On the tradeoff between effectiveness and scalability of measurement-based admission control

András Veres, Zoltán Richárd Turányi

Ericsson Traffic Analysis and Network Performance Laboratory
e-mail: {Andras.Veres, Zoltan.Turanyi}@eth.ericsson.se

July 7, 1999

Abstract

The key problem to be solved in MBAC is to efficiently utilise the information provided by the sources and measured by the network. There is a direct relationship between the amount of information available and the resulting effectiveness of the admission algorithm. In this paper we analyse the tradeoff between complexity and effectiveness using a practial class of AC methods.

1 Introduction

When a network aims to provide any guarantee on the quality of its service it must introduce some form of Admission Control (AC) to protect the network from overload. AC algorithms check incoming requests and decide if they can be served at the requested quality without degrading the quality of other connections. The decision is based on traffic information supplied by the sources, measured by switches or assumed by the AC algorithm.

Recently many papers discussed different AC approaches utilising a different mix of the above three information sources. In [KnigZhan95] the decision is based on a complex, deterministic descriptor, which, if tightly set, can give high utilisation. No on-line measurements are made. Heuristic methods [CKT96, JJB97, JDSZ97, JamShen97] make simple measurements and assume that some statistical properties of the traffic remain constant or predictable. [KWC93] describes a theoretic approach where sources are assumed to be well described by MMPP processes. [KM98] gives effective bandwidth formulae for a class of non Markovian sources, where state transitions are modelled by a markov chain. [GAN91] uses measurements of the average and variance of the rate and assumes Gaussian property of the aggregate. Sally Floyd's method [Flo96] uses the Hoeffding bound to compute a theoretical upper bound on the overflow probability, expects only a peak rate from the sources, measures the aggregate mean rate and assumes the independence and weak stationarity of the sources. Brichet and Simonian [BriSim98] can achieve higher utilisation using token bucket parameters supplied by the sources. In [TurVer98] we proposed an AC algorithm that uses admitted peak rates and token buckets, and measurements of the variance and mean providing a tighter bound. Finally the work of S. Crosby et al. is mentioned [CLL97], their Mosquito algorithm uses detailed per-flow measurements to approximate the SCGF of the flows and assumes weak dependence among consecutive time periods.

The efficiency of these AC methods can be evaluated by their ability to meet the primary and secondary goals of AC. The primary goal is to ensure that the admittance of a flow will not degrade the service level of already admitted traffic. The secondary goal is to admit as many traffic as possible while ensuring the primary goal.

Each information source has its own drawbacks and benefits in meeting these goals:

- *Information provided by the sources.* This type of information can be the most reliable and also this can be enforced at the ingress point by traffic policing or shaping. However in many cases sources can not provide an exact and detailed characterisation of their traffic at the time of flow setup because of the highly varying nature of traffic generated by certain applications e.g. streaming video. As a result these descriptor sets usually can only be fitted very loosely thus limiting the efficiency of admission control algorithms that rely only on the supplied parameters. Furthermore we should not force applications to fit certain complex traffic models rather on the contrary: the traffic descriptors should be general and simple enough to fit all the needs of current and future applications.

- *Information based on assumptions.* Assumptions are very useful in approaching the AC problem with analytical tools. The drawback is that these assumptions can be very misleading and certain –traditionally used– assumptions in telecommunications such as Poission models have been proven to model IP traffic poorly. [PaxFlo95] Thus more complex models are suggested by most authors, but their applicability in AC have not been proven satisfactorily.
- *Information based on measurements.* Measurements are made to increase our knowledge about the traffic. The measured data is then used to fit the model we assumed the traffic fits best. On the other hand measurement has its limits as well and the infromation which can be gained using measurements is also limited due to practical or theoretical reasons. Measurements always have some variance (error) that will cause erroneous AC decisions. This variance can be decreased by measuring a larger number of flows or increasing the measurement period. This is not always possible, as the time interval that would be required for precise measuremnet may be larger than the length of the flows or the timescale during which the process can be considered stationary. There are also practical limitations of using measurements, as obtaining certain complex statistics or measuring a large number of parameters may require an uneconomical amount of resources.

In a previous paper [TurVer98] we argued that an efficient and scalable admission control algorithm should utilise these three information sources in the following way:

1. Information provided by the sources should be simple and easy to police.
2. The required assumptions should be simple and reasonable for example weak stationarity or independence.
3. Measurements should be done only on large traffic aggregates and not per-flow.

In this paper we analyse the tradeoff between the detailedness of measurements and the gain in AC effectiveness, i.e., the tightness of the effective bandwidth calculations based on the parameters. Our conclusion is that measuring the traffic in a small number of groups gives the same performance as per-flow measurements. This results in a scalable AC design. We present an optimal strategy to group the flows and also demonstrate our conclusions using measured real-life TCP/IP traffic.

2 AC model with flow grouping

We assume a class of AC methods, where at flow setup time the sources admit only their peak rates h_k. The router has a FIFO buffer on the outgoing interface with service rate C. For efficient admission control decision, the router makes measurements on the average bit-rate of the traffic. We can examine the tradeoff of efficiency vs. complexity by varying the detailedness of these measurements. The two ends of the spectrum under consideration are:

- Only the aggregate traffic on the link is measured. This is very easy to implement, just a simple counter is required on the interface.
- All packets are classified according to their flow (e.g., they are classified based on the source and destination addresses, port numbers, logical connection identifier) and the rate of each flow is measured separately: m_k. This is much more complex to do as it not only requires to keep per flow states in each router, but also classification is needed for each packet to decide which flow k it belongs to.

The first end of the spectrum scales well but gives only an overall measurement althogh the flows may be very different in their traffic statistics and how they utilize their profiles. Using aggregate measurements, this information is lost, so a conservative effective bandwidth calculation should be used. If per-flow measurements are available then we know much more about the behaviour of the sources and can give tighter bounds for the effective bandwidth.

In between the two extremes we can find methods in which we group a number of flows into a limited number of groups and we make measurements on the groups only. We do not differentiate the flows further within a group. This way we have more information than just the aggregate and also classiffication may be easier to do for a limited number of groups. We predict that the more groups we make the tighter bound can be given. Furthermore, if the grouping is done in an optimal way even higher utilization can be achieved.

In this paper, the presented AC methods are based on the concept of *effective bandwidth* and zero buffer approximation. (A non zero buffer approximation can be found in [TurVer98].) The effective bandwidth of the aggregate traffic is BW if

$$\Pr\left(\sum X_k \geq BW\right) = \epsilon, \tag{1}$$

where ϵ is the saturation probability. A new flow is admitted if $BW \leq C$ where the new flow is included in BW.

The rest of the paper is organized in the following way: in Section 3 we show how to construct an AC method with aggregate measurement, in Section 4 the method based on grouping is derived. Examples for the methods are given in Section 5. Section 6 discusses the way we group flows. In Section 7 a real life example is shown.

3 Effective bandwidths based on aggregate measurement

In [GibKel97] it is shown how the Chernoff bound leads to the Hoeffding bound using a suitable overestimation of the moment generating function. In this section another overestimation of the Chernoff bound can be found which will make it easy to understand the extension we give later in the paper. Using the Chernoff bound the effective bandwidth of the aggregate traffic can be expressed as a function of s ($s > 0$)

$$BW(s) = \frac{1}{s}\sum_{k=1}^{N} \ln\left(1 + \frac{e^{sh_k} - 1}{h_k} m_k\right) + \frac{\gamma}{s} \tag{2}$$

Where N is the number of flows, h_k and m_k are the peak and measured mean rates of flow k respectively. $\gamma = -ln\epsilon$, where ϵ is the probability that the link capacity is exceeded. This inequality holds for any $s > 0$, so for a tight result (2) should be optimized for s.

If we know the individual values of m_k (per-flow measurements) then this expression can be directly used for AC. In this case the m_k of the new flow can be estimated (e.g. by h_k or a token rate if provided). If we can measure only the aggregate then we need to modify this expression to contain only the aggregate mean. First we rewrite this expression as

$$BW(s) = \frac{1}{s}\ln\prod_{k=1}^{N}\left(1 + \frac{e^{sh_k} - 1}{h_k} m_k\right) + \frac{\gamma}{s} \tag{3}$$

Factorizing:

$$\prod_{k=1}^{N}\left(1 + \frac{e^{sh_k} - 1}{h_k} m_k\right) = \prod_{k=1}^{N}\left(\frac{e^{sh_k} - 1}{h_k}\right)\prod_{k=1}^{N}\left(m_k + \frac{h_k}{e^{sh_k} - 1}\right) \tag{4}$$

The geometric mean is smaller than the algebric

$$\sqrt[N]{\prod_{k=1}^{N}\left(m_k + \frac{h_k}{e^{sh_k} - 1}\right)} \leq \frac{1}{N}\sum_{k=1}^{N}\left(m_k + \frac{h_k}{e^{sh_k} - 1}\right) \tag{5}$$

And the effective bandwidth expression is

$$BW(s) \leq BW_{\text{approx}}(s) = \frac{N}{s}\ln\left(\frac{M + \sum_{k=1}^{N}\frac{h_k}{e^{sh_k}-1}}{N}\right) - \frac{1}{s}\sum_{k=1}^{N}\ln\left(\frac{h_k}{e^{sh_k} - 1}\right) + \frac{\gamma}{s} \tag{6}$$

Where $M = \sum m_k$ is the aggregate load measurement. The saturation probability is not exceeded for any value of $s > 0$, for optimal utilization we may find the optimal s where this expression is minimal

$$BW_{\text{approx}} = \min_s (BW_{\text{approx}}(s)) \tag{7}$$

This bound –which is an overestimation of the Chernoff bound– is always tighter than the Hoeffding bound, but cannot be expressed in a closed form. A good approximation for s can be found by using the first few elements of the Taylor series and then differentiating on s. (See [TurVer98].)

$$s_{opt} = \sqrt{\frac{\gamma}{\frac{1}{8}\sum_{j=1}^{N} {h_j}^2 - \frac{1}{2}\cdot\frac{1}{N}\left(M - \frac{1}{2}\cdot\sum_{j=1}^{N} h_j\right)^2}}$$

The closed form result using this approximation is still in most cases tighter than the Hoeffding bound, especially if the traffic is biased to high or low mean-to-peak ratio, which is the typical case for most current applications. (See again [TurVer98].)

Note that during the calculations a major step was (3) where using an approximation we replaced the expression with an other where only the *sum* of means is present. On a network link this means that only a simple counter is needed for the aggregate traffic measuring the transmitted amount of bytes during a given period.

4 Effective bandwidths based on measurements of flow groups

In this section an effective bandwidth expression is given which contains information on load measurements only on a per group basis, where the number of groups can range from 1 (in this case we have aggregate measurement only) to the total number of flows (in this case per flow measurements are done). First we give a solution which is true for arbitrary grouping. Then in section 6 problem of grouping is addressed.

Let's group the N flows into G sets (groups), and denote these sets as $A_i, i = 1..G$ and let $n_i := |A_i|$ the number of flows in group i. Then using the algebric-geometric inequality among the groups we get

$$\prod_{k=1}^{N}\left(m_k + \frac{h_k}{e^{sh_k}-1}\right) = \prod_{i=1}^{G}\left(\sqrt[n_i]{\prod_{k\in A_i}\left(m_k + \frac{h_k}{e^{sh_k}-1}\right)}\right)^{n_i} \leq \prod_{i=1}^{G}\left(\frac{1}{n_i}\sum_{k\in A_i}\left(m_k + \frac{h_k}{e^{sh_k}-1}\right)\right)^{n_i} \tag{8}$$

The full effective bandwidth expression is:

$$BW(s) \leq BW_{\text{groups}}(s) = \frac{1}{s}\sum_{i=1}^{G} n_i \ln\left(\frac{M_i + \sum_{k\in A_i}\frac{h_k}{e^{sh_k}-1}}{n_i}\right) - \frac{1}{s}\sum_{k=1}^{N}\ln\left(\frac{h_k}{e^{sh_k}-1}\right) + \frac{\gamma}{s} \tag{9}$$

Where $M_i = \sum_{k\in A_i} m_k$ is the load measurement for group i.

A good approximation for s can be found similarly as in the previous section

$$s_{opt} = \sqrt{\frac{\gamma}{\frac{1}{8}\sum_{i=1}^{N} {h_i}^2 - \frac{1}{2}\sum_{k=1}^{G}\frac{1}{n_k}\left(M_k - \frac{1}{2}\sum_{i\in A_k} h_i\right)^2}}$$

Clearly the tightness of this expression depends on the choices of:

- the number of groups (G)
- the distribution of flows among the groups (A_i)

Expression (9) can be optimized for both. For the second a simple "rule of thumb" exists: if we select "similar" flows into the same group then the geometric mean will be closer to the algebric in (8) resulting in tighter bandwidth estimation.

The term "similar" here denotes flows with similar $V(k) = m_k + h_k/(e^{sh_k} - 1)$ values. The problem is hard because s is still an unresolved parameter. But in any case flows with similar m_k and h_k are "similar" in this sense. As a practical example we can say that all flows belonging to similar application types (e.g. IP telephony, streaming video, file transfer) are "similar".

5 Is there any practical gain in grouping?

The gain is obvious if the grouping is done efficiently, and there are 'similar' types of flows. We did a test to see the magnitude of this gain. We assumed two typical applications and created two randomized set of flows with similar mean and peak rates within a group. (See Figure 1a.) Type-1 flows have average peak rate 100kbit/s and average mean-to-peak ratio 0.5, while type-2 flows have average peak rate 1Mbit/s and average mean-to-peak ratio 0.1. The first type may model high quality streaming audio with low burstiness and lower peak rates, the second may correspond to MPEG compressed real-time video applications with high peak rates and relatively low mean-to-peak ratios.

On Figure 1b. the gain in the admitted number of flows can be observed if we group the flows into two groups and make load measurements per group instead of measuring the aggregate. The lower line is the admittance region if aggregate load measurement is used (Hoeffding bound based admission control), the upper line shows when we group the flows according to their types (group-1 contains type-1 and group-2 contains type-2 flows).

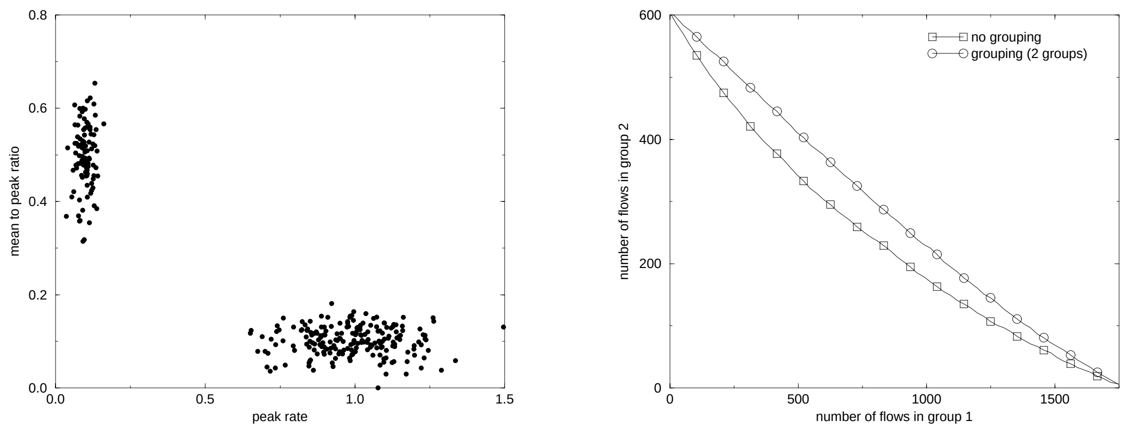

Figure 1: a) Randomly generated test set of flows. b) Admittance region of flow type-1 versus flow type-2.

We can see that in case of distinct flow sets high gain can be achieved when we do correct grouping.

6 How to group flows?

A tight approximation of the per-flow formulae can be done if flows with similar $V(k)$ values are grouped together. On Figure (2) the $V(k)$ values of flows is plotted in ascending order.

After sorting - interestingly - all flows of type 'music' can be found on the left side and all flows of type 'video' on the right side, and there is a significant jump exatcly where the two groups can be separated. This suggests to define the groups according to where such jumps are visible on the $V(k)$ plot.

$V(k)$ is influenced by the transport protocol, the application and the behavior of the user as well. In real networks it can be expected that there will be typical combinations of the above three (e.g. telephony: RTP/UDP, ON/OFF type with peak rate around 10kbps and mean-to-peak ratio close to 0.5) resulting in clusters on the $(m_k, m_k/h_k)$ scatter plot.

In the following example there are more traffic types differing significantly in peak and mean rates. See Figure 3a. The corresponding sorted $V(k)$ plot is also diplayed on Figure 3b. This latter graph also shows the seasonal difference of the $V(k)$ curve which reveals three regions - suggesting the use of three groups.

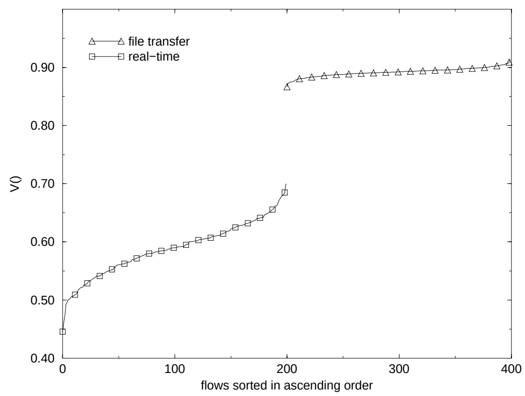

Figure 2: V(k) values for flows (s=1.122).

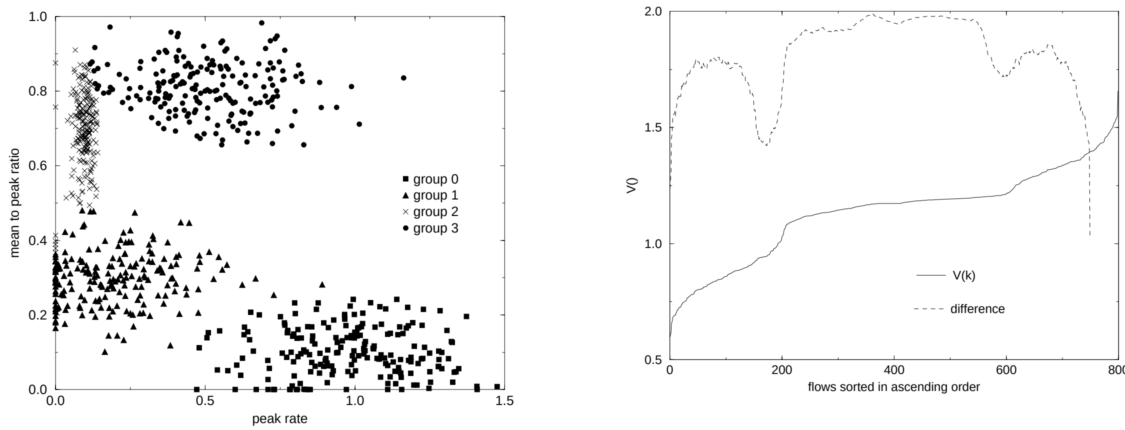

Figure 3: a) Scatter plot of flows belonging to 4 groups. b) $V(k)$ plot of flows and its seasonal difference plot.

7 How many groups are needed?

By increasing the number of groups we gain more knowledge and so we expect tighter bounds on the effective bandwidth. On the other hand if the number of groups is large, more measurements are needed and classification may be more complex which may limit the scalability of the method.

To show this tradeoff a random set of flows are evenly placed on the (peak, mean-to-peak) space (see Figure 4a). Then the number of groups G is increased to see how much gain is achieved with different number of groups. For a certain G the flows are sorted by their $V(k)$ values and are evenly divided into G groups such that N/G flows fall into each group. Figure 4b shows the calculated effective bandwith as a function of G.

The figure suggests that it is enough to use only a very few number of groups, and any further increase in the number of groups (e.g. per-flow measurements) does not give significantly higher gain.

8 Analysis of a dial-in service.

Finally we used measurements of a real dial-in service. Flows were defined as all packets arriving to a specific IP address and with the same source TCP/UDP port number, as we tried to differentiate

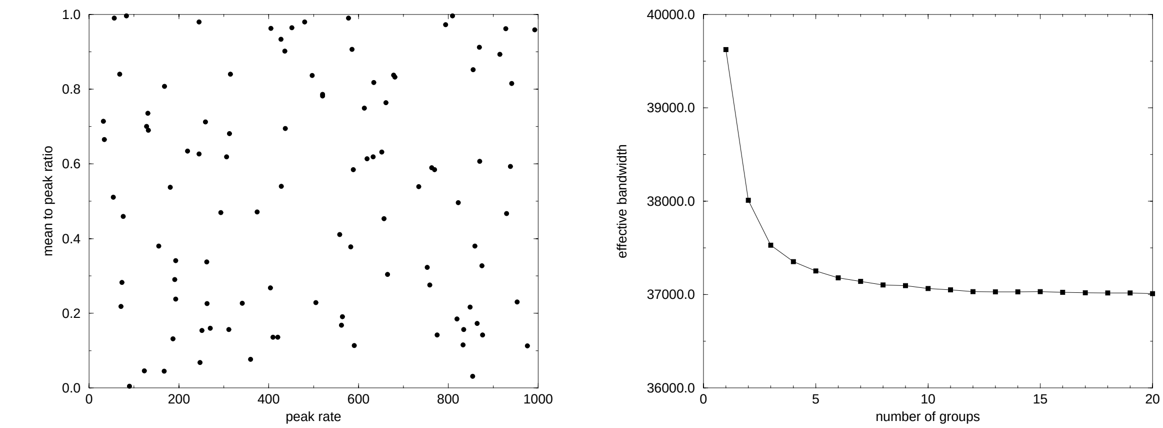

Figure 4: a) Evenly placed random flow set. b) The effective bandwidth vs. the number of groups in case of evenly placed random flows.

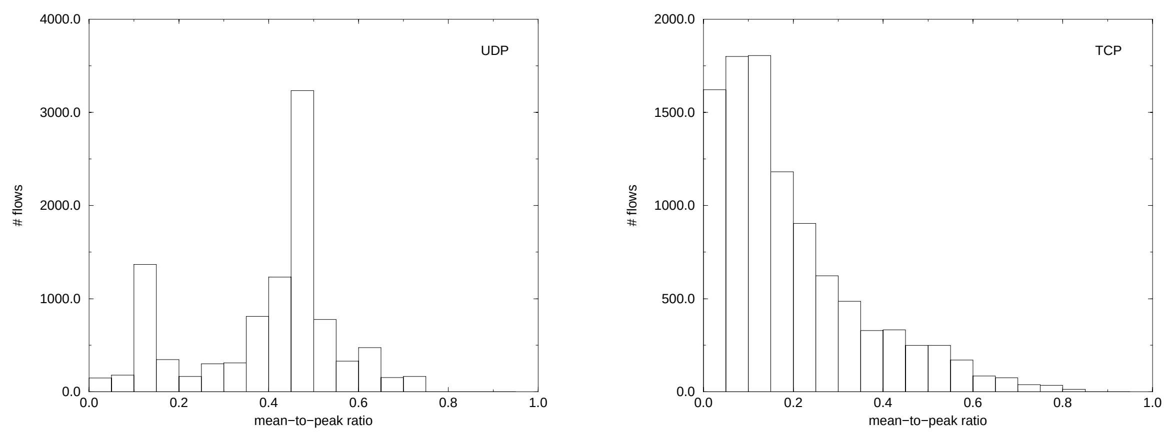

Figure 5: Histogram of measured mean-to-peak ratios of modem flows containing a) UDP and b) TCP packets.

among different application types. We assumed that the peak rates of the flows equaled the modem connection speed which was in most cases 33Kbits/sec. It is likely that the peak rate values provided by the sources will not be more accurate in the near future (especially in the absence of appropriate real-time operating systems and protocols). The observed average rates are displayed on Figure 5a. The TCP flows are spread out widely, but the UDP flows are more concentrated around 0.5 mean-to-peak (approximately 16kbps) which is caused by streaming video/audio applications (e.g. RealAudio). After grouping the flows according to their protocol, the admittance regions are plotted for different TCP-UDP traffic mixes on Figure 6b.

9 Conclusion

We introduced a class of admission control algorithms based on the effective bandwidth concept. We argued that between the two extremes suggested so far for MBAC (i.e., aggregate and per-flow measurements) there exists a range of algorithms based on flow grouping that differ in implementation complexity and in the extent they can exploit the statistical multiplexing gain. The paper

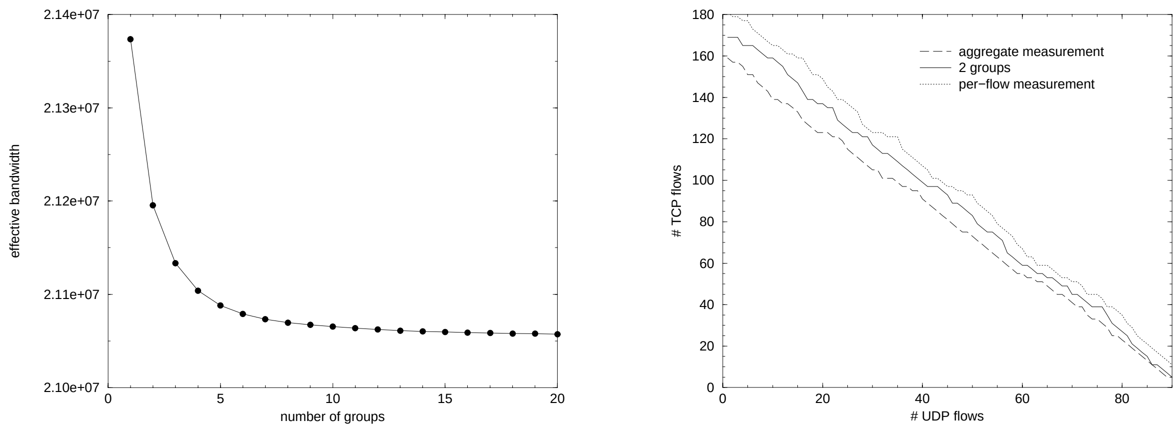

Figure 6: a) The effective bandwidth vs. the number of groups in case of dial-in traffic. b) Admittance region of flow UDP versus flow TCP flows, link capacity $C = 2Mbit/s$.

contains closed form formulae for these AC algorithms. Through examples we demonstrated that by increasing the number of groups the aggregate effective bandwidth decreases but the benefit from this is significant only for very few groups which means that a simple implementation of only a few groups can utilise most of the gain. We also give a heuristics of how to group flows efficiently and finally we demonstrate the efficiency of the method on data collected in a real-life dial-in service.

References

[BriSim98] F. Brichet, A. Simonian, "Conservative Gaussian models applied to Measurement-based Admission Control", *submitted to IEEE INFOCOM '98*, San Francisco, March. 1998

[CKT96] C. Casetti, J. Kurose, D. Towsley, "A New Algorithm for Measurement-based Admission Control in Integrated Services Packet Networks", *Protocols for High Speed Networks '96*, INRIA, Sophia Antipolis, Oct. 1996

[Flo96] S. Floyd, "Comments on Measurement-based Admissions Control for Controlled-Load Service", *unpublished*, available at ftp://ftp.ee.lbl.gov/papers/admit.ps.Z

[GAN91] R. Guerin, H. Ahmadi, M. Naghshineh, "Equivalent capacity and its application to bandwidth allocation in high-spped networks", *IEEE Jurnal on Selected Areas in Communications*, 9(7), pp. 968-981, Sep. 1991

[GibKel97] R. J. Gibbens, F. P. Kelly, "Measurement-Based Connection Admission Control", *International Teletraffic Congress 15*, Jun. 1997

[JamShen97] S. Jamin, S. Shenker, "Measurement-based Admission Control Algorithms for Controlled-load Service: A Structural Examination", *Internal report*, Apr. 1997

[JDSZ97] S. Jamin, P. Danzig, J. Shenker, L. Zhang, "A Measurement-Based Admission Control Algorithm for Integrated Service Packet Networks", *IEEE/ACM Transactions on Networking*, vol. 5. no. 1. Feb. 1997

[JJB97] S. Jamin, C. Jin, L. Breslau, "A Measurement Based Admission Control Algorithm for Controlled-Load Service with a Reference Implementation Framework",*Internet draft*, Nov. 1997

[Kel96] F. P. Kelly, "Notes on Effective Bandwidths", *In F. P. Kelly, S. Zachary and I. B. Ziedins, Stochastic Networks: Theory and Applications, Royal Statistical Society, Lecture Note Series 4, p 141-168, Oxford Univ. Press*

[KWC93] G. Kesidis, J. Walrand, C. Chang, "Effective Bandwidths for Multiclass Markov Fluids and Other ATM Sources", *IEEE Trans. Networking*, Vol. 1, No. 4, pp. 424-428, Aug. 1993.

[KM98] K. Kontovasilis, N. Mitrou, "Effective Bandwidths for a Class of Non Markovian Fluid Sources", *Proc. ACM Sigcomm 98* pp. 263-274, Sept. 1998.

[KnigZhan95] E. Knightly, H. Zhang, "Traffic Characterization and Switch Utilization using a Deterministic Bounding Interval Dependent Traffic Model", in *Proc. IEEE INFOCOM '95*, Boston, April 1995

[TurVer98] Z. R. Turányi, A. Veres, "A Family of Measurement-Based Admission Control Algorithms", *Proc. IFIP Performance of Information and Communication Systems '98*

[CLL97] S.Crosby, I.Leslie, J.T.Lewis, R.Russell, F.Toomey and B.McGurk, "Statistical Properties of a Near-Optimal Measurement-based CAC algorithm", *Proceedings IEEE ATM '97, June 1997, Lisbon*

[PaxFlo95] V. Paxson, S. Floyd, "Wide-area traffic: the failure of Poisson modeling", *IEEE/ACM Transactions on Networking 3:226-244*, 1995

A Finding an approximate value for s

The effective bandwidth formula for grouped flows is the following:

$$BW_{\text{groups}}(s) = \frac{1}{s}\sum_{i=1}^{G}\sum_{j\in A_i} \ln\left[\left(\frac{M_i}{n_i} + \frac{1}{n_i}\sum_{k\in A_i}\frac{h_k}{e^{sh_k}-1}\right)\frac{e^{sh_j}-1}{h_j}\right] + \frac{\gamma}{s}$$

Taking the series expansion of $h_k/(e^{sh_k}-1)$ we get:

$$\frac{h_k}{e^{sh_k}-1} = \frac{1}{s} - \frac{1}{2}h_k + \frac{1}{12}h_k^2 s - \frac{1}{720}h_k^4 s^3 + O(s^5)$$

Substituting the first 3 elements (higher terms would disappear during the next steps) we get

$$BW_{\text{groups}}(s) = \frac{1}{s}\sum_{i=1}^{G}\sum_{j\in A_i} \ln\left[\left(\frac{M_i}{n_i} + \frac{1}{s} - \frac{1}{2n_i}H_i + \frac{1}{12n_i}HH_i s\right) n_i \frac{e^{sh_j}-1}{h_j}\right] + \frac{\gamma}{s}$$

using the notation $H_i = \sum_{k\in A_i} h_k$ and $HH_i = \sum_{k\in A_i} h_k^2$. Taking the series of the logarithm:

$$\ln\left[\left(\frac{M_i}{n_i} + \frac{1}{s} - \frac{1}{2n_i}H_i + \frac{1}{12n_i}HH_i s\right) n_i \frac{e^{sh_j}-1}{h_j}\right]$$
$$= \left(\frac{M_i}{n_i} - \frac{1}{2n_i}H_i + \frac{1}{2}h_j\right)s + \left(\frac{1}{12}\frac{HH_i}{n_i} + \frac{1}{24}h_j^2 - \frac{1}{8}\frac{(2M_i - H_i)^2}{n_i^2}\right)s^2$$

Substituting back to BW_{groups}

$$BW_{\text{groups}}(s) \approx \sum_{i=1}^{G} M_i + \frac{1}{8}HH_i s - \frac{1}{8}\frac{(2M_i-H_i)^2}{n_i}s + \frac{\gamma}{s}$$

By deriving this solving for $s = 0$ we get an approximation

$$s \approx \sqrt{\frac{8\gamma}{\sum_{i=1}^{G} HH_i - \frac{(2M_i-H_i)^2}{n_i}}}$$

and

$$BW_{\text{groups}}(s) \approx \sum_{i=1}^{G} M_i + \sqrt{\frac{1}{2}\gamma\left(\sum_{i=1}^{G} HH_i - \frac{(2M_i-H_i)^2}{n_i}\right)}$$

which is always smaller than the Hoeffding bound and is equal when the group mean-to-peak ratio is 0.5 which is leads to the highest variance in case of on/off sources and is thus worst case in some sense.

$$BW_{\text{Hoeffding}}(s) = \sum_{i=1}^{G} M_i + \sqrt{\frac{1}{2}\gamma\sum_{i=1}^{G} HH_i}$$