Admission Control Placement in Differentiated Services Networks

Stylianos Georgoulas Panos Trimintzios George Pavlou

Centre for Communication Systems Research, University of Surrey
Guildford, Surrey, GU2 7XH, United Kingdom
Email: {S.Georgoulas, P.Trimintzios, G.Pavlou}@eim.surrey.ac.uk

Abstract

The primary role of admission control in Quality of Service enabled networks is to control the amount of traffic injected into the network so that congestion is avoided and certain performance requirements are met. We consider engineered and provisioned IP Differentiated Services networks able to support real-time traffic and we address the placement of admission control at the traffic aggregation points of the network, and the granularity of the admission control logic. Regarding the first issue, we show that sophisticated admission control schemes that take into account statistical multiplexing gains need only be employed at the first traffic aggregation points. Further downstream, peak rate admission control will suffice. With respect to the second issue, we propose a framework for admission control, which involves a combined approach of traffic descriptor and measurement–based techniques.

1. Introduction

It is widely accepted that the today best effort Internet is not able to satisfactory support emerging services and market demands, such as Voice over IP (VoIP) and Videoconference. Differentiated services (Diffserv) are seen as the technology to support Quality of Service (QoS) even for real-time services in IP networks without the scalability problems of Integrated services (Intserv). Admission control is the set of actions taken by the network during the service establishment phase to check whether a service request is to be admitted or rejected. A new service request is admitted when the desired QoS for the new service can be satisfied, without causing any QoS violation to the already established services. An additional role of admission control is to optimize the use of network resources. The challenge is to design simple control functions that improve efficiency under any offered traffic conditions. In Diffserv, the lack of per flow information and signaling in the core network imposes restrictions to the employed admission control schemes, since checking whether sufficient capacity exists end-to-end might not even be feasible in the first place, e.g. when traffic crosses many domains or even intra-domain when an equal cost multi-path routing scheme is employed.

The various admission control approaches differ in the methods they use to decide if there is enough capacity for the new service request and can be divided into three categories: (a) Admission control using *a priori* traffic descriptors, (b) Measurement based admission control (MBAC), and (c) Endpoint admission control (EAC).

In admission control using *a priori* traffic descriptors [1], it is assumed that there exists perfect knowledge of each traffic source type that will be used and, additionally, of the current number of established service instances. This information will enable admission control to compute the total amount of bandwidth required. However, since no traffic measurements are taken into consideration, if the provided traffic descriptors do not depict the actual behavior of the sources, or the appropriate traffic descriptors are not known *a priori*, the performance of this admission control scheme can be very poor.

MBAC shifts the task of traffic specification from the user to the network [2] in order to cope with the inherent problems of admission control using *a priori* traffic descriptors. Instead of users 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 user-specified traffic descriptors can be very simple, e.g. peak rate, which can be easily policed. 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. However, relying on measured only quantities for admission control raises a number of issues that need to be considered, such as the estimation errors [2], system dynamics and memory related issues.

Endpoint admission control is based on some metric applied on probing packets sent by the end host/application along the transmission path [3]. A requirement is for the end-to-end route to be the same for probing packets and flows. Setup delays may be high and, furthermore, simultaneous probing by many sources can lead to a situation known as *thrashing* [3]. That is that even though the number of admitted flows is small, the cumulative level of probe packets prevents further admissions.

In this work we consider admission control for real-time traffic. We define as real-time traffic sources, the ones which have a strict small delay requirement and a bounded, not necessarily too low, packet drop rate (PDR) requirement. In a Diffserv domain we assume that such traffic will be aggregated to form one or more real-time traffic aggregates. The delay requirement of the traffic aggregate has been taken into account in the provisioning stage, i.e. by appropriately setting small queues and by manipulating the routing process to choose appropriate paths. Packets are expected to be lost only at the first point of aggregation (ingress node), where we are going to have the serialization of the various traffic sources. We assume that the interior of the Diffserv domain has been provisioned and engineered in order to support the real-time traffic aggregate. As a result of this provisioning process, and taking into account the routing behavior, at each ingress node we have an estimate of the minimum available bandwidth that is available from that ingress to each of the corresponding egress nodes. This available bandwidth is the basis for our admission control approach, which is employed at the edge (ingress) node of the first Diffserv aggregation point, for accepting a traffic source on behalf of the entire Diffserv domain.

We propose a bandwidth management approach that overcomes the limitations imposed by this lack of per-flow state and signaling in the core network whilst allowing for efficient admission control schemes for real-time traffic. This approach takes into account statistical multiplexing gains due to the aggregation of sources, accounts for issues such as the increase or the decrease in the level of aggregation along the real-time traffic paths, and determines the granularity of the employed admission control schemes with respect to the candidate points of enforcement. Taking the implications of the bandwidth management approach into account, a novel framework for admission control is proposed and its performance is compared against existing approaches.

2. On the Placement of Admission Control

For bandwidth allocation we adopted the effective bandwidth approach. According to [4], when the effect of statistical multiplexing is significant, the distribution of the stationary bit rate can be accurately approximated by a Gaussian distribution. In that case, the effective bandwidth of the multiplexed sources is given by:

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

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 queue overflow probability. In the rest we will denote the function implied by (1), as $\mathit{eff}(S)$, where S is the set of sources under aggregation. When admission control logic takes into account the statistical gains by using (1), there are two issues that need to be considered:

1. whether already statistically multiplexed aggregated traffic, for which the bandwidth allocation was performed using the effective bandwidth approach, allows for further statistical multiplexing gains when merged with similarly statistically multiplexed aggregated traffic.
2. since the effective bandwidth formula is only valid for an aggregation of sources, what happens if, for traffic engineering reasons, the sources that their aggregation gives the value for allocating bandwidth follow different paths before even being aggregated.

These two issues play important role on the placement and granularity of the required admission control functions. The first issue answers the question if we need to perform fined grained, that is per-flow admission control at aggregation points where the received traffic is composed of already aggregated flows. This case arises at the ingress node of a Diffserv domain, which peers with other domains and receives traffic aggregates. As we will show in this section, there is no need to apply any sophisticated algorithms to aggregates for admission control, like applying (1), but simple peak rate admission control schemes will suffice.

Furthermore, we show that the admission control procedures need to take into account statistical gain only at the first point of aggregation, which is close to the access, i.e. concentration metropolitan networks. This result is also important because the resources in such networks are scarce compared to further downstream core networks. The results answering the first issue, point the location for placement of statistical admission control at the edge node of the first aggregation network and peak rate-based admission control at the edge nodes of the subsequent Diffserv domains, making the assumption that admission control is applied only at the edge nodes. This assumption is inline with the Diffserv philosophy, i.e. all complexity is at the edges and the core is as simple as possible.

It is possible the bandwidth allocated to a Diffserv aggregate at an edge node for admission control purposes, to be composed of the total bandwidth of multiple paths, created for engineering reasons (e.g. load-balancing), from that ingress edge to the egress node. In this case it is not clear that admission control solutions that take into account statistical gains, e.g. using (1), will take correct admission decisions. We will show that even in these cases, the admission control decisions based on statistical gain assumptions are generally correct.

2.1. Merging of Traffic Aggregates

We want to test whether we can achieve multiplexing gains when aggregating streams that are composed of already aggregated flows. We used the scenario shown in Fig. 1 in order to test whether the aggregation of aggregated traffic allows for further statistical multiplexing gains when merged.

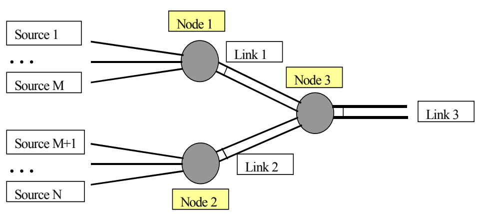

Figure 1. Merging situation.

In this scenario it is assumed that the set of sources, S_i, $i = 1,...,M$, are aggregated at node 1. The required capacity for link 1 was set equal to the value computed using (1). Let C_1 denote the allocated capacity of link 1, then:

$$C_1 = eff(S_1,...,S_M) \quad (2)$$

Similarly the set of sources S_i, $i = M+1,...,N$, are aggregated at node 2 and if C_2 denotes the capacity of link 2, it is:

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

The question we want to answer is if the required capacity of link3, C_3, can be set equal to the effective bandwidth of the total number of sources as a whole, that is:

$$C_3 = eff(S_1,...,S_N) \quad (4)$$

or if we need to set it equal to the sum of the capacities of links 1 and 2, that is:

$$C_3 = C_1 + C_2 \quad (5)$$

in order to achieve the target bound on packet drop rate (PDR) for that traffic aggregate.

The former case suggests that further multiplexing gains can be expected at downstream aggregation points, whereas the latter case suggests that we cannot expect any significant additional gains apart from the ones at the *first point* of aggregation.

For the simulations we assumed that in all cases the target bound packet drop rate was 0.01 and that the same number of sources were aggregated in both links 1 and 2, i.e. $M = N/2$. Regarding the type of traffic sources used, we considered three scenarios: a) all sources are VoIP sources [5] with peak rate 64kbps and exponentially distributed ON and OFF periods with average durations 1.004sec and 1.587 sec respectively, b) all sources are Videoconference sources [6] with mean rate 3.89Mbps and peak rate 10.585Mbps, and c) that we have a mixture of VoIP and Videoconference sources. PDR1 corresponds to the PDR achieved when using (4) and PDR2 to the PDR achieved when using (5). The achieved PDR as a function of the mean aggregate bit rate of the sources was as follows:

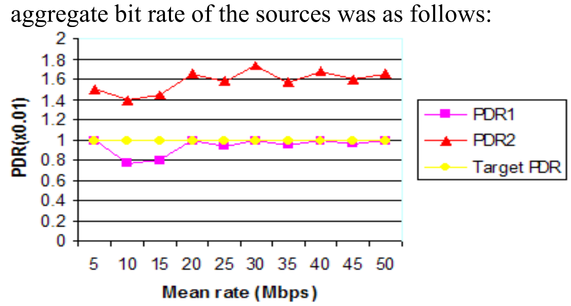

Figure 2. PDR for VoIP sources.

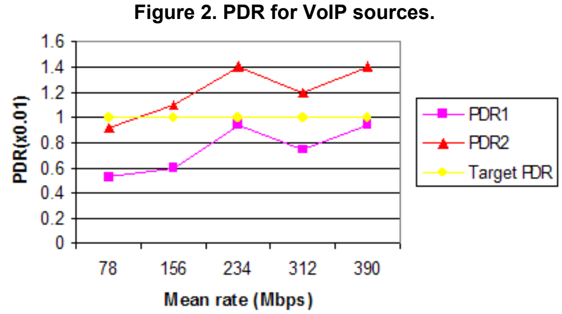

Figure 3. PDR for Videoconference sources.

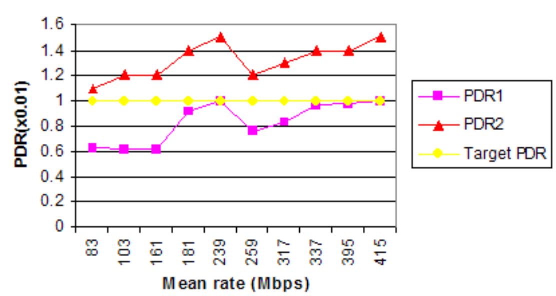

Figure 4. PDR for mixed sources.

The results suggest that we cannot anticipate any further significant gains after the first point of aggregation. An explanation for this is that the small output queues on links 1 and 2 (needed to cope with coincident packet arrivals from distinct flows [7]), combined with output capacities that are significantly lower than the sum of peak rates of individual sources, act as shapers so that the output streams tend to be more of constant bit rate nature, thus not allowing for further gains from aggregation.

The implication of this result is that admission control based on aggregation gain expectations need only to be applied at the first aggregation ingress point of the first Diffserv domain. At domains further downstream, where the traffic under aggregation is composed of aggregates from upstream domains, we cannot expect any significant aggregation gain, which implies that peak rate-based admission control will suffice. Furthermore, this result implies that if we have some information about the expected traffic at the edge nodes, we can provision the interior nodes assuming that the merging of bandwidths of real-time traffic aggregates will be done in a simple additive manner.

2.2. Splitting of Traffic Aggregates

We assume that in a Diffserv domain, admission control is applied only at the edge nodes. The engineering and provisioning of the network may impose some implications on the admission control system. For example, there may be the case that the total bandwidth, considered by the admission control system, for one traffic aggregate, from an ingress to an egress node of the domain, is, for load balancing reasons, composed out of multiple paths. An admission control framework, which takes into account the gains from aggregating sources, e.g. by using (1), is only valid for the aggregation of sources. An important question is what happens if the sources under aggregation follow different paths, before they are even aggregated. This can happen when there are for example multiple paths towards the egress node and the ingress node decides to route the aggregate from different paths. In order to investigate this situation, we used the scenario shown in Fig. 5.

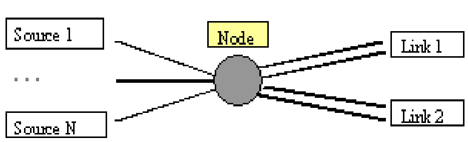

Figure 5. Splitting situation.

Lets assume that a set of sources, S_i, $i = 1,\ldots,N$, all carrying real-time traffic, i.e. belonging to the same traffic aggregate, and having the same egress node, enter the domain at the same ingress node. Therefore, it is assumed that they will be aggregated and that their required bandwidth will be:

$$C_3 = eff(S_1,\ldots,S_N) \quad (6)$$

In the case where all the traffic aggregate follows the same path towards the egress, and the network was provisioned to support this required bandwidth over this path, then an admission decision based on (1) will be correct. But what happens when, for traffic engineering reasons, the total edge-to-edge bandwidth allowed for that traffic aggregate is composed of more than one path towards the egress? This means that the next hop for the various flows, which compose the aggregate, may be different. Let's assume that in our example have two next hops, with allocated capacities for this traffic aggregate C_1 and C_2, so that $C = C_1 + C_2$. The splitting of the traffic aggregate means that the level of gains from aggregation will be lower than when aggregating all N sources, since some sources, say M will be forwarded towards link 1 and the rest, $N - M$, will be forwarded towards link 2 without being aggregated altogether. The question we need to answer is, whether this decrease in aggregation level will impair the desired target bound on the PDR.

The simulation setup was similar to the previous one. We used 0.01 as the target bound on the PDR for all the cases. We experimented with different splitting ratios, $M / N - M$, ranging from 1:1, i.e. each next hop receiving 50% of the incoming traffic, to up to 1:9, i.e. where one receives 10% and the other 90%.

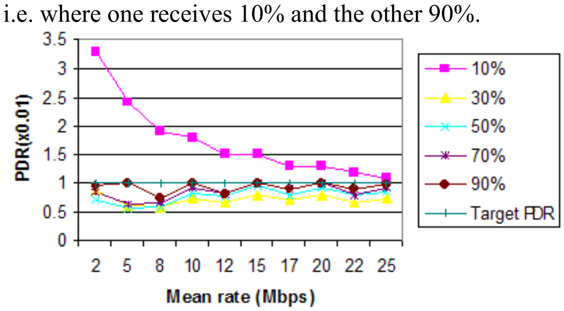

Figure 6. PDR for the links carrying VoIP traffic.

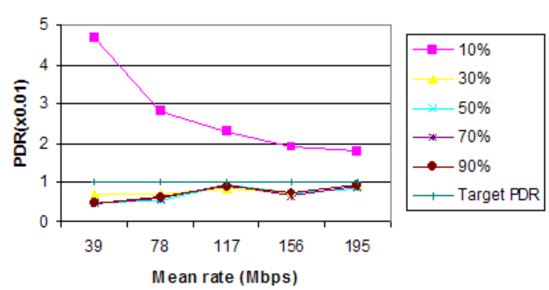

Figure 7. PDR for the links carrying Videoconference traffic.

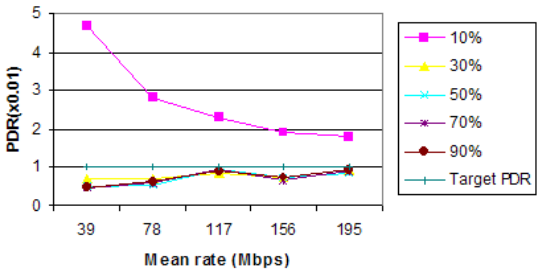

Figure 8. PDR for the links carrying mixed traffic.

The simulations showed that the desired PDR could be achieved in most cases. We only have violation of the PDR bound on the link carrying the small percentage of the split traffic in cases where the mean total rate is small and the individual sources have high rate requirements (Videoconference), resulting in having a small number of sources and thus a small aggregation gain. This set of simulations shows that, apart from cases where the population of sources is very small and the splitting ratio results in very small parts of the aggregate traffic, the effective bandwidth can be considered by admission control in an additive manner. This result is important because it allows the decoupling of admission control, from the engineering and provisioning of a Diffserv domain.

2.3 Implications for Admission Control

The results on merging and splitting traffic aggregates from the previous two sections suggest that the effective bandwidth can be treated in an additive manner. That is, if a value for bandwidth, computed using (1) is given as input to network provisioning functionalities, it can be treated as arithmetic value and still the designated edge-to-edge QoS be delivered, provided that it is manipulated in an additive manner. If the aforementioned approach for deriving and manipulating the traffic demands is employed, it will mean that “fine-grained” admission control, taking into account individual sources’ characteristics, need only be employed at the first points of aggregation since no further degradation of QoS is expected to occur in the core network. Further downstream, e.g. for inter-domain admission control, the aggregated stream need only be treated in a peak rate manner.

3. Admission Control

Our proposed approach for admission control is a combination of admission control using *a priori* traffic descriptors and measurement-based admission control. The term Measurement-based and *a priori* Traffic descriptor Admission Control (MTAC) will refer to our approach. This approach encompasses the positive features of both admission control schemes. It provides a systematic way to derive the bandwidth requirements of the already established flows and the candidate for admission flow using the traffic descriptors and facilitating the bandwidth estimation as given in (1). In addition, we use real-time measurements of the actual load in order to cope with the fact that the traffic descriptors may not depict the exact characteristics of the individual flows. We require measurements to be taken only at the ingress nodes, since we assumed that the network was properly provisioned and engineered as discussed in the previous sections. With reference to the previous sections, our work is positioned at the first aggregation point of sources.

We assume that through provisioning and traffic engineering, C_{total} bandwidth is available edge-to-edge for the real-time traffic aggregate. We use a *reference source*, with mean and standard deviation (m_{ref}, σ_{ref}) as a model source for engineering reasons. We define as *reference trunks* (T_{ref}) the number of simultaneously established reference sources that can fit in C_{total} for a given target bound on packet drop rate.

For more details, regarding the admission control logic, the reader can refer to [8].

4. Performance Evaluation

In order to evaluate the performance of our admission control framework we simulated the dumbbell topology of Fig. 9.

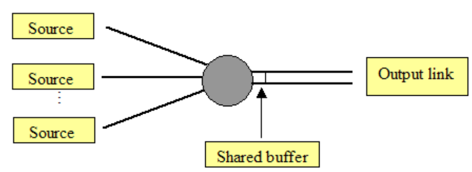

Figure 9. Simulation topology.

We configure the output queue to hold a maximum of 5 packets. We use scenarios with the target bound on packet drop rate for the aggregate real-time traffic being equal to 0.01 and 0.001. We test varying load conditions ranging from 0.5 to 5, where the value 1 corresponds to the *design load*. The latter is the load the system is designed to handle for a desired service rejection probability of 0.01, if we employ a circuit switched approach for *reference sources*.

The output link capacity corresponds to T_{ref} equal to 100. This means that for a service rejection probability equal to 0.01, the *reference source* activation rate would be a Poisson arrival process with average 1000 source activations/hour.

For traffic sources we use VoIP and Videoconference traffic. For VoIP we use the ON/OFF model of [5]. The active time of the VoIP sources is exponentially distributed with an average of 300sec. This model is also used as the *reference source* model. For Videoconference we use an H263 coded trace of the StarWars movie [9]. Its peak rate is 332.8kbps, its mean rate is 64kbps and the standard deviation of the bit rate is 28.8kbps. The active time of the Videoconference sources is exponentially distributed with an average of 180sec. For both VoIP and Videoconference sources, the activation processes are Poisson arrival processes with averages following a ratio 2:1. We consider the cases where the sources requesting admission are mixed Videoconference and VoIP sources.

In order to compare the performance of our framework against other existing proposals, we implemented the scheme described in [10] as Rate Envelope Multiplexing (REM), with adaptive weight factor and no histogram update. We call this MBAC (Measurement-based Admission Control).

The only available traffic descriptor for use in both MTAC and MBAC is the sources' peak rate p_{new}.

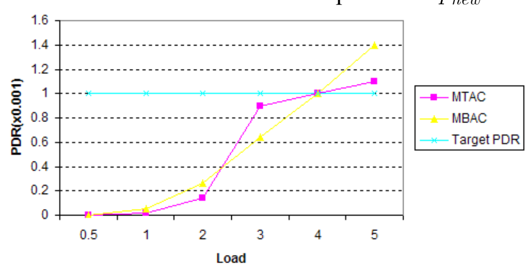

Figure 10. Achieved PDR for mixed traffic and target PDR 0.001.

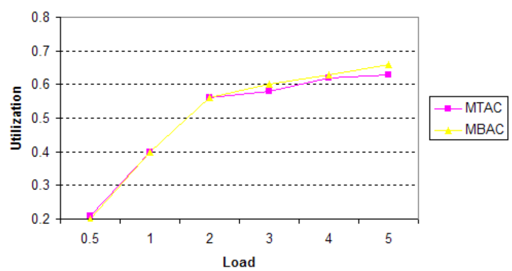

Figure 11. Utilization for mixed traffic and target PDR 0.001.

Both MTAC and MBAC achieve the target PDR 0.01 with MTAC achieving higher utilization. For target PDR 0.001 MTAC violates this PDR for load conditions more than 4 times the *design load.* MBAC also exhibits the same behavior.

5. Conclusions

In this paper we proposed an approach for bandwidth management that restricts the need for *fine-grained*, that is per-flow admission control in the core network. We showed that if a Diffserv network is engineered and provisioned using this method, *fine-grained* admission control need only be employed at the first points of aggregation and no further QoS degradation is expected in the core network as long as further downstream the aggregate streams are treated in a peak rate manner. Regarding admission control at the first aggregation points, we showed through simulation that the performance of our approach is satisfactory for traffic sources with diverse characteristics and bandwidth requirements.

Acknowledgment: This work was undertaken in the context of the Information Society Technologies Management of End-to-end Quality of Service Across the Internet at Large (MESCAL-IST-2001-37961) research project, which is partially funded by the Commission of the European Union.

6. References

[1] T.Lee and M.Zukerman "An Efficiency Study of Different Model-Based and Measurement-Based Connection Admission Control Techniques using Heterogeneous Traffic Sources", Proceedings of IEEE ATM 1999, May 1999.

[2] M.Grossglauser and D.Tse "A Framework for Robust Measurement-Based Admission Control", IEEE/ACM Transactions on Networking, June 1999, Vol. 7, No. 3, pp.293-309.

[3] L.Breslau, E.Knightly, S.Shenker, I.Stoica and Z.Zhang "Endpoint Admission Control: Architectural Issues and Performance", SIGCOMM 2000, Stockholm 2000.

[4] 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, Vol. 9, No. 7, pp.968-981.

[5] C.Chuah, L.Subramarian and R.Katz "Furies: A Scalable Framework for Traffic Policing and Admission Control", U.C Berkeley Technical Report No. UCB/CSD-01-1144, May 2001.

[6] 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, Vol. 36, No. 7, pp.834-843.

[7] T.Bonald, A.Proutiere and J.Roberts "Statistical Performance Guarantees for Streaming Flows using Expedited Forwarding", Proceedings of IEEE INFOCOM 2001, pp.1104-1112.

[8] S.Georgoulas, P.Trimintzios and G.Pavlou 'Joint Measurement- and Traffic Descriptor-based Admission Control at Real-Time Traffic Aggregation Points', To appear in ICC 2004.

[9] http://www-tkn.ee.tu-berlin.de/research/trace/trace.html

[10] T.Lee and M.Zukerman "Admission Control for Bursty Multimedia Traffic", Proceedings of IEEE INFOCOM 2001, pp.478-487.