

Joint Measurement- and Traffic Descriptor-based Admission Control at Real-Time Traffic Aggregation Points

Stylianos Georgoulas, Panos Trimintzios and George Pavlou
Centre for Communication Systems Research, University of Surrey
Guildford, Surrey, GU2 7XH, United Kingdom
e-mail: {S.Georgoulas, P.Trimintzios, [G.Pavlou](mailto:G.Pavlou@eim.surrey.ac.uk)}@eim.surrey.ac.uk

Abstract— The primary role of admission control is to decide on the amount of traffic accepted into the network so that users conforming to their established traffic contracts achieve predefined performance objectives, e.g. bounded packet loss probability, end-to-end delay. We consider Diffserv networks able to support real-time traffic and we propose a novel framework for admission control that involves both traffic descriptor and measurement-based techniques and we compare its performance with existing approaches. Our simulation results show that the performance of our approach is rather insensitive to variations of the traffic sources. Even when the provided traffic descriptors are as simple as a single value denoting the required peak rate, the proposed scheme achieves satisfactory performance.

I. INTRODUCTION

Real-time services have stringent delay and loss requirements. Differentiated services (Diffserv) is seen as the emerging technology to support IP Quality of Service (QoS) in a scalable fashion, without the relevant problems of Integrated services (Intserv), which require per micro-flow state information in the core network. In Diffserv, per flow state is only kept at edge routers while in the core network, traffic with similar QoS requirements is grouped in one of the engineered traffic classes and forwarded in an aggregate fashion. Admission control is the set of actions required at ingress nodes at the service instance establishment phase to check whether a service request is to be admitted or rejected. A new service instance is admitted when the requested QoS can be satisfied without causing QoS violation to the already established service instances. In Diffserv, the violation refers to the QoS perceived by each of the engineered aggregate traffic classes rather than the QoS of individual micro-flows. An additional role of admission control is to optimize the use of network resources in order to achieve realistic network efficiency. The challenge is to design simple control functions that improve efficiency under any offered traffic conditions.

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 coarsely divided into three categories: (a) Admission control using *a priori* traffic descriptors, (b) Measurement based admission control (MBAC), and (c) Endpoint admission control (EAC).

For the first approach [1], admission control is based on the assumption that it has perfect knowledge of each traffic source type that will be used in every link. It also knows the current number of established service instances. This

information will enable admission control to compute the total amount of bandwidth required. Hence, it will only accept a new service request if the minimum amount of bandwidth required by the total number of established service instances, including the new one, is less than the available service rate. It is obvious that this approach is optimal if the accurate traffic descriptors are known and used for admission control decisions. As such, the performance of this approach gives the upper bound for all other admission control approaches, provided of course that sources are conformant. However, since no traffic measurement is taken into consideration, if the provided traffic descriptors do not depict the actual behavior of the sources (for instance that could happen in case of non-conformant non-policed sources) or the appropriate traffic descriptors are not known *a priori* (for instance, only peak rate may be known), the performance of this admission control scheme can be very low.

The second approach (MBAC) tries to avoid the aforementioned problems by shifting the task of traffic specification from the user to the network [2]. 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 ([3] suggests that it is hard to police traffic according to a statistical model). 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.

In the third approach, the end host/application probes the network by sending test packets along the transmission path [4]. Based on some metric then the host decides if the flow can be admitted. Because these methods rely on potentially imprecise end-to-end measurements to guide their admission control decisions, endpoint admission control is primarily intended for soft real-time services, similar to Intserv Controlled Load or Diffserv qualitative services, in which the aggregate load is kept at reasonable levels but no hard guarantees are given to individual flows. Since the end-to-end route might not even be the same for probing packets and flows in high-priority classes, probing does not seem to be an

efficient method. Simultaneous probing by many sources can lead to a situation known as *thrashing* [4]. 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 that have a strict, usually small, delay requirement and a bounded, not necessarily too low, packet loss rate (PLR) requirement. Note that jitter is another QoS parameter for real-time traffic but we do not consider it in this work. In a Diffserv domain we assume that such real-time traffic is aggregated so that traffic from sources composing each traffic aggregate will receive the same treatment over the entire domain. The delay requirement of the traffic aggregate has been taken into account in the provisioning stage, i.e. by appropriately setting small packet 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 the serialization of the various traffic sources takes place. We assume that further downstream; traffic aggregates are treated in a peak rate manner. This is feasible since, as stated in [5], in a common network configuration, backbone links are over-provisioned.

This paper is organized as follows: Section II presents our admission control framework. In Section III we evaluate the performance of our framework against others found in the literature. Finally, in Section IV we conclude, summarizing our findings.

II. FRAMEWORK FOR ADMISSION CONTROL

In Diffserv networks, customer traffic contracts are pre-agreed through SLSeS [6] that include traffic descriptors. Our proposed approach is a combination of Measurement-based and *a priori* Traffic descriptor Admission Control -we will be referring to it as MTAC. As such, it encompasses the positive features of both admission control schemes. It provides a systematic way to derive the bandwidth requirements of the already established flows through measurements and of the candidate for admission flow through its traffic descriptors. We use real-time measurements of the actual load in order to cope with the fact that the traffic descriptors may not depict the actual characteristics of individual flows. In this case, we might end up with either under-utilization, if users overestimate their requirements, or with QoS degradation if users underestimate their requirements and the network is not able to police them efficiently. For bandwidth manipulation and allocation we have adopted the effective bandwidth approach. According to [7], when the effect of statistical multiplexing is significant, the distribution of the stationary bit rate can be accurately approximated by a Gaussian distribution. Regarding the validity of this assumption, in [8] 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. In [9], it is suggested that a number of aggregated sources as low as ten is enough for the Gaussian assumption to hold. In that case, the effective bandwidth of the multiplexed sources is given by:

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

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.

A. Admission control logic

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}) and average duration h_{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 PLR. In our model, we assume that every time a source wants to establish or terminate a service instance, it signals this to the ingress node through a resource reservation protocol; as such, the number of active sources at every point in time is always known.

When a new request arrives, measurements are taken for bandwidth 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 (first aggregation point). We then calculate the number N_m of the reference sources, whose aggregate mean rate is equal to or greater than $M_{measured}$, and the number N_σ of the reference sources, whose aggregate variance is equal to or greater than $\sigma_{measured}^2$. That is, N_m represents the number of active reference sources that would have produced load with average rate at least equal to $M_{measured}$, and N_σ represents the number of active reference sources that would have produced load with variance at least equal to $\sigma_{measured}^2$. They are the integers that satisfy the following relationships:

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

Having estimated N_m and N_σ , we compute their mean value N_{ref} , which represents an estimate of the number of reference sources that produce load with characteristics similar to the ones having been measured.

$$N_{ref} = (N_m + N_\sigma) / 2 \quad (3)$$

The new source requesting admission has declared in the SLS its traffic descriptors (m_{new}, σ_{new}) . In case the only available traffic descriptor is the new source's peak rate p_{new} , the above pair becomes $(m_{new}, \sigma_{new}) = (p_{new}, 0)$. In all cases, having the measurements and the traffic descriptors of the new source, we compute the estimated bandwidth C_{est} as follows:

$$C_{est} = M_{measured} + m_{new} + a'_{PLR} \sqrt{\sigma_{measured}^2 + \sigma_{new}^2} \quad (4)$$

where a'_{PLR} is computed as in (1), based on the target PLR bound of the real-time traffic aggregate. This value C_{est} will be used in the admission control criterion.

B. The Precautious Factor (PF)

Before deriving the admission control criterion, there are two important issues that need to be taken into account. The first one is that the admission control decision needs to be more conservative as the current measurements of the offered load correspond to a number of reference sources N_{ref} that exceeds the number of reference trunks T_{ref} to avoid degradation of performance due to excessive PLR incurred by the overloading of the system.

The second important issue is that the more stringent the target bound on PLR, the more conservative the admission control decision should be. To demonstrate why that should hold, one can consider the following example: the bandwidth that is required by 100 VoIP sources, with peak rate 64kbps and exponentially distributed ON and OFF periods with average durations 1.004sec and 1.587sec, with a bound on queue overflow probability equal to 0.01, using (1), is 3.22608 Mbps. According to (1), the bandwidth required for 101 sources of that type for queue overflow probability 0.01 is 3.355630 Mbps. That means that the additional source gives an increase of bandwidth, $\Delta_{eff}(0.01) = 29.022\text{kbps}$. For queue overflow probability 0.001, the corresponding increase is $\Delta_{eff}(0.001) = 30.182\text{kbps}$. This means that since for decreasing bound on queue overflow probability (which directly translates to a decreasing bound on PLR) the increase in bandwidth requirements for each admitted source is greater, it should be taken into account in the admission control decision.

When we decide whether a new source should be admitted or not, the two issues discussed above need to be taken into account. Therefore we introduce a *Precautious Factor (PF)* before we derive the final expression for our admission control criterion. The more conservative the admission control decision, the greater the value of the precautionous factor.

In order to take into account the first issue, PF should be proportional to the quantity (N_{ref} / T_{ref}) . In order to take into account the second issue, we proceed as follows: given (1), for two different levels of PLR, e.g. ε_1 and ε_2 with $\varepsilon_1 < \varepsilon_2$ and all other parameters (number and characteristics of sources) the same, it can be suggested that this relative increase in additional bandwidth required is proportional to the quantity $\frac{\sqrt{-2 \ln(\varepsilon_1) - \ln(2\pi)}}{\sqrt{-2 \ln(\varepsilon_2) - \ln(2\pi)}}$.

Therefore, the final expression for the precautionous factor that is adopted is:

$$PF = (N_{ref} / T_{ref}) * \frac{\sqrt{-2 \ln(\varepsilon) - \ln(2\pi)}}{\sqrt{-2 \ln(\varepsilon_{ref}) - \ln(2\pi)}} \quad (5)$$

In both expressions, ε_{ref} is a reference PLR level. Also, when it is $N_{ref} < T_{ref}$ then the value $PF = 1$ is adopted.

C. The measurement window

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 similar fashion to [10], we use the following expression for the measurement window:

$$w = \max(DTS, w'), \quad (6)$$

In (6), DTS represents the Dominant Time Scale. DTS is the most probable time scale over which overflow occurs. In [8], a systematic way to derive DTS using real-time measurements is provided with the assumption that the input process to the multiplexing point in the network is Gaussian, which is by definition our assumption when employing (1), and we use this method for estimating DTS. w' represents the mean inter-departure delay [2], and it is defined as follows:

$$w' = \frac{h_{ref}}{N_{active}}, \quad (7)$$

where N_{active} is the number of simultaneously active sources.

We select as measurement window the mean inter-departure delay, that is the time interval within which the system can be considered stationary -no flow departures- unless this time interval is not long enough to capture the packet level dynamics of the aggregate traffic stream. In that case we use DTS as the value of the measurement window.

D. The Admission Control Criterion

Given the allocated bandwidth for the real-time traffic aggregate from edge-to-edge is C_{total} , and having computed the value for C_{est} , employing the precautionous factor and the measurement window, the admission control criterion becomes:

$$\begin{aligned} \text{If } (C_{est} \times PF) \leq C_{total}, & \quad \text{admit} \\ \text{If } (C_{est} \times PF) > C_{total}, & \quad \text{reject} \end{aligned} \quad (8)$$

III. PERFORMANCE EVALUATION

In order to evaluate the performance of the aforementioned admission control framework we used simulations. In our experiments we use the dumbbell topology of Fig. 1.




Figure 1. Simulation topology.

The real-time traffic QoS requirements are as follows. To meet the target bound on the delay of the aggregate, we configure the output queue to hold a maximum of 5 packets, thus enforcing that the upper bound on packet delay will be no more than 5 times the average packet size to the output link capacity. We use scenarios with the target bound on packet loss 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 that the system is designed to handle for a desired service rejection probability of 0.01, if we have employed a circuit switched approach.

The remaining details of the simulation setup are: ϵ_{ref} is set equal to 0.01, the source active time is exponentially distributed with average h_{ref} equal to 300sec. The source activation rate is a Poisson arrival process with average 1000 source activations/hour, which for a service rejection probability equal to 0.01 corresponds to T_{ref} equal to 100.

For traffic sources we use three different models of VoIP traffic. All of them are ON/OFF sources with exponentially distributed ON and OFF times, having a peak rate of 64kbps. The first model [11], has mean duration for the ON and OFF periods of 1.004sec and 1.587sec respectively. This model is also used as the *reference source* model and thus sources that follow this model are considered as conformant. The second model we use is the Brady model [12], which has mean duration for the ON and OFF periods of 1sec and 1.35sec respectively. Sources that follow this model have *worse* characteristics than the reference sources; therefore they are characterized as non-conformant. In the third model [13], the mean durations for the ON and OFF periods are 0.352sec and 0.65sec respectively. Sources that follow this model have better characteristics than the reference sources. The term *mild* sources will refer to this last category of sources.

In order to compare the performance of our framework against other existing proposals, we have implemented two algorithms from the literature. The first one is an admission control scheme using *a priori* traffic descriptors. We call this algorithm TDAC (Traffic Descriptor-based Admission Control). In this scheme, only the declared traffic descriptors are used for bandwidth estimations and admission control. The second algorithm is an implementation of the scheme described in [14] as Rate Envelope Multiplexing (REM), with adaptive weight factor (the packet loss measurement approach) and no histogram update. We call this algorithm MBAC (Measurement-based Admission Control). In this scheme, in addition to measurements, the peak rate of the new candidate for admission flow is used for admission control. Regarding the time scales involved in MBAC, we consider the values 0.05sec, 0.1sec, 0.5sec, 1sec and 2sec. We also adopt the value 0.05sec as the minimum sampling interval involved in our approach in order to derive measurements for the mean rate, the variance of the offered load and additionally estimate the DTS. That means that both the inter-departure delay and

the computed DTS involved in the expression for the measurement window in (6) are considered as integer multiples of that value.

A. Sources that declare their real (m, σ)

In this case the network has been provisioned and engineered using the *reference source* VoIP model. However, the invoked sources may deviate from this model, but they declare their real traffic descriptors. That could correspond to a situation where an ISP has provisioned and dimensioned its network using a hypothetical *reference source* model and has to serve requests that need to send VoIP traffic with different characteristics. We compare the performance of MTAC against TDAC (we don't consider MBAC in this case because it cannot take advantage of other traffic descriptors apart from the peak rate and we want to compare our algorithm with approaches that can take full advantage of the traffic descriptors used by our approach) for all three different types of traffic sources and the simulation results are shown in Fig. 2, 3, 4 and 5, where M, R, W in brackets mean that the relevant curve corresponds to *mild*, *reference* and *worse* (non-conformant) sources respectively.




Figure 2. Achieved PLR for target PLR 0.01.




Figure 3. Utilization for target PLR 0.01.

For target PLR 0.01 and for all traffic sources, MTAC performs reasonably well compared to TDAC, which provides the upper bound in performance. For *mild* sources the target PLR is not violated in any case. For *reference* and *worse* sources, the target PLR is violated only for loading conditions more than 3 times the *design load*, which, in a real network configuration, would correspond to a situation where the network dimensioning and provisioning are inadequate. Additionally, it should be noted that MTAC is prone to measurement errors that affect the measured entities, especially the variance, and the estimation of DTS, since the

computed value for DTS depends on the measured values for mean rate and variance of the offered load. This partly justifies MTAC's performance deficit.

Regarding utilization, MTAC's performance is also satisfactory; taking into account that TDAC provides the upper bound in performance with respect to both the achieved PLR and utilization.




Figure 4. Achieved PLR for target PLR 0.001.




Figure 5. Utilization for target PLR 0.001.

Similar conclusions can be drawn for the case of target PLR 0.001.

B. Peak rate admission control

In this case the network has been provisioned and engineered using the *reference source* VoIP model. However, the invoked sources may deviate from this model and, additionally, the only traffic descriptor they declare is their peak rate p_{new} . We compare the performance of MTAC against MBAC. We don't consider TDAC in this case because its overall performance regarding utilization would be extremely low as a result of the stringent admission control criterion $\sum p_i + p_{new} \leq C_{total}$. The simulation results are shown in Fig. 6, 7, 8 and 9.




Figure 6. Achieved PLR for target PLR 0.01.




Figure 7. Utilization for target PLR 0.01.

MBAC achieves PLR much lower than the 0.01 PLR target for all traffic sources and all loading conditions. For *mild* sources, MTAC achieves the PLR target for all loading conditions. For *reference* and *worse* sources, the target PLR is violated only for loading conditions more than 3 times the *design load*. Regarding utilization, however, MTAC outperforms MBAC in all cases. MBAC's lower utilization is the penalty for achieving PLR significantly lower than the target PLR.




Figure 8. Achieved PLR for target PLR 0.001.




Figure 9. Utilization for target PLR 0.001.

The performance of MBAC is very low for a target PLR equal to 0.001. It violates that target in most cases even though the no histogram update method that we use in our implementation of MBAC, is the more conservative approach among all the other variations presented in [14]. This behavior indicates that for different bounds of the target PLR, some parameters involved in MBAC need to be reconfigured in an ad-hoc fashion. MTAC achieves similar performance to the

case we had target PLR 0.01, without the need for any readjustments of the admission control criterion.

MBAC achieves better utilization than MTAC, but this cannot be considered as an advantage, because the increased utilization is coupled with a violation of the target PLR. This violation nullifies the significance of the increased utilization.

In all cases examined, for both MTAC and MBAC we observe an increase in the achieved PLR for increasing loading conditions. This is anticipated [2] because since they both rely on measurements, every new admission request carries the potential of making a wrong decision. That means that a high source activation rate is expected to have a negative effect on performance.

In the case where only the peak rate is available, MTAC achieves PLR that is slightly decreased compared to the case where the appropriate traffic descriptors are available. This happens because the admission control criterion in the former case is more stringent - in (4) we actually add the peak rate of the new source. However, there is still a violation of target PLR for increasing load conditions, which is partly due to measurement and estimation errors. The utilization is slightly decreased, as a result of the more stringent admission control criterion.

IV. CONCLUSIONS

In this paper we proposed an approach for admission control at the first points of aggregation. This approach is a combination of traffic descriptor- and measurement-based admission control. We showed through simulation that our approach is relatively insensitive to variations of the traffic sources from the traffic source model used for network dimensioning and provisioning, and that it works reasonably well even when the only available traffic descriptor is the peak rate. Our proposed framework can almost match the performance of approaches using *a priori* traffic descriptors and can outperform existing measurement-based approaches for similar simulation setup. In the future we will be looking to extend our admission control framework for different types of real-time sources and for elastic traffic.

V. 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] J.Hui "Resource Allocation for Broadband Networks", IEEE Journal on Selected Areas of Communications, December 1988, Vol. 6, No. 9, pp.1598-1608.
- [4] L.Breslau, E.Knightly, S.Shenker, I.Stoica, and Z.Zhang "Endpoint Admission Control: Architectural Issues and Performance", SIGCOMM 2000, Stockholm 2000.
- [5] G.Iannaccone, M.May, and C.Diot "Aggregate Traffic Performance with Active Queue Management and Drop from Tail", Computer Communications Review, July 2001.
- [6] E.Mykoniati, C.Charalampous, P.Georgatsos, T.Damilatis, D.Goderis, P.Trimintzios, G.Pavlou, and D.Griffin "Admission Control for Providing QoS in Diffserv IP Networks: The TEQUILA Approach",

IEEE Communications Magazine, January 2003, Vol. 41, No. 1, pp. 38-44.

- [7] 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.
- [8] 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, Vol. 11, No. 2, pp. 222-235.
- [9] R.Guerin, and L.Gun "A Unified Approach to Bandwidth Allocation and Access Control in Fast Packet-Switched Networks", INFOCOM 1992, Vol. 1, pp. 1-12.
- [10] S.Belenki "An Enforced Inter-Admission Delay Performance-Driven Connection Admission Control Algorithm", ACM SIGCOMM, April 2002, Vol. 32, No. 2, pp. 31-41.
- [11] 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.
- [12] P.Brady "A Model for Generating On-Off Speech Patterns in Two-Way Conversation", The Bell System Technical Journal, September 1969, pp. 2445-2471
- [13] I.Habib and T.Saadawi "Multimedia Traffic Characteristics in Broadband Networks", IEEE Communications Magazine, page 48, July 1992
- [14] T.Lee and M.Zukerman "Admission Control for Bursty Multimedia Traffic", Proceedings of IEEE INFOCOM 2001, pp.478-487.