

An integrated bandwidth allocation and admission control framework for the support of heterogeneous real-time traffic in class-based IP networks

S. Georgoulas^{a,*}, P. Trimintzios^{b,1}, G. Pavlou^{a,2}, K. Ho^{a,3}

^a Centre for Communication Systems Research, University of Surrey, Guildford, Surrey, GU2 7XH, UK

^b ENISA (European Network and Information Security Agency), EU, P.O. Box 1300, Heraklion, Crete 71001, Greece

Received 4 November 2006; received in revised form 11 October 2007; accepted 14 October 2007

Available online 22 October 2007

Abstract

The support of real-time traffic in class-based IP networks requires the reservation of resources in all the links along the end-to-end paths through appropriate queuing and forwarding mechanisms. This resource allocation should be accompanied by appropriate admission control procedures in order to guarantee that newly admitted real-time traffic flows do not cause any violation to the Quality of Service (QoS) experienced by the already established real-time traffic flows. In this paper we initially aim to highlight certain issues with respect to the areas of bandwidth allocation and admission control for the support of real-time traffic in class-based IP networks. We investigate the implications of topological placement of both the bandwidth allocation and admission control schemes. We show that the performance of bandwidth allocation and admission control schemes depends highly on the location of the employed procedures with respect to the end-users requesting the services and the various network boundaries (access, metro, core, etc.). Based on our results we conclude that the strategies for applying these schemes should be location-aware, because the performance of bandwidth allocation and admission control at different points in a class-based IP network, and for the same traffic load, can be quite different and can deviate greatly from the expected performance. Through simulations we also try to provide a quantitative view of the aforementioned deviations. Taking the implications of this “location-awareness” into account, we subsequently present a new Measurement-based Admission Control (MBAC) scheme for real-time traffic that uses measurements of aggregate bandwidth only, without keeping the state of any per-flow information. In this scheme there is no assumption made on the nature of the traffic characteristics of the real-time traffic flows, which can be of heterogeneous nature. Through simulations we show that the admission control scheme is robust with respect to traffic heterogeneity and measurement errors. We also show that our scheme compares favorably against other admission control schemes in the literature.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Real-time traffic; Bandwidth allocation; Admission control; Class-based IP networks

1. Introduction

Current IP networks provide only one simple service; that is the best-effort service. Such a simple service model

allows routers to be stateless and has allowed the Internet to scale well in both the size of the network and the nature of the applications. However, as the Internet evolves into a global communication infrastructure, there is a need to provide more sophisticated service models in order to support emerging services, such as Voice over IP (VoIP), Videoconference, and real-time traffic in general. Real-time traffic has QoS requirements that the current best-effort Internet cannot provide in a resource and, consequently, cost-effective manner, e.g. without massive overprovisioning.

* Corresponding author. Tel.: +44 1483683641; fax: +44 1483686011.

E-mail addresses: S.Georgoulas@eim.surrey.ac.uk (S. Georgoulas), panagiotis.trimintzios@enisa.europa.eu (P. Trimintzios), G.Pavlou@eim.surrey.ac.uk (G. Pavlou), K.Ho@eim.surrey.ac.uk (K. Ho).

¹ Tel.: +30 2810391280.

² Tel.: +44 1483689480.

³ Tel.: +44 1483689495.

Integrated services (Intserv) was the first attempt towards supporting QoS in the Internet, but the per-flow state information that was required in all routers, both edge and core, resulted in major scalability problems that prevented its wide deployment. Class-based service models, such as the Differentiated services (Diffserv) model, are offering an alternative approach towards QoS in the Internet that alleviates the scalability problems of Intserv. This is achieved by grouping traffic with similar QoS requirements into one of the engineered traffic classes and forwarding it in an aggregate fashion, using resources that have been pre-provisioned for each particular traffic class. By allowing traffic aggregation within the engineered traffic classes, network domains that deploy the class-based service models can take advantage of statistical multiplexing, which allows for efficient use of relevant resources. In order to provide QoS guarantees, a network supporting different classes must also deploy admission control in order to control the amount of traffic injected into the network so as to prevent overload situations that can lead to QoS violations. In class-based networks, the violation refers to the QoS perceived by each of the engineered aggregate traffic classes rather than the QoS of individual flows. That means that upon an overload situation in a given service class, all flows in that class can suffer a potential QoS degradation.

In this work we focus on the support of *real-time* traffic flows in an IP network with class-based service support. We assume that in such a network, there exists end-to-end isolation between the UDP real-time traffic and the TCP data traffic. This is essential for guaranteeing QoS [1] in class-based networks. We can achieve this isolation by using different queues for the two types of traffic. An example of this scenario could be a Diffserv network where the real-time traffic is classified to use a higher priority Assured Forwarding (AF) Per-Hop Behavior (PHB), whereas the TCP controlled traffic is classified, for example, to use a lower priority, possibly best-effort (BE), forwarding PHB.

We define as *real-time* traffic flows, the flows that have strict delay and jitter requirements, and a bounded, not necessarily too small, packet loss rate (PLR) requirement. For services, such as Voice or Video, a certain amount of packet loss can be acceptable [2] without significant quality degradation. Therefore, such services do not need the ‘virtual wire’ (Expedited Forwarding (EF) in the Diffserv model) treatment [3]. Regarding the low delay and low jitter requirements, both are likely to be met in a high-speed core network [2]. Furthermore, certain off-line traffic engineering actions can be taken so that delay and jitter are kept within certain low bounds. For example, the delay requirement can be taken into account at the network provisioning phase by: a) configuring appropriately small packet queues for the real-time traffic class in order to keep the per-hop delay small, and b) controlling the routing process to choose paths with a constrained number of hops. Jitter can remain controlled in successive multiplexing queues as long as the real-time traffic flows are shaped to

their nominal peak rate at the network ingress [4]. In addition, the deployment of non-work conserving scheduling in routers for the real-time traffic class can be beneficial for controlling jitter [5]. Therefore, we assume that the real-time traffic flows can be shaped to their nominal peak rate at the network domain ingress and that the scheduling mechanism for the real-time traffic class is priority scheduling with a strict bandwidth limit and with First-In-First-Out (FIFO) service discipline within the real-time traffic class itself. Taking the above into account, we employ the PLR as the QoS metric that needs to be controlled and we focus on keeping it in values lower than the PLR requirement of real-time traffic.

In this work we first aim to demonstrate how topological placement, that is the location of the employed bandwidth allocation and admission control schemes with respect to the end-users requesting the services and the various network boundaries (access, metro, core, etc.), can affect the performance of these schemes. Then, taking into account the implications of topological placement, we propose an integrated bandwidth allocation and admission control framework that can overcome the effects of topological placement. In the context of this framework, we also present a new Measurement-based Admission Control (MBAC) scheme for real-time traffic and we compare its performance against other admission control schemes found in the literature.

The rest of this paper is organized as follows: in Section 2 we give a brief description of bandwidth allocation schemes that account for statistical multiplexing, being therefore suitable for aggregating real-time traffic flows in class-based IP networks. In Section 3 we also briefly describe the various admission control schemes that can be used for real-time traffic in class-based IP networks, so that later we can refer to each of them and point how the results of our work in the following Sections relate to them. In Section 4 we demonstrate how topological placement, that is the location of the employed bandwidth allocation and admission control schemes with respect to the end-users requesting the services and the various network boundaries (access, metro, core, etc.), can affect the performance of these schemes. In Section 5 we discuss the implications of topological placement and in Section 6, taking into account the implications of topological placement, we present a new Measurement-based admission control scheme for real-time traffic. In Section 7 we compare the performance of our admission control scheme against other admission control schemes found in the literature and in Section 8 we discuss potential limitations and implementation issues of the proposed admission control scheme. Finally, in Section 9, we conclude, summarizing our findings and giving some directions for future work.

2. Bandwidth allocation schemes

Bandwidth allocation schemes can be divided into two main categories. The first category comprises schemes

based on *bufferless statistical multiplexing*. This aims to ensure that the combined arrival rate of the multiplexed sources exceeds the allocated capacity only with very small probability. Examples of such bandwidth allocation schemes can be found in [6,7]. The second category comprises schemes based on *buffered statistical multiplexing*. Contrary to bufferless multiplexing, buffered multiplexing allows an input rate excess, with surplus traffic being temporarily stored in large buffers. Examples of bandwidth allocation schemes that are based on buffered statistical multiplexing can be found in [6,8,9]. Generally speaking, both categories take into account factors such as the number and characteristics of flows (expressed in terms of parameters such as peak bit rate, mean bit rate, variance of bit rate or as token bucket filter parameters), the required loss rate and, in the case of buffered multiplexing, the available buffer size and derive the required capacity (*effective bandwidth*) needed for the actual loss rate to be kept within bounds.

Each of the two categories has its own merits but also its drawbacks. Bufferless multiplexing is attractive for real-time traffic since it ensures that the traffic experiences minimal delay. In addition, the dynamics leading to an overload event in a bufferless system are much simpler than those of a buffered system [10]. The main disadvantage of using a buffer is that the overflow probability depends significantly on assumptions about the flow characteristics [11], and can only be controlled if these characteristics are known a priori without large deviations. Moreover, in this case, provisioning needs to account for statistical variations in the traffic mix as new flows arrive and others terminate. These complications are largely avoided with bufferless multiplexing. On the other hand, buffered multiplexing allows for higher utilization for the same loss rate [7,11,12] but, as stated above, requires more complex traffic management. Also it is not as robust with respect to flow characteristics as bufferless multiplexing and can incur longer delays. We need to stress here that bufferless multiplexing is, of course, just a model abstraction [11]. For packetized traffic, as in IP networks, a small buffer for packet scale queuing is needed to take into account simultaneous packet arrivals from distinct flows [4,13]. However, in this case the buffer is only used for packet scale queuing [13] and not for storing excess traffic and is, therefore, not involved in bandwidth estimation calculations.

3. Admission control schemes

We can broadly divide the various admission control schemes found in the literature into three categories: endpoint admission control (EAC), traffic descriptor-based admission control (TDAC), and measurement-based admission control (MBAC).

Endpoint admission control is based on metrics applied to probing packets sent by the end host/application along the transmission path before the main flow is established [14]. The probing packets can be sent either at the same pri-

ority as the main flow packets (in-band probing) or at a lower priority (out-of-band probing). For in-band probing, the QoS target for the main flow can be directly used as an admission threshold. For out-of-band probing, the probing flow QoS is not directly related to the main flow QoS and it is, therefore, necessary to establish a mechanism to find the optimum admission threshold [15]. A requirement for the end-to-end route is to be the same for probing packets and flows. Setup delays may be high, especially for stringent QoS targets [14], and for reasonably bounded setup delays the metrics do not depict stationary network states but rather snapshots of network status, which can result to an unrealistic picture of the network congestion level. Furthermore, simultaneous probing by many sources can lead to a situation known as thrashing [14]. That is, even though the number of admitted flows is small, the cumulative level of probing packets prevents further admissions, driving the utilization to very low values. For in-band probing, *thrashing* additionally degrades the QoS perceived by the already established flows because of the stress incurred by the probing packets (this situation is referred as *collapse* [14]). For out-of-band probing, since probing packets are sent at a lower priority than the main flow packets, *thrashing* does not degrade the QoS perceived by the already established flows. On the contrary, it can have the opposite effect, since the higher priority main flows traffic class is emptied (this situation is referred as *starvation* [14]). The advantage of EAC is that it does not require core routers to keep per-flow state and to process any per-flow reservation messages [14].

Traffic descriptor-based admission control is based on the assumption that traffic descriptors, either deterministic or stochastic, are provided by the application for each flow prior to its establishment. This approach achieves high utilization when traffic descriptors used by the admission control scheme are appropriate. Nevertheless, in practice, it suffers from several problems [16]. One of them is the inability of the application users to come up with appropriate traffic descriptors before establishing the flows. This is especially so when the bandwidth fluctuates over multiple time scales (long-range dependent traffic). In case users overestimate their requirements, resources will be wasted in the network for the entire flow duration, leading, therefore, to reduced network utilization. On the other hand, if users underestimate their requirements, insufficient resources will be allocated to their flows for the entire flow duration, leading therefore to QoS degradation. Another problem is that the provided traffic descriptors and the associated QoS guarantees define a contract between the application and the network. Therefore, the need to police based on this traffic specification arises, which is difficult, especially for statistical traffic descriptors [16]. Deterministic models, such as token buckets, are easy to police, as they specify the worst-case behavior of traffic on a single time scale, but they fail to provide a sufficient characterization to extract a large fraction of the potential statistical multiplexing gain. While a sequence of token buckets can

approximate such a multiple time-scale characterization, the number of model parameters grows accordingly [17].

Measurement-based admission control tries to avoid the problems of the other approaches by shifting the task of traffic characterization from the application to the network [18]. Instead of applications explicitly specifying their traffic descriptors, the network attempts to “learn” the characteristics of existing flows through real-time measurements. This approach has a number of advantages. First, the specified traffic descriptors can be very simple, e.g. peak rate. Second, an overly conservative specification does not result in over-allocation of resources for the entire duration of the service session. Third, when traffic from different flows is multiplexed, the QoS experienced depends often on their aggregate behavior, the statistics of which are easier to estimate than those of an individual flow (this is a consequence of the law of large numbers). However, relying on measured quantities raises a number of issues. As with any measurement/estimation procedure, there exists the possibility of making errors, which can translate to erroneous flow admission decisions, leading either to QoS degradation, if flows are erroneously granted admission, or reduced network utilization, if flows are erroneously refused admission. The impact of wrong flow admission decisions on performance depends on how long it takes until this error is corrected, that is, on flow departure dynamics. Moreover, the issue of determining a proper amount of memory to use for past information about the flows currently present in the network arises [18].

4. The effects of topological placement

In this Section we will present the first part of our work, which aims to show how topological placement of the respective functionalities can affect the performance of bandwidth allocation and admission control schemes.

For this part of our work, regarding bandwidth allocation, we will initially adopt the normal distribution based bufferless statistical multiplexing approach. According to [6], when the effect of statistical multiplexing is significant, the distribution of the stationary bit rate can be accurately approximated by a Gaussian distribution. In [19,20] it is strongly suggested that the aggregation of even a fairly small number of traffic streams is usually sufficient for the Gaussian characterization of the input process. Furthermore, Gaussian processes cover all second-order long-range dependent or second-order self-similar processes that have been shown to be good models for characterizing actual traffic [19].

In this case, the effective bandwidth of N multiplexed sources is given by [6,8]:

$$C = \sum_{i=1}^N m_i + \sqrt{-\ln(2\pi\varepsilon^2) \sum_{i=1}^N \sigma_i^2} \quad (1)$$

where $\sum_{i=1}^N m_i$ denotes the mean aggregate bit rate, $\sum_{i=1}^N \sigma_i^2$ the variance of the aggregate bit rate, and ε the upper

bound on allowed loss rate. In the rest of this Section we will denote the function of Eq. (1) as $eff(S, PLR)$, where S is the set of sources under aggregation and PLR is the packet loss rate value involved in the calculation of the effective bandwidth C .

We will present this first part of our work with a list of different simulated scenarios using a two-level tree topology, which allows us to illustrate the main points of our study, while, at the same time, being simple enough to suit the nature and the computational demands of the required packet-level simulations. (An initial reduced set of results was originally presented in [21].)

4.1. Scenario I: the effects of aggregate bandwidth allocation

Initially we consider the scenario depicted in Fig. 1.

In this scenario it is assumed that a set of sources, S_i , $i = 1, \dots, M$, are aggregated at node 1 and that another set of sources, S_i , $i = M + 1, \dots, N$, are aggregated at node 2. We assume that the sources connect to nodes 1 and 2 with direct links with negligible congestion, and that all of them will be eventually aggregated in the same traffic class at link 3. The capacity reserved in link 1 for the first set of sources is:

$$C_1 = eff(\{S_1, \dots, S_M\}, PLR_1) \quad (2)$$

where PLR_1 is the packet loss rate budget for the real-time traffic class aggregate in link 1. Similarly, for the second set of sources, the capacity reserved in link 2 is:

$$C_2 = eff(\{S_{M+1}, \dots, S_N\}, PLR_2) \quad (3)$$

where PLR_2 is the allowed packet loss rate budget for the real-time traffic class aggregate in link 2. Since all the sources will be aggregated using the same class at link 3, the required bandwidth to be allocated in link 3 for their aggregation is given by:

$$C_3 = eff(\{S_1, \dots, S_N\}, PLR_3) \quad (4)$$

where PLR_3 is the allowed packet loss rate budget for the real-time traffic class aggregate in link 3.

This scenario could correspond to a situation where end-users (the 1, ..., N sources) connect to the edge routers (nodes 1 and 2), which then connect to the metro/backbone router (node 3) through access links 1 and 2.

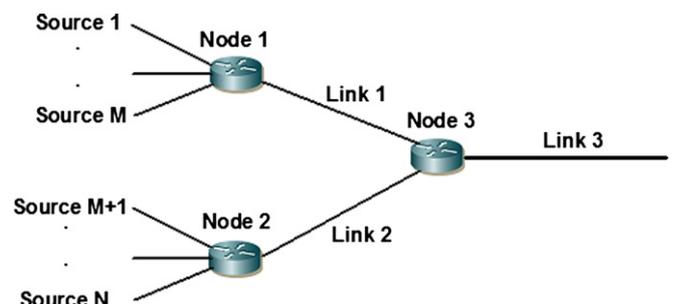


Fig. 1. Topology for assessing the effects of topological placement.

As it can be easily proven [22], packet loss rate parameters are multiplicative. That means that for a set of sources that traverse a sequence of links, l_i , $i = 1, \dots, L$ with packet loss rates PLR_i , the total packet loss rate PLR_{total} can be approximated using the expression:

$$PLR_{total} = 1 - \prod_{i=1}^L (1 - PLR_i) \quad (5)$$

which, in turn becomes additive for low values of PLR_i :

$$PLR_{total} = \sum_{i=1}^L PLR_i \quad (6)$$

Assuming that $PLR_1 = PLR_2 = PLR_3$, that is the allocated capacities at links 1, 2 and 3 for the real-time traffic class are such that allow for the same packet loss rate budget at all links, the expected overall upper bound on end-to-end PLR for the aggregate sources in our topology should be:

$$PLR_{total} = PLR_1 + PLR_3 = PLR_2 + PLR_3 \quad (7)$$

This first part of our work aims to examine whether the actual total packet loss rate experienced by the sources is bounded by the above expression. In order to do so, we run simulations using the network simulator *ns-2* [23]. For the simulations we use two example values for the target link packet loss rate, 0.01 and 0.001. We fix the packet size for the real-time flows to 100 bytes (constant packet size seems to be a reasonable assumption for Voice and Video communications [24]) and we also fix the output queues for packet scale queuing for the real-time traffic class in the order of 500 bytes. We assume, without loss of generality, that the same number of sources is aggregated in both links 1 and 2, i.e. $M = N/2$. This means that the total packet loss rate, end-to-end in our topology, should not exceed 0.02 and 0.002, respectively. We also consider the case where the capacity in link 3 is provisioned so that $PLR_3 = 0$. This happens when $C_3 = C_1 + C_2$, which means that only links 1 and 2 incur losses and in link 3 the real-time traffic aggregates from nodes 1 and 2 are treated using peak rate allocations.

We consider three scenarios for the N traffic sources: (a) all sources are *VoIP* sources with peak rate 64 kbps and exponentially distributed ON and OFF periods with average durations 1.004 s and 1.587 s, respectively (mean rate 24.8 kbps, standard deviation of rate 31.18 kbps) [25], (b) all sources are *Videoconference* sources with mean rate 3.89 Mbps, peak rate 10.585 Mbps and standard deviation of rate 1.725 Mbps [26], and (c) that we have a heterogeneous *mixture of real-time traffic* of both VoIP and Videoconference sources.

Since the real-time traffic class is assumed to be isolated from other classes, we do not consider any best-effort or any other traffic classes in our simulations and we simulate the real-time traffic as being serviced by queues running at the speed of their bandwidth limit. While this is not precisely the behaviour of rate-limited priority queues, it is

rather close [14]. All the results given in this Section are based on averages of simulations for 20 randomly chosen seeds, each for a total of 4100 s, using the first 500 s as a warming-up period. This means that the simulations are effectively run for an hour of simulation time (3600 s) after the warming up period, which is enough to obtain statistically stable results.

In Figs. 2 and 3, PLRa corresponds to the (average) packet loss incurred at links 1 and 2, while PLRb corresponds to the total packet loss for the cases where $PLR_1 = PLR_2 = PLR_3$ and they are given as a function of the mean aggregate bit rate of all sources S_i , $i = 1, \dots, N$ (x -axis).

From these two figures we can see that in all cases, the target packet loss rate in links 1 and 2 (PLRa) is always satisfied. Regarding the total packet loss rate, when the bandwidth in link 3 is set so that link 3 also incurs losses (PLRb) we can see that it is kept below the target total packet loss rate 0.02 (see Fig. 2), but not in the case where the target total packet loss rate is set to 0.002 (see Fig. 3). Since the target packet loss rate in links 1 and 2 is always satisfied, this means that in the latter case link 3 incurs losses that are much higher than the target packet loss rate budget at that link.

These results suggest that even though the original traffic descriptors are valid and appropriate for bandwidth allocation at the first points of aggregation (nodes 1 and 2), they may not be valid in transit nodes such as in node 3. This is because traffic sources become correlated and their characteristics are altered as they traverse a sequence of links and multiplexers. Therefore, using the original traffic descriptors for bandwidth allocation in transit nodes can lead to erroneous bandwidth allocation decisions, indicated in our simulations as excessive values of packet loss rates at link 3. This traffic profile deformation has also been verified in the past [13,27,28] and a number of solutions for dealing with it have been proposed and by means of queuing theory, analytical models for evaluating it for specific types of individual traffic sources and under specific network conditions and assumptions have been presented [13].

One proposed solution is the deployment of traffic-descriptor conserving scheduling disciplines in all links along the end-to-end paths of the real-time traffic flows. Example of such schedulers is the Rate-controlled Static Priority (RCSP) scheduler [29]. This preserves the original traffic descriptors of each individual flow going through it and provides zero packet loss guarantees. However, it requires per-flow queuing and keeping the traffic descriptors of each flow in each node (it induces, therefore, per-flow state in all nodes) and can also be computationally expensive. Therefore, deploying RCSP in a class-based network does not seem feasible (e.g. it violates the basic Diff-serv paradigm with respect to the functionality of the core nodes). Furthermore, since an amount of packet loss can be acceptable for real-time traffic services such as Voice or Video [2], and RCSP is intended to provide zero packet loss guarantees, the deployment of RCSP can lead to unnecessarily conservative use of network resources.

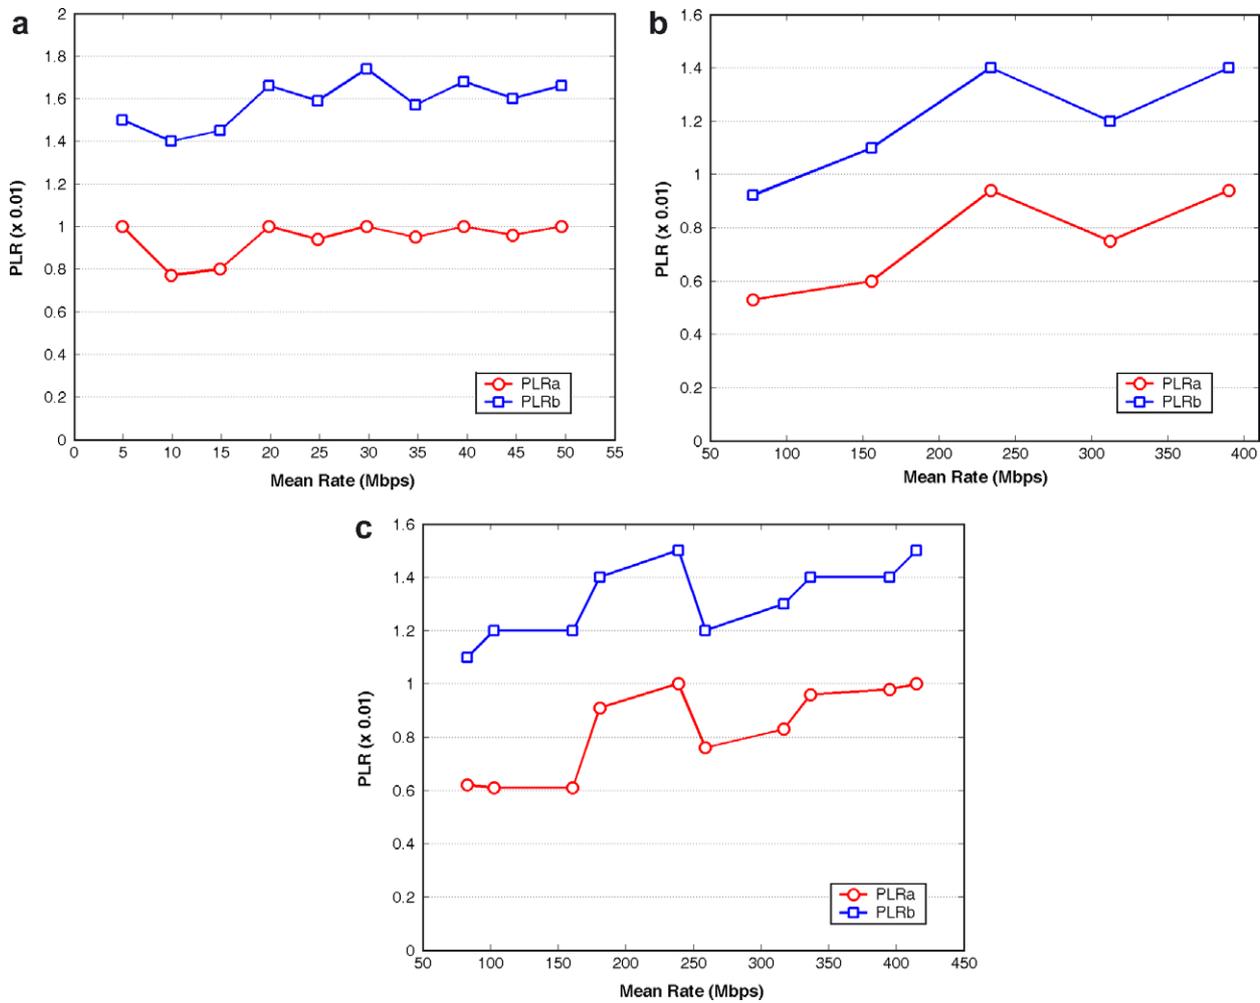


Fig. 2. Incurred PLR for *VoIP* (a), *Videoconference* (b) and *mixed* (c) traffic sources for target link PLR 0.01.

In [13] the issue of traffic profile deformation is discussed in the context of Constant Bit Rate (CBR) sources in Asynchronous Transfer Mode (ATM) networks and a solution for accounting for the traffic profile deformation of individual sources is given based on the estimation of delay variation, which, however, induces per-flow state and, moreover, requires an appropriate method in order to obtain this delay variation estimation. Furthermore, an analytical model for quantifying the effect of traffic profile deformation as a foreground CBR source is multiplexed in a FIFO multiplexer with background traffic and also as a main stream passes a tandem of queues with interfering traffic is given. However, there are two assumptions made with respect to the aforementioned analysis that, as the authors themselves admit “with an overwhelming probability will not be fulfilled in any real applications”. The first assumption is that all queues are loaded to 1 (heavy load) and the second assumption is that the background traffic is a superposition of a number of Bernoulli streams.

In [30], the authors build on the delay variation estimation method presented in [13] and propose the use of probing packets for obtaining the delay variation of individual inter-domain traffic sources (that is sources that cross more

than one domains) in order to update the traffic descriptors of the sources before submitting them to the next downstream domain for use in admission control. However, this requires the cooperation of adjacent domains on a per-flow basis (meaning that a downstream domain must be aware on a per-flow basis of all the flows originating from upstream domains). Furthermore, as in [13], they do not consider generic sources but only one foreground source (modeled as a source with deterministic ON and OFF periods) multiplexed with a background source, which is modeled as a Poisson process.

In this first part of our work, contrary to these aforementioned works, we focus not on the traffic profile deformation of individual and type-specific sources, which, as explained induces per-flow states and can require added functionality in core nodes for the estimation of the delay variation. Instead, we focus on the aggregate traffic profile deformation of generic sources and we try to go one step further and to quantify through simulations the effect that this aggregate traffic profile deformation can have on the incurred packet loss rates and, therefore, on the perceived QoS. We discuss the issues that the effects of the aggregate traffic profile deformation raise with respect to bandwidth

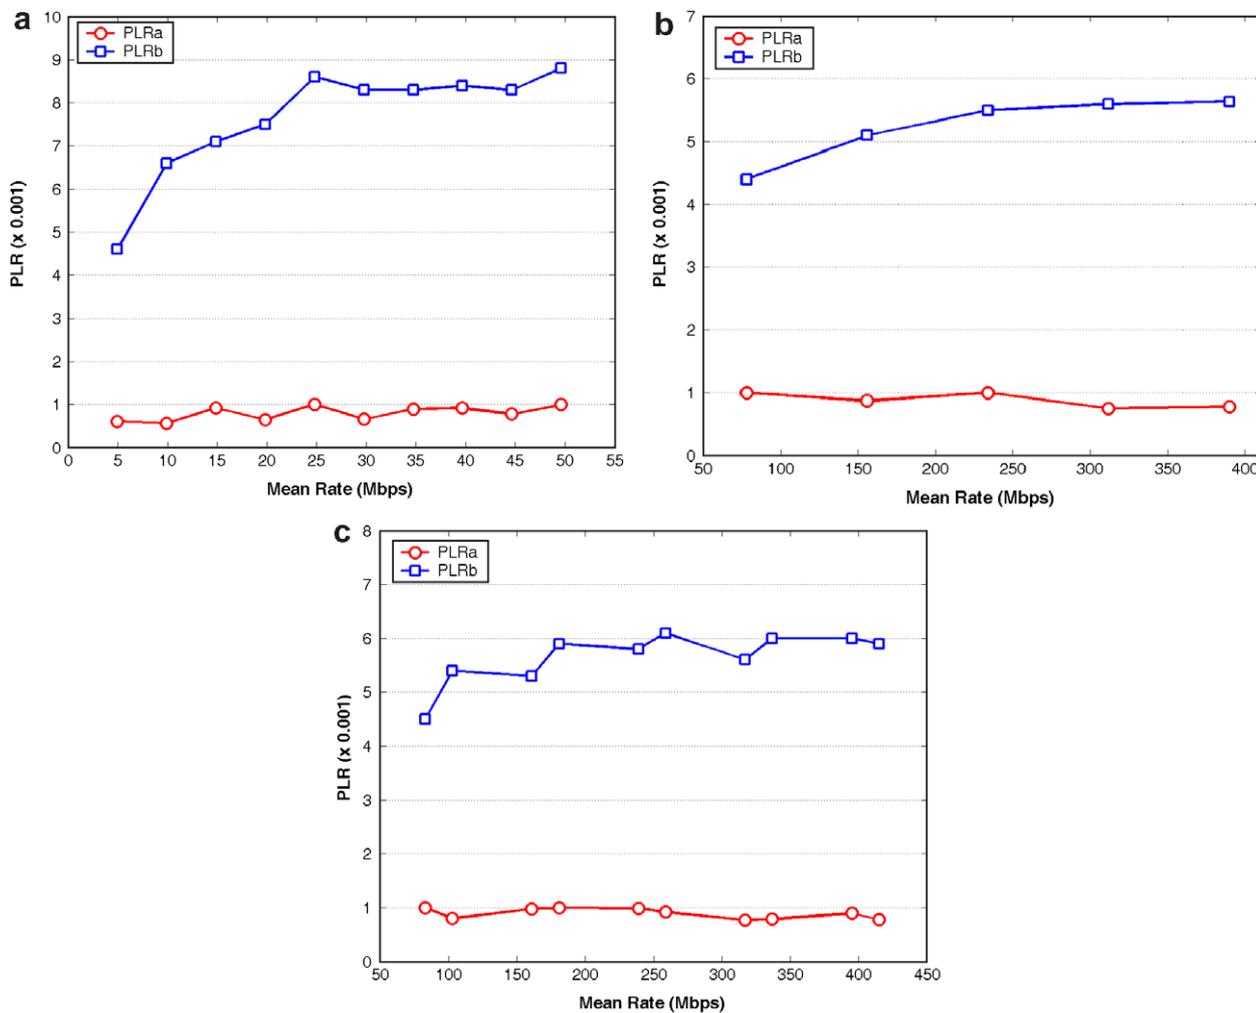


Fig. 3. Incurred PLR for *VoIP* (a), *Videoconference* (b) and *mixed* (c) traffic sources for target link PLR 0.001.

allocation and admission control and how they can be addressed, without, most importantly, imposing the requirement for keeping any per-flow state information in core nodes.

We consider the general cases where short-range dependent (VoIP), long-range dependent (Videoconference), as well as a heterogeneous mixture of short-range and long-range dependent (VoIP and Videoconference) real-time traffic sources are multiplexed in the same traffic class and, therefore, all of them need to be considered and treated as foreground sources. In such cases, obtaining a closed form solution to quantify the effect of the aggregate traffic profile deformation using queuing theory can become very difficult or even intractable. Also, apart from the bufferless statistical multiplexing model we will also examine the case where a buffered statistical multiplexing model is used for aggregate bandwidth allocation.

4.2. Scenario II: quantifying the effects of the traffic profile deformation

In order to quantify the effect of traffic profile deformation further downstream from the first hop node after the

traffic sources are multiplexed, we proceed as follows. We use the same simulated topology, traffic volume and types of traffic sources, as in the previous scenario. We set the capacities allocated to the real-time traffic class in links 1 and 2 equal to the sum of peak rates of the traffic sources that are aggregated in links 1 and 2. This means that links 1 and 2 are transparent to the sources with respect to packet loss (zero packet loss). For link 3 we distinguish two cases. In the first case, we merge the two aggregate demands – one composed of the sources S_i , $i = 1, \dots, M$ and the other composed of the sources S_i , $i = M + 1, \dots, N$, in one bandwidth allocation in link 3, that is:

$$C_3 = \text{eff}(\{S_1, \dots, S_N\}, PLR_3) \quad (8)$$

In the second case we reserve resources for each aggregate demand independently in link 3, even though the sources composing the two aggregate demands will be eventually aggregated in the same traffic class in link 3, (this is referred sometimes as *isolation* [31]), that is:

$$C_3 = \text{eff}(\{S_1, \dots, S_M\}, PLR_3) + \text{eff}(\{S_{M+1}, \dots, S_N\}, PLR_3) \quad (9)$$

In Figs. 4 and 5, the target packet loss rate for link 3 is set to 0.01 and 0.001, respectively. PLRb corresponds to the incurred packet loss rate from link 3 when using Eq. (8) for bandwidth allocation and PLRa corresponds to the incurred packet loss rate when using Eq. (9).

As it can be seen from these two figures, the effect of the traffic profile deformation, as sources become correlated and their characteristics change while they traverse links and multiplexers, even only one hop away from the node where the traffic sources are firstly multiplexed (first point of aggregation), can lead to severe violations of the packet loss rate. Even if the *isolation* method (Eq. (9)) for bandwidth allocation is used, which given the form of the effective bandwidth formula of Eq. (1) leads to more conservative resource reservation compared to the case where the two aggregate demands are merged into one bandwidth allocation (Eq. (8)), the target packet loss rate can be violated by more than one order of magnitude. Furthermore, the increase in the level of aggregation does not improve the situation. On the contrary, it can lead to higher packet loss rate violations. This indicates that the

detrimental effects of the aggregate traffic profile deformation with respect to bandwidth allocation can by far exceed the positive anticipated effects due to the increase in the level of aggregation and the additional statistical multiplexing gain that this increase is expected to have.

4.3. Scenario III: bandwidth allocation with buffered multiplexing models

In order to have a more complete view and not restrict ourselves to the bufferless statistical multiplexing approach, we repeat part of the above simulations using the buffered approach introduced in [8] for bandwidth allocation. According to [8], for a source of type i with average rate m_i , the effective bandwidth is given by:

$$C_i = m_i + \delta\gamma_i/(2B) \tag{10}$$

where B is the buffer size, γ_i is the index of dispersion and $\delta = -\ln(\varepsilon)$ with ε the allowed loss rate. For M different types of traffic sources, with N_i sources of type i the total effective bandwidth is given by:

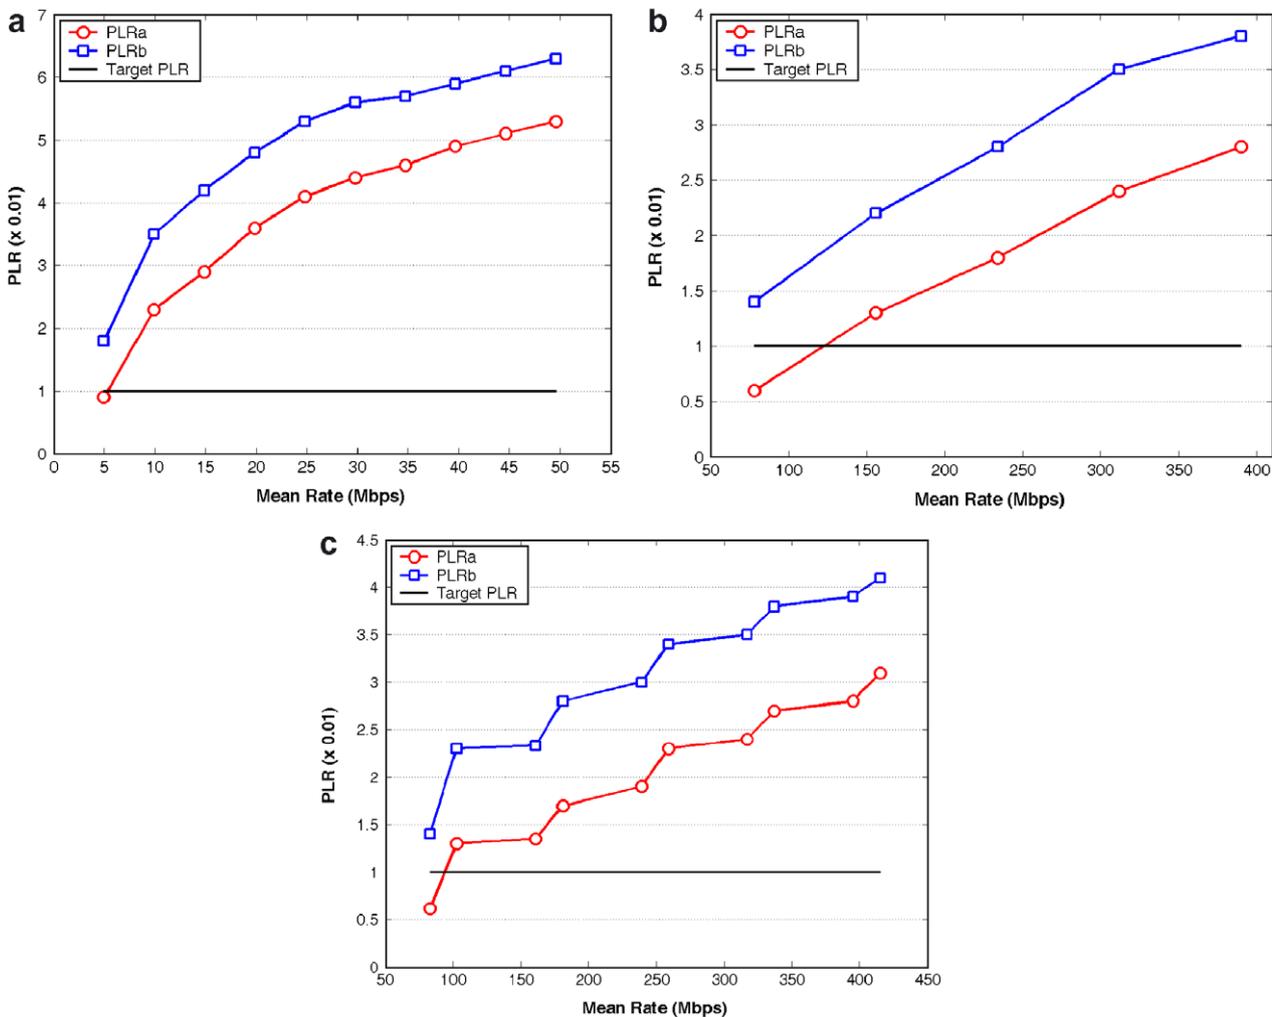


Fig. 4. Incurred PLR for VoIP (a), Videoconference (b) and mixed (c) traffic sources for link 3 with target PLR 0.01.

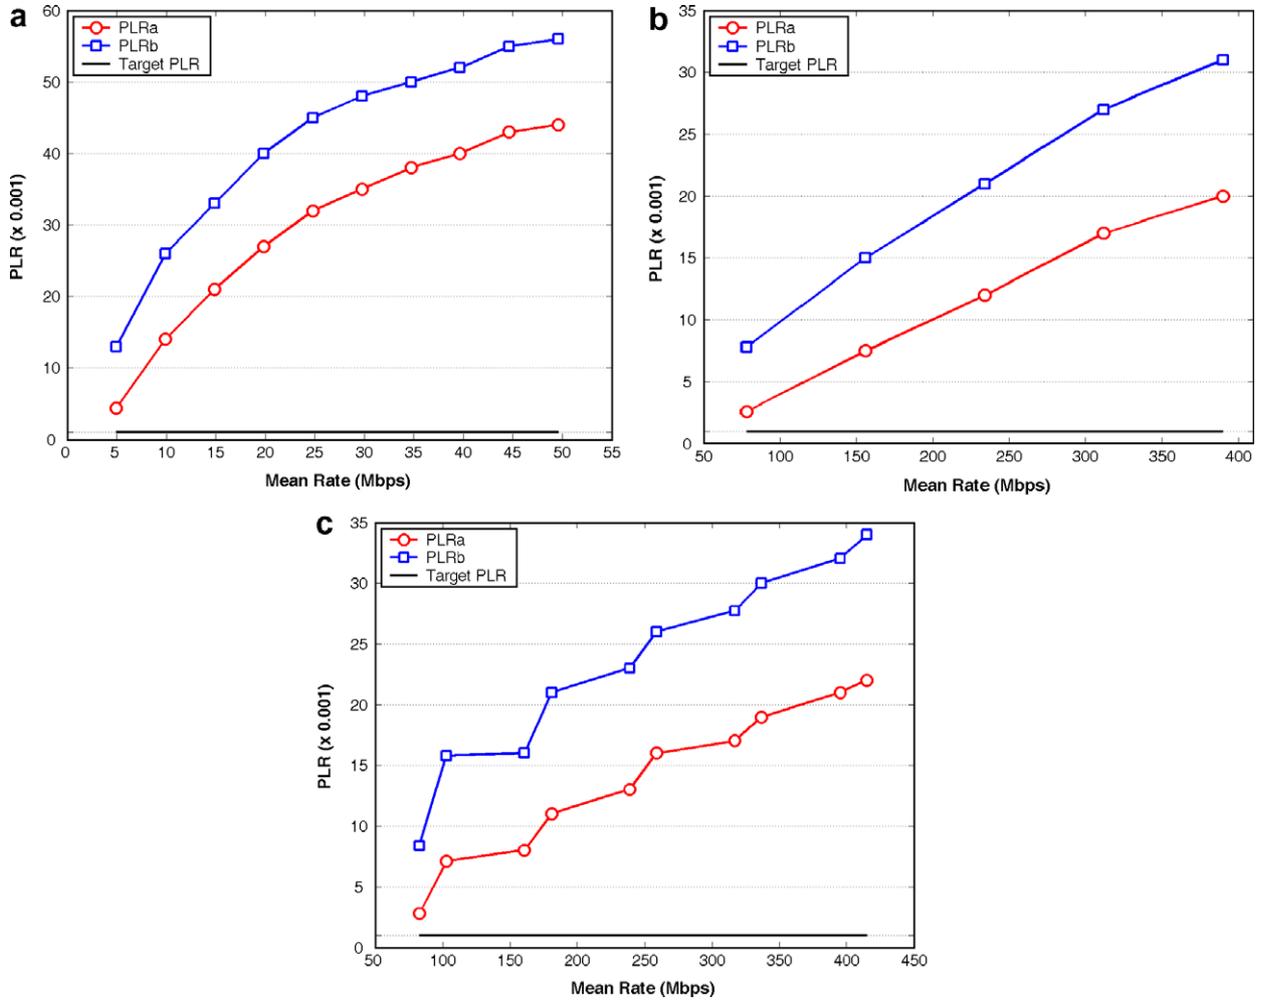


Fig. 5. Incurred PLR for VoIP (a), Videoconference (b) and mixed (c) traffic sources for link 3 with target PLR 0.001.

$$C = \sum_{i=1}^M N_i C_i \tag{11}$$

We use this effective bandwidth formula to estimate the bandwidth for the case of the Videoconference sources. We repeat our experiments for two buffer size levels, 30 kbytes and 50 kbytes, respectively. These levels are set intentionally relatively small since we expect the real-time traffic class to be provisioned with small queues in order to minimize excessive values of per-hop delay.

In Figs. 6 and 7, PLRa corresponds to the average packet loss incurred in links 1 and 2 when the Videoconference sources are aggregated in links 1 and 2 and the target packet loss rate budgets of links 1 and 2 are set equal to 0.01 and 0.001 (similar to PLRa in Scenario I). PLRb corresponds to the packet loss incurred in link 3 in the case where the same number of sources, equal to the number of sources aggregated in link 1 and link 2 for the previous configuration separately, is finally aggregated in link 3, with half of them aggregated through link 1 and half of them aggregated through link 2. For this case the capacities of links 1 and 2 set equal to the sum of peak rates of the sources they are carrying, that means

that links 1 and 2 are transparent to the sources they are carrying with respect to packet loss, and the bandwidth allocated in link 3 is for target link 3 loss rate equal to 0.01 and 0.001.

As it can be seen from PLRa, using Eq. (11) for the sources on links 1 and 2 for target link PLR 0.01 gives the expected results for the incurred packet loss rate on these links, whereas it is quite optimistic for the case of target link PLR 0.001, leading to packet loss rate violations for a small numbers of multiplexed sources. Moreover, applying Eq. (11) to calculate the provisioned effective bandwidth in link 3 (PLRb) can lead to excessive violation of the packet loss rate, especially for link 3 target loss rate 0.001, and for this effective bandwidth formula the violation is bigger for increasing queue size. Furthermore, in contrast to the results in scenario II, the packet loss in link 3 does not increase with the level of aggregation. This is due to the additive nature of the effective bandwidth formula (11), which, in contrast to Eq. (1) that becomes less conservative for increasing levels of aggregation, becomes more conservative and therefore can compensate for the detrimental effects of the aggregate traffic profile deformation.

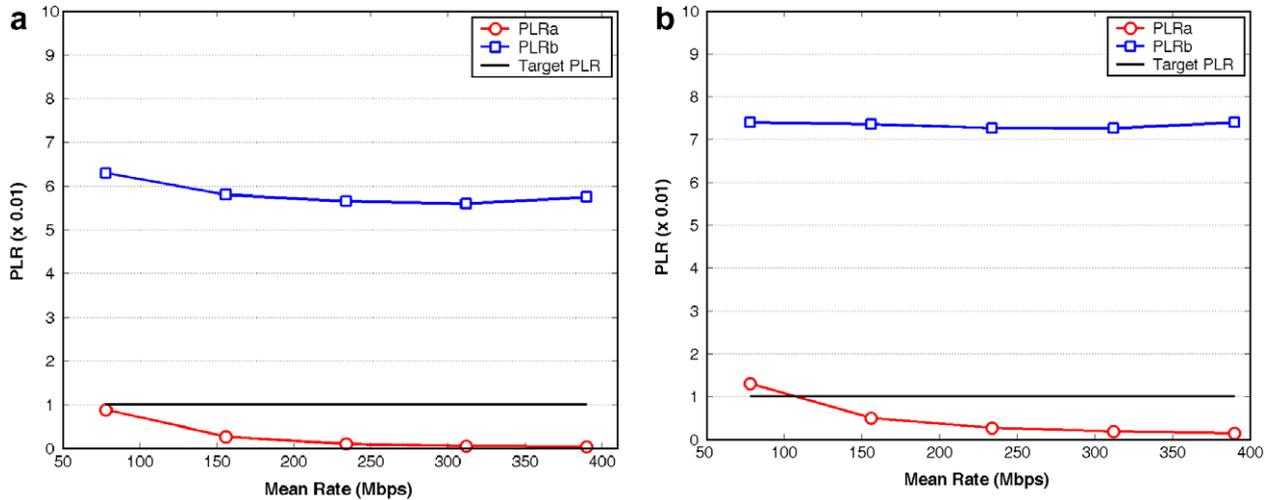


Fig. 6. Incurred PLR for *Videoconference* sources for target link PLR 0.01 and queue size 30 kbytes (a) and 50 kbytes (b).

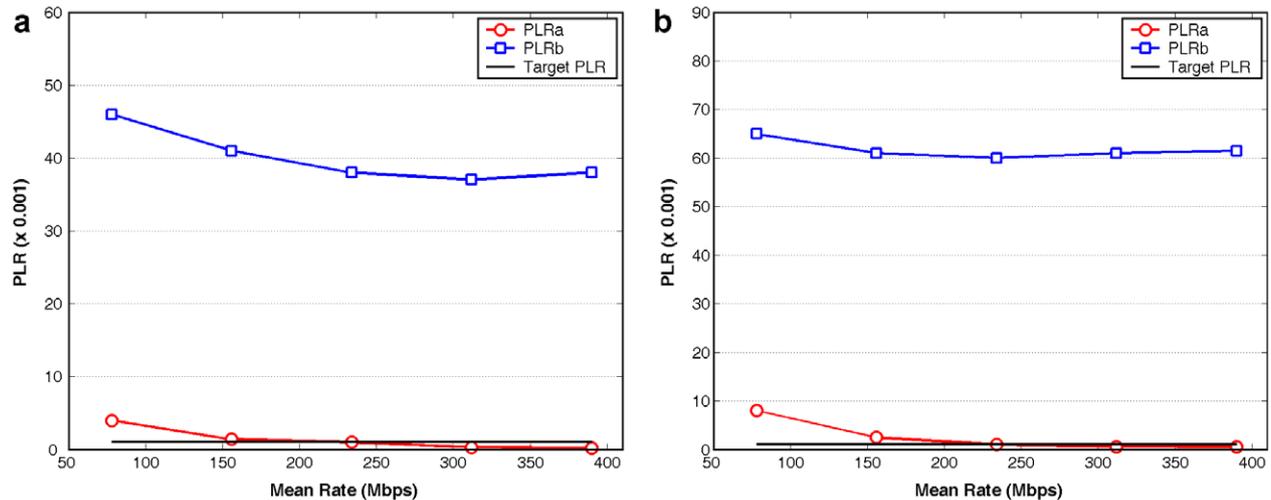


Fig. 7. Incurred PLR for *Videoconference* sources for target link PLR 0.001 and queue size 30 kbytes (a) and 50 kbytes (b).

5. Discussion

In this Section we will first elaborate on the implications of this first part of our work with respect to the aggregate traffic profile deformation and we will point out relevant issues that are raised and how they can affect the performance of bandwidth allocation and admission control schemes. Subsequently, we will present some possible practical traffic engineering solutions for dealing with these issues.

5.1. General implications

The first important implication is that the original traffic descriptors for a given set of real-time traffic sources multiplexed at a network edge, are not only invalid for downstream nodes but also the traffic profile deformation incurred while traversing multiplexers has, in general, a negative effect. That is, the traffic characteristics of a given

set of sources become, on average, worse in downstream nodes, which means that using the original traffic descriptors to depict the behavior of the sources can be an overly optimistic approximation.

The second important implication is that, for a given set of sources, the greatest multiplexing gains are achieved at the network edge, where the sources are uncorrelated. This is clear from the increasing packet loss rates incurred in core links compared to those in edge links for the same bandwidth allocation scheme and for the same number of multiplexed sources.

5.2. Implications for admission control

Regarding admission control, the results have certain implications on the effectiveness of an admission control scheme deployed in core nodes of a class-based IP network. If a TDAC scheme is deployed in a core node, it may fail if it is based on the original traffic descriptors of the traffic

sources. An MBAC scheme in the same case is less likely to fail compared to the TDAC scheme, because it relies on real-time measurements of the aggregate traffic and uses only the traffic descriptors of the source requesting admission. However, the original traffic descriptors declared by the source requesting admission may not depict its behavior in core nodes. It is worth noting that due to multiplexing and buffering, the packets belonging to a specific source may arrive at an interface at a rate exceeding the source's peak rate. That means that even the source's declared peak rate may not depict its worst-case behavior in core nodes [28]. This is also known as the *clumping* effect [13]. Therefore, even a conservative MBAC scheme, which makes the assumption that the source requesting admission will be transmitting at its peak rate, may fail when applied to core nodes.

A similar problem for TDAC and MBAC schemes arises when performing admission control for inter-domain traffic. In this case, if an upstream domain submits the original traffic descriptors of the sources to a downstream domain (without taking into account the traffic profile deformation for the sources within this upstream domain) and the downstream domain performs admission control based on those traffic descriptors, this may lead to QoS violations.

In order for any admission control scheme that uses traffic descriptors to be reliable, when used in nodes other than the first multiplexing point, it should appropriately modify the traffic descriptors to depict the behavior of the sources at that specific multiplexing point. However, this is not trivial, especially for sources whose traffic characteristics cannot be accurately represented by traffic models, since it requires the estimation of delay variation and induces per-flow state [27].

Even if the effects of the traffic profile deformation can be taken into account and appropriate signaling methods exist to learn the sources behavior at downstream nodes, if a TDAC or an MBAC scheme is applied for traffic aggregates on a link-by-link basis, unless the link packet loss rates are set so that the total end-to-end packet loss requirement of the flows traversing the larger number of links is satisfied, this will result in higher flow blocking probabilities for the flows traversing large number of links (long flows). This is because packet loss rate parameters are, for low values—as those required by real-time traffic—additive and this effect is similar to the discrimination against long flows in the case of EAC schemes [14]. However, if the link packet loss requirements are set so that the total packet loss requirement of long flows is satisfied, this can lead to underutilization of the reserved resources for the links carrying flows traversing a small number of links (short flows), because in these links the per-link packet loss rates will be set to lower values than what is actually needed and the lower the target packet loss, the lower the achieved utilization.

The above can be seen in Figs. 8 and 9, which show the utilization of the reserved resources for links 1 and 2 in our topology as a function of the mean rate of the sources mul-

tiplexed in links 1 and 2 for the bufferless statistical multiplexing approach and for target link 1 and 2 packet loss rates 0.01 (Util 0.01) and 0.001 (Util 0.001) (Fig. 8). They also show the utilization of the reserved resources for links 1 and 2 as a function of the mean rate of the sources multiplexed in links 1 and 2 for the buffered statistical multiplexing approach and for target link 1 and 2 packet loss rates 0.01 (Util 0.01) and 0.001 (Util 0.001) (Fig. 9).

From these figures it can be clearly seen that the lower the target packet loss rate at a specific link, the lower the achieved utilization. It is also worth noticing that for the buffered statistical multiplexing approach, the increase in utilization as a function of the aggregation level is much lower than the corresponding increase for the bufferless statistical multiplexing approach. This is, again, due to the additive nature of the effective bandwidth formula of Eq. (11).

5.3. Implications for bandwidth allocation

Regarding bandwidth allocation, the results suggest that if, during the network provisioning phase, the packet loss requirement is translated in a hop count constraint and the bandwidth allocation scheme in core nodes is based on an effective bandwidth formula, even if this is done based on *edge-to-edge isolation*, that is that traffic aggregates that are multiplexed in the same traffic class in the core network are allocated resources on a per ingress-egress pair basis, the consequences may be detrimental. Furthermore, similar to the admission control case, the link capacity dimensioning should be done for link packet loss rates able to satisfy the end-to-end packet loss requirement of the longer flows, which will lead to underutilization of resources on links carrying only short flows.

5.4. Possible practical traffic engineering solutions

Part of the aforementioned issues about bandwidth allocation and admission control schemes (e.g., unfairness against long flows, underutilization of resources for links carrying only short flows) can be overcome if more sub-classes are configured and engineered in order to support the real-time traffic flows. This way, real-time traffic flows (with the same end-to-end packet loss rate requirement) can be aggregated in different sub-classes so that in every link, only flows with similar target link packet loss rate requirements are aggregated. However, this would mean increasing the number of classes that must be engineered and supported in the routers. Apart from the obviously added complexity in relevant network dimensioning, increasing the number of classes the routers must support can lead to dramatically decreased forwarding performance, as demonstrated in [32].

An integrated approach for bandwidth allocation and admission control that can be used in order to overcome the aforementioned problems, including the traffic profile deformation, is to apply admission control only at the net-

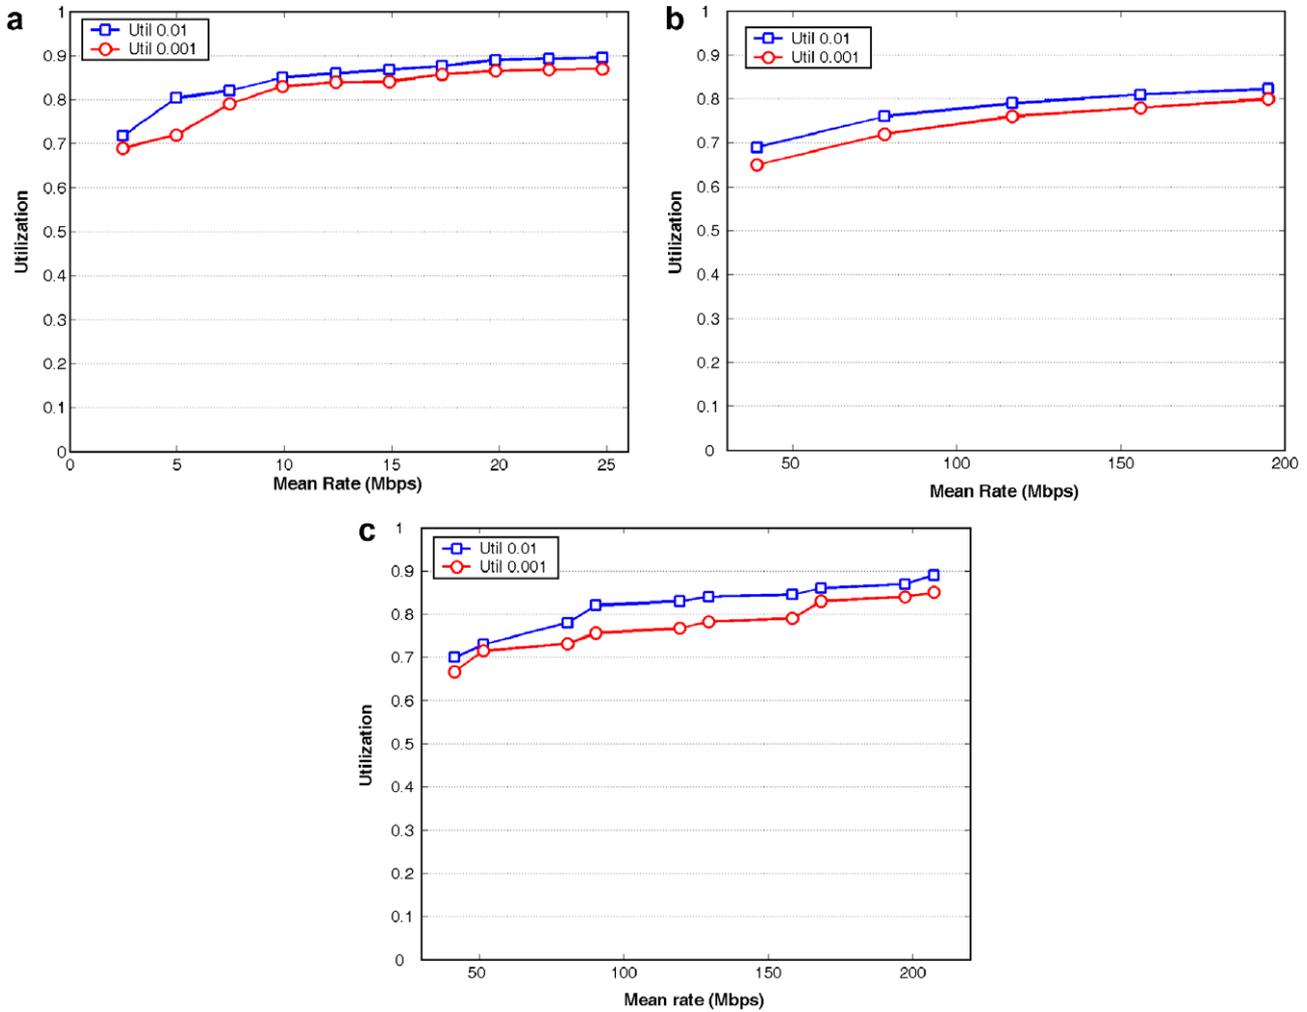


Fig. 8. Average utilization at links 1 and 2 for *VoIP* (a), *Videoconference* (b) and *mixed* (c) traffic sources for target link 1 and 2 PLR 0.01 and 0.001.

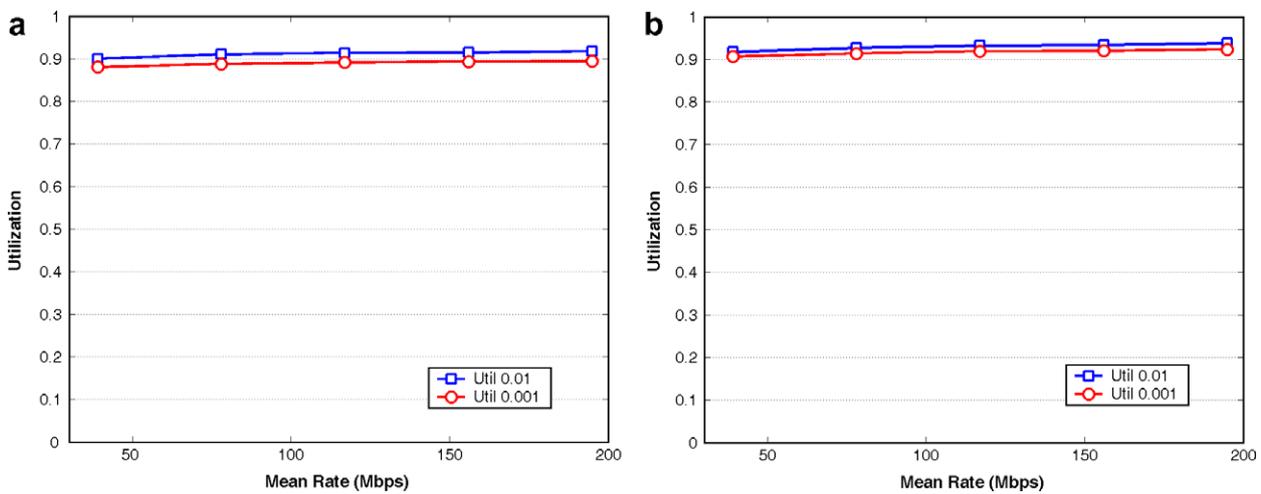


Fig. 9. Average utilization at links 1 and 2 for *Videoconference* sources for target link 1 and 2 PLR 0.01 and 0.001 and queue size 30 kbytes (a) and 50 kbytes (b).

work ingress and further downstream treat the real-time traffic aggregates in a peak rate manner. This is feasible since, as stated in [33], the edge links are currently consid-

ered as the most probable congestion points of a domain, whereas backbone links are overprovisioned [34]. This approach does not induce any additional state in the core

network and does not require core routers to be signaling-aware, which is desired for scalability and resilience reasons, and it is also proven to be a very resource-efficient approach if network dimensioning and provisioning for resilience against network failures is required [35]. Furthermore, as our results suggest, the greatest multiplexing gains are achieved at the network edge anyway, and by employing the aforementioned approach, since packet losses are not incurred by the core network but are restricted to those incurred by the edge links, the target edge link packet loss rates can be set higher, which means increased utilization of these links (e.g. see Figs. 8 and 9). This is important, since the edge links are the most common bottlenecks.

Our measurement-based admission control scheme, which we will present in the next Section, is based on this last integrated approach and makes use of the aforementioned assumptions, being therefore suitable for use in class-based networks where per-flow state in the core nodes cannot be kept. Since we assume that the interior of the class-based network domain has been provisioned and engineered in a peak rate manner in order to support the real-time traffic aggregates after the first points of aggregation, taking into account the routing behavior, at each ingress node we can have an estimate of the minimum bandwidth available for the real-time traffic aggregate from that ingress to each of the corresponding egress nodes. This available bandwidth is the *basis* for our admission control scheme, which is employed at the edge (ingress) nodes of the first aggregation points. It is worth noting that since in our approach, ingress nodes do not compete for resources at the network core, the scheme can be completely distributed, with an instance of it running as an independent module at each ingress node without, however, requiring any cooperation or coordination between the ingress nodes. In distributed admission control schemes that allow sharing of downstream resources between ingress nodes, special actions – such as cooperation between ingress nodes in order to update and ‘share’ among them the state of the shared resources – are needed in order to ensure that admission control decisions made at ingress nodes – without knowledge of the concurrent admission control decisions at other ingress nodes competing for the same downstream resources – do not lead to overacceptance of flows and, therefore, to QoS degradation. (For more details regarding actions that can be taken to handle concurrent admission control decisions made by distributed schemes, the interested reader can refer to [36]).

6. Admission control scheme

As stated in [16,37], in order for an admission control scheme to be successful in practice, it has to fulfill several requirements.

Robustness: An admission control scheme must ensure that the requested QoS is provided. This is not trivial, especially for MBAC schemes, since measurement inevitably has some uncertainty, potentially leading to admission

errors. The QoS should also be robust to traffic heterogeneity, time-scale fluctuations (long-range dependency), as well as to heavy offered loads.

Resource utilization: The secondary goal for admission control schemes is to maximize resource utilization, subject to the QoS constraints for the admitted flows.

Implementation: The cost of deploying an admission control scheme must be smaller than its benefits. In addition, the traffic characteristics required by the scheme should be easily obtained from the traffic sources and the network.

Taking these requirements into account, in this Section we will present our measurement-based admission control scheme, applicable to real-time sources that are able to provide *only* a single traffic descriptor, their peak rate. Given the diversity of Internet-based applications that have real-time requirements, the use of more complex traffic descriptors in admission control, as stated in [38], to accurately characterize source traffic, is neither necessary nor plausible. In addition, since source characteristics depend not only on the applications but also on their use, one cannot make *a priori* characterization of sources with certainty [39]. Therefore, we assume that the only available traffic descriptor to use is the source’s peak rate. This traffic descriptor is easy to police and, even if not available, for sources described by a token bucket filter (r, b) an estimate \hat{p} of it can be derived [38] using the equation:

$$\hat{p} = r + b/U \quad (12)$$

where U is a user-defined averaging period, which defines how conservative the peak rate derivation can be.

Taking into account the benefits of bufferless statistical multiplexing for real-time traffic, as these were described in Section 2, in our scheme we use the normal distribution based bufferless statistical multiplexing approach, which based on Eq. (1), can be rewritten as:

$$C \simeq m + a'\sigma \text{ with } a' = \sqrt{-2 \ln(\varepsilon) - \ln(2\pi)} \quad (13)$$

where m is the mean aggregate bit rate, σ is the standard deviation of the aggregate bit rate and ε is the upper bound on allowed loss probability. The mean aggregate bit rate and the standard deviation of the aggregate bit rate in this case, are derived based on real-time measurements and are not based on traffic descriptors, as in Section 4.

6.1. Algorithm for admission control

In a class-based domain we assume that the real-time traffic aggregate is provisioned and engineered in such a way that at least C_{total} bandwidth is available edge-to-edge. Every time a source wants to establish a service instance, it signals this to the ingress node through some ‘‘RSVP-like’’ signaling protocol [40–42]. A similar assumption can be made for the service termination. If the latter is not explicitly signaled, an alternative option could be to use a timeout period as an indication of the service termination. In any case, at each point in time, the MBAC process at each

ingress point knows the number of active sources at that ingress point.

When a new service request arrives, we need to decide whether or not to allow the source to send traffic using the real-time traffic aggregate resources until the known egress point. Initially, we need to calculate an appropriate time period, the *measurement window*, within which we need to take and use measurements for bandwidth usage estimations. The measured parameters are the mean rate of the offered load, $M_{measured}$, and the variance of the offered load, $\sigma_{measured}^2$, at the output queue of the ingress node. Having the measurements and the peak rate p_{new} of the new source, and by making the worst case assumption that the new source will be transmitting at its peak rate, we compute the estimated bandwidth C_{est} as follows:

$$C_{est} = M_{measured} + p_{new} + a'_{PLR} \sqrt{\sigma_{measured}^2} \quad (14)$$

where a'_{PLR} is computed as in Eq. (13), based on the target PLR bound of the real-time traffic aggregate. This value C_{est} is the estimated bandwidth used in the admission control criterion.

We need to mention here that, as also explained in Section 3, the use of the peak rate of the new source requesting admission in Eq. (14) does not mean that if the source is finally admitted, it will be allocated resources equal to its peak rate. For future admission control decisions, its real traffic contribution will be reflected in the aggregate traffic measurements.

6.2. Measurement window estimation

We define the measurement window w , as the time interval within which the offered load is taken into account for deriving the required measurements. In a similar fashion to [43], we use the following expression for the measurement window:

$$w = \max(DTS, w') \quad (15)$$

In Eq. (15), DTS represents the Dominant Time Scale. DTS is the most probable time scale over which overflow occurs. In [19], the authors describe a systematic way to derive DTS using real-time measurements, with the assumption that the input process to the multiplexing point in the network is Gaussian. This is by definition our assumption when employing (13), therefore we use this method in order to estimate the DTS. DTS, as computed in [19], is a function of the output buffer size. The reader should recall that even though we employ the bufferless multiplexing approach, a small output buffer is still required for packet scale queuing, as explained in previous Section. This value for the output buffer is involved in the estimation of the DTS.

Let w' represent the mean inter-departure delay [18], defined as follows (Little's formula):

$$w' = \frac{h_{avg}}{N_{active}} \quad (16)$$

where N_{active} is the number of simultaneously active sources and h_{avg} is their average duration.

Since we assume that the service establishment and termination is signaled to the ingress nodes, the average duration of the sources can be easily obtained and updated.

Eq. (15) means that we select as measurement window the mean inter-departure delay, i.e., the time interval within which the system can be considered stationary – no flow departures – so as to capture changes in the characteristics of the aggregate traffic stream caused by the past growth or decrease in the number of flows, unless this time interval is not long enough to capture the time-scale fluctuations of the aggregate traffic stream. This can happen in case of long-range dependent traffic. In this case and in order to enable the network to react to these long time-scale traffic fluctuations, we use DTS as the value of the measurement window.

6.3. The admission control criterion

Given that the allocated bandwidth for the real-time traffic aggregate from edge-to-edge is C_{total} , and having computed the estimated bandwidth C_{est} , the admission control criterion in our scheme becomes:

$$\begin{aligned} \text{If } (C_{est} \times APF) \leq C_{total}, & \text{ admit} \\ \text{If } (C_{est} \times APF) > C_{total}, & \text{ reject} \end{aligned} \quad (17)$$

where APF is an Admission Policy Factor we involve in the admission control criterion. The use of APF reflects the provider's policy on how strict the admission control should be. The decisions for setting the APF can be based on simple heuristics or ad hoc engineering methods. In the following Section we describe an example approach for setting APF , in which we take into account two issues that challenge the effectiveness of any MBAC scheme:

- (a) the traffic source heterogeneity, and
- (b) the effect of measurement errors.

6.4. A heuristic for setting admission control policies

The reason for introducing APF is to reflect the provider's policies. This means that appropriately tuning the APF can lead to a more conservative or a more relaxed admission control criterion. In our case we give a heuristic formula for APF with which we address two important issues that need to be taken into account in the admission control decision.

The first issue is that the aggregate traffic stream might have characteristics that do not suit the effective bandwidth formula, as this is expressed by Eq. (13). This, for instance, can happen if the stream is composed of a small number of very bursty connections with high peak rates and low utilizations [6].

To account for this, we use an exponential ON/OFF source, with mean and standard deviation (m_{ref}, σ_{ref}) as a model source for engineering reasons (*reference source*).

The reason for the specific selection is that exponential ON/OFF sources are representative models for VoIP traffic, which is likely to be a big part of the traffic carried by real-time traffic aggregates and their traffic characteristics suit the effective bandwidth formula of Eq. (13). Furthermore, exponential ON/OFF sources are short-range dependent, which means that their traffic characteristics are more easily captured and representative within the given measurement window. We define as *reference trunks* (T_{ref}) the number of simultaneously established reference sources that can fit in C_{total} , according to Eq. (13), for a given bound on packet loss rate.

When a new request arrives, having measured the mean rate $M_{measured}$ and the variance $\sigma_{measured}^2$ of the offered load, we calculate the number N_m of the reference sources, whose aggregate mean rate is equal to or greater than $M_{measured}$. We also calculate the number N_σ of the reference sources, whose aggregate variance is equal to or greater than $\sigma_{measured}^2$. That is, N_m and N_σ satisfy the following relationships:

$$N_m = \left\lceil \frac{M_{measured}}{m_{ref}} \right\rceil \text{ and } N_\sigma = \left\lceil \frac{\sigma_{measured}^2}{\sigma_{ref}^2} \right\rceil \quad (18)$$

Having estimated N_m and N_σ , we compute their mean value N_{ref} :

$$N_{ref} = (N_m + N_\sigma)/2 \quad (19)$$

This value represents a rough estimate of the number of reference sources that produce, within the measurement window, load with mean and variance similar to the ones measured. To compensate for the above, we set APF to be proportional to the quantity (N_{ref}/T_{ref}) .

The second issue that needs to be taken into account with the policy factor is the effect of measurement errors. As shown in [18], assuming that the measured parameters represent real traffic can heavily compromise the performance of an MBAC scheme. Measurement errors tend to compromise QoS, as there exists a fundamental asymmetry associated with the uncertainty of the measured parameters: the negative effect on QoS of an underestimation of parameters – and therefore of an overestimation of the number of permissible flows – far exceeds the positive effect on QoS of an overestimation of parameters [18]. Also, the stringent the PLR requirement, the easier it is to violate it due to measurement errors. In the case where only aggregate bandwidth information is available through measurements, as in our scheme, the degradation in performance can be mainly attributed to errors in the estimation of the variance [16]. With non-negligible probability the variance can be significantly underestimated. To compensate for the measurement uncertainty, we proceed as follows: given Eq. (13), for a specific target PLR, we set APF to

be proportional to the quantity $\frac{\sqrt{-2 \ln(PLR) - \ln(2\pi)}}{\sqrt{-2 \ln(PLR_{ref}) - \ln(2\pi)}}$.

That is, we inflate the part of Eq. (13) that relates to the variance estimation, based on a reference PLR level. By

setting PLR_{ref} to be larger than PLR , we ensure that the more stringent the PLR requirement, the greater the value of this quantity. This reference PLR can be set by policy to adjust the conservativeness of the MBAC scheme.

Combining the two aforementioned quantities, the final expression for the admission policy factor that can be used is:

$$APF = (N_{ref}/T_{ref}) * \frac{\sqrt{-2 \ln(PLR) - \ln(2\pi)}}{\sqrt{-2 \ln(PLR_{ref}) - \ln(2\pi)}} \quad (20)$$

We also set $APF = 1$ whenever the equation above results to APF values less than 1. That means that we use APF in a conservative way in the admission control criterion. The admission policy factor can be considered as a tuning parameter. Even though we derive APF somehow heuristically, based on intuition rather than strict mathematical analysis, one should take into account that all MBACs employ additional admission policy tuning parameters [16,44] because it is not possible to completely decouple performance from traffic characteristics. In addition, difference in performance caused by flow heterogeneity is a matter to be addressed by policy, rather than through algorithmic modifications [44].

7. Performance evaluation

In order to evaluate the performance of our admission control scheme, we run simulations using the network simulator *ns-2*, with the typical, for admission control evaluation purposes, dumbbell topology of Fig. 10.

We assume that the sources connect to the ingress node through links with negligible congestion. Even if the sources originate from local area networks (LANs), this assumption can still be considered valid, since LANs have large enough bandwidth to handle a number of sources, larger than that of wide area networks in most real network situations [45].

We set the reference PLR equal to 0.01 and we fix the output queue for the real-time traffic class to 500 bytes and to be served in a non-work conserving scheduling manner. As a *reference source*, we use an exponential ON/OFF source with a peak rate of 64 kbps and mean durations for the ON and OFF periods 1.004 and 1.587 s, respectively [25]. We use scenarios with the target bound on packet loss rate for the aggregate real-time traffic equal to 0.01 and 0.001. These bounds represent typically acceptable PLR values for the VoIP service and for real-time applications

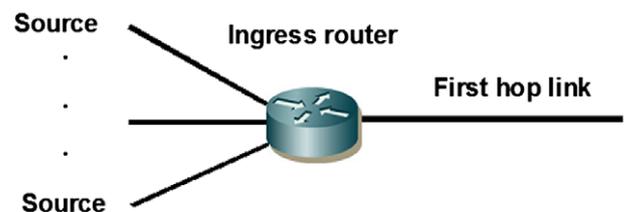


Fig. 10. Simulation topology.

in general, according to [46,47]. We need to stress here that if the real-time applications require different PLR (or QoS in general) targets to be met, then multiplexing them in a single class would require for the most stringent QoS requirement among them to be met. This would lead to severe underutilization of resources. In such a case, applications with different QoS requirements should be multiplexed in different classes based on the value of the requested QoS. For our simulations we assume that the real-time applications have the same PLR requirement and can, therefore, be multiplexed in a single class.

We set the output link capacity allocated to real-time traffic to correspond to T_{ref} equal to 100. This means that the output link capacity is set equal to 3.33 Mbps for the target PLR 0.01 case and equal to 3.56 Mbps for the target PLR 0.001 case. In a real network situation, unused capacity of the real-time traffic class would be fully available to a lower priority, e.g. best-effort, traffic, so there would be no waste incurred by this partitioning. In a similar fashion to the simulations in Section 4, we do not consider any best-effort traffic or the existence of any other traffic classes in our simulations and all the results are based on averages of simulations for 20 randomly chosen seeds, each for a total of 4100 s, using the first 500 s as a warming-up period.

In order to test the *robustness* of the scheme with respect to traffic *heterogeneity* and *long-range dependency*, we use both VoIP and Videoconference traffic sources. For VoIP traffic we use an ON/OFF source model with exponentially distributed ON and OFF times, having a peak rate of 64 kbps. The mean durations for the ON and OFF periods are 0.350 and 0.650 s, respectively [48], meaning that its average rate is 22.4 kbps. For Videoconference traffic we use an H.263 coded trace from [49] with average rate 64 kbps and peak rate 332.8 kbps. The H.263 format has been widely employed to model Videoconference traffic, e.g., see [50,51]. The active time of both VoIP and Videoconference sources is exponentially distributed with an average of 300 s.

In order to test the *robustness* of the scheme with respect to *offered load*, we test varying load conditions ranging from 0.5 to 5, where the value 1 (*reference load*) corresponds to the average load that would be incurred by a source activation rate equal to 1000 VoIP sources/hour. Given the average rates and the durations of the VoIP and Videoconference traffic sources, this value 1 also corresponds to the load incurred by a source activation rate equal to approximately 350 Videoconference sources/hour.

In order to compare the performance of our scheme, which we call MBAC-GEO, against other existing proposals, we implement three other algorithms. The first algorithm is an MBAC scheme described by Zukerman et al in [52] as Rate Envelope Multiplexing (REM), with adaptive weight factor and no histogram update (the packet loss measurement approach). The reasons for the selection of the specific MBAC scheme (we call it MBAC-ZUK)

for comparison with our scheme are that: (a) REM also makes the zero buffer approximation with respect to statistical multiplexing and (b) implementation-wise, in a similar fashion to our scheme, it requires only aggregate bandwidth measurements and the peak rate of the sources requesting admission in order to derive the admission control decision. The time scales involved in this scheme are 0.05, 0.1, 0.5, 1 and 2 s. We also adopt the value 0.05 s as the minimum sampling interval in our approach in order to derive measurements for the mean rate, the variance of the offered load and to additionally estimate the DTS.

The second algorithm is an EAC scheme described by Karlsson et al in [53]. In order to test this scheme (we call it EAC-KAR), since it is an out-of-band probing scheme, we implement, using the built-in class-based queuing mechanisms of *ns-2*, an additional lower priority queue for the probing packets that can store, as in [53], a single probe packet and which is only served when the higher priority real-time traffic queue is empty. As in [53], we set the probing rate equal to the peak rate of the source requesting admission, we consider probe durations of 0.5 s up to 5 s, and we also assume that there is no latency involved between the probing phase completion and the admission control decision.

The third algorithm is a simple TDAC Peak Rate Allocation scheme (we call it TDAC-PRA) that only admits a new source if the following condition is satisfied:

$$\sum p_i + p_{new} \leq C_{total} \quad (21)$$

where $\sum p_i$ is the sum of peak rates of the already established sources. With this scheme, there are no losses, since it does not account for any statistical multiplexing.

As stated in [54], any admission control scheme must address the trade-off between *packet loss*, which is the mostly used QoS index [37], and *utilization*. Therefore for performance evaluation we use these two metrics, together with the average blocking rate.

In our simulations we consider two cases for the mixture of traffic sources that request admission: (a) Videoconference sources only and (b) VoIP sources only.

For TDAC-PRA we do not show any PLR results because the PLR is constantly zero. For MBAC-ZUK, the results shown are for the algorithm parameters set equal to the values used in [52] (for a detailed description of MBAC-ZUK, the interested reader can refer to [52]). For EAC-KAR, the results shown are for probe duration of 3 s for the VoIP sources only case and for probe duration of 2 s for the Videoconference sources only case. These probe durations give the best trade-off between packet loss and utilization for the examined cases. For shorter probing durations we observe violation of the target PLR, whereas for longer probing durations we observe significant *thrashing* effects. (A reduced initial set of the results presented in this section can be found in [55]).

7.1. Videoconference sources

The performance results for Videoconference traffic sources are shown in Figs. 11 and 12.

For videoconference traffic, all schemes satisfy the target PLR for all loading conditions. For target PLR 0.01, all schemes are unnecessarily conservative, which can be partly attributed to the stringent admission control criterion (all schemes make the worst case assumption that the new source will be transmitting at its peak rate) and the high peak rate of the videoconference sources compared to their average rate. Regarding utilization and blocking, the performance of MBAC-GEO is, on average, better than that of EAC-KAR and slightly better than that of MBAC-ZUK.

The reader should recall at this point that the objective is not to achieve the lowest PLR possible, but to keep the incurred PLR within the limits of the target PLR, while maximizing utilization and minimizing blocking. For example, for TDAC-PRA the incurred PLR is zero, but the blocking is significantly higher and the utilization

achieved is also significantly lower than any of the other three algorithms, and because of the high peak rate of the H.263 videoconference sources compared to their average rate, it does not exceed 15%.

7.2. VoIP sources

The performance results for VoIP traffic sources are shown in Figs. 13 and 14.

For VoIP traffic, all schemes satisfy the target PLR 0.01. MBAC-ZUK is more conservative than MBAC-GEO and EAC-KAR, achieving therefore lower utilization and incurring higher blocking. For EAC-KAR we can see that for loading conditions more than 4 times the *reference load* we enter the *thrashing* region, which for an out-of-band probing EAC algorithm is indicated by a decrease in the incurred PLR, accompanied by a decrease in the achieved utilization and an increase in the blocking rate. For target PLR 0.001, MBAC-GEO and EAC-KAR satisfy this PLR for all loading conditions with EAC-KAR being less conservative, achieving, therefore, slightly higher utilization

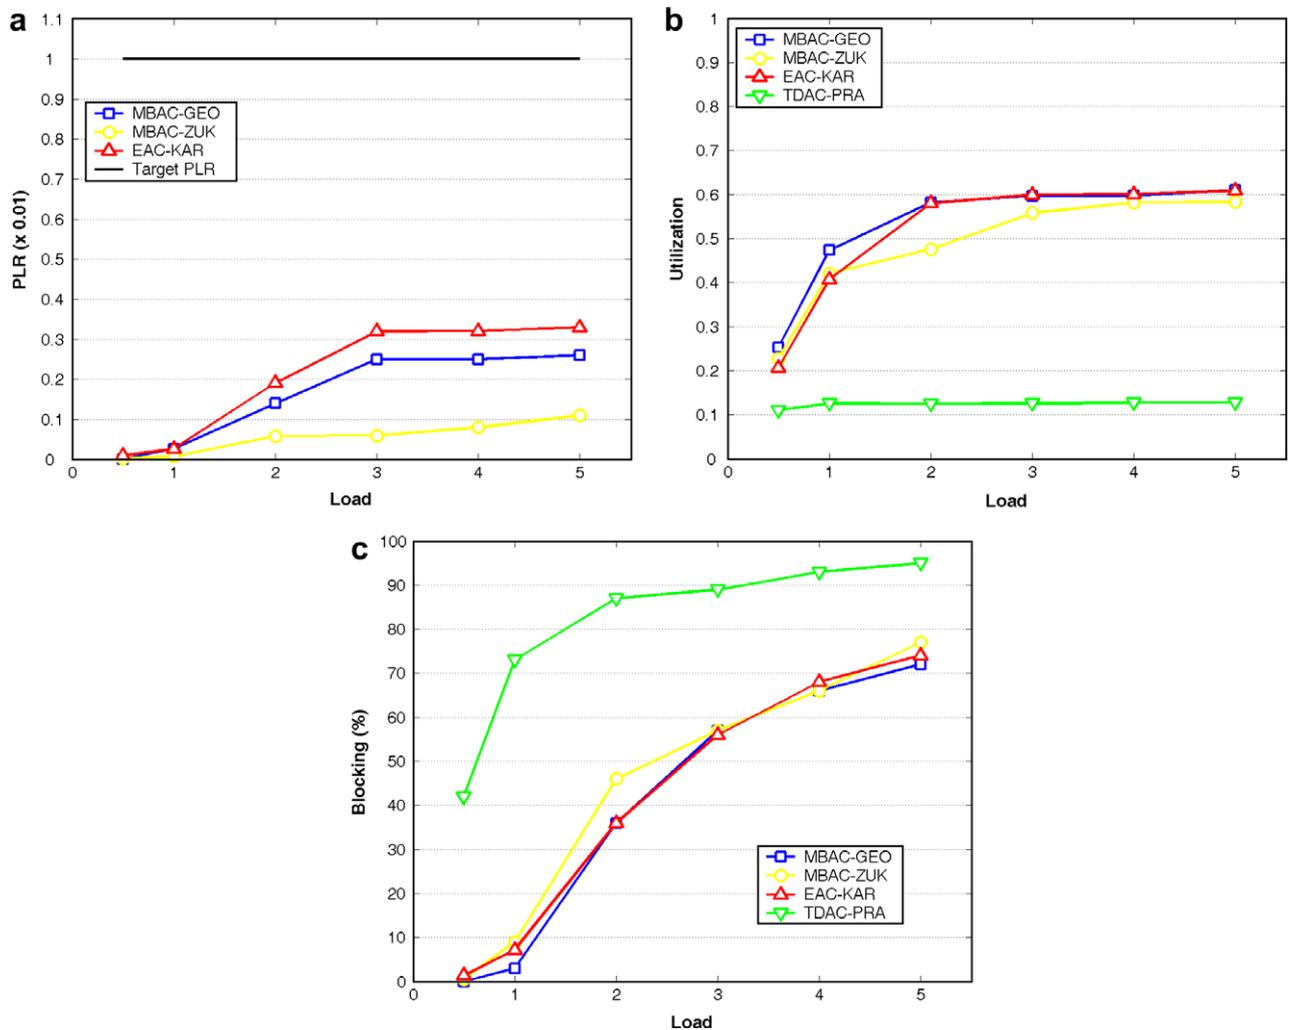


Fig. 11. Incurred PLR (a), utilization (b) and blocking rate (c) for target PLR 0.01.

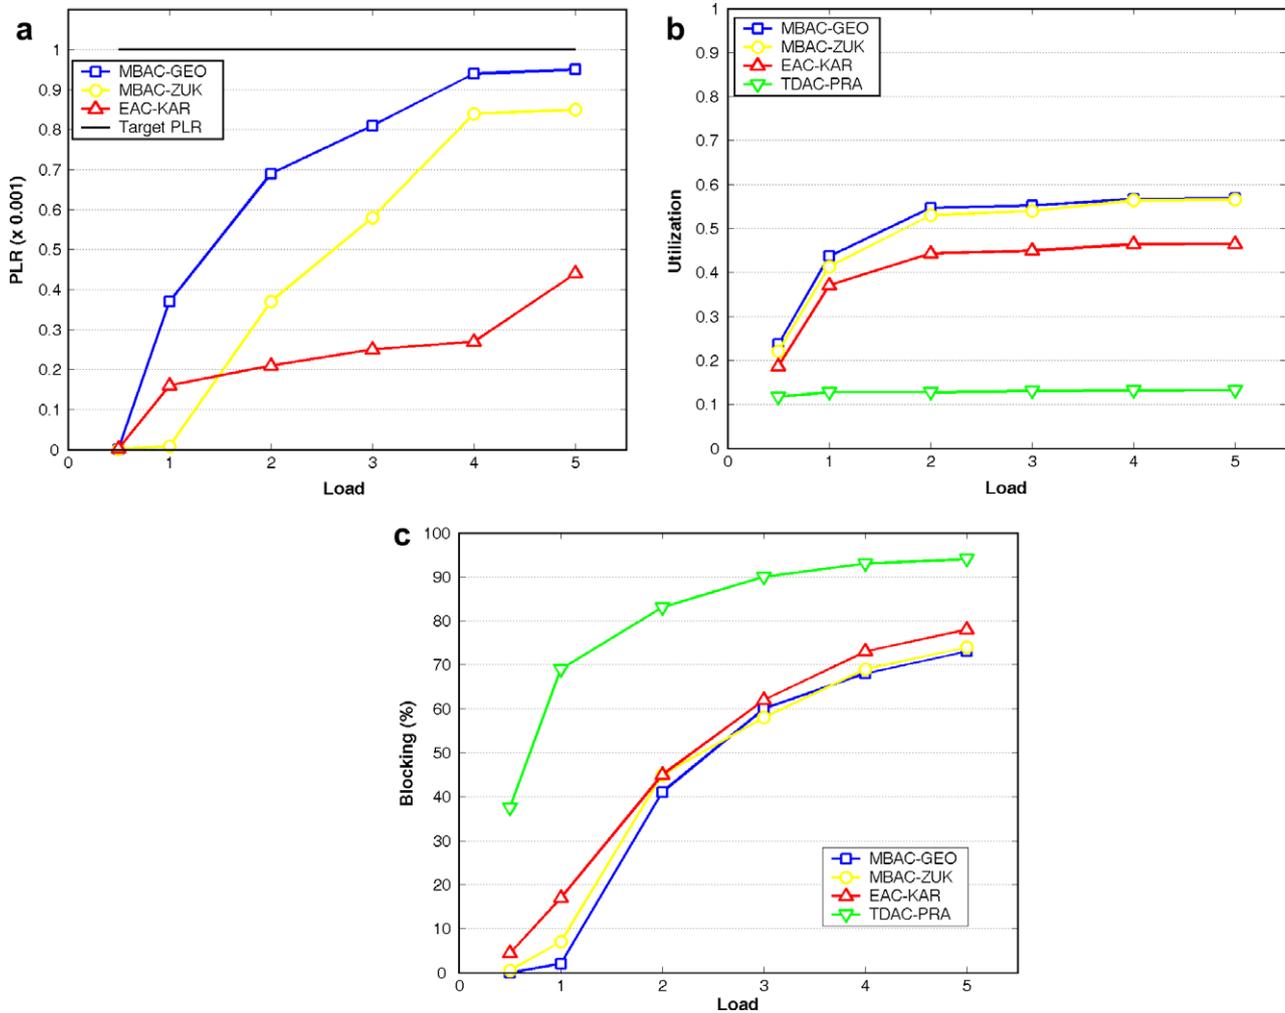


Fig. 12. Incurred PLR (a), utilization (b) and blocking rate (c) for target PLR 0.001.

and incurring slightly lower blocking. EAC-KAR in this case enters the *thrashing* region for loading conditions more than 3 times the *reference load*. MBAC-ZUK violates this PLR for loading conditions more than 3 times the *reference load*, even though the no histogram update method used in our implementation of MBAC-ZUK is the most conservative approach among all the other variations presented in [52]. This means that the tuning parameters involved in MBAC-ZUK should be reconfigured in a trial and error fashion and set to new values -different from the ones used in the Videoconference sources only case and the target PLR 0.01 VoIP sources only case – in order to satisfy the target PLR 0.001 for all loading conditions for VoIP sources only. We need to stress here that EAC-KAR also employs a tuning parameter, which is the probe duration, which we have to vary from 0.5 s up to 5 s in order to find its optimal value (3 s) for the VoIP sources simulated loading conditions. This value is different from the value used for the Videoconference sources only case (2 s). For VoIP sources only and for probe duration 2 s, we observed target PLR violations even higher than these observed for MBAC-ZUK.

TDAC-PRA incurs lower blocking and achieves much higher utilization compared to the previous case, where we have Videoconference sources only, because of the lower peak rate of the VoIP sources compared to their average rate, but still significantly lower than any of the other three schemes.

7.3. Discussion of simulation results

The simulation results show that MBAC-GEO can satisfy the target PLR in all cases (which is the primary objective) without requiring any further reconfiguration of its parameters for individual traffic scenarios and loading conditions and despite the effects of measurement errors. As we show, it is possible without any reconfiguration of its tuning parameters to achieve reasonably good performance regarding utilization and blocking (which is the secondary objective) for a variety of traffic scenarios and loading conditions.

MBAC-ZUK satisfies the target PLR for the target PLR 0.01 case without requiring reconfiguration of its tuning parameters, it fails, however, to satisfy the tar-

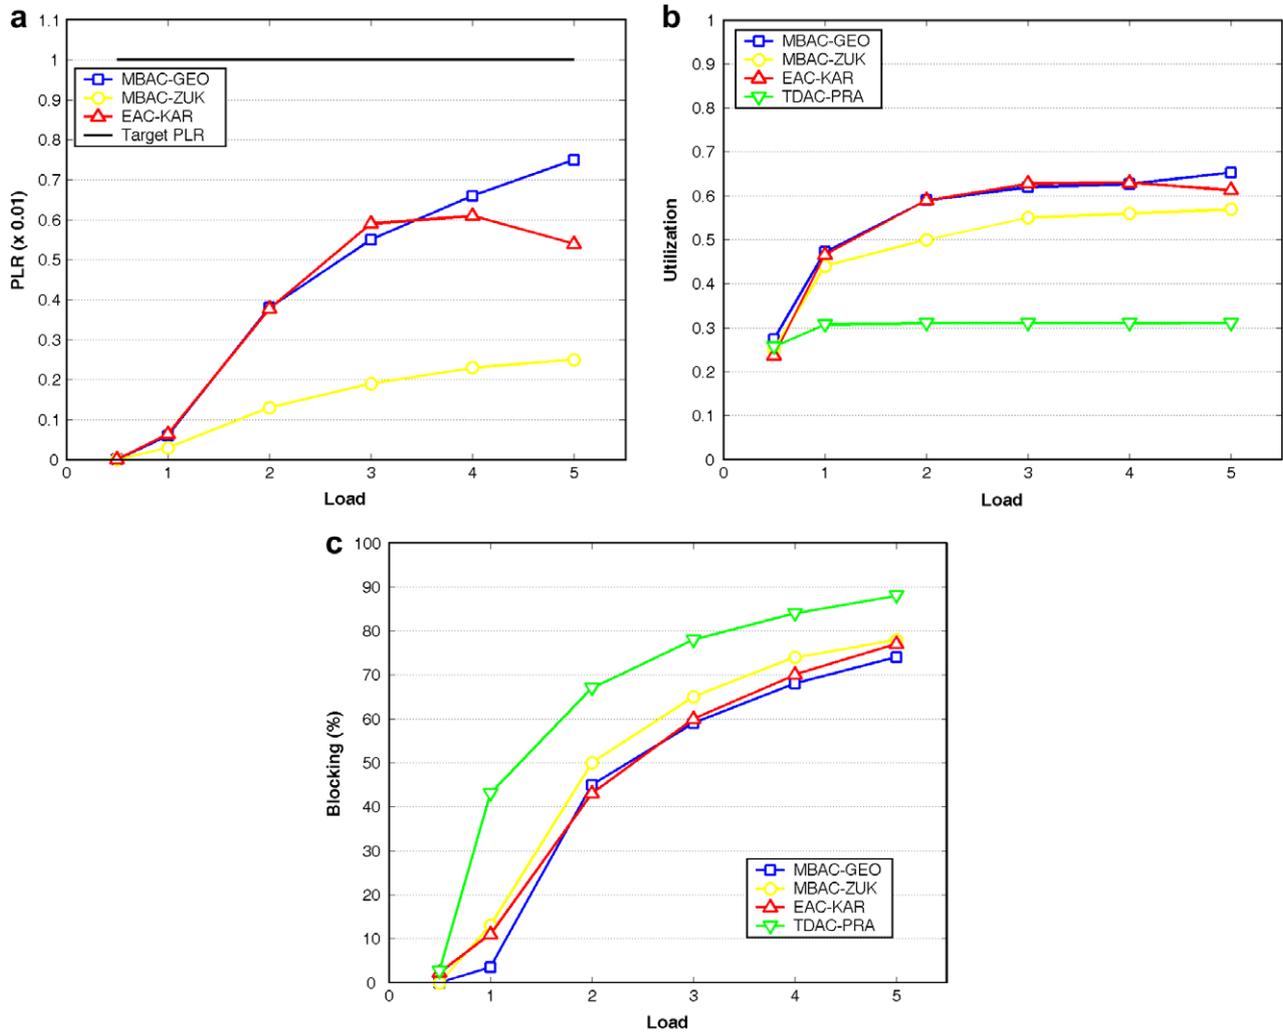


Fig. 13. Incurred PLR (a), utilization (b) and blocking rate (c) for target PLR 0.01.

get PLR 0.001 value for VoIP sources traffic scenarios. For these scenarios new values for its tuning parameters should be sought in a trial-and-error fashion.

EAC-KAR satisfies the target PLR for all simulated cases by appropriately reconfiguring its tuning parameter, following the guidelines in [53]. However, for the VoIP sources only case for both target PLR values and for high loading conditions, this is achieved while entering the thrashing region, which is a region an EAC scheme should avoid entering, which means that the target PLR is not actually satisfied due to the ability of EAC-KAR to adapt to the traffic conditions.

In Fig. 15 the averages of utilization and blocking rate for all simulated cases and loading conditions for all four schemes, including the scenarios and loading conditions where MBAC-ZUK violates the target PLR, are shown.

As it can be seen, MBAC-GEO achieves, on average, the highest utilization and incurs the lowest blocking rate among all schemes, while satisfying the target PLR in all

cases and without requiring any reconfiguration of its parameters. It is followed by MBAC-ZUK, which, however, violates the target PLR in certain cases. EAC-KAR achieves the third best performance regarding blocking and utilization (which can be attributed to the fact that it ‘needs’ to enter the thrashing region in order to satisfy the target PLR), while satisfying the target PLR, but, as shown, this requires the adjustment of its tuning parameter (probe duration). As expected, TDAC-PRA is the worst performer with respect to utilization and blocking.

In all cases examined, for both MBAC-GEO and MBAC-ZUK we observe an increase in the incurred PLR for higher loading conditions. This is anticipated [18] because they both rely on measurements, so every new admission request has the potential of being a wrong decision. This means that a high source activation rate (which directly translates to high loading conditions) is expected to have a negative effect on performance. The same holds for EAC-KAR but only for the cases where we have not entered the *thrashing* region.

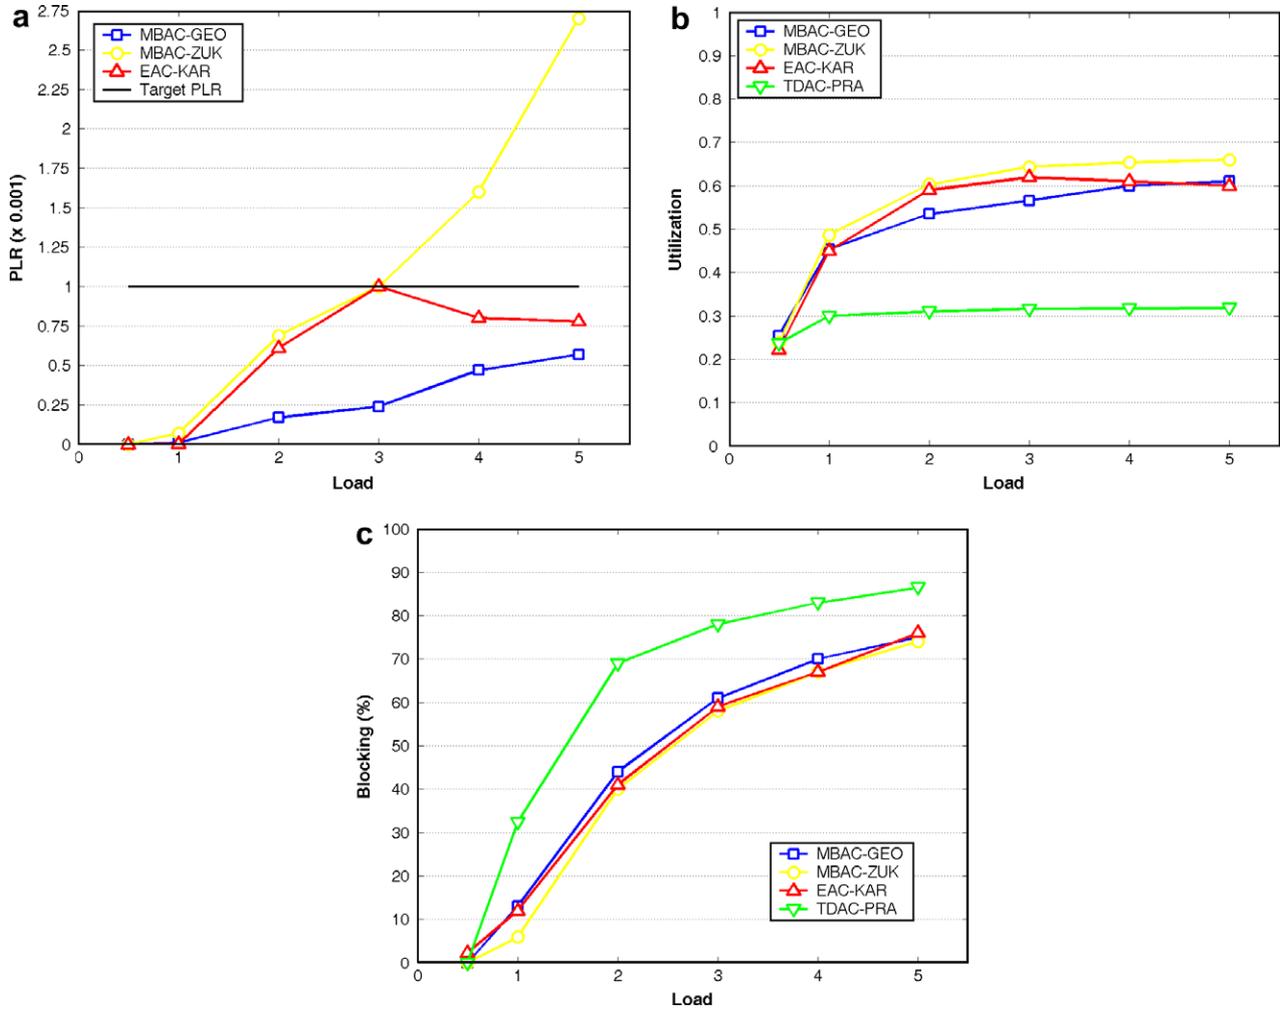


Fig. 14. Incurred PLR (a), utilization (b) and blocking rate (c) for target PLR 0.001.

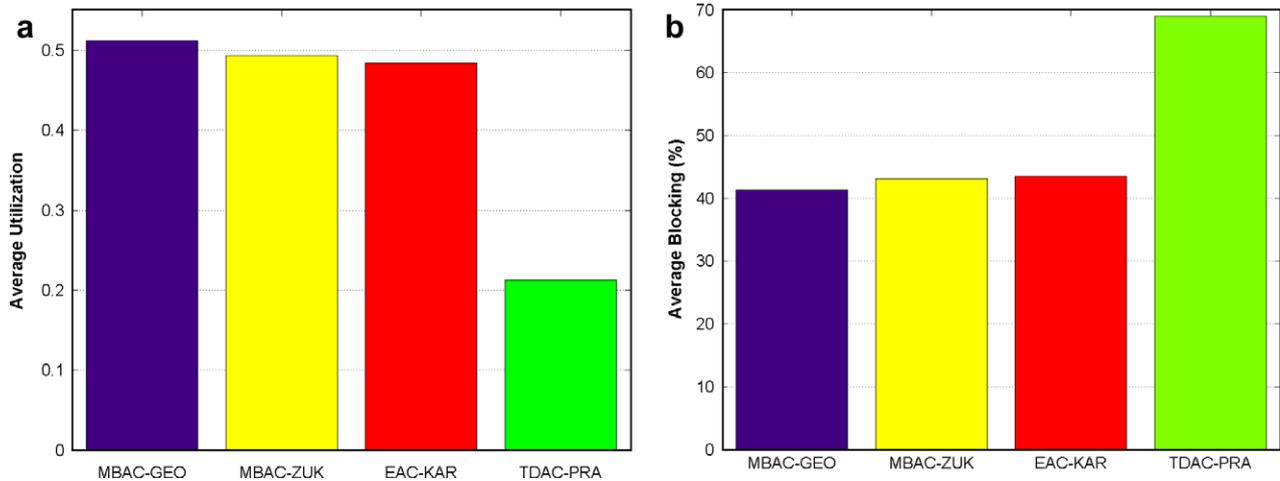


Fig. 15. Average utilization (a) and blocking rate (b).

It is worth mentioning that we also considered the scenario where mixed VoIP and Videoconference sources request admission, however the simulation results for this scenario are very similar to the VoIP sources only and Videoconference sources only scenarios and, therefore, are not shown here.

8. Limitations and implementation issues of the admission control scheme

In this Section we will discuss the limitations and the potential implementation and deployment scenarios of our admission control scheme, which is based on our inte-

grated approach for bandwidth allocation and admission control.

8.1. Limitations

The limitations of any admission control scheme are a consequence of the fundamental assumptions it makes and on which it builds its logic. These assumptions inevitably limit the applicability and validity of any admission control scheme in cases where they are not met. In that respect, our admission control scheme makes two key assumptions.

The first one, which relates to the network configuration, is that either the network core is overprovisioned, or the real-time traffic class is very carefully pre-provisioned. Either of these assumptions affects our admission control scheme in the same way, i.e. the only place where we need to apply admission control is at the network domain ingress nodes since treating real-time traffic aggregates further downstream in a peak rate manner is feasible. This approach also implicitly guarantees that no further QoS degradation (incurred PLR to be more precise) will take place further downstream from the first multiplexing point (ingress node). As already mentioned, the ‘overprovisioned core’ assumption is not uncommon in the literature, e.g. see [33,34]. Also, careful and strategic pre-provisioning and engineering for real-time traffic with Diffserv has been an area that has attracted a lot of research interest in the recent past, so such practices can be expected in the future. However, in cases where a network domain core is neither overprovisioned nor pre-provisioned for real-time traffic, the admission control scheme would not be valid as it is. It would need to be extended to take into account monitored information from all the potential bottlenecks within a domain and also concurrency-related issues, i.e. competition between admission control decision points for the same downstream resources, as well as account for the traffic profile deformation of the flows requesting admission at each admission controlled multiplexing point. We believe that this would be a very challenging task, and at a first glance it would be more cost effective to opt for a simple admission control scheme, like ours, and spend some effort and resources to overprovision or properly engineer the domains.

The second assumption, relates to the modeling of the traffic aggregate at the ingress nodes using the Gaussian assumption. This assumption is required in order to be able to use Eq. (13) to derive the effective bandwidth for the real-time traffic aggregates at the ingress nodes. The Gaussian assumption also needs to hold for the method introduced in [19] to derive the Dominant Time Scale value (DTS) which is used in our approach to derive the expression for the measurement window. As already mentioned, the Gaussian assumption is common in the literature, e.g., see [6,8,19,20] and has been proven to be a good approximation of real traffic. But it does not hold when the number of multiplexed sources is very

small, especially when these sources are very bursty with high peak rates and low utilizations [6]. In our approach we try to partly account for this limitation by introducing the Admission Policy Factor (*APF*), however if the number of multiplexed sources is extremely small then a different traffic modeling approach should be sought as well as a different method to estimate the expression for DTS.

Another limitation that needs to be mentioned is that our admission control scheme employs a tuning parameter, the Admission Policy Factor (*APF*), which is derived based on heuristics rather than mathematical analysis. That means that in cases where the *APF* is not appropriately tuned, this may lead to either a too conservative admission control criterion or to a rather loose one. In future work we will attempt a more analytical approach for setting the *APF*, so as to make our admission control scheme less reliant on heuristics.

8.2. Implementation issues

In Section 5.4 we described our admission control scheme as a distributed system running as an independent module at each of the ingress nodes in a domain. Since there is no competition between ingress nodes for resources at the network core, either because of overprovisioning or proper engineering, no cooperation or coordination among ingress nodes is required, therefore the only signaling necessary is from the end-users to the ingress nodes they are connected to. This signaling can be performed through an “RSVP-like” signaling protocol [40–42]. As in [40], one could, for example, consider the use of RSVP in a ‘lightweight’ mode where the RSVP PATH message from the end-user conveys the service request to the admission control module at the ingress node and the RSVP RESV message from the admission control module at the ingress node conveys the admission control decision for the service request back to the end-users.

In addition to supporting real-time service request/response operations, the admission control module at each of the ingress routers must have monitoring information regarding the aggregate rate of real-time traffic at each of the output interfaces of the ingress nodes. If the admission control module at a certain ingress router is implemented as an internal part of the routers’ software suite, then the monitoring information about the state of the output queues can be easily collected through monitoring software operating within the ingress node. If the admission control module is operating on a separate control station attached to the ingress node, then in order to obtain the required monitoring information the ingress router should support remote monitoring capabilities. Most of the modern routers do have such functionality. For example they can expose monitoring information through the Simple Network Management Protocol (SNMP) [56] or through other means, such as IP Flow Information Export (IPFIX) [57] or NetFlow records [58].

The deployment decision of having the admission control module running onto a control station directly attached to the ingress router rather than being integrated into the router's software, offers flexibility, especially when the router's software does not support the required admission control functionality or is not open for customization. On the other hand, though, it adds a small delay to the responsiveness of the admission control scheme.

9. Conclusions and future work

In this paper we first highlight several issues with respect to bandwidth allocation and admission control for the support of real-time traffic in class-based IP networks. We discuss the implications of topological placement, that is the location of the employed bandwidth allocation and admission control schemes with respect to the end-users requesting the services and the various network boundaries (access, metro, core, etc.) and we show that the performance of bandwidth allocation and admission control schemes can be heavily dependent on it. This means that their performance at different points in a class-based IP network and for the same traffic load conditions can be quite different and can deviate significantly from the expected targets. Through simulation we also try to provide a quantitative view of the aforementioned issues for a variety of heterogeneous traffic scenarios.

Taking the implications of topological placement into account, we then propose a measurement-based admission control scheme for heterogeneous real-time traffic. We assume that an instance of our MBAC scheme runs at every ingress node in a distributed manner. We demonstrate through simulation that the scheme is *robust* to traffic heterogeneity, time-scale fluctuations and heavy offered loads. The scheme can satisfy the packet loss rate requirement in all cases despite the effect of measurement errors and without requiring any reconfiguration of its parameters. Furthermore, the scheme achieves satisfactory *utilization* and compares well against existing admission control approaches for the same simulation setup. Finally, we have to mention that our scheme is also easy to *implement*. It only relies on aggregate bandwidth information and does not require any per-flow information state. The required traffic characteristics are the peak rate of the traffic source requesting admission and the mean rate and the variance of the aggregate real-time traffic load at the output queues of the ingress node where the MBAC instance runs. These can be easily obtained. In addition, the scheme does not require any cooperation between the ingress nodes and requires the use of signaling only from the sources to the ingress nodes, but not further downstream within the domain, since it is based on provisioned information.

In our current research efforts, we are focusing on extending our MBAC scheme for inter-domain traffic; that is traffic that crosses more than one domains. For

the inter-domain traffic case, since peering links at the border routers between neighboring domains are often bottlenecks [59–61] and also the source of some of the greatest costs for network operators [61], they cannot be widely overprovisioned. Therefore, MBAC should also take into account the state of these links before deriving the admission control decision. We are investigating what type of information needs to be measured in this case to depict the state of the peering links and how this information can be incorporated in our MBAC scheme, taking again into account the implications of topological placement.

Acknowledgement

This work was undertaken in the context of the IST AGAVE and IST ENTHRONE projects.

References

- [1] W. Sun, P. Bhaniramka, R. Jain, Quality of Service using Traffic Engineering over MPLS: An Analysis, IEEE LCN, 2000.
- [2] G. Schollmeier, C. Winkler, Providing Sustainable QoS in Next-Generation Networks, IEEE Commun. Mag., June 2004.
- [3] J. Lakkakorpi, O. Strandberg and J. Salonen, Adaptive Connection Admission control for Differentiated Services Access Networks, IEEE J. Sel. Area Commun., October 2005.
- [4] T. Bonald, A. Proutiere, J. Roberts, Statistical Performance Guarantees for Streaming Flows using Expedited Forwarding, IEEE INFOCOM, 2001.
- [5] M. Mowbray, G. Karlsson, T. Kohler, Capacity Reservation for Multimedia Traffics, Distr. Syst. Eng., 1998.
- [6] R. Guerin, H. Ahmadi, and M. Naghshieh, Equivalent Capacity and its Application to Bandwidth Allocation in High-Speed Networks, IEEE J. Sel. Area Commun., September 1991.
- [7] P. Joos, W. Verbiest, A statistical bandwidth allocation and usage monitoring algorithm for ATM networks, IEEE ICC, 1989.
- [8] C. Courcoubetis, G. Fouskas, R. Weber, On the performance of an effective bandwidths formula, International Teletraffic Congress, 1994.
- [9] A. Elwalid, D. Mitra and R. Wentworth, A new approach for allocating buffers and bandwidth to heterogeneous regulated traffic in an ATM node, IEEE J. Sel. Area Commun., August 1995.
- [10] D. Tse, M. Grossglauser, Measurement-based call admission control: analysis and simulation, IEEE INFOCOM, 1997.
- [11] T. Bonald, S. Oueslati-Boulahia, J. Roberts, IP traffic and QoS control: the need for a flow-aware architecture, in: World Telecommunications Congress, September 2002.
- [12] E. Knightly, N. Shroff, Admission control for statistical QoS: theory and practice, IEEE Network, March 1999.
- [13] J. Roberts, U. Mocci and J. Virtamo, Broadband network teletraffic, Final report of action COST 242, Springer 1996.
- [14] L. Breslau, E. Knightly, S. Shenker, I. Stoica, Z. Zhang, Endpoint admission control: architectural issues and performance, SIGCOMM, 2000.
- [15] K. Mase, Toward scalable admission control for VoIP networks, IEEE Commun. Mag., July 2004.
- [16] M. Grossglauser, D. Tse, A time-scale decomposition approach to measurement-based admission control, IEEE/ACM Trans. Networking, August 2003.
- [17] E. Knightly, H. Zhang, Traffic characterization and switch utilization using a deterministic bounding interval dependent traffic model, IEEE INFOCOM, 2005.

- [18] M. Grossglauser, D. Tse, A framework for robust measurement-based admission control, *IEEE/ACM Trans. Networking*, June 1999.
- [19] D. Eun, N. Shroff, A measurement-analytic approach for QoS estimation in a network based on the dominant time scale, *IEEE/ACM Trans. Networking*, April 2003.
- [20] R. Guerin, L. Gun, A unified approach to bandwidth allocation and access control in fast packet-switched networks, *IEEE INFOCOM*, 1992.
- [21] S. Georgoulas, P. Trimintzios, G. Pavlou, K. Ho, On the location-awareness of bandwidth allocation and admission control for the support of real-time traffic in class-based IP networks, *IEEE/IFIP MMNS*, 2006.
- [22] S. Lima, P. Carvalho and V. Freitas, Distributed admission control for QoS and SLS management, *J. Network Syst. Manage*, September 2004.
- [23] K. Fall, K. Varadhan, The *ns* manual, (http://www.isi.edu/nsnam/ns/ns_doc.pdf).
- [24] I. Mas, V. Fodor, G. Karlsson, The performance of endpoint admission control based on packet loss, *QoFIS*, 2003.
- [25] C. Chuah, L. Subramanian, R. Katz, *Furies*: a scalable framework for traffic policing and admission control, UC Berkeley Technical Report, May 2001.
- [26] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, J. Robbins, Performance models of statistical multiplexing in packet video communications, *IEEE Trans. Commun.*, July 1988.
- [27] K. Shiomoto, N. Yamanaka, T. Takahashi, Overview of measurement-based connection admission control methods in ATM networks, *IEEE Commun. Surveys*, 1999.
- [28] H. Perros, K. Elsayed, Call admission control schemes: a review, *IEEE Commun. Mag.*, November 1996.
- [29] H. Zhang, D. Ferrari, Rate-controlled static-priority queuing, *IEEE INFOCOM*, 1993.
- [30] M. Dabrowski, A. Beben, W. Burakowski, On inter-domain admission control supported by measurements in multi-domain IP QoS networks, *IEEE IPS*, 2004.
- [31] H. Fu, E. Knightly, Aggregation and scalable QoS: a performance study, *IEEE IWQoS*, 2001.
- [32] M. Torneus, Testbed for measurement based traffic control, Master's Thesis, KTH IMIT, Sweden, June 2000.
- [33] V.N. Padmanabhan, L. Qiu, H.J. Wang, Server-based inference of Internet Link Lossiness, *IEEE INFOCOM*, 2003.
- [34] G. Iannaccone, M. May, and C. Diot, Aggregate traffic performance with active queue management and drop from tail, *Comput. Commun. Rev.*, July 2001.
- [35] M. Menth, Efficient admission control and routing for resilient communication networks, PhD Thesis, Univ. of Wurzburg, July 2004.
- [36] S.R. Lima, P. Carvalho and V. Freitas, Distributed admission control in multiservice IP networks: concurrency issues, *J. Commun.*, June 2006.
- [37] G. Mao, D. Habibi, Loss performance analysis for heterogeneous ON-OFF sources with application to connection admission control, *IEEE/ACM Trans. Networking*, February 2002.
- [38] S. Floyd, Comments on measurement-based admission control for controlled-load services, Lawrence Berkeley Laboratory Technical Report, July 1996.
- [39] S. Jamin, S. Shenker, P. Danzig, Comparison of measurement-based admission control algorithms for controlled-load service, *IEEE INFOCOM*, 1997.
- [40] M. Karsten, J. Schmitt, N. Berier, R. Steinmetz, On the feasibility of RSVP as general signalling interface, *QoFIS*, 2001.
- [41] X. Fu, C. Kappler, Towards RSVP lite: light-weight RSVP for generic signalling, *IEEE AINA*, 2003.
- [42] M. Karsten, J. Schmitt, R. Steinmetz, Implementation and evaluation of the KOM RSVP engine, *IEEE INFOCOM*, 2001.
- [43] S. Belenki, An enforced inter-admission delay performance-driven connection admission control algorithm, *ACM SIGCOMM*, 2002.
- [44] L. Breslau, S. Jamin, S. Shenker, Comments on the performance of measurement-based admission control algorithms, *IEEE INFOCOM*, 2000.
- [45] A. Bilhaj, K. Mase, Endpoint admission control enhanced systems for VoIP networks, *IEEE SAINT*, 2004.
- [46] I. Frigui, Services and performance requirements for broadband fixed wireless access, *IEEE P802.16 Broadband Wireless Access Working Group*.
- [47] T. Chaded, IP QoS Parameters, TF-NGN, November 2000.
- [48] I. Habib, T. Saadawi, Multimedia traffic characteristics in broadband networks, in: *IEEE Communications Magazine*, July 1992.
- [49] <http://www-tnk.ee.tu-berlin.de/research/trace/trace.html>.
- [50] S. Atallah, O. Layaïda, N. Palma, D. Hagimont, Dynamic configuration of multimedia applications, *IFIP/IEEE MMNS*, 2003.
- [51] M. Diaz, D. Garduno, T. Gayraud, S. Owezarski, Multimedia multicast protocols based on multimedia models, *MMM*, 2001.
- [52] T. Lee, M. Zukerman, Admission control for bursty multimedia traffic, *IEEE INFOCOM*, 2001.
- [53] V. Elek, G. Karlsson, R. Ronngren, Admission control based on end-to-end measurements, *IEEE INFOCOM*, 2000.
- [54] R. Gibbens, F. Kelly, Measurement-based connection admission control, in: *15th International Teletraffic Congress*, June 1997.
- [55] S. Georgoulas, P. Trimintzios, G. Pavlou, K. Ho, Heterogeneous real-time traffic admission control in differentiated services domains, *IEEE GLOBECOM*, 2005.
- [56] J. Case, M. Fedor, M. SchoffStall, J. Davin, Simple network management protocol, *Internet RFC 1157*, May 1990.
- [57] J. Quittek, T. Zseby, B. Claise, S. Zander, Requirements for IP flow information export, *Internet RFC 3719*, October 2004.
- [58] B. Claise, Cisco Systems NetFlow Services Export Version 9, October 2004, *Internet RFC*, 3954.
- [59] A. Akella, S. Seshan, A. Shaikh, An empirical evaluation of wide-area internet bottlenecks, in: *ACM Internet Measurement Conference*, 2003.
- [60] T. Bressoud, R. Rastogi, Optimal Configuration for BGP Route Selection, *IEEE INFOCOM*, 2003.
- [61] Akamai, Internet bottlenecks: the case for edge delivery services, technical report, Akamai White Paper, 1999.



Stylianos Georgoulas was admitted to the Department of Electrical and Computer Engineering at the University of Patras in Greece and received his Diploma in 2001. He subsequently got his PhD degree in Networking at the Department of Electronic Engineering, University of Surrey, UK, in 2006. He is currently a post-doctoral Research Fellow at the Center for Communication Systems Research (CCRS), University of Surrey. During his time at Surrey he has been involved in the EU IST MESCAL, AGAVE and ENTHRONE projects. His research interests are in the area of admission control and bandwidth management/ traffic engineering for QoS-enabled IP networks.



Panos Trimintzios was a Research Fellow at CCSR, University of Surrey and now is employed at the European Network and Information Security Agency (ENISA). He received a B.Sc in Computer Science and an M.Sc in Computer Networks from the University of Crete, Greece and a Ph.D in Networking from the Department of Electronic Engineering, University of Surrey, UK. His interests include network control and management, network monitoring and security, traffic engineering and quality of service.



George Pavlou is a Professor of Communication and Information Systems at the CCSR, University of Surrey, UK, where he leads the activities of the Networks Research Group. He holds a MEng in Electrical and Mechanical Engineering from the National Technical University of Athens, Greece, and M.Sc and Ph.D degrees in Computer Science from University College London, UK. He is responsible for a number of European and UK research projects and industrial collaborations. His research

interests include network planning and dimensioning, traffic engineering, multicast, mobile ad hoc networks, multimedia service control, programmable networks and communications middleware. He has published approximately 170 papers in fully refereed international conferences and

journals and has contributed to standardization activities in ISO, ITU-T, TMF, OMG and IETF.



Kin-Hon Ho is a post-doctoral Research Fellow at the CCSR, University of Surrey, UK. He holds a BSc degree in Computer Studies from the City University of Hong Kong, an M.Sc degree in Data Communications from the University of Sheffield, UK, and a Ph.D degree in Networking from the Department of Electronic Engineering, University of Surrey, UK. During his time at Surrey he has been involved in the EU IST MESCAL and AGAVE projects. His research interests include Internet traffic engineering and network optimization.