On the Location-Awareness of Bandwidth Allocation and Admission Control for the Support of Real-Time Traffic in Class-Based IP Networks*

Stylianos Georgoulas[1], Panos Trimintzios[1,2], George Pavlou[1], and Kin-Hon Ho[1]

[1] Centre for Communication Systems Research, University of Surrey
Guildford, Surrey, GU2 7XH, United Kingdom
[2] ENISA, PO Box 1309, GR-71001, Heraklion, Crete, Greece

Abstract. The support of real-time traffic in class-based IP networks requires reservation of resources, accompanied by admission control in order to guarantee that newly admitted real-time traffic flows do not cause any violation to the Quality of Service (QoS) perceived by the already established ones. In this paper we highlight certain issues with respect to bandwidth allocation and admission control for supporting real-time traffic in class-based IP networks. We investigate the implications of topological placement of both bandwidth allocation and admission control schemes. We show that their performance depends highly on the location of the employed procedures with respect to the end-users and the various network boundaries. We conclude that the strategies for applying these schemes should be location-aware, because their performance at different points in a class-based IP network can be different and can deviate from the expected performance. Through simulations we also provide a quantitative view of these deviations.

1 Introduction

Class-based service models, such as the Differentiated services (Diffserv) model, offer a scalable approach towards QoS. This is achieved by grouping traffic with similar QoS requirements into one of the engineered traffic classes and forwarding it in an aggregate fashion. By allowing traffic aggregation, networks that deploy class-based service models can take advantage of statistical multiplexing, which allows for efficient use of the resources. In order to provide QoS guarantees, a network supporting different classes must also deploy admission control to control the amount of traffic injected into the network so as to prevent overloads that can lead to QoS violations.

In this work we focus on the support of *real-time* traffic 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 controlled data traffic. This is essential for guaranteeing QoS [1]. 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 UDP real-time traffic is classified to use a higher priority Assured Forwarding (AF) Per-Hop Behavior (PHB), whereas the TCP controlled data traffic is classified to use a lower priority, possibly best-effort, forwarding PHB.

* This work was undertaken in the context of the FP6 Information Society Technologies AGAVE (IST-027609) project, which is partially funded by the Commission of the European Union.

A. Helmy et al. (Eds.): MMNS 2006, LNCS 4267, pp. 195 – 207, 2006.

We define as *real-time* traffic flows, flows that have low delay and jitter requirements, and a bounded packet loss rate (PLR) requirement and we focus on the 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]. The low delay and jitter requirements are likely to be met in a high-speed network [2] and offline traffic engineering actions can be additionally taken so that delay and jitter are kept within low bounds. For example, the delay requirement can be taken into account at the network provisioning phase by a) configuring small queues for the real-time traffic so as 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 as long as the flows are shaped to their peak rate at the network ingress [4]. Also, 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 peak rate at the network 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.

In this work we 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 and the various network boundaries, can affect their performance. Before proceeding, we will briefly describe the bandwidth allocation schemes suitable for aggregating real-time traffic and also describe the admission control schemes that can be used for real-time traffic in class-based IP networks, so that later we can refer to them and point out how the results of our work relate to them.

2 Bandwidth Allocation Schemes

Bandwidth allocation schemes can be divided in two main categories. The first category comprises schemes based on *bufferless statistical multiplexing*. Bufferless statistical multiplexing 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 stored in large buffers. Examples of such bandwidth allocation schemes can be found in [6, 7, 8]. Generally speaking, both categories take into account factors such as the number and characteristics of flows, the required loss rate and, in case of buffered multiplexing, the available buffer size and derive the required capacity (*effective bandwidth*) needed for the loss rate to be kept below the required threshold.

Each of the two categories has its merits but also its drawbacks. The dynamics leading to an overload event in a bufferless system are much simpler than those of a buffered system [9]. On the other hand, buffered multiplexing allows for higher utilization for the same loss rate [9]. We need to stress though, that bufferless multiplexing is just a model abstraction [9]. For packetized traffic, a small buffer for packet scale queuing is needed to account for simultaneous packet arrivals from distinct flows [8]. However, in this case the buffer is not used for storing significant amounts of excess traffic and is, therefore, not involved in bandwidth estimations.

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).

EAC is based on some metrics applied to probing packets sent along the transmission path before the flow is established [10]. A requirement is for the end-to-end route to be the same for probing packets and flows. For reasonably bounded setup delays the metrics do not depict stationary network states but snapshots of network status, which can result to a quite unrealistic picture of the network congestion and, furthermore, simultaneous probing by many sources can lead to a situation known as thrashing [10]. That is, even though the number of admitted flows is small, the cumulative level of probing packets prevents further admissions.

TDAC is based on the assumption that traffic descriptors are provided for each flow prior to its establishment. This approach achieves high utilization when the traffic descriptors used by the scheme are appropriate. Nevertheless, in practice, it suffers from several problems [11]. One is the inability to come up with appropriate traffic descriptors before establishing the flow. Another problem is that the traffic descriptors and the associated QoS guarantee define a contract between the flow and the network. Therefore, the need to police based on this traffic specification arises, which is difficult for statistical traffic descriptors [11]. Deterministic models, such as token buckets, are easy to police but they fail to provide a sufficient characterization to extract a large fraction of the potential statistical multiplexing gain [11].

MBAC tries to avoid the problems of the other approaches by shifting the task of traffic characterization to the network [11]. That means that 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, which is easy to police. Second, a 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 on their aggregate behavior, the statistics of which are easier to estimate than those of an individual flow. However, relying on measured quantities raises certain issues, such as estimation errors, flow-level dynamics and memory related issues [11].

4 The Effects of Topological Placement

In this section we will show how the topological placement of the functionalities can affect the performance of bandwidth allocation and admission control schemes.

For the needs of our work 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 [12] it is suggested that the aggregation of even a fairly small number of traffic streams is usually sufficient for the Gaussian characterization of the input process.

In this case, the effective bandwidth of N multiplexed sources is given by [6, 7]:

$$C \simeq m + a'\sigma \text{ with } a' = \sqrt{-2\ln(\varepsilon) - \ln(2\pi)} \tag{1}$$

where $m = \sum_{i=1}^{N} m_i$ is the mean aggregate bit rate, $\sigma = \sqrt{\sum_{i=1}^{N} \sigma_i^2}$ is the standard deviation of the aggregate bit rate, and ε the upper bound on allowed loss rate. We will denote the function of eq. 1 as $\mathit{eff}(S, PLR)$, where S is the set of aggregated sources and PLR is the PLR value used in the calculation of the effective bandwidth C.

We will present our study with a list of scenarios using a two-level tree topology, which allows us to illustrate the mains points of our study, while being simple enough to suit the nature and computational demands of the required packet-level simulations.

4.1 Scenario I: The Effects of Aggregate Bandwidth Allocation

Initially we consider the scenario depicted in Fig. 1.

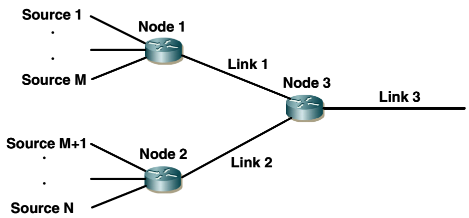

Fig. 1. Topology for assessing the effects of topological placement

In this scenario it is assumed that a set of sources, $S_i, i = 1,...,M$, are aggregated at node 1 and that another set of sources, $S_i, i = M+1,...,N$, 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:

$$C_1 = \mathit{eff}(\{S_1,...S_M\}, PLR_1) \tag{2}$$

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:

$$C_2 = \mathit{eff}(\{S_{M+1},...S_N\}, PLR_2) \tag{3}$$

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, the required bandwidth to be allocated in link 3 for their aggregation is given by:

$$C_3 = \mathit{eff}(\{S_1,...S_N\}, PLR_3) \tag{4}$$

where PLR_3 is the allowed packet loss rate budget for the real-time traffic 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.

As it can be easily proven [13], packet loss rate parameters are multiplicative. That means that for a set of sources that traverse a sequence of links, l_i, $i = 1,...,L$, with packet loss rates PLR_i, the total packet loss rate PLR_{total} can be approximated by:

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

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 total, end-to-end in our topology, packet loss rate for the aggregate sources should be:

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

Our study initially aims to examine whether the actual total packet loss rate is bounded by the above expression. In order to do so, we run simulations using the network simulator *ns-2* [14]. For the simulations we use two example values for the target link packet loss rates, 0.01 and 0.001. 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 64kbps and exponentially distributed ON and OFF periods with average durations 1.004sec and 1.587sec respectively (mean rate 24.8kbps, standard deviation of rate 31.18kbps) [15], b) all sources are *Videoconference* sources with mean rate 3.89Mbps, peak rate 10.585Mbps and standard deviation of rate 1.725Mbps [16], and c) that we have a *mixture* of both VoIP and Videoconference sources. We fix the packet size to 100bytes (constant packet size seems to be a reasonable assumption for Voice and Video communications [17]) and 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 and we simulate the real-time traffic class as being serviced by queues running at the speed of their bandwidth limit.

In the following figures, 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, $i = 1,...,N$ (x-axis).

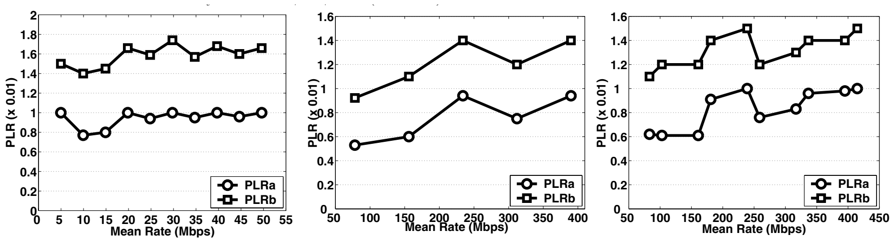

Fig. 2. Incurred PLR for *VoIP* (left), *Videoconference* (center) and *mixed* (right) traffic sources for target link PLR 0.01

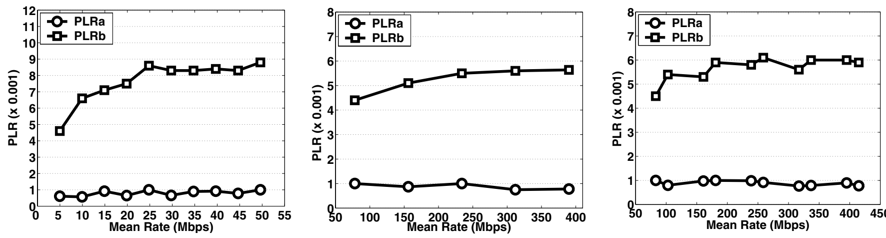

Fig. 3. Incurred PLR for *VoIP* (left), *Videoconference* (center) and *mixed* (right) traffic sources for target link PLR 0.001

From the above figures we can see that in all cases, the target packet loss rate in links 1 and 2 (PLRa) is 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, that 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 at the first points of aggregation, they may not be valid in transit nodes such as in node 3. This is because sources become correlated and their characteristics are altered as they traverse links and multiplexers. Therefore, using the original traffic descriptors in transit nodes can lead to erroneous bandwidth allocation decisions. This traffic profile deformation has been verified in the past [8, 18, 19] and analytical models for evaluating it for specific types of individual sources and under specific network conditions and assumptions have been presented [8].

In this work, we try to go one step further and quantify the effects that this aggregate traffic profile deformation can have on the incurred packet loss and, therefore, on the perceived QoS when short-range dependent (VoIP), long-range dependent (Videoconference), as well as a mixture of short-range and long-range dependent real-time traffic sources are aggregated in the same traffic class. In such general cases, obtaining a closed form solution to quantify its effects can become very difficult.

4.2 Scenario II: Quantifying the Effects of Traffic Profile Deformation

In order to quantify the effects of traffic profile deformation further downstream from the first hop node, we proceed as follows. We use the same simulated topology, traffic volume and types of 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 sources aggregated in links 1 and 2 (links 1 and 2 are transparent to the sources with respect to packet loss). For link 3 we distinguish two cases. In the first case, we merge the two demands –one composed of the sources S_i, $i = 1,...,M$ and the other composed of the sources S_i, $i = M+1,...,N$ – in one bandwidth allocation in link 3:

$$C_3 = eff(\{S_1,...S_N\}, PLR_3) \tag{8}$$

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

$$C_3 = eff(\{S_1,...S_M\}, PLR_3) + eff(\{S_{M+1},...S_N\}, PLR_3) \tag{9}$$

In the following figures, the target packet loss rate for link 3 is set to 0.01 and 0.001. PLRb corresponds to the incurred packet loss rate from link 3 when using eq. 8 and PLRa corresponds to the incurred packet loss rate when using eq. 9.

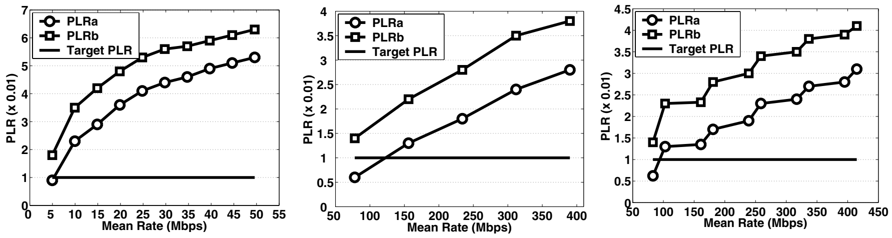

Fig. 4. Incurred PLR for *VoIP* (left), *Videoconference* (center) and *mixed* (right) traffic sources for link 3 with target PLR 0.01

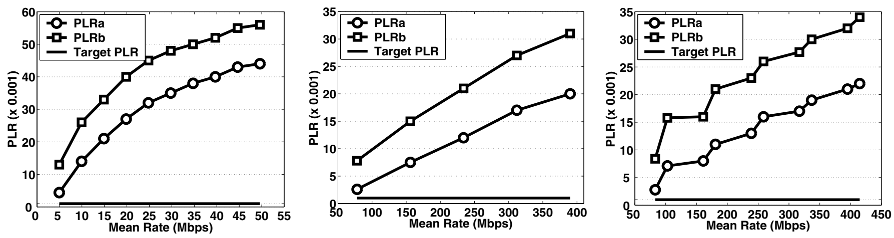

Fig. 5. Incurred PLR for *VoIP* (left), *Videoconference* (center) and *mixed* (right) traffic sources for link 3 with target PLR 0.001

As it can be seen, the effects of the traffic profile deformation, even one hop away from the node where the sources are firstly multiplexed, can lead to severe violations of the packet loss rate. Even if the *isolation* method (eq. 9) is used, which, given the form of the effective bandwidth formula of eq. 1, leads to more conservative resources reservation compared to when 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 can lead to higher packet loss rate violations. This indicates that the detrimental effects of the aggregate traffic profile deformation can far exceed the positive effects due to the additional multiplexing gains that this increase is expected to have.

4.3 Scenario III: Bandwidth Allocation with Buffered Multiplexing Models

In order not to restrict ourselves to the bufferless approach, we repeat part of the above simulations using the buffered approach introduced in [7]. According to [7], 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:

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

We use this effective bandwidth formula to estimate the bandwidth for the case of Videoconference sources and we repeat our experiments for two buffer size levels, set intentionally to relatively small values, 30kbytes and 50kbytes respectively.

In the following figures, PLRa corresponds to the average packet loss incurred in links 1 and 2 when the sources are aggregated in links 1 and 2 and the target packet loss rates of links 1 and 2 are set equal to 0.01 and 0.001. PLRb corresponds to the packet loss incurred in link 3 when the same number of sources is aggregated in link 3 with the capacities of links 1 and 2 set equal to the sum of peak rates of the sources they are carrying -links 1 and 2 are transparent with respect to packet loss- and the bandwidth allocated in link 3 set for target link 3 loss rate equal to 0.01 and 0.001.

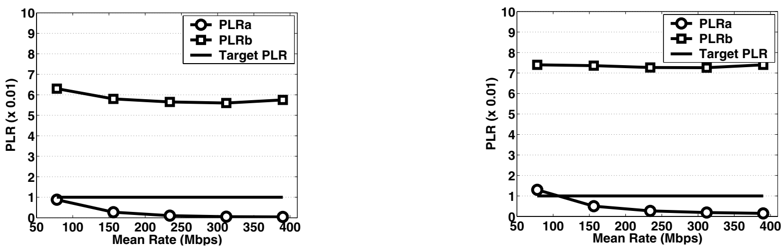

Fig. 6. Incurred PLR for *Videoconference* sources for link PLR 0.01 and queue size 30kbytes (left) and 50kbytes (right)

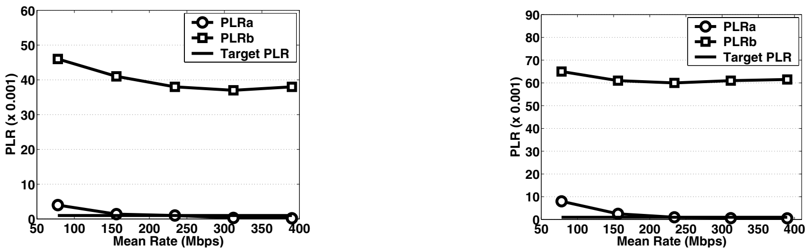

Fig. 7. Incurred PLR for *Videoconference* sources for link PLR 0.001 and queue size 30kbytes (left) and 50kbytes (right)

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 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. Furthermore, contrary to scenario II, the packet loss in link 3 does not increase with the level of aggregation. This is due to the additive nature of eq. 11, which becomes more conservative for increasing levels of aggregation and, therefore, can partly compensate for the detrimental effects of the aggregate traffic profile deformation.

5 Discussion

In this section we will first elaborate on the implications of our study and we will point out relevant issues that are raised. Subsequently, we will present some possible practical traffic engineering solutions for dealing with these issues.

5.1 General Implications

The first implication of our study is that the original traffic descriptors for a set of real-time traffic sources multiplexed at a network edge, are not only invalid at downstream nodes but, also, the incurred traffic profile deformation has, in general, a negative effect. That is, the traffic characteristics of a 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 implication is that for a 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 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 is less likely to fail 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. Due to multiplexing and buffering, the packets belonging to a source may arrive at an interface at a rate even 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 [19]. Therefore, even a conservative MBAC scheme, which while deriving the admission control decision, assumes that the source requesting admission will be transmitting at its peak rate, may fail when applied in core nodes.

A similar problem for TDAC and MBAC schemes arises when performing admission control for inter-domain traffic; that is traffic that traverses transit domains. 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 some kind of 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, since it requires the estimation of delay variation and induces per flow states [18].

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 MBAC scheme is applied on a link-by-link basis, if the link packet loss rates are not 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 effect is similar to the discrimination against long flows in the case of EAC schemes [10]. 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 resources for the links carrying short real-time traffic flows only, 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 [9].

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 in core nodes is based on an effective bandwidth formula, even if done based on *edge-to-edge isolation*, that is that traffic aggregates multiplexed in the same traffic class in the core network are allocated resources on an ingress-egress pair basis, the consequences may be detrimental. Furthermore, similar to the admission control case, the 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 regarding 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 classes so that in every link, only flows with similar target link packet loss rate requirements are aggregated in the same class. However, this would mean increasing the number of classes that must be engineered and supported in the routers. Apart from the added complexity in network dimensioning, increasing the number of classes that routers must support can lead to decreased forwarding performance [20].

A unified approach for bandwidth allocation and admission control that can be used to overcome the aforementioned problems, including the traffic profile deformation, is to apply admission control only at the network ingress and further downstream treat the real-time traffic aggregates in a peak rate manner. This simplifies network dimensioning, since it removes the packet loss related hop count constraint, and is feasible since the edge links are the most probable congestion points of a domain [21], whereas backbone links are overprovisioned [22]. This approach does not induce any states in the core network and does not require core routers to be aware of any signaling, which is desired for scalability and resilience reasons, and it is also proven to be very resource-efficient if resilience against network failures is required [23]. Furthermore, as our results suggest, the greatest multiplexing gains are achieved at the network edge, and by employing this approach, since 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.

6 Related Work

As aforementioned, the problem of traffic profile deformation has been verified in the past [8, 18, 19] and a number of solutions for dealing with it have been proposed.

One solution is the use of traffic-descriptor conserving scheduling disciplines in all links along the end-to-end paths. Example of such schedulers is the Rate-controlled Static Priority (RCSP) scheduler [24]. This preserves the original traffic descriptors of each 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. Furthermore, since RCSP is intended to provide zero packet loss guarantees, the deployment of RCSP can lead to unnecessarily conservative usage of network resources.

In [8] the issue of traffic profile deformation is discussed in the context of Constant Bit Rate sources in Asynchronous Transfer Mode 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 states and requires an appropriate method in order to obtain this delay variation estimation (e.g. in [25] the delay variation is estimated by employing a probing procedure). Similar to RCSP, this method may not be feasible in class-based networks since it can impose the requirement for adding functionality and keeping per-flow state in core nodes.

Contrary to these works, we focus not on the traffic profile deformation of individual sources, which induces per-flow states and can require added functionality in core nodes for the delay variation estimation. Instead, we focus on quantifying the aggregate traffic profile deformation for a variety of real-time traffic sources multiplexed in the same class. We discuss issues that are raised and how they can be addressed, without imposing the requirement for keeping per-flow state information in core nodes.

7 Concluding Remarks

In this paper we 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 and the various network boundaries (access, core, etc.), and we show that their performance depends on it; that is their performance at different points in a class-based IP network and for the same traffic load can be different and deviate greatly from the expected targets. Through simulations we also provide a quantitative view of these deviations. Finally, we propose a unified approach for bandwidth allocation and admission control that can overcome the detrimental effects of this "location-awareness".

References

1. W. Sun, P. Bhaniramka and R. Jain, "Quality of Service using Traffic Engineering over MPLS: An Analysis", IEEE LCN 2000.
2. G. Schollmeier and C. Winkler, "Providing Sustainable QoS in Next-Generation Networks", IEEE Communications Magazine, June 2004.
3. J. Lakkakorpi, O. Strandberg and J. Salonen, "Adaptive Connection Admission control for Differentiated Services Access Networks", IEEE Journal on Selected Areas in Communications, October 2005.
4. T. Bonald, A. Proutiere and J. Roberts, "Statistical Performance Guarantees for Streaming Flows using Expedited Forwarding", IEEE INFOCOM 2001.
5. M. Mowbray, G. Karlsson and 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 Journal on Selected Areas in Communications, September 1991.
7. C. Courcoubetis G. Fouskas and R. Weber, "On the Performance of an Effective Bandwidths Formula", International Teletraffic Congress 1994.
8. J. Roberts, U. Mocci and J. Virtamo, "Broadband Network Teletraffic", Final report of action COST 242, Springer 1996.
9. T. Bonald, S. Oueslati-Boulahia and J. Roberts, "IP traffic and QoS control: the need for a flow-aware architecture", World Telecommunications Congress, September 2002.
10. L. Breslau, E. Knightly, S. Shenker, I. Stoica and Z. Zhang, "Endpoint Admission Control: Architectural Issues and Performance", SIGCOMM 2000.
11. M. Grossglauser and D. Tse, "A Framework for Robust Measurement-Based Admission Control", IEEE/ACM Transactions on Networking, June 1999.

12. D. Eun and N. Shroff, "A Measurement-Analytic Approach for QoS Estimation in a Network Based on the Dominant Time Scale", IEEE/ACM Transactions on Networking, April 2003.
13. S. Lima, P. Carvalho and V. Freitas, "Distributed Admission Control for QoS and SLS Management", Journal of Network and Systems Management, September 2004.
14. K. Fall and K. Varadhan, "The *ns* manual" (www.isi.edu/nsnam/ns/ns_doc.pdf).
15. C. Chuah, L. Subramarian and R. Katz, "Furies: A Scalable Framework for Traffic Policing and Admission Control", U.C Berkeley Technical Report, May 2001.
16. B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson and J. Robbins, "Performance Models of Statistical Multiplexing in Packet Video Communications", IEEE Transactions on Communications, July 1988.
17. I. Mas, V. Fodor and G. Karlsson, "The Performance of Endpoint Admission Control Based on Packet Loss", QoFIS 2003.
18. K. Shiomoto, N. Yamanaka and T. Takahashi, "Overview of Measurement-based Connection Admission Control Methods in ATM Networks", IEEE Communication Surveys 1999.
19. H. Perros and K. Elsayed, "Call Admission Control Schemes: A Review", IEEE Communications Magazine, November 1996.
20. M. Torneus, "Testbed for Measurement Based Traffic Control", Master's Thesis, KTH IMIT, Sweden, June 2000.
21. V. N. Padmanabhan, L. Qiu and H. J. Wang, "Server-based inference of Internet Link Lossiness", IEEE INFOCOM 2003.
22. G. Iannaccone, M. May, and C. Diot, "Aggregate Traffic Performance with Active Queue Management and Drop from Tail", Computer Communications Review, July 2001.
23. M. Menth, "Efficient Admission Control and Routing for Resilient Communication Networks", PhD Thesis, Univ. of Wurzburg, July 2004.
24. H. Zhang and D. Ferrari, "Rate-controlled Static-Priority queuing", IEEE INFOCOM 1993.
25. M. Dabrowski, A. Beben and W. Burakowski, "On Inter-domain Admission Control Supported by Measurements in Multi-domain IP QoS Networks", IEEE IPS 2004.