Available online at www.sciencedirect.com
SCIENCE DIRECT®
Computer Networks 50 (2006) 2088–2111
Computer Networks
www.elsevier.com/locate/comnet

On-line bandwidth control for quality of service mapping over satellite independent service access points

Mario Marchese *, Maurizio Mongelli

DIST—Department of Communication, Computer and System Sciences, University of Genoa, Via Opera Pia 13, 16145 Genoa, Italy

Available online 8 November 2005

Abstract

The *Quality of Service* (QoS) provision requires the cooperation of all network layers from bottom to top. More specifically, for *Broadband Satellite Multimedia* (BSM) communications, the physical layers (strictly satellite dependent) are isolated from the rest by a *Satellite Independent Service Access Point* (SI-SAP), which should offer specific QoS to IP and the upper layers. The structure of SI-SAP and the BSM protocol model "opens the door" to the problem of mapping the performance requests of *Satellite Independent* layers over *Satellite Dependent* technology. In such a context, a *QoS mapping* problem arises when different encapsulation formats are employed along the protocol stack of the SI-SAP interface, for instance, when IP packets are transferred over an ATM, DVB or MPLS core network. In this perspective, we investigate here a novel control algorithm to estimate the bandwidth shift required to keep the same performance guarantees, independent of the technology change. By exploiting *Infinitesimal Perturbation Analysis* to capture the network performance sensitivity, we obtain an adaptive control law suitable for on-line control on the basis of traffic samples acquired during the network evolution. Owing to the generality of the mathematical framework under investigation, our control mechanism can be generalized for other network scenarios and functional costs.

Keywords: QoS interworking architecture; Satellite communications; QoS mapping; Bandwidth tuning; On-line network control; Fading countermeasures

1. Introduction

The definition of "heterogeneous network" may assume different aspects. Network portions may use different transmission means such as cable, satellite, radio, and may implement different protocols. From a user perspective, an end-to-end QoS in a satellite/wireless/terrestrial network depends on the QoS achieved at each layer of the network based on functions performed at the layer interfaces whose characteristics take a topical role. ETSI-BSM (*Broadband Satellite Multimedia*) QoS architecture (defined in Refs. [1–4]) is a good example and it is the basis of this paper, which concerns satellite networks.

QoS requires the cooperation of all network layers from bottom to top. At each layer, the user performance requirements are achieved by using efficient technologies and counteracting any factor affecting performance degradation. The idea is creating a bandwidth pipe to be provided at the layers'

* Corresponding author. Tel.: +39 010 3532985; fax: +39 010 3532154.
E-mail addresses: mario.marchese@unige.it, mopa@dist.unige.it, dama@dist.unige.it (M. Marchese).

1389-1286/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2005.10.006

interface (called *Service Access Point*) and efficiently used by the upper layers. In this view, the protocol stack of the gateway at the entrance of the BSM network should be seen as a *Protocol Optimization Through Layer* because it involves each specific protocol, which cooperates to guarantee a specific QoS level over the satellite network portion. The considered protocol stack separates the layers identified as *Satellite Dependent* (SD) and *Satellite Independent* (SI). SD layers depend on the physical implementations and are often covered by industrial copyright. SI layers are composed of IP and upper layers. The interface between SI and SD is defined by ETSI through SI-SAPs (*Satellite Independent-Service Access Points*). QoS requirements must flow through SI-SAPs and be implemented at the SD layers.

In this perspective, the paper introduces a control scheme that allows matching QoS requirements between the SI and SD layers, without any closed form-formula for the performance metric under investigation. The key point is measuring the *equivalent bandwidth* shift imposed by the possible change in the encapsulation formats throughout the layers (e.g., from IP implemented at SI layers and ATM, used as the SD transport technology). The main idea is to have a control algorithm that can correctly dimension the necessary bandwidth shift but does not need any closed-form formula of the performance metric and is able to react under variable system conditions. We estimate the bandwidth shift needed to support IP over ATM service by assuming *IP Packet Loss* as the performance metric.

Some studies report that a possible evolution over the next few years will include a core network composed of ATM (or MPLS) technology connecting border IP portions. It is a widespread view that "...capital expenditure constraints in both service providers and enterprises will mean that MPLS will evolve in the carrier core network first, with ATM remaining for some time to come as the primary technology for multiservice delivery in bandwidth-limited edge and access networks" [5]. In such a context, the BSM architecture, taken as a reference in this paper, can therefore be considered as representative for the incoming network protocol infrastructures of future satellite communications.

The remainder of the paper is organized as follows. Section 2 addresses the state of the art of optimization techniques for telecommunication networks and the need to study adaptive control techniques, suitable for on-line network control. Section 3 summarizes the role of the information transport technologies and describes the BSM protocol architecture. The solution to map QoS over SD technology is proposed in Section 4, taking ATM as an example solution at layer 2 [6,7]. Section 5 is devoted to presenting our control methodology to map the QoS between the SI and SD layers of the BSM architecture. Performance evaluation is the object of Section 6 and, in Section 7, we conclude by summarizing the obtained results and emphasizing directions for future research.

2. State of the art in network optimization and control

Optimization techniques for telecommunication networks usually follow the minimization of a proper functional cost capturing the system performance:

$$^{\mathrm{Opt}}\boldsymbol{\theta} = \arg\min_{\boldsymbol{\theta}} \underset{\omega}{E}\{L[\boldsymbol{\theta}, \omega]\}, \tag{1}$$

where $L[\boldsymbol{\theta}, \omega]$ is the performance index of interest (e.g., blocking probability of connection requests, packet loss probability, packets' mean delay or delay jitter) and $\boldsymbol{\theta}$ is the vector of the decision variables related to the available resources (e.g., service rate, buffer size). The expectation $\underset{\omega}{E}\{\cdot\}$ is provided over all the feasible sample paths ω of the system (i.e., the set of all possible realizations of the stochastic processes involved in the system). Such problems are often solved by means of centralized approaches, in which the control systems are strictly based on closed-form expressions for the performance measure (see, for example, [8–11] and references therein, for what concerns *call admission* and *bandwidth control*, *routing* and *pricing* issues). For instance, in [10], the *Tsybakov–Georganas* formula for the *packet loss probability* in the presence of self-similar traffic is used and in [9,11] the *blocking probability* of call requests is computed through the *Erlang B* formula. Such optimization approaches act according to a *parameter adaptive certainty equivalent control* (PACEC) [12], where a mapping between the current statistical behaviour of the system and the parameters of the functional costs must be periodically performed on line, in order to maintain good performance of the resource allocation algorithms. The main drawback of these approaches is that the conditions for the applicability of closed-form functional costs are hardly matched in real contexts. Not only are closed forms for important performance parameters (e.g., mean delay and delay jitter of the packets) not always available (for example in the presence of self-similar

traffic), but also "... even under Markovian assumptions for processes of queuing systems, there are only limited cases where closed-form expressions can be obtained" [13]. In general, it is difficult to assure that the necessary hypotheses are verified in a real application. PACEC also requires a strong on-line consumption of computing power, due to the continuous on-line minimization of a global cost through the adoption of a proper mathematical programming algorithm. *Dynamic Programming* (DP) is often essential [10,14], but it gives rise to the *curse of dimensionality* phenomenon whose related computational burden severely affects the application itself of DP in real time environments. On the other hand, when the stochastic environment does not allow the derivation of closed-form expressions of $L[\boldsymbol{\theta}, \omega]$, the situation is complicated by the need for estimating $\underset{\omega}{E}\{L[\boldsymbol{\theta}, \omega]\}$. It generally requires *Monte Carlo* simulation approaches that are not realistic if adaptive reactions are pursued in real time.

Control theory helps provide active measurement techniques suitable for on-line network management, thus avoiding analytical models of network phenomena. In the last few years, a wide range of control techniques have emerged to optimize the network performance (see, e.g., [15] and references therein), especially in wireless environments, where the available bandwidth is a more valuable resource than for wired ones [16]. Traditional PID (*proportional integral derivative*) controllers, despite the relatively simple approach (almost primitive, in eyes of some control theorists), often yield good performance, in particular if we look at the *Active Queue Management* and TCP control research fields (see, e.g., [17,18]). On the other hand, more sophisticated techniques (e.g., fuzzy tuning, neural networks, identification methodologies, learning theory) have been studied and applied to obtain optimized performance (see, for instance, [16,19,20] and references therein). In this work, we follow the principles of *stochastic approximation* [21] by using a *sensitivity estimation* technique, recently developed in the field of *Perturbation Analysis* (PA) for *Discrete Event Systems* (DESes). More specifically, we tackle the *QoS mapping* problem arising at the SI-SAP interface of the satellite BSM architecture by means of a novel control methodology, driven by a gradient estimation procedure taken from [22]. Emphasis is put on on-line, quick reactions to non-stationary conditions involving time-varying channel degradation and traffic variability. The key idea is to develop a resource allocation algorithm, suited for the SI-SAP interface management, that is *light* (requiring a small on-line computational effort), *adaptive* (able to "learn" the statistical changes of the system), and *non-parametric* (able to optimize the system performance without any performance metric in closed form).

3. Broadband satellite multimedia protocol architecture

As mentioned before, a QoS-based service derives from reliable physical layers (including, in this case, layer 2 and layer 1) that can offer specific services to the upper layers. In general, the idea is that QoS provision requires the cooperation of all network layers. The services offered by the lower layers should provide the necessary QoS mapping at the higher layers. Fig. 1 shows a graphical model of the relation between lower (physical) layers and higher layers. The connections (or bundles of them) are forwarded down to a physical interface that transports the information along a channel. More specifically for BSM communications, where the reference is the support of services offered by the UDP-TCP/IP suite over a satellite bearer service, the protocol architecture is defined in Refs. [1,2] and is reproduced in Fig. 2.

The physical layers (i.e., satellite physical, MAC and Link Control, strictly satellite dependent) are isolated from the rest by a *Satellite Independent Service Access Point* (SI-SAP), which should offer specific QoS services to the upper layers. The SI-SAP acts over BSM bearer services and decouples SI from SD *adaptation functions*. Concerning BSM QoS, it is important to consider the restrictions imposed by the air interface and to define robust and efficient QoS mechanisms. User traffic (actually the traffic classes) should be connected with an IP QoS mechanism and, in sequence, mapped into a BSM bearer service with a suitable QoS guarantee.

The service offered by SI-SAP should be accurately defined to assure a QoS-based service. The definition is not yet complete. Basic documents are Refs. [23,24]. QoS is explicitly considered within the BSM architecture through the *Queue Identifier* (QID) that is a SI-SAP parameter specification for an abstract queue. In practice, SI layers can ask for a specifically guaranteed service in terms performance parameters and SD layers have the role of mapping the QoS request over the real transport technology. Primitives allow opening, modifying, and closing the resource reservation.

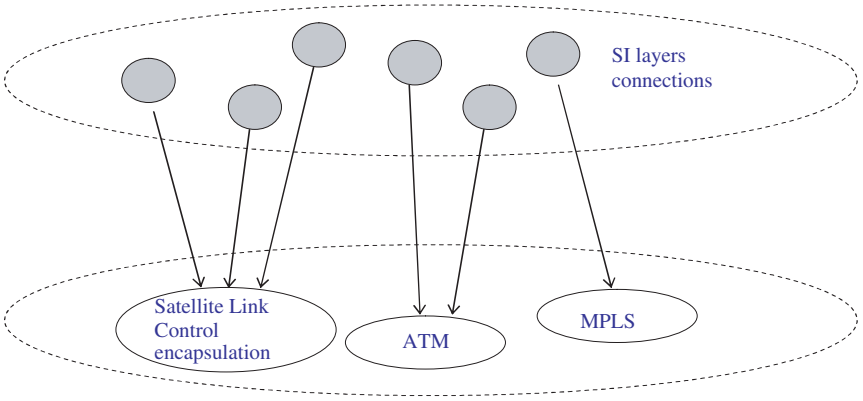

Fig. 1. Higher *Satellite Independent* (SI) layers over some possible *Satellite Dependent* (SD) transport technologies.

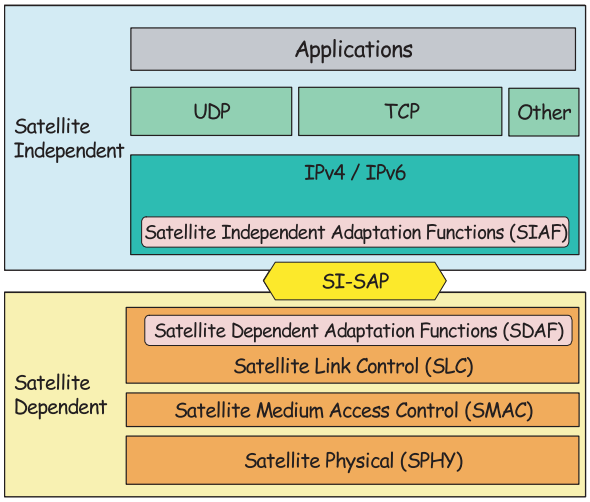

Fig. 2. ETSI BSM Architecture: the protocol stack [1,2].

4. QoS mapping over BSM technology

In such a context, there are two main problems: the change of information unit (encapsulation) and the different features offered by the technology with the consequent need for traffic aggregation. An example concerning the former is the transport of IP traffic over a BSM portion that uses ATM as the SD layer transport technology (as often implemented in industrial systems [7]). Concerning the latter, as clearly indicated in [4], "it is accepted in the BSM industry that at the IP level (above the SI-SAP interface) between 4 and 16 queues are manageable for different IP classes. Below the SI-SAP these classes can further be mapped into the satellite dependent priorities within the BSM which can be from 2 to 4 generally". So, even if, at the IP layer, BSM queues can be defined by QoS specific parameters, including IntServ flowSpecs, Path labels and DiffServ marking [23], in practice, the DiffServ approach is the business and implementation choice at the SI layer. Moreover, the association of IP QoS classes to SD transfer capabilities (e.g., ATM queues, MPLS queues, DVB-RCS request categories) is strongly limited by hardware implementation constraints.

A queue model describing the SI–SD mapping for the SI-SAP interface, similarly used in Ref. [6], is given in Fig. 3. The first block (IP traffic classifier) assigns the arriving IP packet to a queue depending on the DSCP (*DiffServ Code Point*) value. The example shows one EF class, four AF classes and one BE class. Below the DiffServ scheduler that picks the packets up from the queues, an AAL5 adapter block encapsulates the IP packet within the AAL5 frame and creates the ATM cells. The ATM portion is satellite dependent and, in this example, implements three ATM queues acting at MAC layer: one queue dedicated to EF traffic, one for all the AF and one for BE. The problem is how much bandwidth must to be assigned to each ATM queue so that the IP-based SI SLA (*Service Level Agreement*, i.e., the performance expected) is guaranteed. SLA is expressed in terms of objective metrics such as loss, delay and jitter. The control algorithm studied in this work considers the *IP Packet Loss Probability* (PLP) as the only performance metric, both for SI and SD layers, but can be also generalized for other metrics.

In brief, there are two components to evaluate: the "change of encapsulation format" and the "need to aggregate traffic". Only the former is the object of study in this work. The latter is currently the subject of ongoing research, following the same optimization methodology investigated here.

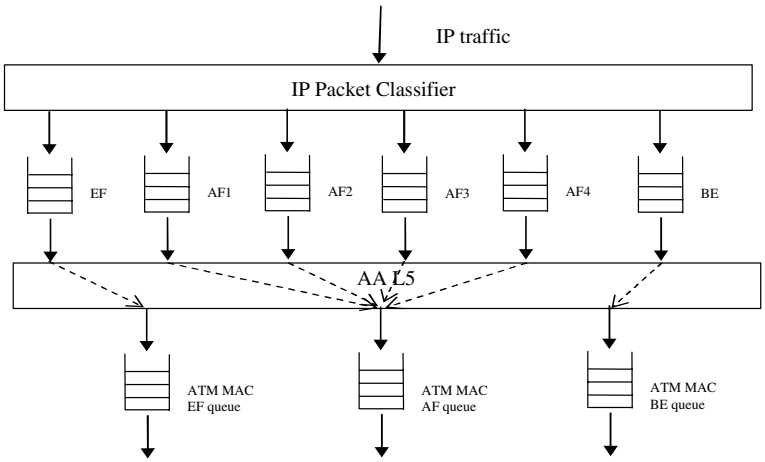

Fig. 3. SI-SAP interface. SI layer (IP DiffServ-based) over SD layer (ATM-based) [6,7].

In the QoS mapping scenario we are considering, an IP packet (coming from the SI layer) is lost if at least one of the ATM cells (of the SD layer) that compose the packet is lost. In this context, two issues arise: the first one concerns how the *ATM Cell Loss Probability* (CLP) can influence, in the SD core, the PLP. The second one concerns how much bandwidth must be reserved in the ATM tunnel to preserve the SI-based SLA. An example may help. If a DiffServ flow at the SI layer receives a bandwidth allocation of, e.g., 1.0 Mbps in order to guarantee a PLP of 1×10^{-2}, when the flow is mapped in the SD core, we need to answer the following questions:

- What is the CLP* that must be guaranteed in the ATM tunnel to keep PLP $\leqslant 10^{-2}$?
- How much bandwidth must be reserved for the ATM tunnel to get CLP*?

The control algorithm we are introducing avoids decoupling the problem and provides the bandwidth that must be allocated to the ATM SD layer to keep the performance metric (the SLA that is the IP PLP, in this case) under a given threshold by using only IP PLP losses measured on line at the SD layer. The performance threshold may be time variant. Actually, the control algorithm may follow a given limit (e.g., 10^{-2}), but may also chase a dynamic QoS threshold. To stress this dynamic characteristic, many tests have been performed by imposing the packet losses measured at the IP layer as the performance threshold of the ATM SD layer.

4.1. The QoS mapping concept: state of the art

The ATM over IP environment, still relevant and timely due to the widespread diffusion of ATM backbones, may be extended to other encapsulation formats. The *Internet Engineering Task Force* (IETF) is currently developing several techniques to manage the interworking between ATM, IP and MPLS flows (see, e.g., [25,26]). In particular, the *Pseudo Wire Emulation Edge-to-Edge* (PWE3) IETF group is working on protocols for transporting layer 2 services (which, in the PWE3 view, also includes ATM) over IP-based network. This includes not only the transport of ATM over MPLS, but frame relay, circuit emulation, synchronous optical network (SONET), and Ethernet over MPLS, too. More recently, in the BSM environment, *Digital Video Broadcasting* (DVB) may be a possible implementation choice for the SD layer but still involves standardization actions.

Refs. [27,28] of IETF and [29] of the ATM Forum are the most important proposed standards for the mapping of QoS declarations between different QoS technologies. They mainly concern IP to ATM mapping operations and investigate which ATM service categories must be chosen to support different QoS IP service classes. In Ref. [30] the mapping of IntServ over ATM is studied, showing that the ATM nrt-VBR service class gives more bandwidth saving than the CBR service class. Ref. [31] shows by simulations that the IntServ-ATM mapping of [28] causes an excessive cell loss rate. These works have in common the investigation of the alteration of the QoS parameters when a flow

changes its physical transport technology. In particular, [32] highlights the problem of different encapsulation formats in IP and ATM. This can have a serious effect on the bandwidth requirement. It also improves the efficiency of the system by slightly modifying the AAL5 structure. Nevertheless, the problem of tracking the equivalent bandwidth shift due to the change in the encapsulation format is still an open issue.

4.2. The notion of equivalent bandwidth and its application at the SI-SAP interface

Equivalent bandwidth is defined as the minimum amount of network resources to be allocated to guarantee a fixed degree of performance. Applied to the SI-SAP context, it is the minimum bandwidth necessary at the SD layer to keep the PLP below the threshold. Traditionally, equivalent bandwidth techniques are based on the statistical characterization of the traffic, described in terms of descriptors (*peak rate*, *mean rate*, *maximum burst size*). Unfortunately, in the SI-SAP case (Fig. 2) the ingoing flow to SD buffers is the outgoing flow from SI buffers and the aforementioned traffic descriptors may hardly be applied. If such IP traffic statistics were perfectly known (or perfectly estimated), heuristics could be applied to establish a proper SD bandwidth allocation but, as said, traffic estimation in this case is a hard task and the necessary approximations could lead to allocation errors. For this reason, we investigate an adaptive control scheme that does not use any knowledge of the IP traffic statistics and is based only on active measurements (or a priori reference values) for the loss volumes. These are the only information the proposed algorithm needs.

We also suppose that no information is available about the buffer dimension at the ATM layer 2, since it is a SD layer and its implementation is often covered by industrial copyright. In practice, the bandwidth assigned to the SD queues is allocated dependent on the packet loss measured in the queues themselves. As already said, the method may be applied comparing the measure of SD losses either with a loss volume derived from a reference value as, for example, the loss imposed by the SI SLA, or with another measure, as the loss measured at the SI buffer whose traffic enters the SD layer. The latter is prefered in this paper because it allows highlighting the algorithm's dynamic behaviour but results in both cases will be reported in the performance analysis section.

5. The control mechanism

In order to manage the bandwidth provisioning at the SI-SAP interface with unknown IP traffic statistics and unavailable closed-form formulas for the mapping of PLP over CLP, we proceed by formulating a control scheme driven by *Infinitesimal Perturbation Analysis* (IPA). IPA is a sensitivity estimation technique for *Discrete Event Systems* (DESes) [13,22,33]. It is based on the observation of the sample paths followed by the stochastic processes of a DES and gives an estimation of the derivative of the chosen performance metric. As shown in [22,33], IPA algorithms applied to a *stochastic fluid model* (SFM) of a packet network are only based on monitoring the buffer status and do not need any explicit feedback concerning the number of active sources and their statistical behaviour.

As said before, we propose to measure the IP packet loss volume within the SD core and within the SI buffers or, as an alternative to this, to have an off-line performance threshold value, and to get an SI reference loss volume over time as shown in the next section. Given such on-line measurements, the algorithm proposed adapts the bandwidth assigned at the SD layer so as to keep the SD loss volume below the SI loss volume. QoS is guaranteed, independently of the change in the transfer mode.

5.1. The derivative estimation of the loss performance metric

To do this, we recall a derivative estimation procedure for the performance index (the packet loss) defined in the SLA. With a notation that slightly differs from [22], we adopt a SFM for each of the modeling buffers, shown in Fig. 3. Each buffer has a finite-capacity buffer of fixed size c. Fig. 4 shows the formal model of a buffer.

The stochastic processes associated with this model and essential for the optimization procedure are: $\alpha(t)$, the input flow rate (*inflow*) process into the SFM; $x(t)$, the buffer occupancy process (the

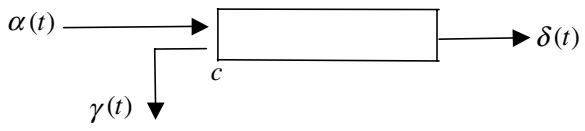

Fig. 4. *Stochastic Fluid Model* (SFM) of the SI-SAP interface traffic buffers.

fluid portion already in service is not included in SFMs); $\delta(\theta,t)$, the fluid *outflow* rate process, namely

$$\delta(\alpha(t),\theta(t),t)=\begin{cases}\theta(t) & \text{if } x(t)>0,\\ \alpha(t) & \text{if } x(t)=0,\end{cases}$$

where $\theta(t)$ is the buffer service rate; and $\gamma(\theta(t),t)$, the loss rate (*overflow*) process

$$\gamma(\alpha(t),\theta(t),t)=\begin{cases}\alpha(t)-\theta(t) & \text{if } x(t)=c,\\ 0 & \text{otherwise.}\end{cases}$$

Only the time dependency will be explicitly reported in the remainder of the paper for the stochastic processes to avoid excessive notational burden.

Since our aim is to get a control scheme for the loss probabilities of the SI packets on the SD core, we make use of the following IPA performance measure: the *loss volume* $L_V(\cdot)$ over a time interval $[0,T]$. It is defined as

$$L_V(\theta)=\int_0^T \gamma(\theta,t)\,\mathrm{d}t. \tag{2}$$

Our first goal is to obtain a derivative estimate of the performance metric, $L_V(\theta)$, with respect to the service rate $\theta\in\Re^+$. Let B_κ be an "active" period of the buffer between two instants of bandwidth reallocation, namely, a period of time in which the buffer is not empty ($x(t)\neq 0$). Let ξ_κ be the starting point of B_κ and v_κ be the instant of time when the last loss occurs during B_κ. For every θ, it can be shown [22] that

$$\frac{\partial L_V^\kappa(\theta)}{\partial\theta}=-(v_\kappa(\theta)-\xi_\kappa(\theta)). \tag{3}$$

The contribution to the derivative in each active period B_κ, during which some losses occur, is the length of the time interval elapsed from the beginning of B_κ to the last instant in B_κ where the buffer is full. Denoting with N_B the number of active periods during an observation window (for instance between two consecutive service rate reallocations of the buffer), an estimation of the derivative performance can be obtained as (see Fig. 5):

$$\frac{\partial L_V(\theta)}{\partial\theta}=\sum_{\kappa=1}^{N_B}\frac{\partial L_V^\kappa(\theta)}{\partial\theta}. \tag{4}$$

Fig. 5 illustrates a possible buffer sample path and related IPA estimator computation. The IPA derivative estimator (3) is also known as "non-parametric", since it is computable directly from an observed sample path of the system (denoted by ω) without any further knowledge concerning the probability distributions of the stochastic processes involved in the system. "The form of the IPA estimators is obtained by analyzing the system as a SFM, but the associated values are based on real data" [33]. This implies that (4) is applicable not only in off-line simulations (aimed, for example, at planning telecommunication networks), but also in on-line network management and control. We proceed according to this direction in Section 5.3.

5.2. The fading effect

Before detailing the formulation of the bandwidth control at the SI-SAP interface, it is necessary to recall the channel degradation effect, potentially affecting the SD layer.

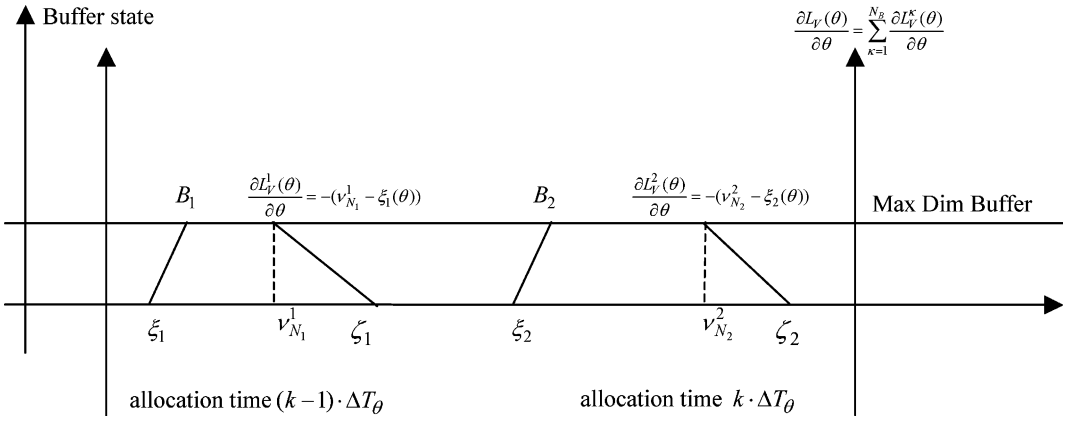

Fig. 5. IPA derivative estimation (4) $\frac{\partial L_V(\theta)}{\partial\theta}=\sum_{\kappa=1}^{N_{B_i}}\frac{\partial L_V^\kappa(\theta)}{\partial\theta}$, looking at the length of the busy periods of the buffer [22,33]. ΔT_θ is the time interval between two consecutive bandwidth reallocations of the service rate θ.

In satellite networks, variable channel fading conditions can heavily affect the transmission quality, especially in the *Ka band*, where the effect of rain over the quality of the transmission is evident. Error recovery techniques, such as *Automatic Repeat Request* (ARQ) and *Forward Error Correction* (FEC), are employed to face adverse channel conditions. ARQ techniques are not recommended, since real-time traffic requires stringent latency constraints. On the other hand, FEC mechanisms allow recovering erroneous packets, despite channel degradation, but because of their overhead, cause bandwidth reduction with consequent additional packet loss in the network. Modelling this behaviour is very important.

Let $\theta^{SD}(t)$ be the service rate allocated to a traffic buffer at the SD layer at time t. The effect of fading is modeled here as a reduction in the bandwidth actually "seen" by the buffer. The reduction is represented by the stochastic process $\phi(t)$. At time t, the "real" $\hat{\theta}^{SD}(t)$ is

$$\hat{\theta}^{SD}(t) = \theta^{SD}(t) \cdot \phi(t); \quad \phi(t) \in [0,1]. \tag{5}$$

In presence of fade countermeasures located at the physical layer, and applying FEC codes, $\theta^{SD}(t) \cdot \phi(t)$ may be considered the bandwidth portion dedicated to the transmission of information bits, while $\theta^{SD}(t) \cdot (1 - \phi(t))$ will transport redundancy bits. Whenever the fading effect causes errors over the packets, an adaptive control monitors the C/N (*Carrier/Noise Power*) factor and, based on this measure, increases the redundancy of the packets through a FEC scheme. In this way, the bandwidth $\hat{\theta}^{SD}(t)$ is reduced because some bits are devoted to carry redundancy and the outflow rate is modified by the fading effect.

The fading effect will be considered within the control algorithm, acting at the SI-SAP interface to map QoS, which is addressed below.

5.3. Problem formulation

Let $\alpha^{SI}(t)$ be the inflow process entering the SI layer at time t. More than one queue is involved in the $\alpha^{SI}(t)$ process. As mentioned in Section 4, the interaction of multiple SI queues is related to the differentiation of the traffic classes and the consequent need for traffic aggregation at the SD layer. This subject is left open for future research. Here, for the sake of simplicity, we consider the case of a single $\alpha^{SI}(t)$ process entering one buffer (with service rate $\theta^{SI}(t)$) at the SI layer and conveyed to a single SD buffer (with service rate $\theta^{SD}(t)$) at the SD layer.

Let $\alpha^{SD}(t)$ be the inflow process of a buffer at the SD layer at time t. The process $\alpha^{SD}(t)$ comes from the outflow processes of the SI buffer or directly from the $\alpha^{SI}(t)$ process, if no buffering is applied at the SI layer. In any case, a change in the encapsulation format is applied when $\alpha^{SI}(t)$ is conveyed to the SD layer (for example applying LLC-SNAP together with AAL5). We indicate with $L_V^{SD}(\cdot, \theta^{SD}(t))$ the loss volume at the SD layer. It is a function of the following elements: the SD inflow process $\alpha^{SD}(t)$ (depending on: the SI inflow process $\alpha^{SI}(t)$ and the transport technology change), the fading process $\phi(t)$ and the SD bandwidth allocation $\theta^{SD}(t)$.

We suppose that resource allocation for SI can satisfy the required QoS at the SI level. Here, the key problem is to "equalize" the QoS measured at the SD layer in dependence on the QoS imposed at the SI layer. To capture this concept, it is useful to think of a penalty cost function, whose values can be interpreted as an indication of the current inability of the SD layer to guarantee the required QoS. In practice, $L_V^{SD}(\cdot)$ can chase both the performance measured at the SI layer and a threshold value (fixed or time variant) that flows from SI to SD layer and represents the SLA. In both cases, the performance reference of SI is described through $L_V^{SI}(\cdot)$: in the first case, it is just the measure of the buffer loss over time; in the second case, it is the product of the required PLP, (denoted by PLP^*), and the SD inflow volume over a given time horizon $[t, t + \forall t]$, i.e., $L_V^{SI}(\cdot) = PLP^* \cdot \int_t^{t+\forall t} \alpha^{SD}(t) \cdot dt$.

QoS Mapping Optimization (QoSMO) Problem: it finds the optimal bandwidth allocation ${}^{Opt}\theta^{SD}(t)$ so that the cost function $J(\cdot, \theta^{SD}(t))$ is minimized:

$$
\begin{aligned}
&{}^{Opt}\theta^{SD}(t) = \arg\min_{\theta^{SD}(t)} J(\cdot, \theta^{SD}(t));\\
&J(\phi(t), \alpha^{SI}(t), \alpha^{SD}(t), \theta^{SI}(t), \theta^{SD}(t))\\
&\quad = \underset{\omega \in \Theta}{E} L_{\Delta V}(\phi(t), \alpha^{SI}(t), \alpha^{SD}(t), \theta^{SI}(t), \theta^{SD}(t))\\
&L_{\Delta V}(\phi(t), \alpha^{SI}(t), \alpha^{SD}(t), \theta^{SI}(t), \theta^{SD}(t))\\
&\quad = \left[L_V^{SI}(\alpha^{SI}(t), \theta^{SI}(t)) - L_V^{SD}(\alpha^{SD}(t), \theta^{SD}(t) \cdot \phi(t))\right]^2
\end{aligned}
\tag{6}
$$

We denote by ω a *sample path* of the system, i.e., a realization of the stochastic processes involved in the problem $(\phi(t), \alpha^{SI}(t), \alpha^{SD}(t))$ according to the statistical behaviour of the IP sources and to the channel degradation and with $\underset{\omega\in\Theta}{E}(\cdot)$ the mean over the set of all the possible sample paths Θ. The QoSMO problem (6) describes *an optimization* problem. The optimal $\theta^{SD}(t)$ is aimed at guaranteeing the SI-based SLA at the SD layer. The solution of the optimization problem through analytical tools is a very hard task. Numerical approximations are strongly recommended. On the other hand, even if a closed-form formula for the cost function $J(\cdot, \theta^{SD}(t))$ were available, it would require a priori assumptions on the traffic sources that we try to avoid. So, bypassing any "certainty equivalent" assumption, we investigate a way to spread the solution of QoSMO over time, by exploiting a (computationally light) sensitivity estimation procedure "capturing" the current bandwidth need of the SD buffer. We also take care that our approach does not prevent a quick on-line convergence toward the solution of (6).

To establish an optimization procedure aimed at approximating the optimal solution of (6), a derivative of the cost function $L_{\Delta V}(\cdot, \theta^{SD}(t))$ is needed. To get it we capture the temporal behaviour of $L_V^{SD}(\hat{\theta}^{SD}(t))$ through on-line measures ($\hat{\theta}^{SD}(t)$ is the real bandwidth, including fading effect, "seen" by the buffer). The mentioned derivative can be obtained from (7):

$$\frac{\partial L_{\Delta V}(\theta^{SD}(t))}{\partial \theta^{SD}(t)} = 2 \cdot \frac{\partial L_V^{SD}(\hat{\theta}^{SD}(t))}{\partial \hat{\theta}^{SD}(t)} \times [L_V^{SD}(\hat{\theta}^{SD}(t)) - L_V^{SI}(\theta^{SI}(t))] \quad \forall t, \tag{7}$$

where the $\frac{\partial L_V^{SD}(\hat{\theta}^{SD}(t))}{\partial \hat{\theta}^{SD}(t)}$ component in (7) is computed according to the IPA formulae (3) and (4), with v_κ (i.e., the instant of time when the last loss occurs during the buffer busy period B_κ), being the time when the first SD (ATM) packet (belonging to the last SI (IP) packet lost during the busy period B_κ) is lost.

Let $\Delta T_{\theta^{SD}}$ be the time interval between two consecutive bandwidth reallocations at the SD buffer and $k \cdot \Delta T_{\theta^{SD}}$ the instant when the kth reallocation is performed. The gradient estimation (computed, for instance, at instant k) is performed with respect to the "real" service rate $\hat{\theta}^{SD}$, thus "capturing" both the effects of the previous bandwidth allocation $\theta^{SD}((k-1)\cdot\Delta T_{\theta^{SD}})$ and of the realization of the stochastic process $\phi(t)$ in the time interval $[(k-1)\cdot\Delta T_{\theta^{SD}}; k\cdot\Delta T_{\theta^{SD}}]$.

The proposed optimization algorithm (called *Reference Chaser Bandwidth Controller*) is based on the gradient method, whose descent step is ruled

> Reference Chaser Bandwidth Controller (RCBC) tracks over time the optimal solution ${}^{Opt}\theta^{SD}(t)$ of the QoSMO problem (6) through the on-line gradient descent:
>
> $$\theta^{SD}(k+1) = \theta^{SD}(k) - \eta_k \cdot \frac{\partial L_{\Delta V}(\theta^{SD}(k))}{\partial \theta^{SD}(k)}; \quad k = 1, 2, \ldots \theta^{SD}(0) = \theta^{SD}_{init} \tag{8}$$

by (8). We denote with η_k the gradient stepsize. The gradient descent (8) is initialized with θ^{SD}_{Init}, whose definition is detailed in Section 5.6.

No statistical knowledge (i.e., mean, variance) about the IP traffic flow $\alpha^{SI}(t)$ is necessary to apply the gradient descent. All the variables in (7) can be computed by monitoring the buffer behaviour at the SI-SAP interface.

In [34], a similar *Perturbation Analysis* (PA) technique is employed for the optimization of the *blocking probability* in a circuit-switched network. In [35], the aforementioned IPA estimator is used to minimize the *packet loss probability* in a satellite network. To the best of our knowledge, this is the first time that such technique is adopted to adjust the equivalent bandwidth dimensioning.

5.4. Proof of convergence

The employed gradient-based algorithm is a standard stochastic approximation scheme [21], driven by the IPA-based derivative estimator (7). Four technical conditions are required to establish convergence to a global optimum.

1. The first one is related to a decreasing behaviour of the gradient stepsize η_k. The form $\eta_k = \frac{c_1}{c_2+k}$, $c_1, c_2 > 0$ is usually employed. However, (as also in [33–35]) we verified in our tests that a fixed gradient step size does not affect the convergence and simplifies the control implementation.

2. The second one involves *unbiasedness* and *consistency* of the derivative estimator (7). A *Perturbation Analysis* (PA) estimator is defined to be unbiased if the derivative operator can be replaced with the expectation operator and vice versa, i.e.

$$\frac{\partial}{\partial \theta^{SD}} \underset{\omega}{E}[L_{\Delta V}(\cdot, \theta^{SD})] = \underset{\omega}{E}\left[\frac{\partial L_{\Delta V}(\cdot, \theta^{SD})}{\partial \theta^{SD}}\right]. \quad (9)$$

Interchanging the expectation with the limit is necessary to build an estimator, which is a function of the current sample path ω. The two conditions ensuring the unbiasedness of IPA derivative are [22]: (2a) for every $\theta^{SD} \in \Re^+$, $\frac{\partial L_{\Delta V}(\cdot, \theta^{SD})}{\partial \theta^{SD}}$ exists, *with probability* 1 (w.p.1); (b) the function $L_{\Delta V}(\cdot, \theta^{SD})$ is Lipschitz-continuous and the Lipschitz constant has a finite first moment.In general, due to the discontinuities of the loss volume for several DES, traditional PA techniques give biased derivative estimators that do not satisfy (9) (see, e.g., [13] for an overview concerning this topic). In our case, since the function $L_{\Delta V}(\cdot)$ is the result of a linear composition of buffer loss volumes, the aforementioned conditions are satisfied. In particular, we will show in the following that $L_{\Delta V}(\cdot, \theta^{SD})$ is a continuous and differentiable function of θ^{SD}.

Consistency is related to the convergence capability of the proposed estimator with respect to the "real" (and unknown) gradient of the functional cost $\frac{\partial}{\partial \theta^{SD}} \underset{\omega}{E} L_{\Delta V}(\cdot, \theta^{SD})$ [13]. Consistency requires that as the length of the observed sample path ω increases, the estimator $\frac{\partial L_{\Delta V}(\cdot, \theta^{SD})}{\partial \theta^{SD}}$ converges to $\frac{\partial}{\partial \theta^{SD}} \underset{\omega}{E} L_{\Delta V}(\cdot, \theta^{SD})$ w.p.1. In practice, the hope is that, as the number of observations increases, the accuracy of the estimator increases, too. This observation regards the ergodicity of the stochastic processes involved in the system, which, here, are assumed to be non-stationary. Hence, it is hard to prove consistency. However, as also pointed out in [33], "we concentrate on obtaining reliable shorter-term sensitivity information tracking the behaviour of the network and seeking to continuously improve its performance". Thus, the need of consistent gradient estimators is disregarded.

3. The third condition requires that $\sup_{\theta^{SD} \in \Re^+} \| \frac{\partial L_{\Delta V}(\theta^{SD})}{\partial \theta^{SD}} \| < \infty$ (where $\|\cdot\|$ denotes the standard Euclidean norm). Namely, the derivative estimation should be defined and limited in $\theta^{SD} \in \Re^+$. For every buffer size and service capacity allocation, it is easily observable that the lengths of the busy periods of the SD buffer (where the *RCBC* is applied) are bounded by the dimension of every reallocation time period $\Delta T_{\theta_{SD}}$. This condition is therefore satisfied.
4. The final condition pertains to the cost function to guarantee a global optimum. It is easily observable that a global optimum exists. This is the subject of the following subsection.

5.5. Counteracting non-stationary stochastic processes

An adaptive optimal operation point of the system can be found despite the change of the encapsulation format and the variability of traffic statistics. To highlight the behaviour of *RCBC*, we depict the components of the cost function $L_{\Delta V}(\cdot)$ in Fig. 6.

The loss volume of a traffic queue being $L_V^{SD}(\theta^{SD})$, it can be reasonably assumed to be continuous, differentiable, with a negative derivative in the service rate θ^{SD} and, as a consequence, the penalty cost function, $L_{\Delta V}(\cdot)$, is also continuous and differentiable with a unique minimum (see Fig. 6). ${}^{Opt}\theta^{SD}$ is the optimal operation point equalizing the loss volumes $L_V^{SD}(\cdot)$ and $L_V^{SI}(\cdot)$.

The challenge is to automatically adapt ${}^{Opt}\theta^{SD}$ to counteract variable system conditions (i.e., either when the performance thresholds or the traffic statistics change (so redefining $L_V^{SI}(\theta^{SI})$), thus varying the slope of $L_V^{SD}(\cdot)$), and without using any closed-form formula of $L_V^{SD}(\cdot)$. Due to the regularity of the penalty cost function $L_{\Delta V}(\cdot)$ in (6), the aforementioned gradient-based algorithm should track the ${}^{Opt}\theta^{SD}$ value efficiently. The major concern is related to the stationarity of the involved stochastic processes. For this reason, it is necessary to assume that the convergence of the algorithm toward a new operation point is faster than the changes in the stationarity of the stochastic environment.

We assume stationarity pertains to: statistics of the traffic sources, call (source) arrival rate, peak and mean packet rate; packet size distribution (at the SI layer) and current fading level conditions. In the results reported in the following, we show how the control algorithm is able to quickly

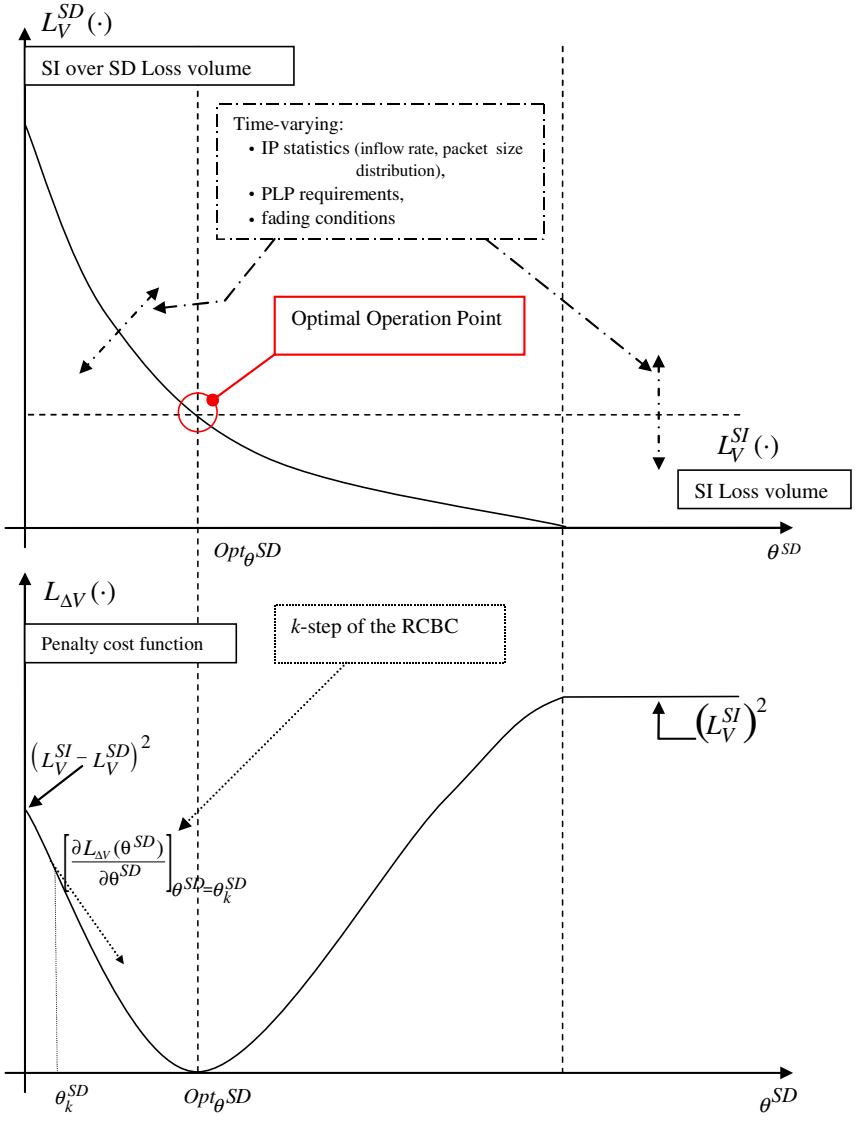

Fig. 6. Functional cost and optimal operation point.

counteract variable traffic and fading level conditions. The key idea is to properly initialize both the gradient descent and the employed gradient stepsize when a change in the stochastic environment takes place.

5.6. Heuristic initialization to maximize convergence speed

As summarized above, the *RCBC* converges to the minimum ATM bandwidth allocation (at the SD layer) to preserve the QoS required by the SI layer. The major concern is related to the transient period required by the gradient descent (8) to reach optimality. It could severely degrade the on-line performance of the proposed algorithm. The transient behaviour can be reduced by properly dimensioning the stepsize η_k, but, unfortunately, this operation can introduce high sensitivity in the control scheme to traffic changes, giving rise to possible bandwidth oscillations and unsatisfactory system performance. Hence, to tune the parameter η_k, a trade-off between convergence speed and sensitivity of the scheme must be investigated in detail.

The other point of intervention to speed up the convergence is to properly dimension the initial point of the gradient descent (8). A possibility is employing the following heuristic.

The increase in the bandwidth allocation necessary at the SD layer can be foreseen by means of the *CellTax* effect introduced by the AAL5 with LLC-SNAP encapsulation and by the ATM header. Since, during the generation of the ATM frame, two octets (for the LLC-SNAP overhead) need to be added to each IP packet, the average number of ATM cells for each IP packet is

$$\overline{\#ATMCells} = \frac{\overline{DimIPPacket} + 2}{48}, \tag{10}$$

where $\overline{DimIPPacket}$ denotes the average IP packet's size in bytes and 48 is the payload of an ATM cell in bytes. Hence, it is possible to compute the overall overhead due to the encapsulation format of the SD frame and the average percentage bandwidth increase necessary in the SD core of the network, identified in the following as *CellTax*:

$$CellTax = \frac{\overline{\#ATMCells} \cdot 53 - \overline{DimIPPacket}}{\overline{DimIPPacket}}, \tag{11}$$

where 53 is the overall size of an ATM cell in bytes. A possible forecast for the SD bandwidth allocation is determined by the CellTaxAllocation heuristic that approximates the solution of the QoSMO problem (2) with the following bandwidth allocation law (12):

$${}^{\mathrm{CellTax}}\theta^{\mathrm{SD}} = (1 + CellTax) \cdot \theta^{\mathrm{SI}} \tag{12}$$

A similar method can be employed when the transport technologies of interest are different, e.g., in case of a QoS mapping involving either IPv6 over IPv4 or IPv4 over MPLS at the SI-SAP interface. The only concern is related to the knowledge of the precise encapsulation formats.

The heuristic in (12) can be used to initialize the gradient descent of the *RCBC* in order to speed up convergence, thus guaranteeing fast reactions to variable system conditions. In practice, each time a change in the stochastic environment is detected (e.g., when the number of sources, input traffic statistic or the fading level change), the gradient descent (8) of the *RCBC* restarts from

$$\theta^{\mathrm{SD}}(0) = \theta^{\mathrm{SD}}_{\mathrm{init}} = (1 + CellTax) \cdot \theta^{\mathrm{SI}}. \tag{13}$$

In this way, the gradient descent (8) can be seen as an active measurement, self-tuning procedure around the heuristic allocation (12), capable of tracking the optimal operation point with a small number of suboptimal reallocation steps. The limit of (13) is that the knowledge of θ^{SI} (not required by RCBC) is needed. The CellTax approximation (12) of QoSMO problem is also used as a comparison solution in the next section.

6. Performance analysis

To test *RCBC* and evaluate the bandwidth provision at the SI-SAP interface according to different traffic scenarios, we developed a C++ simulator where AAL5 is based on LLC-SNAP encapsulation. The scenario presented in Fig. 3 is considered for simulation, but only one IP queue and one ATM queue (e.g., EF traffic over ATM) is taken into account. *RCBC* faces only the encapsulation problem as described above.

We first start evaluating the convergence behaviour of the proposed control algorithm, together with its capability of tracking a time variant bandwidth requirement in different traffic conditions. The bandwidth assignment of *RCBC* is compared with the allocation computed by the *CellTaxAllocation* (12), also highlighting the equivalent bandwidth shift due to the change of the encapsulation format for different SD buffer sizes. To show the *RCBC* behaviour both chasing the losses measured at SI layer and the loss volume computed from a fixed SLA threshold, we use VoIP traffic. The reaction of the algorithm to channel status variations is tested over the same VoIP scenario through the application of real fading levels.

6.1. Convergence behaviour and tracking capability of the control algorithm

In this scenario, the IP traffic flow entering the SI-SAP interface is composed of 10 aggregated sources, each of them characterized by the *Trimodal* distribution concerning the packet size. According to this, the packet size can assume three different values: a with probability p_a, b with probability p_b and c with probability $1 - p_a - p_b$. The Trimodal distribution is widely used for current Internet traffic. Measures collected by the *Politecnico di Torino* telecommunication research group have shown that assuming $a = 48$ bytes, $b = 576$ bytes, $c = 1500$ bytes, $p_a = 0.559$, $p_b = 0.2$ (denoted by the notation "Trimodal$(48, 576, 1500, 0.559, 0.2)$") accurately approximates the traffic traces collected during the first 13 days of the year 2001 by the gateway router of the *Politecnico di Torino* [36]. We use the Trimodal$(48, 576, 1500, 0.559, 0.2)$ distribution for each active IP source in the following.

In order to verify the adaptive capability of the *RCBC*, an increase in the IP sources' rate is applied after 1 min of simulation. We suppose that the PLP guaranteed at the SI layer is 3×10^{-2}. The bandwidth θ^{SI} for the IP flow, shown in Table 1, is computed by simulation inspection. Both the SI buffer and SD buffer sizes are set to 150,000 bytes. We also assume that no fading degradation affects the satellite channel, for now.

A sample path of the PLP measured at the buffers is depicted versus time in Fig. 7. *RCBC* is applied by setting the time between two SD reallocations $\Delta T_{\theta^{SD}}$ to 0.2 s and the gradient stepsize $\eta = \eta_k$ to 1×10^3 $\forall k$. Figs. 8 and 9 show, respectively, the values of the ATM bandwidth allocations and of the functional cost derivative (7) versus time. It is easy to see that the proposed control algorithm is able to equalize the PLP of the two buffers after a transient period of about 10.0 s.

It is also worth noting that, in the case presented, no bandwidth is provided to the SD buffer at the traffic load change instant. It means that the tracking of the bandwidth needed is started from the null value. The convergence time is much reduced if the bandwidth allocated to SD buffer "starts" from heuristic values, as should be clear from the results reported in Section 6.4. ${}^{Opt}\theta^{ATM}$ is time-dependent according to the time-varying IP sources' behaviour. Fig. 10 shows the PLP in the first minute of simulation after reaching the ${}^{Opt}\theta^{ATM}$ steady state. The bandwidth necessary for SD layer is obtained and a proper "QoS mapping" is applied to guarantee the required SLA, despite the changes in the transport technology and in the sources' inflow rate. *RCBC* tracks a time variable threshold (SI layer PLP) and can satisfy the required SLA almost constantly, except for sporadic instants, due to quick traffic (or threshold) variations.

Table 1
Convergence behaviour scenario

Time interval (s)	0.0–60.0	60.0–120.0
IP sources' rate (Mbps)	1.0	2.0
θ^{SI} for PLP $\leqslant 3 \times 10^{-2}$ at the SI layer (Mbps)	9.55	19.0

IP sources' rate and IP equivalent bandwidth.

6.2. Comparison with CellTaxAllocation

We now compare *RCBC* with the *CellTaxAllocation* heuristic. The latter exploits a perfect knowledge about the bandwidth assignment at the SI layer and about the SI packet size distribution. We recall that *CellTax* is obtained as in (11). We will use percentages in the following as shown in (14).

$$\%CellTax = CellTax \times 100. \tag{14}$$

The mean number of ATM cells $\overline{\#ATMCells}$, generically indicated in (10), is generated by an IP source that produces n different packet's size $DimIPPacket_i$ (in bytes), $i = 1, \ldots, n$, each of which with probability p_i, $(\sum_{i=1}^{n} p_i = 1)$.

$$\overline{\#ATMCells} = \sum_{i=1}^{n} \left\lceil \frac{DimIPPacket_i + 2}{48} \cdot p_i \right\rceil. \tag{15}$$

We compare *RCBC* with *CellTaxAllocation* by varying the traffic conditions. Interarrival time between IP packets are modeled by a *Pareto* distri-

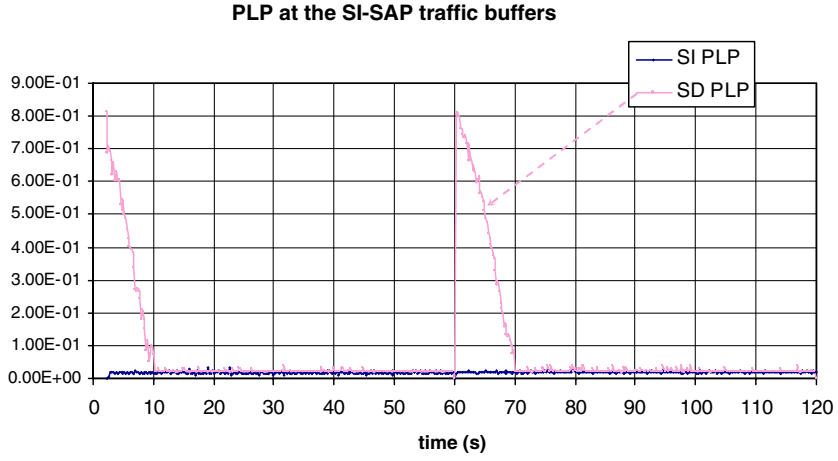

Fig. 7. Convergence behaviour scenario. PLP at the both the SI and SD buffers of the SI-SAP interface.

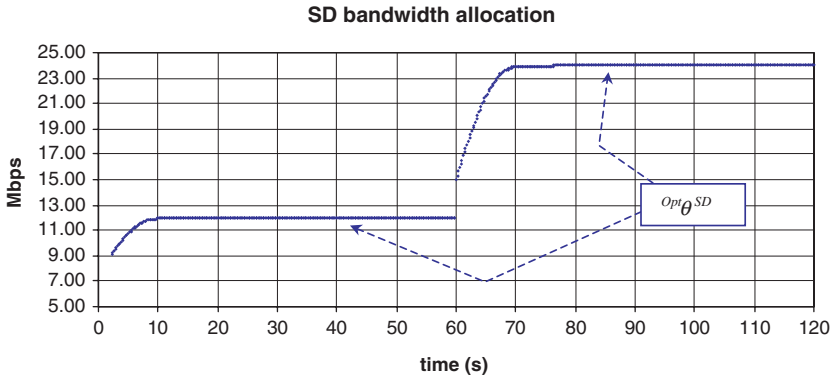

Fig. 8. Convergence behaviour scenario. Steady states of the SD bandwidth allocation using *RCBC*.

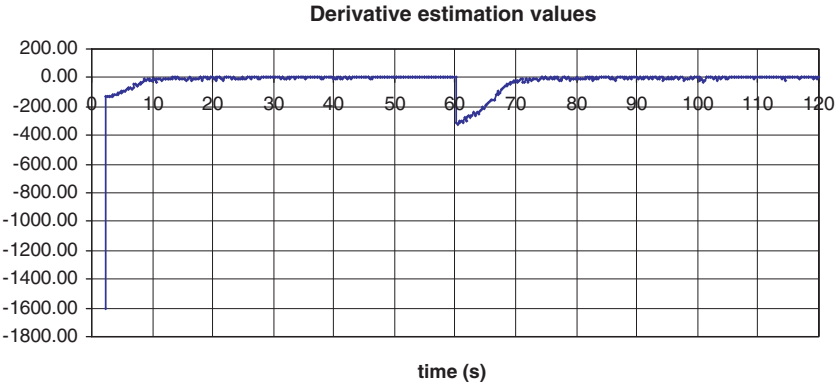

Fig. 9. Convergence behaviour scenario. Derivative values of the penalty cost function $L_{\Delta V}(\cdot, \theta^{SD})$.

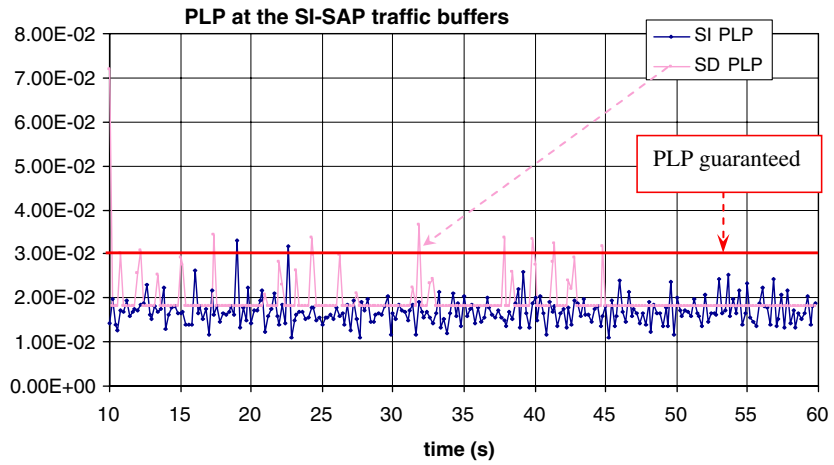

Fig. 10. RCBC behaviour over time. PLP after a stationary $^{Opt}\theta^{SD}(t)$ has been reached.

bution. The mean interarrival times used are 10 ms and 100 ms and the number of connections in the IP flow is 1, 20, and 100. As in the previous subsection, the packet size distribution is Trimodal-$(48, 576, 1500, 0.559, 0.2)$. The SI buffer size is set to 150,000 bytes. Bandwidth allocation θ^{SI} guarantees a PLP $\leqslant 1 \times 10^{-2}$. The comparison is shown in Table 2 for SD buffer size = SI buffer size = 2830

Table 2
Comparison with *CellTaxAllocation*

Arr-Time IP packets [S]	#Conn	θ^{SI} [Mbps]	% *CellTax*	${}^{CellTax}\theta^{SD}$ [Mbps]	${}^{Opt}\theta^{SD}$ [Mbps]	% *RCBC CellTax*	*CellTax* difference (%)
0.01	1	0.28	20.44	0.337	0.352	25.82	4.38
0.01	20	5.700	20.44	6.865	7.024	23.22	2.31
0.01	100	28.500	20.44	34.325	34.944	22.61	1.80
0.1	1	0.039	20.44	0.047	0.048	23.66	2.19
0.1	20	0.78	20.44	0.939	0.964	23.61	2.62
0.1	100	4.100	20.44	4.938	4.957	20.90	0.46

SD buffer size = 2830 ATM cells.

Table 3
Comparison with *CellTaxAllocation*

Arr-Time IP packets [S]	#Conn	θ^{SI} [Mbps]	% *CellTax*	${}^{CellTax}\theta^{SD}$ [Mbps]	${}^{Opt}\theta^{SD}$ [Mbps]	% *RCBC CellTax*	*CellTax* difference (%)
0.01	1	0.29	20.44	0.349	0.413	42.56	18.37
0.01	20	5.750	20.44	6.925	8.075	40.44	16.60
0.01	100	28.750	20.44	34.627	41.657	44.89	20.30
0.1	1	0.038	20.44	0.046	0.06	57.64	30.89
0.1	20	0.78	20.44	0.939	1.130	44.93	20.33
0.1	100	3.900	20.44	4.697	5.556	42.45	18.28

SD buffer size = 200 ATM cells.

ATM cells and in Table 3 for SD buffer size = 200 ATM cells. The values reported are, from left to right: mean interarrival time of the IP packets (*Arr-Time IP packets*); number of connections (sources) in the flow (*#Conn*); SI bandwidth allocation (θ^{IP}); %*CellTax* as in (14); SD bandwidth allocation scheme (${}^{CellTax}\theta^{SD}$) computed by the *CellTaxAllocation*; SD bandwidth allocation computed by *RCBC* (${}^{Opt}\theta^{SD}$); the percentage additional bandwidth in steady state, computed through *RCBC* and called %*RCBC CellTax* (16), required by the SD layer to chase SI PLP.

$$\%RCBC\ CellTax = \frac{{}^{Opt}\theta^{SD} - \theta^{SI}}{\theta^{SI}} \times 100. \qquad (16)$$

The last column contains the difference (*CellTax difference*) between the *RCBC CellTax* (16) and the *CellTax* (14). The width of the confidence interval over the simulated loss measures is less then 1% for 95% of the cases.

As shown in the previous subsection, ${}^{Opt}\theta^{SD}$ (obtained after the convergence of the *RCBC*) is the optimal value to dimension the bandwidth necessary at the SD core and it should be taken as the target value for the following comparison.

From these results, it is clear that *CellTaxAllocation* produces good results only if the buffers have the same size. In this case, the difference between the %*CellTax* (14) and %*RCBC CellTax* (16), is below 5% (see *CellTaxDifference* column in Table 2). If the SD buffer has only 200 cells (Table 3), *CellTaxAllocation* underestimates the bandwidth needs. In more detail, if the mean interarrival time is 0.01 s, the *CellTaxDifference* is around 20% with a minimum of 18% for one connection in the flow; if the mean interarrival time is 0.1 s the, *CellTaxDifference* is larger and reaches a maximum of about 31% for one source. The explanation of this behaviour comes from the fact that with increasing mean interarrival time with a small number of connections, the rate variability of the flow increases. It has a strong impact on *CellTaxAllocation* performance.

On the other hand, *RCBC* shows itself as a promising tool for planning the SI-SAP interface. It is worth noting that *RCBC*, without any a priori knowledge about statistical behaviour of the traffic sources, encapsulation format and buffer dimension, guarantees the correct QoS mapping, despite the changes in traffic flow. Only information about the amount of losses is necessary to apply the algorithm. It means that, in theory, the notion of buffer might also be dropped without affecting the validity of the control mechanism.

Fig. 11 extends the numerical results presented in Tables 2 and 3 and shows the bandwidth shift in percentage at the SI-SAP interface (i.e., %*RCBC CellTax* in (16)) as a function of the SD buffer size. The SI buffer is again fixed at 150,000 bytes. Each

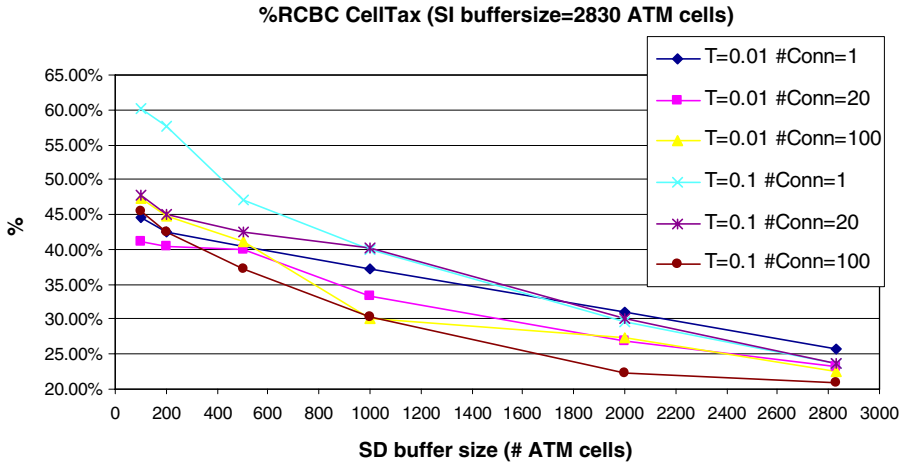

Fig. 11. Bandwidth shift measured as *%RCBC CellTax* by varying the SD buffer size. "T": mean interarrival time of SI packets (Pareto distributed); "#Conn": number of connections in the flow.

curve is related to a combination of IP packet mean interarrival time and number of connections in the flow. Fig. 11 confirms the observations previously reported: when the SD buffer size is reduced by an order of magnitude with respect to the SI buffer size, SI bandwidth provisioning θ^{SI} has to be increased up to 60% at the SD layer (the "$T = 0.1$, #Conn = 1" case with SD buffer size = 100 ATM cells). On the other hand, when the buffers are of equal size, the necessary bandwidth increase is around 20–25%.

6.3. Remarks about convergence

The transient period of *RCBC* to reach the optimal SD allocation can be efficiently reduced not only by changing the initial value, but also by means of a proper tuning of $\Delta T_{\theta^{SD}}$ (the time interval between two SD bandwidth reallocations) and η (the gradient stepsize). Clearly, if the values of the latter are too large, instability in the SD allocation arises; if, on the contrary, they are too small, a long time is necessary to reach the optimal steady state. So, a proper trade-off between the two parameters must be investigated. For instance, a satisfying combination of $\Delta T_{\theta^{ATM}}$ and η is: $\Delta T_{\theta^{ATM}} = 0.1$ s and $\eta = 3 \times 10^3$; using these values for the previous simulation tests, the transient period is reduced approximately to 4 s.

6.4. Voice over IP traffic scenario

The aim is to show how to speed-up *RCBC* convergence and to highlight *RCBC* behaviour with a fixed SLA threshold to get the loss volume (as indicated in Section 5.3), which is a common situation in real BSM networks where the SI layer passes the SLA threshold to the SD layer. Additionally, we will test *RCBC* in the presence of time-variant fading. We consider the case of *Voice over IP* (VoIP) traffic, coming from a single EF queue at the SI layer, and conveyed to a single ATM buffer at SD layer. Each VoIP source is modeled (as in [37]) as an exponentially modulated on–off process, with mean on and off times equal to 1.008 s and 1.587 s, respectively [38]; during on periods it generates 16.0 kbps over the RTP/UDP/IP stack. The VoIP packet size is 80 bytes. The required SLA performance threshold for VoIP flows is set to 2% of PLP. The SLA is reported in Table 4. We suppose that the SI bandwidth provision θ^{SI} has been already dimensioned to guarantee the required PLP of SI layer. As stated above, two implementation schemes of *RCBC* are considered: (1) the PLP at the SD buffer chases the packet losses measured at the SI buffer, as in the results presented up to

Table 4
VoIP over ATM

Service level agreement	Range
Premium VBR for *Voice over IP*	*Variable bit rate* (VBR)
Traffic description and conformance testing of VoIP	• Peak rate: 16 kbps • Mean rate: 14.87 kbps • Packet size: 80 bytes
Performance guarantees	Packet loss probability: 2%

VoIP *Service Level Agreement* [37,38].

now, (2) the SD PLP chases the performance level defined in the SLA (2%). The two methods are denoted as "feedback from measured SI losses" and "feedback from SLA reference level", respectively.

We compare *RCBC* with *CellTaxAllocation* strategy by setting the same size of the SI and SD buffers (fixed at 20 VoIP packets corresponding to 31 ATM cells). Buffer sizes are short to assure low propagation delays at each SI-SAP, thus also limiting *Packet Transfer Delay* (PTD) values, which are not part of the SLA for now. We progressively increase the number of VoIP sources in the flow from 70 to 110. The increasing step is 10 sources to stress the working conditions and the reaction of the algorithm. The heuristic introduced in (13) is used to initialize the *RCBC* gradient descent in order to speed up the convergence and to assure fast reactions to time-varying system conditions. The RCBC's bandwidth initialization is performed when the number of VoIP source changes (every 3000 s). A new SD bandwidth allocation $\theta^{SD}(k)$ is performed through the gradient descent (8) every $\Delta T_{\theta^{SD}} = 30$ s. The gradient stepsize $\eta_k = \eta\ \forall k$, is fixed at 7.0 in the "feedback from measured SI losses" case and at 11.0 in the "feedback from SLA reference level" case. They are the best values (found by simulation inspection) to maximize the convergence speed and to avoid strong SD rate oscillations.

Fig. 12 shows the performance of the *CellTaxAllocation*, which often fails to guarantee the required PLP over time. On the other hand, the performance of *RCBC* (reported in Figs. 13a and 13b) is very satisfying both for "feedback from measured SI losses" and for "feedback from SLA reference level". Table 5 reports the average PLP over the entire simulation horizon. It is easily observable that the *RCBC* that uses measured SI losses yields a PLP very close to

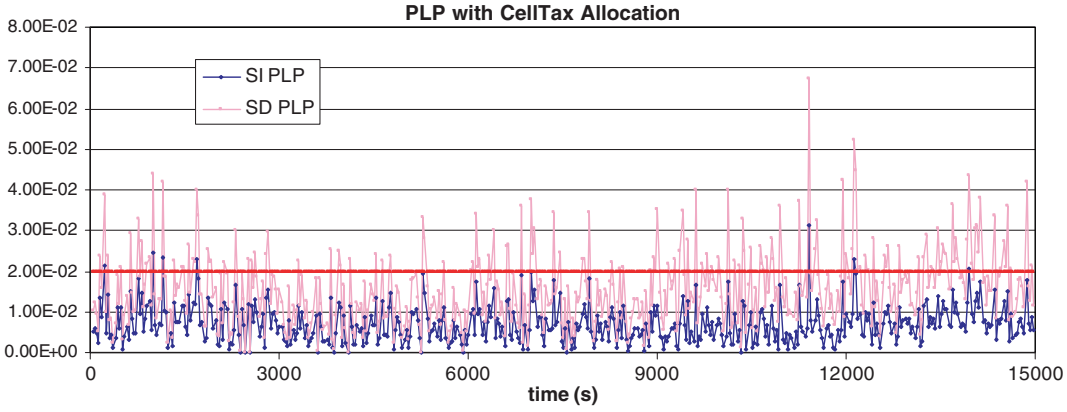

Fig. 12. VoIP over ATM. PLP at the SI and SD layer applying *CellTaxAllocation (12)*.

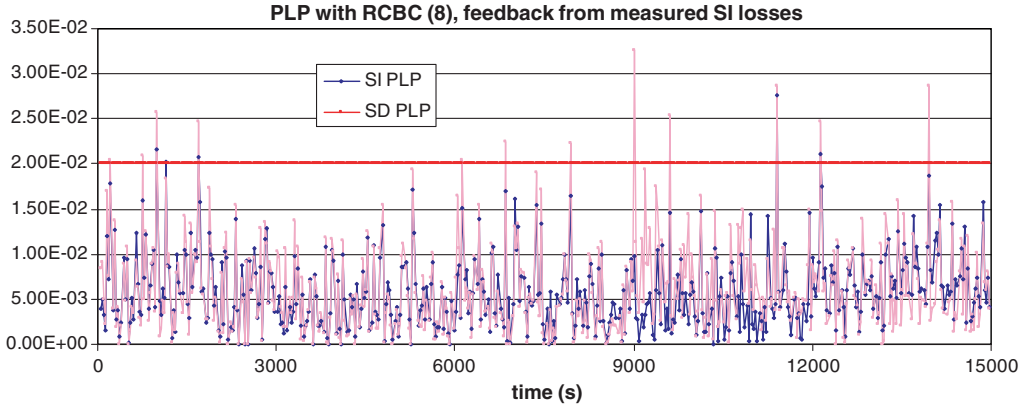

Fig. 13a. VoIP over ATM. PLP at the SI and SD layer applying *RCBC (8)*, feedback from measured SI losses.

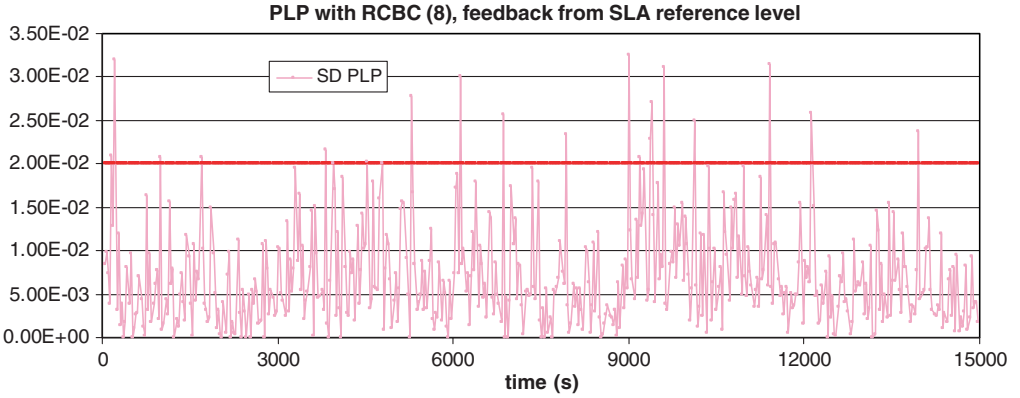

Fig. 13b. VoIP over ATM. PLP at SD layer applying *RCBC (8)*, feedback from SLA reference level.

Table 5
VoIP over ATM

SI layer	*CellTax*	RCBC, feedback: SI losses	RCBC, feedback: SLA
5.78×10^{-3}	1.55×10^{-2}	6.16×10^{-3}	7.03×10^{-3}

VoIP. Average PLPs over the simulation horizon.

the IP one. *RCBC* chasing SLA threshold satisfies the required performance and allows reaching an average PLP well below the value obtained through *CellTaxAllocation*, which, even if it assures an average value below the threshold, has a very "jumpy" behaviour not completely satisfactory over many periods (Fig. 12).

The bandwidth provision of *RCBC* and *CellTaxAllocation* are compared in Fig. 14 for the measured SI losses case. The SLA reference level case has approximately the same performance and it is not reported. *CellTaxAllocation* underestimates the required SD bandwidth to carry the VoIP flows correctly.

Fig. 14 also gives a first indication about the improvement of the convergence speed that is worth highlighting. The price to pay to minimize the convergence speed of the *RCBC* is to assume some knowledge about both SI packet size and SI bandwidth allocation so that (13) can be applied through (10) and (11) for the initialization of the gradient descent (8). In case of a variable packet size, if off-line information is not available, an on-line estimation of the mean packet size may be used to approximate the real value of the term $\overline{\#ATMCells}$ in (15).

RCBC quick reaction to traffic changes is clearer in Fig. 15, where the values of the cost function $L_{\Delta V}(\cdot, \theta_k^{SD})$ are shown, and in Fig. 16, where *%CellTax* computed by the *CellTaxAllocation* and

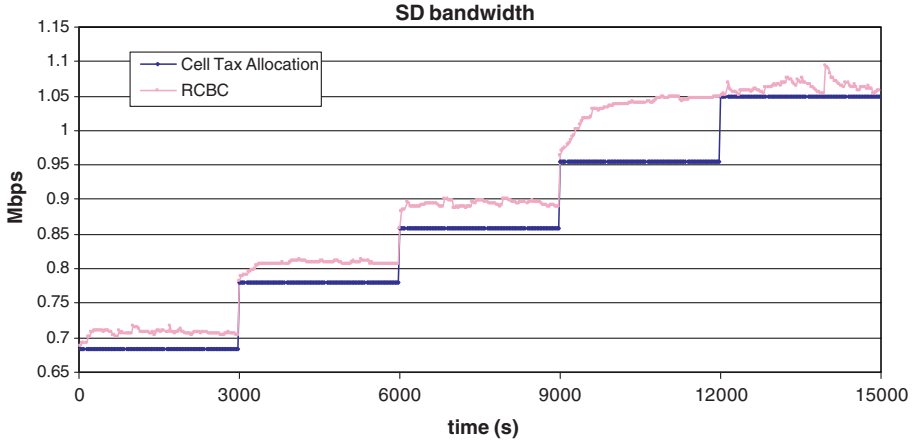

Fig. 14. VoIP over ATM. SD bandwidth, *CellTaxAllocation* versus *RCBC*.

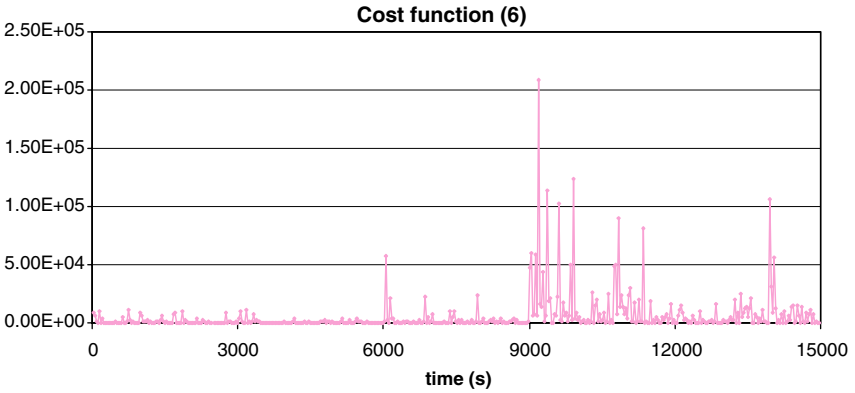

Fig. 15. VoIP over ATM. Cost function (6) minimized through *RCBC*.

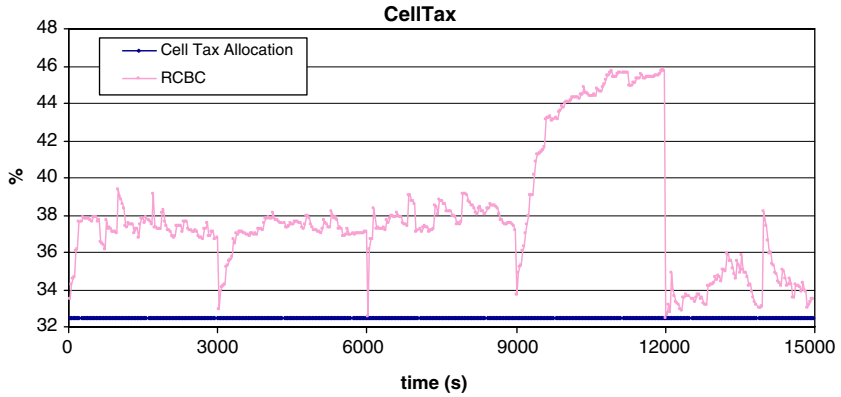

Fig. 16. VoIP over ATM. *%RCBC CellTax* (16) versus *%CellTax* (14) (*CellTaxAllocation*).

%RCBC CellTax(k) (given by (17)), are compared for each reallocation step:

$$\%RCBC\ CellTax(k) = \frac{\theta^{SD}(k) - \theta^{SI}}{\theta^{SI}} \times 100; \quad k = 1, 2, \ldots \tag{17}$$

The convergence speed is outstanding. Despite the change in the transport technology, *RCBC* assures QoS mapping by tracking the best operation point of the system. The bandwidth provisioning at both the SI and SD layers is compared in Fig. 17 to highlight the bandwidth shift computed by *RCBC* when the encapsulation format changes.

As far as the end-to-end delay of VoIP flows is concerned, not being included in the SLA and not being computed by RCBC for now, it may be considered by limiting the maximum number of hops of the end-to-end path, as often done in the literature.

Similar performance (not shown here) can be obtained by imposing a random interarrival time of *VoIP* calls. In this case, the traffic variability step of calls is a single source. The sensitivity of $L_{\Delta V}(\cdot)$ is much lower than the one presented in Fig. 15, thus helping *RCBC* convergence. The gradient descent can be initialized without any knowledge of the *CellTax* as above. To achieve quick on-line reactions, it is sufficient to assign $\theta^{SD}(0) = \theta^{SI}$ each time a *VoIP* call enters or exits the network. This choice was validated by simulation results not reported here. It implies some knowledge about SI layer status. We finally must note that similar results have been obtained for many examples involving other traffic types (e.g., video streaming) and traffic intensities.

6.5. Counteracting variable fading levels

We now consider the effect of the fading phenomenon. The time horizon of the simulation scenario is 133.0 min. Again, an aggregate trunk of 50 VoIP on-off sources composes the inflow process at the

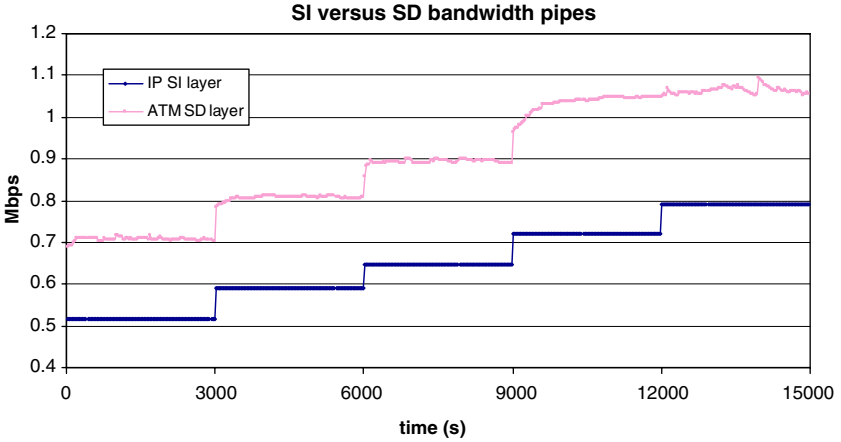

Fig. 17. VoIP over ATM. SI and SD layer bandwidth pipes.

traffic buffers at the SI-SAP interface. The buffer size is set to 1600 bytes (20 VoIP packets) for the SI layer and 3770 bytes (70 ATM) cells for the SD layer. The gradient stepsize η is fixed to 1.0 and the time interval between two consecutive SD rate reallocations $\Delta T_{\theta^{SD}}$ is 1 min. *RCBC* is initialized with $\theta^{SD}(0) = \theta^{SI}$ at the beginning of the simulation. This time, no traffic change takes place, but time-varying channel degradation affects the SD buffer service capacity.

The employed fading process has been taken from [10], where real attenuation samples are extracted from an experimental data set carried out in the *Ka band* on the *Olympus* satellite by the CSTS (*Centro Studi sulle Telecomunicazioni Spaziali*) Institute (Milan, Italy), on behalf of the *Italian Space Agency*. The up-link (30 GHz) and down-link (20 GHz) samples consider averages (over 1 s), expressed in dB, of the signal power attenuation with respect to clear sky conditions. The *Carrier/Noise Power* (C/N_0) factor is monitored at each station and, on the basis of its values, different bit and coding rates are applied to keep the BER below a chosen threshold of 10^{-7}. Six different fading classes are defined, corresponding to combinations of channel bit and coding rate that give rise to redundancy factors $\xi_{\text{level}}(t)$, $level = 1, \ldots, 6$ $(\xi_{\text{level}}(t) \geqslant 1.0)$. $\xi_{\text{level}}(t)$ represents the ratio between the *Information Bit Rate* (IBR) in clear sky and the IBR in specific working conditions. The corresponding bandwidth reduction factor is: $\phi(t) = \frac{1}{\xi_{\text{level}}(t)}$. With the data adopted in [10], we have

$$\phi(t) \in \{0.0, 0.15625, 0.3125, 0.625, 0.8333, 1.0\}. \tag{18}$$

The bandwidth reduction can be computed as $\hat{\theta}^{SD}(t) = \phi(t) \cdot \theta^{SD}(t)$; with $\phi(t) = \frac{1}{\xi_{\text{level}}(t)}$. The value $\phi(t) = 0$ corresponds to the outage condition. As shown in Fig. 18, the fading process results in strong channel degradation, especially in the time interval $[4800, 6000]$. The PLP measured at the SD layer is

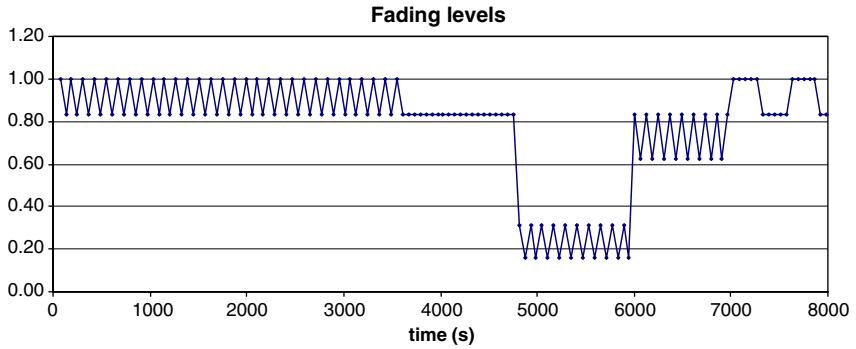

Fig. 18. Fading scenario. Reduction factors [Cel03].

depicted in Fig. 19. *RCBC* chases the losses at the IP layer, which, obviously, is not affected by fading. The required PLP is guaranteed almost all the time. Only four peaks of performance degradation (confirmed by the values of the cost function $L_{\Delta V}(\cdot, \theta_k^{SD})$ reported in Fig. 20) happen when fading changes. Actually, performance degradation corresponds to 5 min of "over threshold behaviour" (PLP 5×10^{-2} instead of 2×10^{-2}) within the entire simulation period. It is important to note that, if the fading level changes, the initial value of bandwidth when *RCBC* starts minimizing is set again to $\theta^{SD}(0) = \theta^{SI}$. Fading, as evident in the trace of Fig. 18, stresses the algorithm. Its success in dealing with these fades is clear in Fig. 19. A potential poor response to transients is avoided either by using heuristics or by memorizing the bandwidth allocation assigned during similar fading levels from in past. This capability is already part of the features of the current implementation of the scheme. The tracking capability of *RCBC* is easily observable in Figs. 21 and 22, where the rate allocation of the

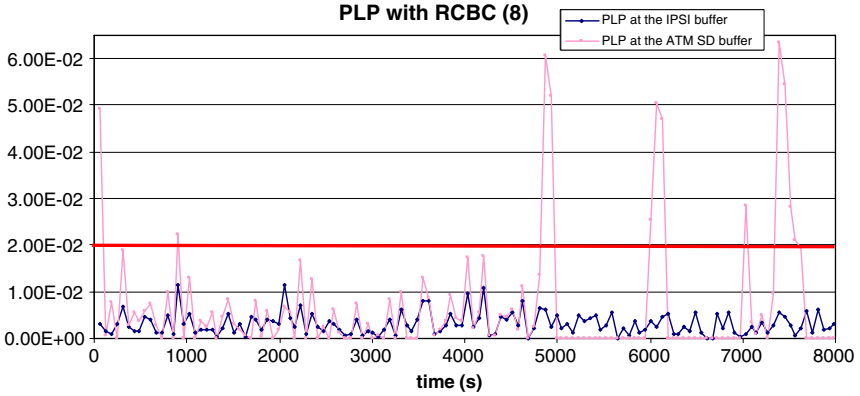

Fig. 19. Fading scenario. PLP with *RCBC* at SI and SD layers.

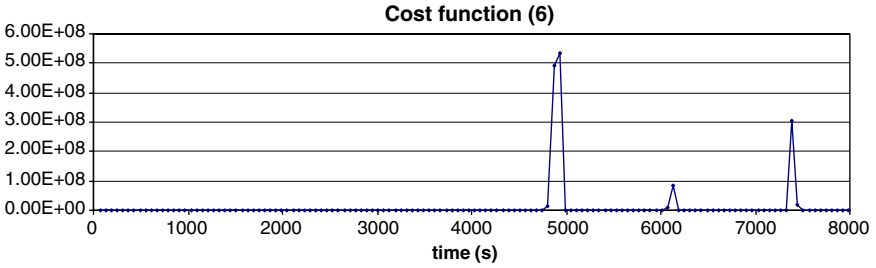

Fig. 20. Fading scenario. Cost function (6) minimized through *RCBC*.

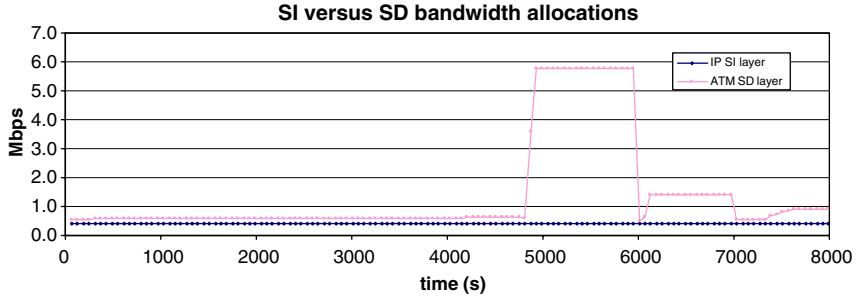

Fig. 21. Fading scenario. Rate allocations, including the additional bandwidth necessary to implement fading counteraction.

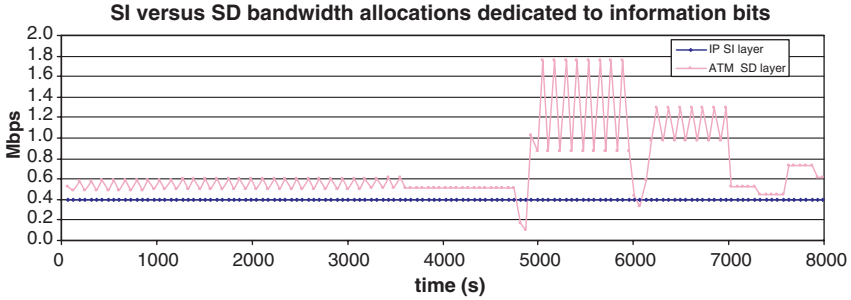

Fig. 22. Fading scenario. Rate allocations, not including the additional bandwidth necessary to implement fading counteraction.

SD layer is compared to that of SI. Fig. 21 depicts the *RCBC* rate allocation including the additional bandwidth necessary to transport redundancy information to implement fading. Fig. 22 shows the "real" service rate $\hat{\theta}^{SD}(t) = \theta^{SD}(t) \cdot \phi(t)$ dedicated to information bits, not considering redundancy. *RCBC* provides the optimal SD bandwidth allocation to counteract time-varying fading levels. Also in this case, the quick convergence assures good performance in real time. It is important to note that *RCBC* finds out the necessary bandwidth to offer the requested SLA also in case of severe fading values, as in the time interval $[4800, 6000]$), when it assigns up to 5.79 Mbps, corresponding to an "information bit" service rate around 0.9 Mbps. The residual bandwidth is used to counteract the channel degradation: $\theta^{SD} = 5.79$ Mbps, $\phi = 0.15625$, $\hat{\theta}^{SD} = \theta^{SD} \cdot \phi \cong 0.90$ Mbps.

7. Conclusions and future work

The challenging problem of the *Broadband Satellite Multimedia* architecture (currently under investigation by the ETSI SES BSM standardization group) is the interaction between the *Satellite Independent* (SI) and *Satellite Dependent* (SD) layers. The SD stack should offer a QoS service to the upper layers independent of a *Service Level Agreement* fixed at the SI layer. The SI–SD interface encompasses the current trend in the interconnection and interworking of network infrastructures, thus giving rise to complex network interdependence where QoS is the final goal.

In this light, a novel control scheme (based on an active measurement technique) has been studied to allow bandwidth adaptation and consequent tracking of the chosen performance metric (the *Packet Loss Probability*, in this case), despite possible changes in the stochastic environment. Extensive simulation results show the good efficiency of the proposed control approach.

The protocol architecture and the control technique investigated in our work can be generalized for other network environments, such as terrestrial and radio broadband systems. In particular, the proposed control scheme is currently under investigation to tackle the aggregation problem mentioned in Section 4. The idea is to deploy a bandwidth control mechanism for the QoS mapping problem arising when different traffic classes are mixed together in an aggregated trunk.

References

[1] ETSI, Satellite Earth Stations and Systems (SES), Broadband Satellite Multimedia, Services and Architectures, ETSI Technical Report, TR 101 984 V1.1.1, November 2002.
[2] ETSI, Satellite Earth Stations and Systems (SES), Broadband Satellite Multimedia, IP over Satellite, ETSI Technical Report, TR 101 985 V1.1.2, November 2002.
[3] ETSI, Satellite Earth Stations and Systems (SES), Broadband Satellite Multimedia, Service and Architectures; functional architecture for IP interworking with BSM networks, ETSI Technical Specification, TS 102 292 V1.1.1, February 2004.
[4] ETSI, Satellite Earth Stations and Systems (SES), Broadband Satellite Multimedia, Services and Architectures; BSM Traffic Classes, ETSI Technical Specification, TS 102 295 V1.1.1, February 2004.
[5] M. Bocci, J. Guillet, ATM in MPLS-based converged core data networks, IEEE Comm. Mag. 41 (1) (2003) 139–145.
[6] ETSI, Technical Committee Satellite Earth Stations and Systems (TC-SES), Standardization issues for QoS control over BSM, Broadband Satellite Multimedia Working Group. Meeting n. 19, Sophia Antipolis, France, Material source: ESA project "Integrated QoS and resource management in DVB-RCS networks" (ASP, BT, EMS Technologies), June 2004.
[7] E. Lutz, H. Bischl, J. Bostic, C. Delucchi, H. Ernst, M. Holtzbock, A. Jahn, M. Werner, ATM-based multimedia communication via satellite, Eur. Trans. Telecommun. 10 (6) (1999) 623–636.

[8] C.M. Barnhart, J.E. Wieselthier, A. Ephremides, Admission control policies for multi-hop wireless networks, Wireless Networks 1 (4) (1995) 373–387.
[9] R. Bolla, F. Davoli, M. Marchese, Bandwidth allocation and admission control in ATM networks with service separation, IEEE Comm. Mag. 35 (5) (1997) 130–137.
[10] N. Celandroni, F. Davoli, E. Ferro, Static and dynamic resource allocation in a multiservice satellite network with fading, Int. J. Satellite Commun. Networking 21 (4–5) (2003) 469–487.
[11] N.J. Keon, G. Anandalingam, Optimal pricing for multiple services in telecommunications networks offering quality-of-serive guarantees, IEEE/ACM Trans. Networking 11 (1) (2003) 680.
[12] D. Bertsekas, Dynamic Programming and Optimal Control, second ed., Athena Scientific, Belmont, MA, 2001.
[13] C.G. Cassandras, S. Lafortune, Introduction to Discrete Event Systems, Kluwer Academic Publisher, Boston, MA, 1999.
[14] K. Ross, Multiservice Loss models for Broadband Telecommunication Networks, Springer Verlag, Berlin, 1995.
[15] C.M. Lagoa, H. Che, B.A. Movsichoff, Adaptive control algorithms for decentralized optimal traffic engineering in the Internet, IEEE/ACM Trans. Networking 12 (3) (2004) 415–428.
[16] P.R. Kumar, New technological vistas for systems and control: the example of wireless networks, IEEE Control Syst. Mag. 21 (1) (2001) 24–37.
[17] J. Aweya, M. Ouelette, D.Y. Montuno, A self-regulating TCP acknowledgement (ACK) pacing scheme, Int. J. Network Manage. 12 (1) (2002) 145–163.
[18] Y. Gao, J.C. Hou, A state feedback control approach to stabilizing queues for ECN-enabled TCP connections, in: Proc. IEEE Infocom 2003, San Francisco, CA, March–April 2003, pp. 2301–2311.
[19] M. Baglietto, T. Parisini, R. Zoppoli, Distributed-information neural control: the case of dynamic routing in traffic networks, IEEE Trans. Neural Networks 12 (3) (2001) 485–502.
[20] F. Delli Priscoli, D. Pompili, G. Sette, A control based traffic controller in wireless and satellite downlink channels, European Control Conference (ECC) 2003, Cambridge, UK, September 2003.
[21] H.J. Kushner, G.G. Yin, Stochastic Approximation Algorithms and Applications, Springer-Verlag, New York, NY, 1997.
[22] Y. Wardi, B. Melamed, C.G. Cassandras, C.G. Panayiotou, Online IPA gradient estimators in stochastic continuous fluid models, J. Optim. Theory Appl. 115 (2) (2002) 369–405.
[23] ETSI, Satellite Earth Stations and Systems (SES), Broadband Satellite Multimedia, Guidelines for the Satellite Independent Service Access Point (SI-SAP), ETSI Technical Report, Draft TR 1XX XXX V⟨0.2.0⟩, June 2004.
[24] ETSI, Satellite Earth Stations and Systems (SES), Broadband Satellite Multimedia, Common Air Interface Specification, Satellite Independent Service Access Point (SI-SAP), TS 102 357 V0.2.5, February 2005.
[25] J. Fischer, PWE3: ATM Service Description, draft-fischer-pwe3-atm-service-03, Internet draft, IETF, March 2002.
[26] L. Martini, Encapsulation Methods for Transport of ATM Cells/Frame over IP and MPLS Networks, Internet draft draft-martini-atm-encap-mpls-01.txt, Internet draft, IETF, June 2002.
[27] S. Ayandeh, A. Krishnamurthy, A. Malis, Mapping to ATM classes of service for differentiated service architecture, draft-ayandeh-diffserv-atm-00.txt, Internet Draft, IETF, April 2000.
[28] M. Garret, M. Borden, Interoperation of controlled-load service and guaranteed service with ATM, RFC 2381, IETF, August 1998.
[29] ATM Forum Tech. Comm., Addendum to traffic management 4.1: Differentiated UBR, July 2000.
[30] P. Giacomazzi, L. Musumeci, Transport of IP controlled-load service over ATM networks, IEEE Network 13 (1) (1999) 36–47.
[31] P.F. Cobley, H. Davies, Performance implications of QoS mapping in heterogeneous networks involving ATM, in: Proc. 1st IEEE Conf. ATM (Icatm) 1998, Colmar, France, June 1998, pp. 529–535.
[32] J. Schmitt, Translation of specification units between IP and ATM quality of service declarations, Int. J. Commun. Syst. 16 (4) (2003) 291–310.
[33] C.G. Cassandras, G. Sun, C.G. Panayiotou, Y. Wardi, Perturbation analysis and control of two-class stochastic fluid models for communication networks, IEEE Trans. Automat. Contr. 48 (5) (2003) 23–32.
[34] K. Gokbayrak, C.G. Cassandras, Adaptive call admission control in circuit-switched networks, IEEE Trans. Automat. Contr. 47 (6) (2002) 1234–1248.
[35] F. Davoli, M. Marchese, M. Mongelli, Optimal resource allocation in satellite networks: certainty equivalent approach versus sensitivity estimation algorithms, Int. J. Commun. Syst. 17 (10) (2004) 935–962.
[36] M. Mellia, A. Carpani, R. Lo Cigno, Measuring IP and TCP behavior on an edge node, in: Proc. IEEE Globecom 2002, Taipei, Taiwan, November 2002, pp. 2540–2544.
[37] K. Byungsuk, I. Sebüktekin, An integrated IP QoS architecture—performance, in: Proc. IEEE Military Commun. Conf. (MILCOM 2002), Anaheim, CA, October 2002, pp. 1189–1193.
[38] ITU-T, Artificial Conversational Speech, ITU-T Recommendation P.59, March 1993.

Mario Marchese (S'94-M'97-SM'04) received a Laurea degree (cum laude) from the University of Genoa, Italy, in 1992, and qualified as a Professional Engineer that same year. He obtained his Ph.D. in Telecommunications from the University of Genoa in 1996. From 1999 to 2004 he worked with the University of Genoa Research Unit, Italian National Consortium of Telecommunications (CNIT), where he was Head of Research. Since 2005 he has been an associate professor in the Department of Communication, Computer and Systems Science (DIST), University of Genoa. He founded and is still technically responsible for the CNIT/DIST Satellite Communication and Networking Laboratory (SCNL), University of Genoa, which contains high-value devices and tools. In this position, he is also responsible for managing different units of specialized scientific and technical personnel. He is Vice-Chair of the IEEE Satellite and

Space Communications Technical Committee. He is the author and coauthor of more than 80 scientific papers in international magazines, international conferences, and book chapters.

His main research interests include satellite networks, transport layer over satellite and wireless networks, quality of service over ATM, IP, and MPLS, and data transport over heterogeneous networks.

Maurizio Mongelli was born in Savona, Italy in 1975. He received his "Laurea" degree (cum laude) from the University of Genoa, Italy in 2000. He obtained his Ph.D. degree in "Electronic and Computer Engineering" from the University of Genoa in 2004. He worked for both Selex Communications S.p.A and the Italian Consortium of Telecommunications (CNIT), at the University of Genoa Research Unit from 2000 to 2004. He is now a member of the research staff of the Telecommunication Networking Research Group at the University of Genoa, with a post-doctoral scholarship sponsored by Selex Communications S.p.A. His main research interests include: QoS protocol architectures and management functions, optimization algorithms and pricing for telecommunication systems.