MBAC algorithm for streaming flows in Cross-protect

A. Kortebi, S. Oueslati and J. Roberts
France Telecom R&D

May 24, 2004

Abstract

Measurement based admission control (MBAC) algorithms constitute an attractive way to provide streaming flows with statistical QoS guarantees, namely controlled jitter and loss. The Cross-protect router architecture requires MBAC to be performed with minimal information on flow traffic. In this paper, we propose a simple and efficient MBAC approach based on a fluid model that assumes a Poisson distribution of the number of active bursts. New flows are admitted if the measured aggregate load is less than a predefined threshold, which is determined using a minimal traffic descriptor: a common maximum peak rate for all flows. The algorithm is evaluated by means of ns simulations, in order to derive packet-level performance, considering a variety of traffic scenarios including video and real traces. Our results show that high link utilization can be achieved while maintaining low packet delay and loss rate.

1 Introduction

Flow-aware networking (FAN), based on the Cross-protect router mechanisms, allows performance guarantees for streaming and elastic flows while preserving the user-network interface of the best effort Internet [18], [21]. In this architecture user-defined flows are identified on the fly. Per-flow scheduling using Priority Fair Queueing (PFQ) ensures max-min bandwidth sharing and gives head of line priority to packets of flows emitting at a rate less than the current fair rate.

Rather than distinguishing explicitly between streaming flows, requiring negligible packet loss and delay, and elastic flows, requiring 'fast as possible' document transfer, FAN distinguishes flows that are under or over the fair rate. Admission control is employed to maintain the fair rate sufficiently high to provide streaming like quality for flows of peak rate up to a chosen threshold, i.e., the targeted streaming flows are always in the 'under' category. This admission control is necessarily measurement-based since there is no prior knowledge of flow traffic characteristics.

Admission control is also necessary to keep the fair rate high even in situations of overload. Previous work on the definition of admission control for elastic traffic [1] has clarified that the target threshold should be as high as 1% of the link capacity, even though this may be much higher than what most users would consider acceptable. Given a reliable measurement of the current fair rate (provided by the PFQ scheduler), it is relatively easy to meet this objective. Moreover, as long as the elastic flows constitute a significant part of overall traffic, the same criterion automatically preserves the quality of streaming flows whose packets are handled with priority. This is because the load due to streaming traffic is then low enough to ensure negligible packet loss and delay.

With Cross-protect, some elastic flows also enter into the 'under' category as their peak rate (i.e., their rate limitation due to causes other than the considered link) is less than the current fair rate. Indeed, in very high speed backbone links, the vast majority of flows are in this category since the user's access rate (1 Mbit/s, say) is much less than the limiting fair rate of 1% of link capacity (25 Mbit/s on an OC 48, say). It is necessary in this case to apply an admission control algorithm capable of ensuring negligible packet loss and delay for streaming flows when they are indistinguishable from the bulk of elastic traffic. The objective of this paper is to discuss the definition of such an algorithm.

We first state the nature of the problem, then we present briefly some existing MBAC algorithms. Our approach is described in section 4, simulation results are given in section 5 and finally we conclude.

2 Problem statement

Admission control must be based on very limited information about the traffic characteristics of individual flows. The only fact we know is that their peak rate is less than a given threshold p (some flows may have a peak rate greater than p but, in the event of a rate overload, their packets will be spaced by the scheduler so that they have no impact on the delay of the considered class of streaming flows (*to be verified*). The number of flows in progress is unknown though an estimate could be obtained from the protected flow list (or lists). Flow departures are not signalled and are only deduced after the expiry of a time-out allows their removal from the protected flow list. Flow arrivals are known locally to each admission controller but, typically, the admission control function is distributed over several line cards [18].

While the available information is very limited, the considered context is highly favourable to efficient statistical multiplexing. By assumption, the link rate is much higher than the maximum peak rate. The latter is necessarily less than the 1% (or so) fair rate threshold and typically even smaller (e.g., 1 Mbit/s DSL access lines sharing a 2.5 Gbit/s OC 48 backbone link). High efficiency is thus possible while maintaining excellent per flow performance. In addition, the economics of backbone provision are such that high efficiency is generally not a necessity. A target average utilization at 'saturation' (i.e., when new flows would be subject to rejection) of 80% may be perfectly acceptable (*to be verified*).

The measurements provided by the PFQ scheduler are load counts in successive short time intervals. We assume there is no limit on the length of this interval. We also assume there is no processing limit on the number of successive measurements taken into account or on the complexity of the functions derived from them (e.g., means, maxima, higher moments). However, the algorithm evaluated in this preliminary study relies just on a running estimate of the mean load per time interval.

3 Existing MBAC approaches

Many MBAC algorithms have been proposed in the literature (e.g., [10], [12], [15], [20], [13], [5], [19], [6], [7]). In the present section, we briefly review some of these algorithms with respect to their applicability to the FAN architecture.

It is first useful to distinguish between bufferless and buffered multiplexing. Buffered multiplexing is performed when a large buffer is used to absorb a prolonged burst level arrival rate excess. A major problem with this kind of multiplexing is that performance depends on detailed traffic characteristics which are unknown a priori.

Bufferless multiplexing is simpler to control [22]. The combined rate of ongoing flows is maintained less than link capacity with high probability by admission control. A small buffer is used to absorb simultaneous packet arrivals from two or more different flows. A major advantage of this scheme is that the loss probability (which can be used as QoS target) depends only on the stationary distribution of the input rate. Bufferless multiplexing is efficient when the maximum flow peak rate is less than around 1% of link capacity otherwise.

Cross-protect essentially realizes bufferless multiplexing for flows of peak rate less than the fair rate. In the following, we discuss briefly the suitability for our purposes of some of the most well-known MBAC algorithms.

Jamin et al [16] proposed a simple MBAC algorithm called *measured sum (MS)*. A new flow is admitted if the sum of its nominal rate (r) and the estimated rate of aggregate flows ($\widehat{v}$) is less than a utilization target (u) times the link bandwidth (C), i.e., when:

$$\widehat{v} + r \leq uC \tag{1}$$

The estimated rate of aggregate flows is derived using a time window estimator. The performance of this algorithm depends on the choice of the window estimator length; Casetti et al. [5] proposed an algorithm to adapt the measurement window length. This algorithm is clearly very simple and takes no account of the nature of flow traffic.

Gibbens et al. [10] propose a decision theoretic approach, where a new flow is admitted if the current aggregate load is less than an appropriate threshold. The flow peak rate is the only flow information taken into account in the simplest method proposed (a more complex model also requires the number of flows in progress). This method thus corresponds with present requirements and presented results provide useful confirmation that efficient multiplexing is possible even in this restricted framework. Unfortunately, to compute the admission threshold requires prior assumptions about the level of offered load and flow burstiness.

Gibbens and Kelly [9] present a family of algorithms based on tangents to an equivalent bandwidth curve computed from the Chernoff Bounds. Different choices of tangent imply knowledge of different traffic characteristics. Flows are assumed to belong to a given number of classes. The least complex approach requires knowledge of the instantaneous combined load (as in [10]) together with the number of flows of each class and their peak rate. This method derives from the use of Hoeffding bounds proposed by Floyd [8]. The alternatives identified in [9] all require additional information that is not available in Cross-protect.

The methods proposed by Gibbens at al [10], Gibbens and Kelly [9] and Floyd [8] all propose to use a "back-off" strategy: if a flow is blocked, no further flow is accepted until one of the flows in progress finishes. This avoids the problem of multiple new flows being admitted in heavy traffic as the result of a single low utilization measurement. Unfortunately, this device is not possible with Cross-protect since flow terminations are not explicitly known.

Tse and Grossglauser take account of the impact of measurement errors on the performance of an MBAC algorithm [11], [12]. In their earlier work, it is assumed that individual flow characteristics (rate mean and variance) are measured. The second paper [12] provides methods that are closer to our present requirements being based on overall measured traffic. Their method is based on a Gaussian approximation of the aggregate bandwidth. The admission condition is derived to satisfy a target loss probability assuming bufferless multiplexing. They introduce the notion of critical time scale ($\tilde{T_h}$) broadly equal to the time scale over which the impact of an admission decision persists. They establish that $\tilde{T_h} = \frac{T_h}{\sqrt{n}}$ where T_h is the average flow holding time and n is the maximum number of active flows the link can carry.

A new flow is admitted if the following condition is met:

$$C - (N_t + 1)\widehat{\mu}_t > \alpha_q \widehat{\sigma}_t^H \sqrt{N_t + 1} \tag{2}$$

where C is the link capacity, N_t is the number of flows in the system at time t, $\widehat{\mu}_t$ is the measured per flow rate, $\alpha_q = Q^{-1}(\epsilon)$ with ϵ is the target loss probability, $Q(.)$ is the complementary cdf of a $N(0,1)$ Gaussian random variable and $\widehat{\sigma}_t^H$ is the estimated standard deviation of per flow rate. We cannot use this method notably because the number of flows in progress N_t is not known in Cross-protect.

Grossglauser and Tse propose an alternative method when the number of flows is unknown. A new flow is admitted if:

$$C - A_t^\lambda - p > \alpha_q \widehat{\sigma}_t^{AH} \tag{3}$$

where A_t^λ is the measured aggregate load, p is the flow peak rate and $\widehat{\sigma}_t^{AH}$ is the measured standard deviation of the aggregate load. When a flow is admitted its peak rate is added to A_t^λ. This prevents momentary overloads under heavy traffic (their assumption) due to flow admissions following a low estimate A_t^λ in some period t. The instants of flow arrival may only be known partially in Cross-protect since admission control is implemented in a distributed way over multiple line cards.

Despite the obvious differences between the different MBACs proposed, it turns out that they all seem to result in the same trade-off between utilization and perceived per-flow performance. This is illustrated in [4] and [24] where performance is measured in terms of a packet loss rate. The algorithms differ in terms of their predictive capability: the choice of parameter values for the admission condition resulting in a given performance target. In fact, [4] shows that all algorithms tested are quite poor at predicting performance which depends significantly on traffic characteristics not explicitly taken into account.

4 An MBAC for Cross-protect

As described in [18], Cross-protect admission control has a dual role: maintain the fair rate sufficiently high in overload to maintain the efficiency of elastic traffic and ensure that streaming flows of peak rate less than p experience negligible packet loss and delay. In this paper we consider the second objective. We seek to define an algorithm using just the measured load due to 'under' traffic and the assumption that this traffic is due to flows whose peak rate is less than a certain threshold p. The fair rate available to backlogged flows is measured over a time scale of hundreds of milliseconds or seconds and is not directly useful to control packet delays that are significant at a much smaller time scale. This section presents our preliminary research.

4.1 Traffic model

We assume traffic constitutes a stationary process constituted by user sessions initiated according to a Poisson process. Each session is a finite alternating succession of flows and 'think times'. All we know about the arrival process of packets within a flow is that the peak rate is less than p bits/s. We further assume packets are of constant size equal to the maximum transfer unit, MTU bits. This simplification constitutes a worst case for multiplexing performance (see [3], *to be verified*).

The arrival process for any flow may be viewed as an on/off process with bursts of MTU bits occurring at rate P, for any $P \geq p$, with a minimal inter-burst interval (i.e., the interval between two successive leading edges) of MTU/p seconds. Under the assumption of Poisson session arrivals, in the absence of admission control, the number of such bursts in progress has a Poisson distribution. To see this, consider a session as a customer in a stochastic network with two infinite servers. The customer visits the first server to emit each burst with alternating visits to the second server for the duration of the inter-burst interval. It is known that the number of customers in each station of this network is as if they received Poisson arrivals for whatever traffic characteristics relating to flow rates and durations [17].

In particular, the number of active bursts, equal to the number of packets arriving in an interval of length MTU/P, is Poisson. The mean number of arrivals is $\lambda_p \mathrm{MTU}/P$ where λ_p is the packet arrival rate.

The above reasoning is valid in the absence of admission control. To simplify, we assume the distribution of the number of packets in an interval is also Poisson when admission control is applied. In other words, we suppose the admission control thins the arrival process without changing its main stochastic property (*needs justifying*).

The admission control algorithm we evaluate below aims to ensure the probability of an arrival excess in an interval of length MTU/p is less than some low target probability ϵ. For instance, assuming the peak rate p is less than 1% of the link rate C, the number of packets at saturation is at least 100 (i.e., C/p) and we would seek to limit load so that Pr[more than 100 packets]$< \epsilon$.

We perform exponential smoothing to derive the load estimate from traffic measured in successive intervals:

$$\text{new } = \alpha \times \text{ old } + (1-\alpha) \times \text{ new measurement.}$$

The admission decision is to reject new flows whenever the estimated load is greater than a threshold depending on the values of C/p and ϵ. Table 1 indicates some typical values. Note that the admission decision typically changes rapidly from one interval to the next, particularly when the actual peak rate of flows is much less than the assumed maximum value p. The average number of new flows arriving in each interval is equal to the average number of packets (less than 100 in the example above) divided by the average number of packets per flow. If the result of this division is too high (leading to large bursts of flow rejects and unstable oscillatory performance) it may be preferable to reduce the size of the interval used to compute the load estimate (i.e., increase the value of P).

C/p	100	100	1000	1000
ϵ	0.01	0.001	0.01	0.001
Threshold	0.79	0.73	0.93	0.91

Table 1: Admission thresholds

5 Simulation results

We have performed a large number of simulation experiments using either hypothetical traffic composed of homogeneous on-off flows or traffic derived from video and Internet traces. The simulated topology illustrated by figure 1 consists of 4 sources, 4 destinations and a bottleneck link where admission control is performed. All flow arrivals are Poisson.

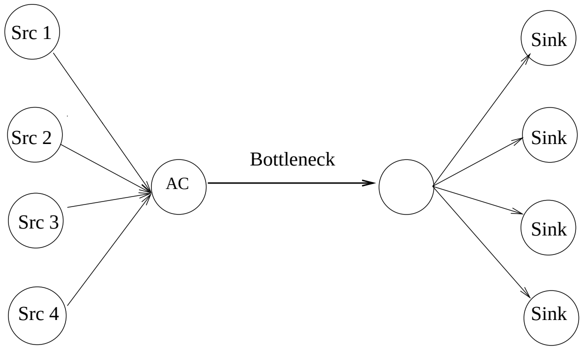

Figure 1: Simulated network topology

5.1 On-off traffic

All flows are of on-off type with a common peak rate and the same exponential distribution for both silence and burst lengths with mean 500 ms. Packets are of constant size 1000 bytes. The link rate is 10 Mbit/s and the peak rate is either 64 Kbit/s or 32 Kbit/s. Flows have an exponential duration with mean 60 seconds.

We first apply the method described in Section 4 assuming a peak rate of 64 Kbit/s and use a measurement interval of 125 ms (8000/64K). Load is averaged using filter parameter α of 0.9. The admission control blocks new flows in the next interval whenever the current load estimate exceeds 0.88 (corresponding to an overflow probability of $5*10^{-2}$).

Figure 2 shows the empirical histogram of the number of packets arriving in one interval when the offered load is 0.8 for traffic with peak rate 64 Kbit/s and 32 Kbit/s, respectively. In this case the probability of blocking is very small. Both histograms closely approximate a Poisson distribution.

Figure 3 shows corresponding results when the load is equal to 1.1. In this case there is a high degree of blocking (approximately 20 % in both cases). The distribution is again similar in form for the two cases with a tendency for slightly higher packet counts for the 32 Kbit/s traffic (due in part to some imprecision in the interval length in the program).

Figures 4 and 5 show how realized performance at packet level depends on link utilization. In Figure 4 the probability of overflow (a number of packets greater than the available capacity of 156 packets per 125 ms interval) is shown to be closely controlled even when the utilization is high (corresponding to a significant overload). Packet delay, represented in Figure 5 by the probability of delay exceeding 3ms, is roughly the same for the two peak rates.

As illustrated by Figure 6, blocking depends on offered load rather than detailed traffic characteristics.

Figure 7 shows results for 32 Kbit/s traffic when the interval length is calculated for this peak rate (equal to 250 ms) and admission threshold is 0.91 corresponding to $\epsilon = 5*10^{-2}$. We conclude that higher utilization can be achieved inducing an increase in packet delays.

5.2 Video trace traffic

Flows are generated according to video traces from [23] coded in H263 VBR (variable bit rate).

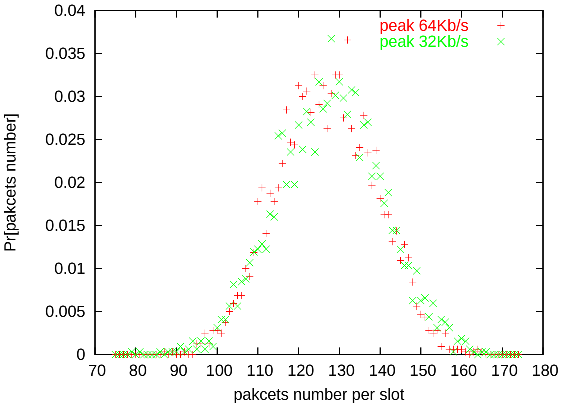

Figure 2: Packets number distribution, $\rho = 0.8$

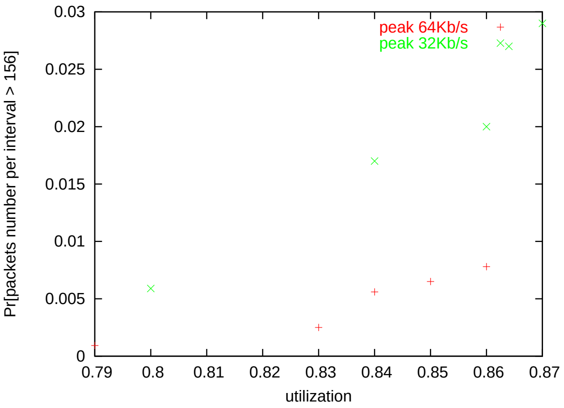

Figure 4: Overflow probability vs. utilization

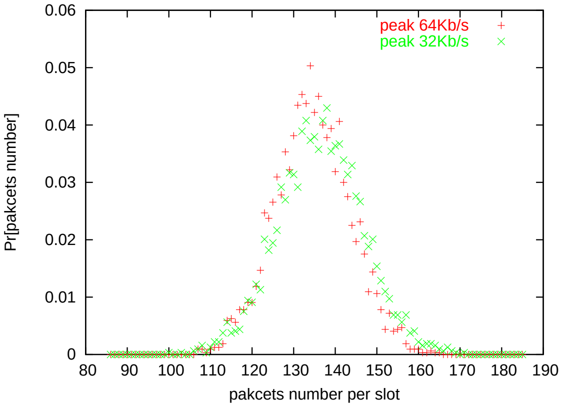

Figure 3: Packets number distribution, $\rho = 1.1$

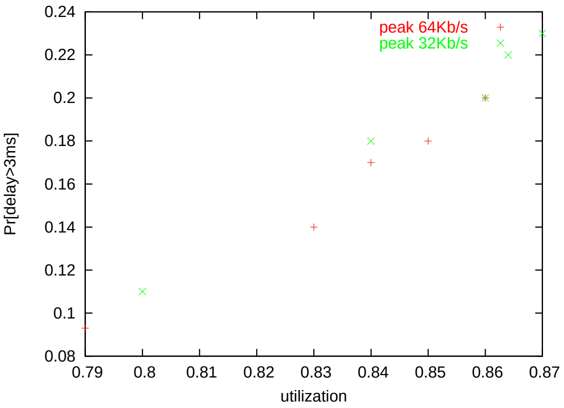

Figure 5: Delays vs. utilization, admission threshold $= 0.88$

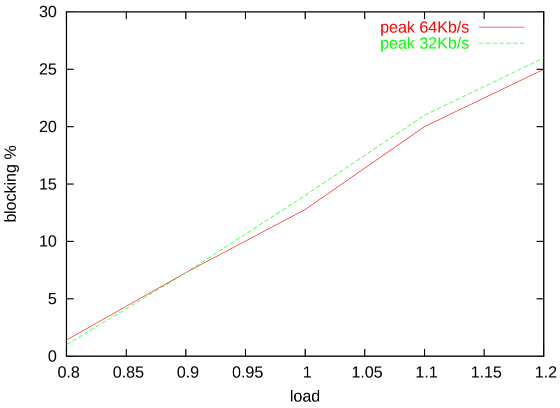

Figure 6: Blocking vs. offered load

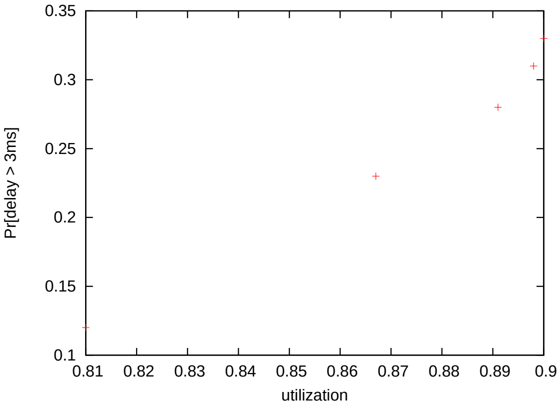

Figure 7: Delays vs. utilization, admission threshold = 0.91

Bottleneck capacity is 320 Mb/s and buffer size 3200 packets. We used four video traces in order to have heterogeneous traffic with different peak rates:

1. Formula 1 video with peak rate = 2.8 Mb/s, mean rate = 533 Kb/s and 30 minutes duration.
2. Star Wars video with peak rate = 1.3 Mb/s, mean rate = 115 Kb/s and 45 minutes duration.
3. Star Trek video with peak rate = 2.2 Mb/s, mean rate = 170 Kb/s and 60 minutes duration
4. South Park video with peak rate = 2.5 Mb/s, mean rate = 310 Kb/s and 20 minutes duration.

Each new flow begins at a random point from a video trace and lasts for 1 minute on average. The traces consist in a list of frame sizes in bytes. We convert this to packet traffic with packets of maximum size 1000 bytes. When a frame length is greater than 1000 bytes, it is split into different packets by the network layer. We generate these packets with either a constant peak rate or with a variable bit rate resulting in two types of video traffic:

1. type-1: packets of fragmented frames are generated with peak rate equal to 2.8 Mb/s.
2. type-2: packets are generated with a variable bit rate equal to (frame length)/(interframe time). In this case the maximum peak rate considered is 2.8 Mb/s corresponding to the F1 video.

The admission threshold was set to 0.8 (corresponding to $\epsilon = 10^{-2}$) and we simulate an offered load of $\rho = 1$. For type-1 videos, the load measurement interval was $cl = 2.5ms$; while for type-2 videos, $cl = 2.8ms$.

In both cases, without admission control streaming packets suffer high delays and a high proportion of them are lost. Admission control restores good performance. Type-1 video flows experience higher blocking than type-2. Performance indicators are reported in Tables 2 and 3.

5.3 Internet trace traffic

We performed simulations using real traces derived from ADSL data collected on an operational 155

Performance indicators	Poisson	no CAC
Utilization	0.74	0.99
Average delay (ms)	0.26	38.1
Maximum delay (ms)	2.1	64.9
Blocking	24%	-

Table 2: ρ=1, video type-1 traffic

Performance indicators	Poisson	no CAC
Utilization	0.78	1
Average delay (ms)	0.14	38.2
Maximum delay (ms)	1.5	66
Blocking	20%	-

Table 3: ρ=1, video type-2 traffic

Mb/s download link shared by several hundred ADSL subscribers with peak rate of 512 or 1024 Kbit/s. Elastic flows included in the trace are constrained by their access rates; therefore they behave like streaming flows in that their rate never exceeds a small fraction of the link capacity. We used four traces to generate a load of 1.05 over 100 seconds simulation duration. Buffer size was set to 1550 packets.

The admission threshold was set to 0.83 based on a peak rate of 1 Mb/s (corresponding to $\epsilon = 10^{-2}$). The load was measured over intervals of length $cl = 10$ ms. When a flow is blocked, we remove all its packets from the trace.

As expected, without admission control performance is very bad. Packets undergo significant delays (34 ms in average) and losses (8%). Applying admission control ensures negligible delays (only 0.1% of packets suffer more than 10ms delay) and no losses while attaining 82% link utilization as plotted in Figure 8. Performance indicators are reported in Table 4.

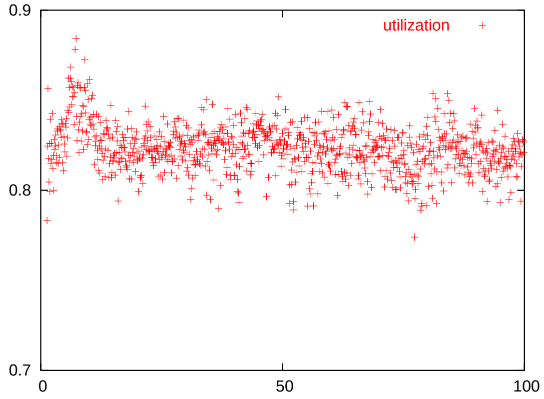

Figure 8: $\rho = 1.05$, real traces traffic with admission control, utilization variations vs. time

Performance indicators	Poisson MBAC	no CAC
Utilization	0.82	1
Average delay (ms)	0.28	34
Maximum delay (ms)	16	38
Blocking	27%	-

Table 4: ρ=1.05, Real traces traffic, capacity = 155 Mb/s

6 Conclusion

Flow-aware networking, or FAN, based on the Cross-protect mechanisms allows appropriate service guarantees for streaming and elastic traffic without the complication of class of service discrimination. It remains, however, to design and calibrate the necessary measurement-based admission control algorithms. Previous work suggests a simple algorithm based an the estimated fair rate is sufficient when a significant proportion of traffic is composed of elastic flows that are bottlenecked at the link in question. However, in a backbone network the majority of flows are bottlenecked elsewhere, notably by the user access links whose rate is usually much smaller than link capacity.

In the latter case it is necessary to perform bufferless statistical multiplexing and to employ an MBAC capable of ensuring the overall incoming rate remains less than link capacity with high probability. Given the limited information available about the traffic of individual flows (only a bound on their peak rate is known) and the distributed nature of the required admission control, MBACs already proposed in the literature are either not adapted or require extension.

We have performed some preliminary investigations by simulation and these suggest a simple approach based on the Poisson distribution of the number of packets arriving in a small interval may be sufficient. The idea is to apply admission control or not in a short time interval based on the current estimation of offered load.

It is emphasized that this evaluation is preliminary. Current work is focused on extending the range of traffic models used, notably with real traffic traces. A further subject of investigation is the definition of an adaptive algorithm where the admission threshold is increased or decreased depending on realized medium to long term performance. Finally, it is necessary to more completely understand the impact of the size of the measurement interval on the stability of the applied control. It may be necessary, for example, to apply a back-off policy as introduced in previous works to prevent a large number of arrivals occurring in a short interval of time and leading to overflow.

References

[1] N. Benameur, S. Ben Fredj, S. Oueslati and J. Roberts. Quality of service and flow-aware admission control in the Internet. In Computer Networks, Vol 40, pages 57-71, 2002.

[2] S. Ben Fredj, T. Bonald, A. Proutière, G. Régnié and J.W. Roberts. Statistical bandwidth sharing: A study of congestion at flow level, In Proc. of ACM SIGCOMM 2001.

[3] T. Bonald, A. Proutière, J. Roberts. Statistical performance guarantees for streaming flows using expedited forwarding, In Proc. of IEEE INFOCOM 2001

[4] L. Breslau, S. Jamin and S. Shenker. Comments on the Performance of Measurement-Based Admission Control Algorithms. In IEEE INFOCOM 2000, Tel Aviv, Israel, March 2000.

[5] C. Casetti, J. Kurose and D. Towsley. An Adaptive Algorithm for Measurement-based Admission Control in Integrated Services Packet Networks. UMASS Technical Report TR 96-76, IFIP Workshop on high speed networks, October 1996.

[6] S. Crosby, I. Leslie, B. McGurk, J.T. Lewis, R. Russell and F. Toomey. Statistical Properties of a Near Optimal Measurement-based CAC Algorithm. In Proceedings IEEE Workshop on ATM, Lisbon, June 1997.

[7] Z. Dziong, M. Juda and L. Mason. A Framework for Bandwidth Management in ATM Networks - Aggregate Equivalent Bandwidth Estimation Approach. IEEE/ACM Transactions on Networking 5, 1, February 1997.

[8] S. Floyd. Comments on Measurement-based Admissions Control for Controlled-Load Services. Lawrence Berkeley Laboratory, Tech. Rep., July 1996.

[9] R.J. Gibbens and F.P. Kelly. Measurement-based connection admission control. In 15th

International Teletraffic Congress Proceedings, June 1997.

[10] R.J. Gibbens, F.P. Kelly and P.B. Key. A Decision-Theoretic Approach to Call Admission Control in ATM Networks. In IEEE Journal on Selected Areas in Communications, Vol 13, No 6, pages 1101-1114, August 1995. oban1

[11] M. Grossglauser and D. Tse. A Framework for Robust Measurement-Based Admission Control. IEEE/ACM transactions on Networking, 5 (1), February 1997.

[12] M. Grossglauser and D. Tse. A Time-Scale Decomposition Approach to Measurement-Based Admission Control. IEEE/ACM Transactions on Networking, 2003.

[13] J.M. Hyman, A.A. Lazar, and G. Pacifici. A Separation Principle Between Scheduling and Admission Control for Broadband Switching. IEEE Journal on Selected Areas in Communications, 11(4):605-616, May 1993.

[14] S. Jamin, S. Shenker. Measurement-based Admission Control Algorithms for Controlled-load Service: A Structural Examination. University of Michigan CSE-TR-333-97, April. 1997

[15] S. Jamin, P.B. Danzig, S.J. Shenker, and L.Zhang. A Measurement-based Admission Control Algorithm for Integrated Services Packet Switched Networks (Extended Version). IEEE/ACM transactions on networking, February 1997.

[16] S. Jamin, S.J. Shenker and P.B. Danzig. Comparison of Measurement-Based Admission Control Algorithms for Controlled-Load Service. IEEE Infocom 1997.

[17] F.P. Kelly. Reversibility and stochastic networks. J. Wiley & Sons, 1979.

[18] A. Kortebi, S. Oueslati and J. Roberts. Cross-protect: implicit service differentiation and admission control. Proceedings of HPSR 04, Phoenix, 2004.

[19] J. Qiu and E.W. Knightly. Measurement-Based Admission Control with Aggregate Traffic Envelopes. IEEE/ACM transactions on networking, vol. 9, NO. 2, April 2001.

[20] M. Reisslein. Measurement-Based Admission Control: A Large Deviations Approach for Bufferless Multiplexers. In Proceedings of IEEE Symposium on Computers and Communications , pages 462-467, Antibes, France, July 2000.

[21] J. Roberts. Internet Traffic, QoS and Pricing. Proceedings of IEEE, to appear, 2004.

[22] J. Roberts, U. Mocci and J. Virtamo. Broadband Network Teletraffic (Final Report of COST 242). LNCS 1155 Springer Verlag, 1996.

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

[24] H. van den Berg and M. Mandjes. Measurement-based admission controle: algorithms, evaluation, and research perspectives. COST 257,1999.