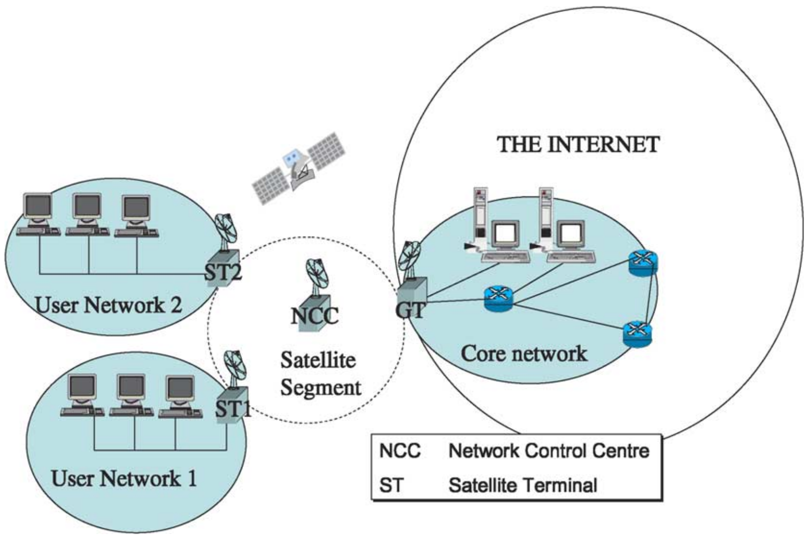

Fig. 1. SATIP6 scenario.

devices implement IP routing and have IP interfaces on the satellite segment.

From the QoS provision perspective, two fundamental functions have to be implemented:

1. *Traffic Control*, whose aim is to perform classification, regulation and scheduling of IP traffic to transmit in the satellite segment in order to implement differentiated bandwidth allocation and delay control to IP flows.
2. *Access Control*, whose aim is to regulate access to the uplink bandwidth contended by the STs in a same spot-beam in order to assure a fair allocation and efficient exploitation of bandwidth.

The first function is performed in both STs and GWs, while *access control* is part of the STs only.

The *traffic control* module performs differentiated management of IP service classes for QoS provision [7–10] to enforce delay as well as bandwidth requirements to IP traffic streams separately. The fundamental component to implement differentiated policy of traffic flows is the scheduler. There are a number of simple scheduling algorithms designed to provide some differentiation in bandwidth allocation and delay control. Fair queuing algorithms, based on the Generalized Processor Sharing (GPS) model, aim at minimization of the Worst-case Fair Index (WFI) [11]. This is the only performance metric which is used to assess scheduler performance and often is not suitable to satisfy link-sharing requirements in practice. In hierarchical link-sharing [12] more flexible mechanisms than WFI-based ones to manage excess bandwidth are envisioned. These allow repartition of excess bandwidth first among applications of a same user and then between different users on a hierarchical basis, which appear more practical and nearer to real link-sharing need. However, the hierarchical link-sharing approach appears not adequate to model requirements of emerging applications which need to be guaranteed on bandwidth allocation and delay control independently. QoS provision is offered through allocation of different bandwidth to users and applications. In the H-FSC algorithm [13] delay and bandwidth requirements cannot be satisfied independently while allowing hierarchical link-sharing. Independency of bandwidth and delay guarantees allows efficient support of applications such a file transfer, with loose delay requirements and high broadband requirements, or web-browsing and telnet, with stringent delay requirements and low bandwidth requirements. Traffic shaping can be used to enforce traffic isolation and improve performance of scheduling algorithms and can allow enforcing delay and bandwidth guarantees, sometimes independently of each other.

In Section 3.1 we propose a traffic control approach for the satellite access based on the Earliest Deadline First (EDF) scheduling with Dual Leaky Buckets (DLB), adapted to support IP flows. The resulting Rate Controlled-EDF (RC-EDF) [14] traffic control, uses DLBs to implement independent control of delay and bandwidth allocation, and the EDF scheduling policy to differentiate the flows. The major advantages of this approach are the minimal computational cost (the EDF algorithm has a logarithmic cost) and tight control of delay and bandwidth.

The *access control* protocols and algorithms provide ST and GWs traffic with the access to the SATIP6 network at Medium Access Control (MAC) layer. SATIP6 supports IP DiffServ classes: *Expedited Forwarding* (EF), which has bandwidth, delay and jitter requirements, *Assured Forwarding* (AF), characterized by bandwidth and loose delay requirements and *Best Effort* (BE), which has no requirements. The IP classes are mapped onto two MAC priority classes: EF packets are mapped onto the Real Time (RT) MAC class, whereas AF and BE packets are mapped onto the Non Real Time (NRT) MAC class; notice that, as clarified in the following, the differentiation between NRT and BE packets is performed at IP layer only. In particular, the *access control* protocols specify the Bandwidth-on-Demand (BoD) mechanism, which allows to dynamically assign the uplink bandwidth on the basis of capacity requests.

The debate on satellite BoD algorithms is becoming a central point even in the framework of standardization bodies such as the Internet Engineering Task Force (IETF) [15] and the ETSI

BSM [6,16]. Ref. [17] presents a performance evaluation of a 'burst targeted' BoD access scheme, based on the prediction of traffic bursts; the drawback of the scheme is that it is suited to supporting ON/OFF type traffic only, i.e., traffic which can be identified by a series of bursts and gaps. The approach in [18] is based on adaptive predictive control methodologies: an adaptive predictor forecasts future input traffic flows, whether bandwidth requests are computed according to a receding-horizon policy, via solution of an optimal control problem. In [19], the global resource-sharing problem is formulated as an optimization problem involving all the terminals, all the uplinks and all the downlinks with a large number of coupled constraints; this approach is troublesome since it is not scalable and requires the NCC to perform the whole resource management task. In [20], the BoD algorithm, developed by using control-theoretic methodologies, computes the capacity requests on the basis of the on-line traffic measures; the resulting BoD scheme, improved in [21], acts so that the queue length in the satellite terminal buffers tracks an appropriate reference value; the scheme is not based on traffic models, so that it is suitable for any traffic sources, and is easy to implement. The proposed SATIP6 BoD algorithm, based on the one presented in [21], is described in Section 3.2.

The support of TCP (Transmission Control Protocol) flows with BoD mechanisms is a critical issue, since the TCP algorithm is affected by the high latency introduced by the satellite network, which is further increased by the BoD. In the literature, some analysis of the TCP behaviour can be found [3,4,22] on the basis of simulation results. A key contribution of this paper is that the presented TCP on BoD results have been obtained via the SATIP6 demonstration platform.

3. Satellite terminal QoS procedures

3.1. Traffic control algorithm

The traffic control module performs buffer and bandwidth management of downstream IP packets. Namely, as soon as the IP forwarding process dispatches a packet to the satellite outgoing interface, a classification process is called. This classifies the packet on the basis of control information in the IP header as belonging to one of the following traffic classes: EF, AF and BE.

- Flows belonging to the EF class have stringent requirements on both queuing delay and jitter, need to be allocated bandwidth on a peak rate basis, are served with highest priority and are mapped onto statically allocated satellite bearers.
- Flows belonging to the AF class have no particularly stringent delay requirements, are allocated a minimum bandwidth on a mean rate basis and are served partially with statically allocated satellite bearers and bandwidth acquired dynamically through the BoD mechanism.
- All traffic which does not require QoS provision is classified as BE and is transmitted using dynamically allocated bandwidth, obtained via the BoD mechanism.

EF and AF flows undergo a regulation process implemented through a set of DLBs, which enforce minimum bandwidth allocation, flow isolation and fairness among flows through the definition of three parameters: average rate, r; peak rate, p; bucket size, b. Namely, when a DLB(r, p, b) is assigned to a flow the following equations define the traffic envelop to apply to the flow:

$$B^{\mathrm{off}}(t_1, t_2) = \int_{t_1}^{t_2} R^{\mathrm{off}}(t) \cdot \mathrm{d}t, \tag{1}$$

$$A(t_2 - t_1) = A(\tau) = \min(p \cdot \tau, b + r \cdot \tau), \tag{2}$$

$$B^{\mathrm{adm}}(t_1, t_2) = \min(B^{\mathrm{off}}(t_1, t_2), A(t_2 - t_1)), \tag{3}$$

where $R^{\mathrm{off}}(t)$ is the instantaneous bit rate offered by the flow to regulate, $B^{\mathrm{off}}(t_1, t_2)$ is the amount of data which is offered in the interval of time $[t_1, t_2]$, $A(\tau)$ is the so-called traffic envelope associated to the DLB(r, p, b) which upper bounds the admitted data by the DLB for any interval of time τ and $B^{\mathrm{adm}}(t_1, t_2)$ is the admitted data by the DLB in the interval of time $[t_1, t_2]$.

Two separate buffers are present at Medium Access Control (MAC) layer: a RT buffer, which avails of statically assigned capacity, and a NRT buffer, which feeds capacity acquired dynamically through the BoD mechanism (see the next section). Then, the scheduler process maps regulated traffic from EF, AF and BE flows applying the following criteria:

- Regulated traffic from EF flows is reordered using an EDF discipline and then directly passed to the MAC queue for RT traffic which avails of statically allocated capacity. Although always enough capacity is available to transmit EF traffic, the EDF discipline has been found useful to correctly pace the different EF flows which have to be multiplexed on the uplink.
- Regulated traffic from AF flows is reordered using an EDF discipline and then passed to the NRT MAC queue whenever the EF queue is empty (alternatively, to prevent packet starvation, a minimum rate could be granted to AF traffic).
- BE traffic is simply passed to the NRT MAC queue whenever IP EF and AF queues are completely emptied (alternatively, to prevent packet starvation, a minimum rate could be granted to BE traffic).

Fig. 2 shows the proposed traffic control scheme, where the Priority Queuing (PQ) is in charge of prioritizing EF, AF and BE traffic according to the above-mentioned criteria.

3.2. Access control algorithm

The access to the uplink is allowed by three kinds of capacity assignments:

- *Static capacity assignments*. At connection setup the NCC assigns a certain number of timeslots to a connection for its lifetime. Since this pre-assigned capacity is always available to the connection, low delays can be granted; on the other hand, the capacity availability is independent of the current bandwidth requirements, so that the network resources might be underutilized.

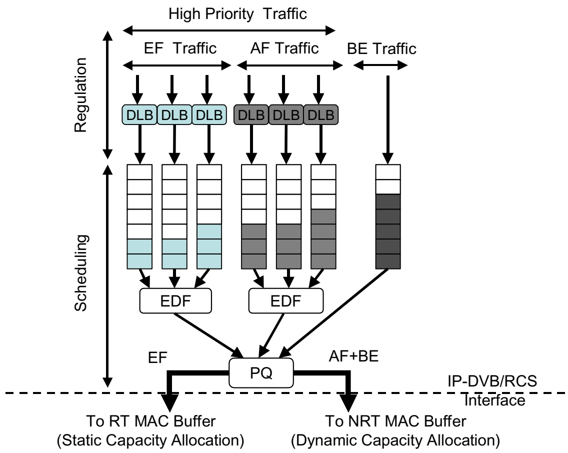

Fig. 2. Traffic control.

- *Dynamic capacity assignments*. The NCC assigns the time-slots to a certain connection on the basis of the requested capacity. Since the source STs have to send explicit capacity requests to receive capacity assignments, the latency of this approach is large due to the intrinsic propagation delay of the satellite links (the round-trip delay is about 0.5 s); however, the capacity assigned to the connection reflects the current bandwidth requirements and is granted so that the network capacity be well exploited. The BoD protocol defines the rules by which the *capacity requests* are computed.
- *Free Capacity Assignments* (FCA). The NCC assigns spare (i.e., unrequested) time-slots to the STs on the basis of a certain fairness criterion. FCA allows reducing the delay caused by the request-assignment cycle.

BoD algorithms are based on a capacity request-assignment cycle. The capacity requests are regularly transmitted by the STs to the NCC every superframe, and are computed on the basis of the sum of the queue lengths in the NRT MAC buffers, denoted with $q(t)$, and of the rate of the IP packets feeding the NRT MAC queues, denoted with $r_{\mathrm{IN}}(t)$. The time interval between the transmission of a capacity request by the terrestrial terminal and the associated capacity allocation is equal to the round-trip time between the ST and the NCC, T_{RTD}, plus the NCC algorithm runtime

T_{NCC}. Since both the considered delays are constant, $T_{\mathrm{RTD}} + T_{\mathrm{NCC}}$ constitutes the constant feedback delay of the system, denoted with T_{FB}.

The objective of the dynamic access control procedure is twofold: (i) STs should be capable of utilizing the whole requested capacity (*full link utilization*); this means that a proper amount of packets should be accumulated in the terminal buffer: as a matter of fact, if the terminal has no packets in the buffer queue when the capacity allocation is received, the unused allocated capacity is wasted; (ii) congestions must be recovered, to avoid packet starvation; (ii) FCA should be efficiently exploited.

3.2.1. Satellite network model

The system model consists of four parts:

(i) The source terminal, which is modeled by an input bit rate, $r_{\mathrm{IN}}(t)$, which, obviously, is non-negative:

$$r_{\mathrm{IN}}(t) \geqslant 0 \quad \forall t. \tag{4}$$

(ii) The NRT buffer in the source terminal, in which the packets wait for transmission in the uplink, is modeled by an integrator. Let $q(t)$ denote the queue length in this buffer; the variation of the queue length is given by the input rate $r_{\mathrm{IN}}(t)$ minus the output rate $r_{\mathrm{OUT}}(t)$, i.e.:

$$\dot{q}(t) = r_{\mathrm{IN}}(t) - r_{\mathrm{OUT}}(t) \tag{5}$$

(for the sake of simplicity, buffer underflows will be neglected; see [21] for a detailed analysis).

(iii) The controller $C(t)$ in the source terminal, which computes the capacity requests $r_{\mathrm{REQ}}(t)$ on the basis of the available measures: input rate $r_{\mathrm{IN}}(t)$ and queue length $q(t)$.

(iv) The NCC, which assigns the capacity on the basis of the requests and of the available link capacity. If the network is not congested, the assigned rate is equal to the requested bit rate, and thus $r_{\mathrm{OUT}}(t) = r_{\mathrm{req}}(t - T_{\mathrm{FB}})$. Conversely, if the network is congested, the control centre assigns less capacity: $r_{\mathrm{OUT}}(t) < r_{\mathrm{req}}(t - T_{\mathrm{FB}})$. Finally, if there is some leftover available capacity, the NCC can assign it as FCA. Thus, the NCC together with the transmission delay can be modeled as a delay block cascaded to an additive disturbance $d(t)$, defined as follows:

$$d(t) = r_{\mathrm{REQ}}(t - T_{\mathrm{FB}}) - r_{\mathrm{OUT}}(t). \tag{6}$$

If the network is congested, $d(t) > 0$; otherwise, $d(t) \leqslant 0$.

Fig. 3 shows the network model.

3.2.2. Model-based control scheme

The idea is to define an appropriate reference queue $q_{\mathrm{REF}}(t)$ and a controller $C(s)$ aimed at driving the actual queue length $q(t)$ to 0. Now, the first problem is to design $q_{\mathrm{REF}}(t)$ with the aim of *full link utilization*; obviously, this objective cannot be accomplished if the NCC assigns non-requested capacity: thus, the *full link utilization* must be achieved when no FCA is present. The appropriate reference queue length in question is the following:

$$q_{\mathrm{REF}}(t) = \int_{t-T_{\mathrm{FB}}}^{t} r_{\mathrm{IN}}(\tau)\mathrm{d}\tau. \tag{7}$$

In fact, Eq. (7) is straightforwardly obtained by noting that if the requested rate is equal to the input rate (i.e., $r_{\mathrm{REQ}}(t) = r_{\mathrm{IN}}(t)$), then the buffer accumulates the packets received during the request-assignment cycle T_{FB}.

It follows that a simple rate-based request $r_{\mathrm{REQ}}(t) = r_{\mathrm{IN}}(t)$ is capable of achieving the *full link utilization* objective. However, we must consider that the network can be also congested and that FCA can be available: in these cases, the $q(t)$ becomes larger and smaller than $q_{\mathrm{REF}}(t)$, respectively. Then, the controller $C(s)$ has two tasks: (i) it must react to congestion by increasing the rate-based requests on the basis of the excess

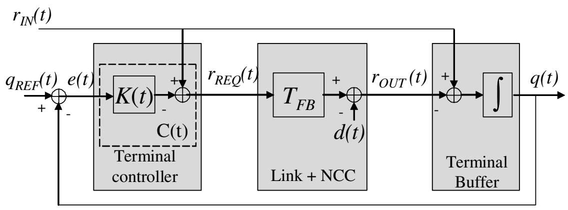

Fig. 3. Model of the BoD control system.

packets accumulated in the buffer; (ii) it must not decrease the rate-based requests if FCA are available, in order to reduce the queuing delay.

In view of these considerations, the proposed controller $C(s)$, shown in Fig. 4, computes the capacity requests as the sum of a rate-based part $r_{\mathrm{IN}}(t)$ and of a queue-based part $r_Q(t)$. By inspecting Figs. 3 and 4, it results that the request $r_{\mathrm{REQ}}(t)$ is given by the following equation ($\otimes$ indicates the convolution product):

$$\begin{aligned} r_{\mathrm{REQ}}(t) &= r_{\mathrm{IN}}(t) - r_Q(t) \\ &= r_{\mathrm{IN}}(t) - \min(0, K(t) \otimes [q_{\mathrm{REF}}(t) - q(t)]), \end{aligned} \tag{8}$$

where the $r_Q(t)$ is computed via $K(t)$, which is a sub-controller driven by the queue length error, defined as $e(t) = q_{\mathrm{REF}}(t) - q(t)$.

The meaning of the rate-based part of the request (8) is straightforward: in order to avoid packet starvation, the requested transmission rate should equal the input rate. The second term of Eq. (8) deals with congestions since it aims at driving the queue length error to 0 by reducing $r_{\mathrm{REQ}}(t)$, i.e., at driving the actual queue length to the reference one. Furthermore, the minimum operation (represented by the saturation block σ in Fig. 4) does not allow $r_{\mathrm{REQ}}(t)$ to be lower than $r_{\mathrm{IN}}(t)$, so that FCA can be exploited.

Since the process model $P(s) = \mathrm{e}^{-sT_{\mathrm{FB}}}/s$ contains a delay (see Fig. 3), and since the satellite delay is constant and known, the controller $K(s)$ has been developed by utilizing Smith's principle [23], which is a control-theoretic method to overcome the stability problems caused by the feedback delay:

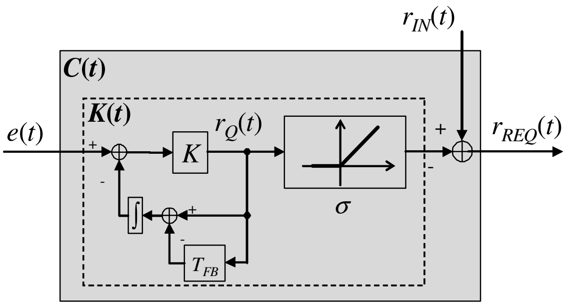

Fig. 4. ST BoD controller.

$$\begin{aligned} K(s) &= \frac{r_Q(s)}{e(s)} = \frac{C_0(s)}{1 + C_0(s)[P_0(s) - P(s)]} \\ &= \frac{K}{1 + K\left[\frac{1}{s} - \frac{\mathrm{e}^{-sT_{\mathrm{FB}}}}{s}\right]}, \end{aligned} \tag{9}$$

where $C_0(s) = K$ is the proportional primary controller and $P_0(s) = 1/s$ is the process model without delay. From Eq. (9), it follows that $r_Q(t)$ is then computed as

$$r_Q(t) = -K\left[q(t) - q_{\mathrm{REF}}(t) - \int_{t-T_{\mathrm{FB}}}^{t} r_Q(\tau)\mathrm{d}\tau\right]. \tag{10}$$

The following properties of *reference tracking*, *congestion recovery* and *full link utilization* hold; note that, as above-mentioned, these properties are relevant only if FCA are not present; thus, in the following, $d(t)$ is considered non-negative and the saturation block σ is neglected.

Property 1. *By setting $K > 0$, $q(t)$ tracks $q_{\mathrm{REF}}(t)$ exponentially, with time constant $\tau = 1/K$ (reference tracking).*

Proof. Since the system is linear, the part of $q(t)$ due to $q_{\mathrm{REF}}(t)$ can be computed by considering $r_{\mathrm{IN}}(t) \equiv d(t) \equiv 0$:

$$\frac{q(s)}{q_{\mathrm{REF}}(s)} = \frac{s}{s+K}\mathrm{e}^{-sT_{\mathrm{FB}}}. \tag{11}$$

The proof derives from the fact that (11) is a first-order transfer function, with pole $p = 1/K$, followed by a delay. □

Property 2. *By setting $K > 0$, if at a time t_C a congestion terminates (i.e., $d(t) > 0$ for $t < t_C$ and $d(t) = 0$ for $t \geqslant t_C$), the queue length part due to the congestion is exponentially driven to 0, with time-constant $\tau = 1/K$ (congestion recovery).*

Proof. Since the system is linear, the part of $q(t)$ due to the congestion, i.e., due to $d(t)$, can be computed by considering $r_{\mathrm{IN}}(t) \equiv q_{\mathrm{REF}}(t) \equiv 0$:

$$\frac{q(s)}{d(s)} = \frac{1 - \mathrm{e}^{-sT_{\mathrm{FB}}}}{s} + \frac{\mathrm{e}^{-sT_{\mathrm{FB}}}}{s+K}. \tag{12}$$

The inverse Laplace transform of (12) is

$$q(t) = \int_{t-T_{\mathrm{FB}}}^{t} d(\tau)\mathrm{d}\tau + \int_{0}^{t} \mathrm{e}^{-K(t-T_{\mathrm{FB}}-\tau)} d(\tau)\mathrm{d}\tau \qquad (13)$$

which, considering that $K > 0$ and $d(t) \geqslant 0\ \forall t$, proves Property 2. □

Theorem 1. *By setting* $K > 0$, *the control scheme meets the full link utilization objective (i.e.,* $q(t) \geqslant 0\ \forall t$) *if no FCA is provided (i.e.,* $d(t) \geqslant 0\ \forall t$).

Proof. Eq. (7) shows that $q_{\mathrm{REF}}(t)$ is calculated from $r_{\mathrm{IN}}(t)$. Thus, since the system is linear, $q(t)$ can be computed as the sum of two terms: the first one is due to $d(t)$, and its contribution is given by Eq. (13); the second one is due to $r_{\mathrm{IN}}(t)$ and can be computed by considering $d(t) \equiv 0$:

$$\frac{q(s)}{r_{\mathrm{IN}}(s)} = \frac{1 - \mathrm{e}^{-sT_{\mathrm{FB}}}}{s}. \qquad (14)$$

The inverse Laplace transform of Eq. (14) is

$$q(t) = \int_{t-T_{\mathrm{FB}}}^{t} r_{\mathrm{IN}}(\tau)\mathrm{d}\tau. \qquad (15)$$

The proof follows from Property 2, stating that Eq. (13) is non-negative, and from Eq. (4), entailing that Eq. (15) is non-negative.

3.3. Interaction between traffic control and access control

The interaction between the IP and MAC layer aims at calculating the right amount of bandwidth to be requested without resource over-provisioning while allowing service differentiation. Fig. 5 shows the overall architecture.

As shown by the 'data path' in Fig. 5, the IP flows are classified into the defined IP Classes of Service (CoS), regulated, stored in the proper buffer and scheduled at IP level (as specified in Section 3.1). The classified and regulated IP flows are then segmented, by the Segmentation and Reassembly (SAR) module, into ATM cells, which are stored into two MAC-layer buffers: the RT ATM cells stored in the RT MAC buffer are transmitted via the statically allocated capacity, while

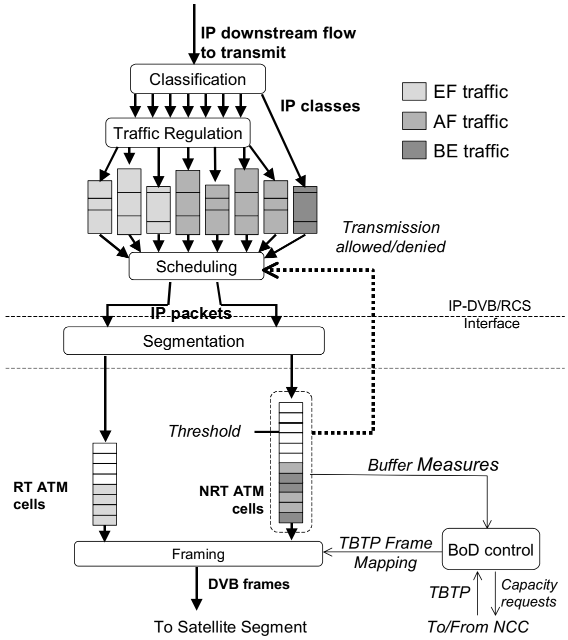

Fig. 5. QoS architecture in the ST.

the ATM cells stored in the NRT MAC buffer are transmitted via the dynamically allocated capacity; NRT packets are allowed to use a statically allocated time-slot if the RT MAC queue is empty. Finally, the framing entity maps the ATM cells onto the uplink frame.

The BoD controller, located at the MAC layer of the STs and shown in Fig. 5, is in charge of controlling the access to the satellite link. The BoD controller computes the capacity requests on the basis of measures concerning the NRT queue length and the rate of the packets feeding the NRT MAC buffer. With the aim of avoiding MAC-layer buffer overflows, if the queue length exceeds a threshold level, the scheduler is prevented to feed the NRT MAC queue. This is the only interlayer signalling between IP and MAC layer: thanks to this simplicity, the presented architecture has been selected for implementation in the SATIP6 demonstrator. Note that, to obtain the full advantages offered by the proposed architecture, the BoD controller should compute the capacity requests on the basis of measures concerning the sum of the AF and BE queue lengths at IP layer, and on the basis of the rate of the packets feeding the AF and BE queue lengths at IP layer. However, this solution has been postponed for the next version of the demonstrator, since it implies an additional implementation complexity: traffic measures have to be computed over several queues, and the overhead due to the addition of header and padding in the segmentation process must be considered. The impairment due to this simplification will be examined in the following.

On the basis of the NRT MAC buffer queue length, the BoD controller computes the bandwidth requests for the NRT traffic and sends them to the NCC (as specified in Section 3.2); the BoD controller then regularly receives the time-slot assignments from the NCC via the so-called Terminal Burst Time Plan (TBTP) messages [20].

Compared to a traditional architecture, in which the regulated IP packets are collected in CoS buffers, segmented into ATM cells, mapped onto the DVB traffic priority classes, re-regulated and, finally, re-scheduled, the proposed architecture claims various advantages:

(i) it avoids the duplication of the scheduling, shaping and policing functionalities in the IP and MAC layers;
(ii) it avoids MAC buffer overflows and hence partial transmission of an IP packet, which causes a waste of transmission capacity;
(iii) when a congestion state occurs packets are accumulated in the IP queues and not in the MAC queues, so that when the congestion ends the scheduler is able to use all its discrimination capability, selecting the IP packets with the most stringent delay requirements;
(iv) moreover, if the BoD controller is based on IP queue measures, the threshold level of the MAC NRT buffer can be set to a very low value: the MAC queues would be used only to store the packets that will be sent in the next uplink frame and, consequently, the delay introduced by the MAC layer queues becomes negligible.

Conversely, in the actually implemented BoD controller shown in Fig. 5, for the sake of implementation simplicity, the BoD requests are computed based on MAC queue measures, as shown in Fig. 5. In this case, the threshold level of the MAC NRT buffer has to be selected in order to allow the BoD controller to compute the capacity request: considering that the rate-based part of the capacity request is based on the measure of the rate of the NRT packets feeding the MAC NRT buffer, and that the request-assignment cycle duration is equal to T_{FB}, the threshold should allow to collect the packets transmitted at full rate for T_{FB} seconds: by denoting with r_{MAX} the maximum transmission rate of the ST, the threshold must be set equal to $r_{\mathrm{MAX}} \cdot T_{\mathrm{FB}}$. In so doing, even if advantage (i)–(iii) are preserved, the advantage (iv) is lost.

To partially overcome this problem, an additional parameter, α, has been introduced, which is used to lower the reference queue length in the following manner:

$$q_{\mathrm{REF}}(t) = \alpha \int_{t-T_{\mathrm{FB}}}^{t} r_{\mathrm{IN}}(\tau)\mathrm{d}\tau \tag{16}$$

with $\alpha \in [0, 1]$. If $\alpha < 1$, since the NRT MAC buffer queue length tracks a lower reference queue, the queue length and, consequently, the queuing delay are reduced, at the price of some underutilization of the assigned capacity; thus, there is a trade-off between queuing delay and utilization efficiency. The parameter α has been tuned via the demonstrator. Note, however, that the delay reduction is fully achieved only when the actual queue reaches the reference one. The exponential tracking dynamic is driven by the pole $p = -K$ of Eq. (11) (i.e., by the time constant $\tau = 1/K$).

4. Experimental results

The architecture and the algorithms proposed in this paper have been implemented in a demonstration platform. The platform, shown in Fig. 6, is composed of four main parts: two uplink spotbeams are emulated by two user LANs (LAN1, LAN2), shown on the left-hand side of the figure; in the middle, another LAN (SATLAN) emulates satellite system; on the left-hand side, a Gateway, connected to the IPv6 network, is emulated by LAN3. Each user LAN (1 and 2) includes two main Work Stations (WSs) and one SubLAN con-

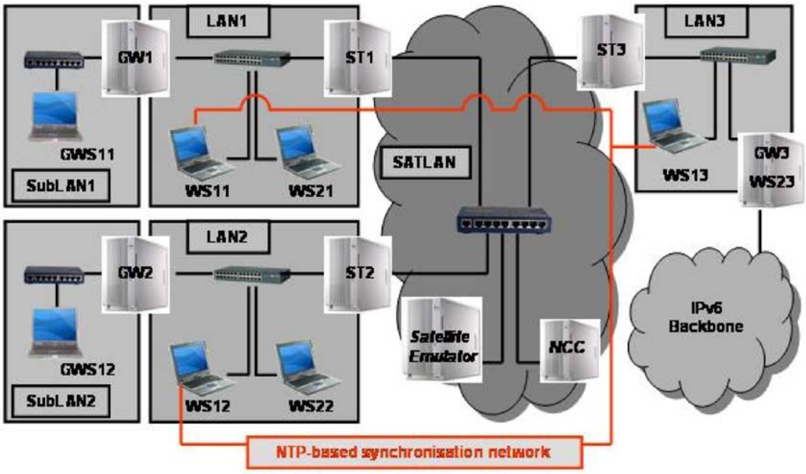

Fig. 6. SATIP6 testbed.

nected to one work station, named GWS; the GWS works as a home Agent and/or Access Router, depending on the studied scenario (as suggested in [24]).

The satellite network is emulated using an Ethernet network connecting the STs (i.e., LAN1 and LAN2), the NCC, and the Satellite Emulator. Finally, LAN3 is used as ISP (Internet Service Provider) network, providing the access to a native IPv6 network, or as a headquarter, used to test multicast applications. A Network Time Protocol (NTP)-based network linking the work stations provides network synchronization.

Thanks to its modular design and implementation, the platform is able to emulate in a realistic and flexible way a complete DVB–RCS system. It is then possible to configure the platform to simulate a system using a regenerative satellite with an on-board switching matrix, as in SATIP6, or to simulate a transparent DVB–RCS system dimensioned around a single Hub, only by changing some configuration files. A complete DVB–RCS protocol stack is implemented, and the modulation/coding part is simulated in real time thanks pre-calculated Bit Error Rate (BER) profiles.

Table 1 shows the software used within the platform testbed.

Each ST avails of a statically assigned bandwidth share, while the remaining transmission capacity can be consumed by dynamic requests: each ST sends its bandwidth request to the NCC according to the BoD algorithm, then the NCC divides the available bandwidth among the STs on the basis of the received requests and sends the bandwidth allocations. If the total amount of the requests is greater than the available bandwidth, the rate assigned to a certain ST is likely to be less than the requested one (*congestion state*); on the contrary, if it is lower, the leftover available bandwidth is equally distributed among the STs as FCA. Obviously, the availability of FCA decreases with the uplink load, while the occurrence of *congestion states* grows.

Table 1
Platform software

Station	Operating system	Software
WS	RedHat 9.0	Gnomemeeting, VideoLanClient, vsftp, Apache, Mozilla
GWS	Windows XP with MIPv6	VideoLanClient, TTCP, ping6
GW	Windows XP with MIPv6	TCP Performance Enhancing Proxy and MIPv6 Home Agent
NCC, SE and ST	RedHat 9.0	SATIP6-developed satellite emulation softwares

The uplink capacity is divided into frames, whose duration is 0.05 s. A super frame is composed of 10 internal frames. The capacity requests are computed every superframe, and the assignments are received every superframe. This means that the BoD scheme of Section 3.2 is sampled with sampling time $T_C = 0.5$ s, equal to the superframe length. In the sampled version of the BoD algorithm, according to stability considerations, the constant K of the capacity request (10) is set equal to $0.5T_C^{-1} = 1\ \mathrm{s}^{-1}$ (see [21] for a detailed discussion).

With these parameters, if α is set equal to 1, the average MAC queuing delay of the NRT packets is about 1250 ms.[1] Even if some advantages of the proposed architecture are still achieved (avoidance of function duplication and avoidance of MAC buffer overflows), the end-to-end delay performances are almost equivalent to the ones obtained with a traditional architecture. Thus, this configuration has been selected as the comparison system. Through a simulation campaign, the parameter α has been set equal to 0.5, in order to effectively reduce the delays while guaranteeing, on average, more than 95% of capacity utilization efficiency.

Recalling that the exponential tracking dynamic is driven by the time constant $\tau = 1/K$, and considering that $K = 1\ \mathrm{s}^{-1}$, the reference tracking is practically achieved only after about $4\tau = 4$ s. This means that the delays of the first packets of the traffic bursts cannot be reduced, and that sporadic transmissions suffer from high queuing delays. The queuing delays of these first packets can be reduced via FCA (as discussed in [17]). Appropriate access mechanisms are under research to cope with this problem, which affects also short-living TCP flows.

Further delay reduction is obtained since the actual capacity requests are a little bit larger than the ones computed by Eqs. (8)–(10), for the reason that the capacity requests have to be expressed in time-slots per frame: if the NCC assigns the full requested capacity, the effect is like availing of some 'implicit' FCA.

The codecs used to generate multimedia traffic are summarized in Table 2. In particular, VideoLan[2] [25] was used to broadcast the video, with the following codecs: Xvid[3] [26] for DivX-like High Quality (HQ) video; DivX 3 Low-Motion and MPEG (Moving Picture Experts Group) for Low Quality (LQ) video.

Figs. 7 and 8 show one of the scenarios used to tune α; in this scenario, the considered input traffic is given by a DivX codec, with average bit rate of 400 kbps (Fig. 7(a)). Fig. 8(a) and (b) show the Cumulative Density Functions (CDF) of the end-to-end packet delays with FCA equal to 0 and 40 kbps (which is 10% of the average bitrate), respectively, and with α set to 1, 0.5 and 0. The figures show that the proposed architecture with the selected $\alpha = 0.5$ significantly enhances the delay performances with respect to the case of $\alpha = 1$, which are representative of the performance obtainable with a traditional architecture, and that further reduction of α are not appreciably effective; in fact, the average delay: for instance, the average delay with $\alpha = 0$ and with $\alpha = 0.5$ during the simulation run is about 1 s, whereas the average delay with $\alpha = 1$ is about 1.7 s. Moreover, the figures show the effects of FCA availability. Fig. 7(b) shows the end-to-end packet delay with FCA = 40 kbps and $\alpha = 0.5$. Other scenarios with different traffic sources (DivX, GSM, Web Browsing, FTP, etc.) and traffic mixes showed analogous results.

Several tests have been executed in order to validate the proposed QoS architecture.

[1] 500 ms of the request-assignment cycle, 250 ms of the transport delay plus, on average, 250 ms that ATM cells have to wait in the NRT MAC buffer for the request opportunity, plus, on average, 250 ms that ATM cells have to wait in the NRT MAC buffer for the transmission opportunity within the superframe. The worst case occurs when the ATM cells wait the whole superframe in the buffer for both request and transmission opportunities: the maximum delay is therefore 1750 ms.

[2] The VideoLAN project targets multimedia streaming of MPEG and DivX files, DVDs, digital satellite channels, digital terrestial television channels and live videos, on a high-bandwidth IPv4 or IPv6 network, in unicast or multicast.

[3] XviD is the first open-source MPEG-4 codec, released under the GPL license.

Table 2
Multimedia codecs

Codec	MPEG-1	Xvid	DivX 3 Low-Motion
Width × Height	320 × 240	672 × 288	360 × 288
Frame Rate (fps)	30	24	25
Bitrate mode	Constant Bit Rate	Variable Bit Rate	Variable Bit Rate
Bitrate (kbps)	1035	1274 (average)	297 (average)
Codec audio	MPEG-1 Layer 2	MPEG-1 Layer 3	PCM Audio
Sampling rate (Hz)	44	44	44
Bitrate (kbps)	128	319	353

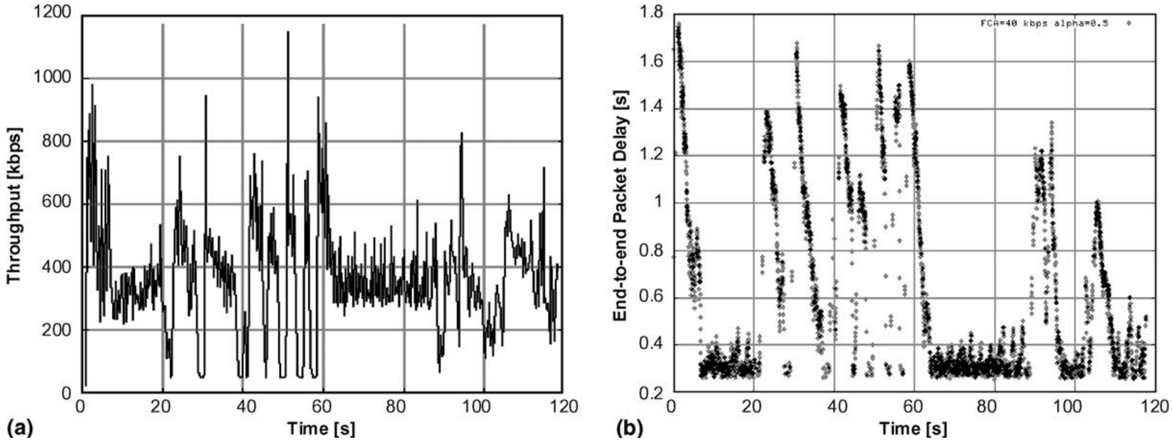

Fig. 7. (a) DivX traffic feeding the satellite terminal (average bit rate = 400 kbps); (b) end-to-end packet delays with $\alpha = 0.5$ and FCA = 40 kbps.

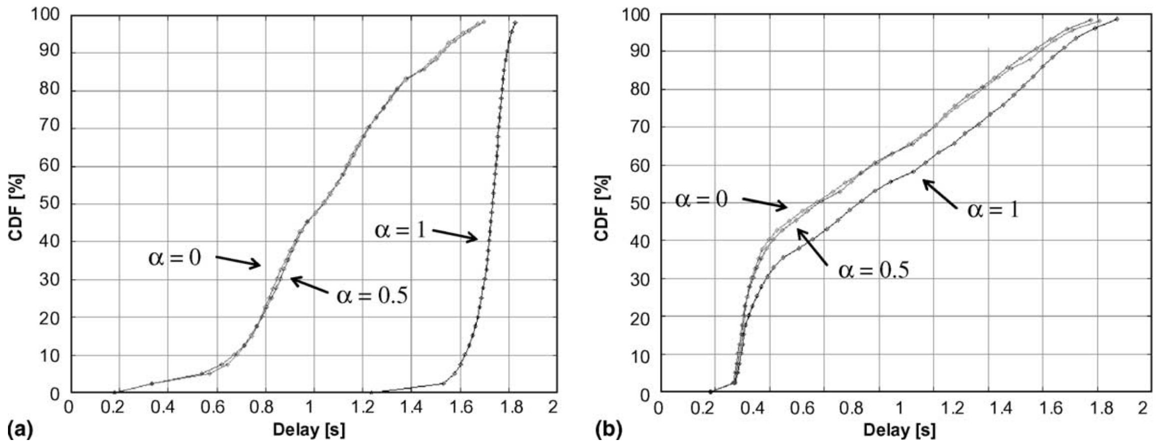

Fig. 8. (a) Cumulative Density Function (CDF) of end-to-end delays with FCA = 0 kbps; (b) Cumulative Density Function (CDF) of end-to-end delays with FCA = 40 kbps.

The first presented test is aimed at evaluating the TCP on BoD behaviour; the test consists in a transferring a file of about 1 Gb using TCP. To cope with the high latency of the satellite network, further increased by the BoD mechanism, TCP parameters have been set as shown in Table 3. Fig. 9(a) and (b) show the throughput of a TCP connection supported by WS11 and the associated capacity requests, respectively.

In the test in question, the acknowledgements (ACK) packets—constituting the feedback required by the TCP to regulate its transmission rate—are transmitted by the receiver ST (i.e., the ST which receives the TCP packets and sends ACKs via the return channel) using available FCA. If no FCA is available, other tests show that the TCP connection is not capable of reaching an acceptable throughput if the ACKs are transmitted via the BoD and no other NRT traffic is present in the destination ST: as a matter of fact, the ACK-generated traffic is sporadic, so that the transmission delay of each ACK packet on the return channel is about 1250 ms. This problem is not present if the receiver ST has other active NRT flows: in this case, ACK packets are multiplexed in the NRT buffer with the other flows and, as discussed above, the queuing delay is effectively reduced by the reference queue (16).

At time t_1 = 100 s and t_2 = 200 s, other WSs log in, so that the available capacity is limited to about 350 kbps. The figure reveals that the TCP connection supported by the BoD algorithm has been able to catch the mean available capacity during the congestion period. After the end of congestion (time t_3 = 420 s, when the other WSs log off), the figure shows that, by exploiting FCA, the BoD scheme achieve low queuing delays: the obtained average Round Trip Time (RTT) of less than 0.6 s, allowing the TCP to transmit almost approximately at 900 kbps. Note that the associated capacity requests perfectly track the TCP throughput (since the payload of an ATM cell is 384 bits, given that the superframe duration is 0.5 s, 1.3 ATM cell per frame equals 1 kbps).

The second test is aimed at showing that dynamic capacity allocation can be fruitfully used also to support video streaming. VideoLan was

Table 3
TCP parameter setting

Maximum segment size (bytes)	1460
Receive buffer (bytes)	128,000
Delayed ACK mechanism	Segment/clock based
Maximum ACK delay (s)	0.2
Slow-start initial count (MSS)	2
Fast retransmit	Enabled
Fast recovery	Enabled
Window scaling	Enabled
Selective ACK (SACK)	Enabled
Nagle's SWS avoidance	Disabled
Karn's algorithm	Enabled
Retransmission thresholds	Attempt based
Initial RTO (s)	2
Minimum RTO (s)	1
Maximum RTO (s)	30
RTT gain	0.125
Deviation gain	0.25
RTT deviation coefficient	4

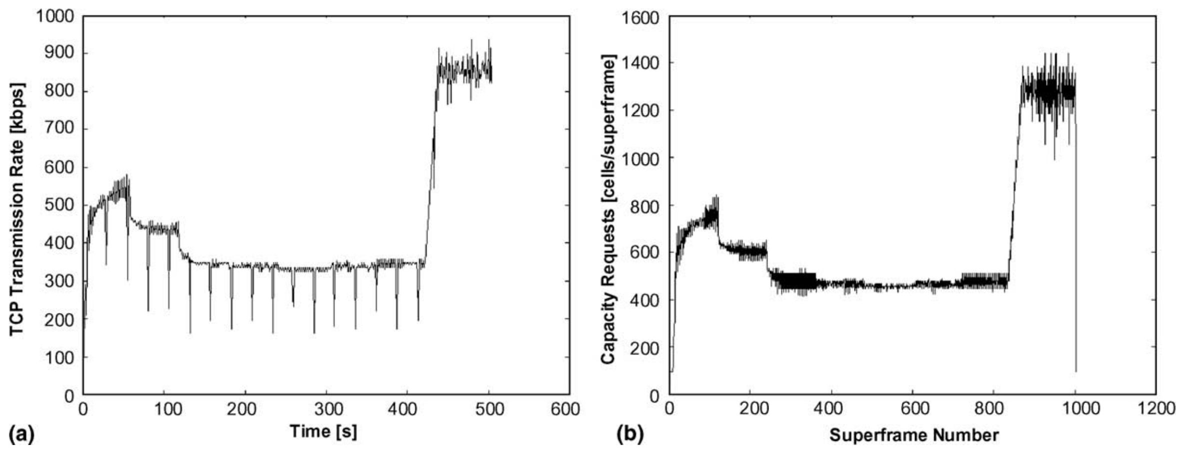

Fig. 9. (a) Measured TCP throughput during congestion; (b) associated capacity requests.

used to broadcast the video, with the XVid codec for DivX-like HQ video, with throughput between 250 Kbits/s and 3.5 Mbits/s, an the MPEG-1 codec for LQ video, with throughput around 1 Mbits/s.

If the network load is negligible (Fig. 10(a)), the HQ video transmission avails of about 300 kbps of FCA; the resulting average end-to-end delay is about 0.5 s and the delay always lower than 1.2 s. If the network load is higher (Fig. 10(b)), no FCA is available, and, consequently, the average delay increases up to almost 1.1 s, with peaks of about 1.6 s. Even if, in both cases, the delay variations are considerable, they are acceptable for Video streaming. In fact, since the required interactivity level is lower than the one required, for instance, for audio conference applications, by using a play out buffer of 2 s the quality will be almost perfect even without FCA.

As shown by Fig. 11, LQ videos experience a lower delay with respect to HQ videos; the reason is that the throughput of the MPEG stream is rather regular with respect to the highly variable XVid throughput. As above discussed, the delay of the first packets suffer from high delay, evident in the beginning of the plots of Fig. 11; the delay of the following packets rapidly decrease due to the effect of the parameter α and of FCA. The final average delay is lowered down to about 300 ms. The effect of FCA is to accelerate the delay reduction, as revealed by the comparison of Fig. 11(a),

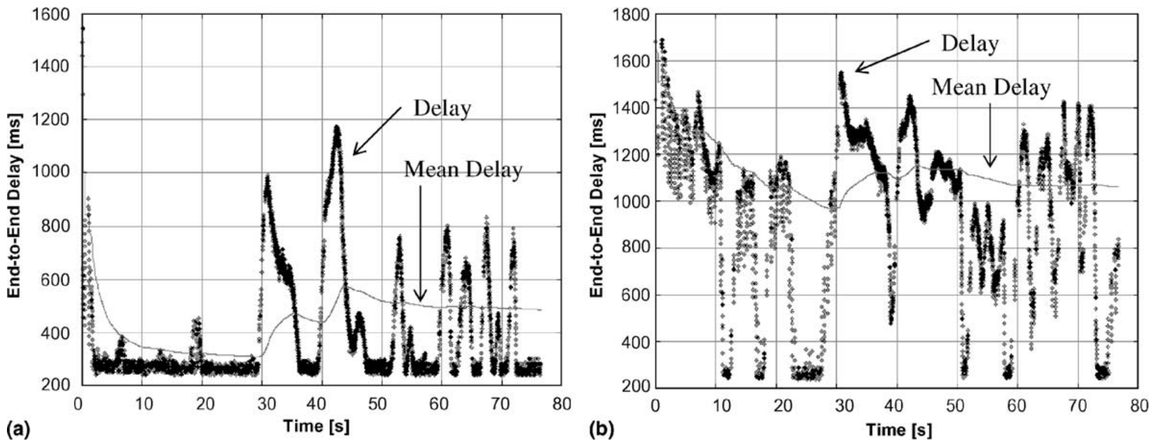

Fig. 10. End-to-end delay (dotted) and mean delay (solid) graphs for Xvid video profile: (a) FCA = 300 kbps, (b) FCA = 0 kbps.

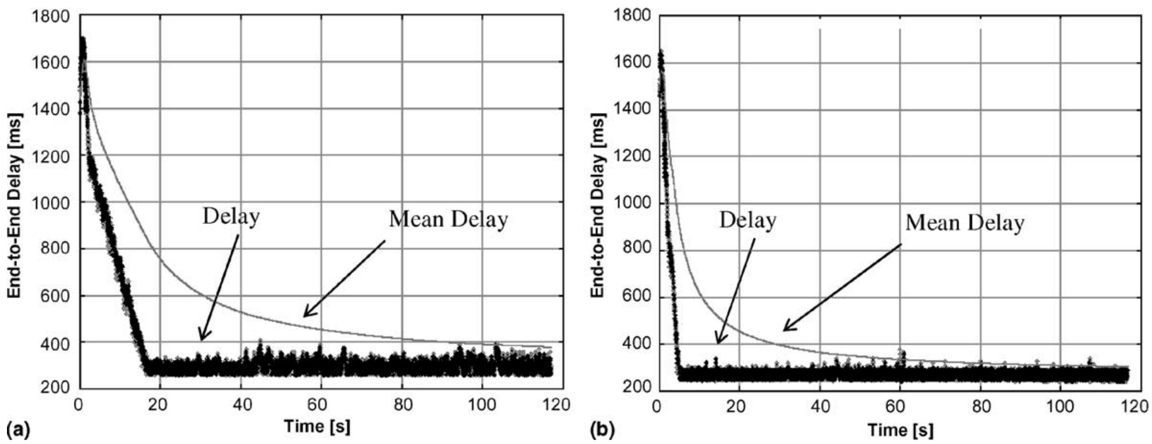

Fig. 11. End-to-end delay (dotted) and mean delay (plain) graphs for MPEG video profile: (a) FCA = 0 kbps, (b) FCA = 300 kbps.

Table 4
Loss percentage comparison with different service classes

Network load [%]	Average delay (Jitter) [ms]			Losses [%]		
	BE	AF	EF	BE	AF	EF
25	293 (26)	293 (26)	283 (23)	0	0	0
50	291 (23)	291 (24)	283 (23)	0	0	0
75	290 (24)	289 (24)	283 (23)	0	0	0
100	919 (23)	948 (24)	283 (23)	0	0	0
125	6753 (23)	1783 (24)	283 (23)	33	0	0
150	6755 (28)	1783 (24)	283 (23)	37	0	0

where the available FCA is only the 'implicit' FCA given by the approximation of the capacity requests, and Fig. 11(b), where 300 kbps of FCA are available.

In conclusion, LQ video transmission is well suited for video conference (because of the low end to end delay and jitter), whereas the second case is more adapted for video diffusion.

Note that in the described tests of LQ and HQ video on BoD no congestions were considered: even in the case with FCA = 0 kbps, the requested capacity was always assigned to the connection. In order to provide QoS guarantees even in presence of congestions, the proposed architecture is capable of differentiating the flow behaviour via the *Traffic Control* algorithms.

With the purpose of testing the performance of the *Traffic Control* procedures, a single reference ST with several different traffic flows (GSM based on the GSM 06.10 codec, DivX 3 Low-Motion LQ video, FTP, Web Browsing) is considered. The heterogeneous traffic flows have been mapped onto three different services: EF, AF and BE. EF provides loss and delays guarantees and is supported by *static bandwidth assignments*; AF offers loss guarantees and is supported by *dynamic bandwidth assignments*, as well as BE traffic, which, however, has no guarantees. The differences among the QoS achieved by the three services are summarized by Table 4.

Table 4 shows that the average delays of the BE and AF classes increase with the network load, while the delay of the EF service remains constant and with low jitter even if the network is overloaded; moreover, Table 4 shows that EF and AF services are effectively protected from packet losses (given that their combined load is less than 100%), while BE service losses increase with the network congestion.[4] Note that the jitter is rather limited also for the BE and AF traffic: the reason is that the network load is high and the queues are always filled up, so that the effect of traffic bursts (higher delay for the first packets with respect to the following ones) is mitigated.

In conclusion, Table 4 shows that the proposed architecture allows a proper differentiation among the services without requiring complex implementation.

5. Conclusions

In this paper, an architecture suitable for QoS provision in broadband TCP/IP over DVB satellites has been proposed. The presented work has been carried out within the IST project SATIP6.

On the one hand, the *access control* aims at requesting to the satellite network the proper amount of time-slots on the DVB frame on the basis of measures of the traffic entering the satellite system. On the other hand, *traffic control* aims at guaranteeing delay and bandwidth requirements to IP flows and users on the basis of the bandwidth allocated to DVB classes by means of the *access control* function. The paper proposes an overall integrated architecture aimed at efficiently coordinating their operation.

[4] Note that the loss percentage of the BE traffic of the considered ST depends also on the BE flows of the other STs: since the NCC follows the *max–min* fairness criterion [27], it assigns the same amount of capacity to the congested flows; thus, small BE flows are less affected by congestions than bigger ones.

The main contribution of this paper is that the proposed QoS architecture has been implemented by the SATIP6 consortium, so that all the proposed algorithms have been tested in a realistic scenario with real applications. The results of the tests performed through the demonstration platform showed that, thanks to the integrated traffic control procedures, the proposed architecture is capable of correctly differentiating the QoS experienced by the IP flows, while, thanks to the efficiency of the resource management procedures allowing dynamic and static capacity assignments, it obtains high exploitation of the available resources.

Further research is needed in the field of TCP on BoD interactions, in order to cope with the sporadic-nature of the acknowledgements and with short-living flows. Moreover, the performance of the implemented architecture can be enhanced by introducing a tighter integration between the *traffic control* and the *access control* modules.

Acknowledgment

The authors wish to thank all the partners of the SATIP6 consortium: Alcatel Space Industries (France), which is the project coordinator, Telecom Italia Lab (Italy), France Telecom SA (France), University of Rome "La Sapienza" (Italy), Sintef (Norway), LAAS/CNRS (France), Alliance Qualite Logiciel (France).

References

[1] A. Jamalipour, M. Marchese, H.S. Cruickshank, J. Neale, S.N. Verma, A.M. Bush, IEEE Journal on Selected Areas in Communications, Broadband IP Networks via Satellites—Part I 22 (2) (2004).

[2] A. Jamalipour, M. Marchese, H.S. Cruickshank, J. Neale, S.N. Verma, A.M. Bush, IEEE Journal on Selected Areas in Communications, Broadband IP Networks via Satellites—Part II 22 (3) (2004).

[3] M. Karaliopoulos, R. Tafazolli, B.G. Evans, Providing differentiated service to TCP flows over bandwidth on demand geostationary satellite networks, IEEE Journal on Selected Areas in Communications 22 (2) (2004) 333–347.

[4] M. Sooriyabandara, G. Fairhurst, Dynamics of TCP over BoD satellite networks, International Journal of Satellite Communications: Special Issue on IP QoS for Satellite IP 21 (2002) 427–449.

[5] A. Iera, A. Molinaro, S. Marano, IP with QoS guarantees via Geo satellite channels: performance issues, IEEE Personal Communications 8 (3) (2001) 14–19.

[6] ETSI, SES Broadband Satellite Multimedia (BSM) Group; <http://www.etsi.org/>.

[7] E. Rosen, A. Viswanathan, R. Callon, Multiprotocol label switching architecture, RFC 3031, January 2001.

[8] R. Braden, D. Clark, S. Shenker, Integrated services in the internet architecture: an overview, Internet Engineering Task Force, RFC 1633, June 1994.

[9] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, An architecture for differentiated services, RFC 2475, December 1998.

[10] L. Georgiadis, R. Guerin, V. Peris, K.N. Sivarajan, Efficient network QoS provisioning based on per node traffic shaping, IEEE/ACM Transactions on Networking 4 (4) (1996) 482–501.

[11] A. Demers, S. Keshav, S. Shenker, Analysis and simulations of a fair queuing algorithm, Journal of Interworking: Research and Experience 1 (January) (1990) 3–26.

[12] S. Floyd, V. Jacobson, Link-sharing and resource management models for packets network, IEEE/ACM Transactions on Networking 3 (4) (1995) 365–386.

[13] Ion Stoica, Hui Zhang, T.S. Eugene Ng, A hierarchical fair service curve algorithm for link-sharing, real-time and priority service, in: Proceedings of SIGCOMM'97, 1997.

[14] V. Firoiu, J. Kurose, D. Towsley, Efficient admission control for EDF schedulers, in: Proceedings of IEEE INFOCOM '97, April 1997, pp. 310–317.

[15] Internet Engineering Task Force-TCP Over Satellite Working Group, <http://tcpsat.grc.nasa.gov/tcpsat/>.

[16] DVB–RCS: Digital Video Broadcasting (DVB); Interaction Channel for Satellite Distribution Systems, ETSI EN 301 790 v.1.3.1 March 2003. Available from: <http://www.etsi.org/>.

[17] P.D. Mitchell, D. Grace, T.C. Tozer, Burst targeted demand assignment multiple-access for broadband internet service delivery over geostationary satellite, IEEE Journal on Selected Areas in Communications 22 (3) (2004) 546–558.

[18] L. Chisci, R. Fantacci, T. Pecorella, Predictive bandwidth control for GEO satellite networks, in: Proceedings of the IEEE International Conference on Communication (ICC), vol. 7, 20–24 June 2004, pp. 3958–3962.

[19] G. Açar, C. Rosenberg, Weighted fair bandwidth-on-demand (WFBoD) for geostationary satellite networks with on-board processing, Computer Networks 39 (1) (2002) 5–20.

[20] F. DelliPriscoli, A. Pietrabissa, Resource management for ATM-based geostationary satellite networks with on-board processing, Computer Networks 39 (1) (2002) 43–60.

[21] F. Delli Priscoli, A. Pietrabissa, Design of a bandwidth-on-demand (BoD) protocol for satellite networks modeled as time-delay systems, Automatica 40 (5) (2004) 729–741.

[22] M. Karaliopoulos, R. Tafazolli, B. Evans, On the interaction of TCP with BoD in GEO broadband satellite networks, in: Proceedings of IEEE GLOBECOM'02, vol. 3, 17–21 November 2002, pp. 2945–2949.
[23] O.J. Smith, A controller to overcome dead time, ISA Journal 6 (2) (1959) 28–33.
[24] S. Combes, S. Josset, P. Very, I. Buret, T. Zein, Network scenario for IP-dedicated satellite access scheme, in: Proceedings of the 8th Ka-Band Utilization Conference, Baveno, Italy, 25–27 September 2002.
[25] VideoLAN (Free Software and Open Source video streaming solution), <http://www.videolan.org>.
[26] XviD (ISO MPEG-4 compliant video codec), <http://www.xvid.org>.
[27] P. Gevros, J. Crowcroft, P. Kirstein, S. Bhatti, Congestion control mechanisms and the best-effort service model, IEEE Network 15 (3) (2001) 16–26.

Antonio Pietrabissa graduated in Electronic Engineering from the University of Rome "La Sapienza", in 2000, where he received his Ph.D. in System Engineering in 2004. Currently, he holds a research contract with the University of Rome "La Sapienza" and works in European Union research projects within the fifth and sixth framework research programmes in the area of IST (Information Society and Technologies). His research focus is the application of system and control theory methodologies to telecommunication networks, with specific interest to the design of connection admission control procedures, congestion control mechanisms and medium access control protocols for multimedia broadband satellite systems, wireless networks and next generation heterogeneous networks. He is the author of more than 20 papers on these topics.

Tiziano Inzerilli received his Ph.D. in Computer Engineering at the University of Rome "La Sapienza" in the year 2004 defending a thesis on efficient traffic control in wireless TCP/IP networks. Since year 2000 he has carried out research EU projects, within the fifth and sixth framework research programmes in the area of IST (Information Society and Technologies), on WLAN and broadband satellite access for TCP/IP applications, TLC middleware and network support for pervasive systems over integrated wireless technologies. He has published in major IEEE magazines and conferences on telecommunication, as well as on several European conferences.

Olivier Alphand received his Engineering Degree from the Ecole Nationale Supérieure d'Ingénieurs de Constructions Aéronautiques (ENSICA), Toulouse, France in 2002. He is currently working towards his Ph.D. at LAAS/CNRS, France, in collaboration with Alcatel Space. After preliminary performance studies on wireless technologies (especially IEEE 802.11), he has been primarily involved in the European IST SATIP6 project. His research interests are in the areas of QoS provisioning in next generation satellite networks and the support of multimedia traffic in QoS-enabled wireless networks.

Pascal Berthou received his Ph.D. in Computer Science and Telecommunication from National Polytechnic Institute of Toulouse in 2001. He is an associate professor at the Toulouse University of Science. He joins the Laboratory for Analysis and Architecture of Systems of the French National Center for Scientific Research (LAAS/CNRS) in 1998 as a research staff member, where he works in the area of high-speed networks and protocols and multimedia communications. Since this date he is working especially in the domain of satellite communications. He served as technical co-chair for the Broadband Satellite Communication Systems workshop of the IFIP 18th World Computer Congress.

Thierry Gayraud received his Ph.D. in Electrical Engineering in 1992 from the University Paul Sabatier, Toulouse, France, where he is currently an associate professor, researcher in Laboratory for Analysis and Architecture of Systems (LAAS/CNRS). His current research interests are in design, implementation and deployment of transport protocols offering guaranteed Quality of Service (QoS) to distributed multimedia applications such as tele-teaching and Computer Supported Co-operative Work (CSCW) over wireless heterogeneous networks (from sensor to satellite networks).
He has more than 30 publications in international journals and conferences. He was co-chair of WCC2004 BSCS Workshop held last year in Toulouse. He was involved in several research projects [European (GCAP) or French (@IRS++)] and the scientific leader in LAAS in the national French RNRT project Amarrage and in the European IST project SATIP6.

Michel Mazzella is a graduate of the Ecole Speciale de Mecanique et d'Electricite in Paris. He has been with Alcatel since 1972, when he joined the Switching Division, leading to his appointment as engineering manager for time-division telephony exchanges. Since 1982 he has been with the Space Division where his work began with ground systems design for telecommunication networks in programmes such as the Data Relay Satellite, IOC, HELIOS, HERMES, and SPOT4. He moved to the research group for new systems in 1991 and focused on network aspects for mobile satellite communications. From 1993 he has been responsible for management of mobile satellite projects, notably as project manager of the SAINT project in RACEII, SINUS in ACTS and BRAHMS and SATIP6 in IST.

Eddy Fromentin received his Ph.D. from University of Rennes 1, France in 1996 on Distributed Automatic Debugging topic. After working two years in University of Rennes as an assistant professor on fault tolerance subject, he moved to industry as a system consultant for three years. He is now an employee of SILICOMP-AQL, France as a network consultant. His current main interests are IPv6 networks and Quality of Service.

Fabrice Lucas received his Computer Science degree in Software Engineering in 1995 from the Institut de Formation Supérieure en Informatique et télé-Communications (IFSIC Rennes). He has been working in SILICOMP-AQL as a project manager in various fields among which wireless networking, embedded computing and multimedia systems. He also contributed in several IST projects dealing with QoS, such as WINE (wireless IP-based architectures), 6HOP (heterogeneous multihop wireless networks), and SATIP6.