Journal of High Speed Networks 13 (2004) 87–102
IOS Press

An implementation-based comparison of Measurement-Based Admission Control algorithms

Andrew W. Moore
University of Cambridge, Computer Laboratory, William Gates Building, J.J. Thomson Ave., Cambridge CB3 0FD, UK
E-mail: andrew.moore@cl.cam.ac.uk

Abstract. In this paper we present an implementation-based comparison of Measurement-based Admission Control algorithms. Through the use of a special purpose environment, a performance and behaviour comparison is made. The results of this paper illustrate the independence of traffic from admission control behaviour in the homogeneous traffic environment. While illustrating the impact the admission decision will make upon heterogenous traffic systems. These results highlight the importance of estimators being robust to statistical variation in measurements and offering calibrated controls. While based upon a comparison of measurement-based admission control algorithms — these conclusions are applicable to any application of measurement-based estimation.

Keywords: Measurement, network resource control, admission control, measurement-based admission control

1. Introduction

This paper reports on a comparison of Measurement-Based Admission Control (MBAC) algorithms. The comparison is conducted using implementations in a purpose-built test environment, subjecting the algorithms to realistic traffic and network conditions, while being constrained to practical limits on memory, computational resource and access to measurements.

Admission Control (AC) is a mechanism for traffic management, which consists of admitting a new traffic source if and only if the network can accommodate the new flow while still supporting existing commitments made to sources already accepted. An AC procedure is employed to maintain a high utilisation of network resources while preserving the guarantees made to existing flows. This is done by balancing higher network utilisation through increased multiplexing against the satisfaction of commitments for existing clients. Such an AC scheme relies on being able to accurately establish the resource requirements of current flows, along with a prediction of the impact a new flow will have upon existing traffic sources. Commonly an AC scheme requires that a new flow declares parameters that can be used to calculate its resource requirements and, therefore, its impact on pre-existing flows.

A traditional AC algorithm requires new flows to supply accurate characterisation of themselves at the time of admission. The characterisation of the new flow must be performed prior to admission. In contrast, measurement-based AC (MBAC) algorithms compute an estimate of the resource requirements of current flows by measurement of such metrics as line utilisation or buffer-loss. Thus MBAC algorithms permit the declared traffic requirements of the new connection to be minimal or incorrect without impacting upon the performance of the ongoing characterisation derived from measurements. Thus an MBAC algorithm can allow better utilisation of network resources with flows admitted with minimal prior characterisation.

AC algorithms have long been of interest in the management of fixed-load and integrated services such as voice and video. In the Internet community, interest in ACs has, until recently, been limited to the controlled-load services of IP under Integrated Services (INTSERV) [4,29]. However, work on admission and control in networks based upon differentiated-services (DIFFSERV) [3,24,25], at network ingress [27] and egress [8,26], has seen a resurgence of work in MBAC techniques. However, there have been many proposed MBAC algorithms. This has motivated the need for an implementation-based comparison in order to identify an ideal MBAC algorithm. In the context of defining criteria by which the ideal MBAC algorithm may be identified, this paper discusses a number of the comparisons that have been used in the past.

The results presented from this work serve to identify further directions for research into MBAC algorithms while noting criteria that have served to confuse past comparison AC work. Additionally, given that a Measurement-Based Estimator (MBE) is fundamental to each MBAC algorithm, lessons learned from this comparison work may be applied equally easily to MBEs used in other network management tasks such as long-term capacity planning or improving current QoS-routing approaches [14,21].

1.1. A note on effective bandwidth

A formal definition of *effective bandwidth* is provided by Kelly [18]. Such a definition serves to capture the subtleties of the traffic multiplex and network QoS (loss, delay and throughput) constraints in combination with the buffer and service-capacity resources made available by the network. Kelly [18] may be interpreted to define the *effective bandwidth* of any individual traffic source as the total bandwidth required to satisfy the QoS constraints of the total traffic multiplex for a given buffer resource when divided among the number of traffic sources present in the multiplex. As it applies to a single source, it is this definition of *effective bandwidth* that is used throughout this paper.

2. Theory

This section presents a number of different MBAC algorithms, noting the fundamental premise upon which each are based. A number of the MBAC algorithms of this study have their basis in the solution or approximation of the *Chernoff Bounds*, while others approach the estimation problem from different theoretical backgrounds such as large-deviation theory or statistical analysis.

Firstly, a simple MBAC algorithm is introduced: AC-ST, with its admission decision based upon a single instantaneous utilisation measurement. For the AC decision, the measurement is compared with a pre-defined threshold value. AC-ST is a valuable template estimator that, while not expected to perform particularly well, forms a useful base-reference. The computation of appropriate thresholding values can be considered the major work of any MBAC algorithm and some authors [13,19] have attempted to tackle this issue directly using the approach of AC-AR.

Driven by a measurement of current line utilisation, the approach of AC-AR uses an acceptance region: a range of values of utilisation where combinations of incoming flows would be admissible. The acceptance region is computed to maximise line utilisation for a nominated packet loss, given a set of flows with a known declaration of peak and mean rates.

First of the algorithms based upon the *Chernoff Bounds*, AC-CB is the instantiation of one of the family of algorithms proposed in [12]. The approach of this algorithm is to estimate the bound to a curve of *effective bandwidth* versus mean rate. Four techniques are presented, each estimating the bound based upon different information about the curve of *effective bandwidth* versus mean rate.

In contrast, the measured sum algorithm, AC-MS, computes an estimate of *effective bandwidth* based upon regular sampling of measured aggregate loads [17]. This algorithm takes a much simpler approach with little theoretical foundation combining a local-maximum prediction with a control over the level of line utilisation.

The AC-MPFE algorithm is based upon the theory of large deviations [10]. Large deviation theory allows the quantification of rare events such as packet loss in a computer network due to traffic interactions. This approach uses the premise that the logarithm of the loss-ratio versus the buffer size may be bound by a straight line. This approximation allows description of the large deviation rate function and in turn this allows description of a Scaled-Cumulative Generating Function (SCGF). An estimate of the *effective bandwidth* becomes the slope of the SCGF for a particular set of (traffic) measurements constrained by a set of buffer characteristics (loss-ratio and buffer size).

The AC-MAE algorithm is derived from AC-MPFE [23]. While imposing a significantly different (reduced) demand on the measurement system, this algorithm is based upon the same foundation theory as AC-MPFE. Both AC-MPFE and AC-MAE operate through the estimation of bandwidth requirements based directly upon available buffer-space and desired packet loss-ratio.

A family of algorithms based upon statistical information derived directly from the measurement of line utilisation is introduced next. AC-MVE is an algorithm based upon the estimation of bandwidth requirements using a combination of mean and variance over one time-scale. This technique incorporates a correction to the variance multiplier to account for potential increased variability when the number of samples is small. The idea of such a correction causes this algorithm to share much in common with the estimator described in Duffield et al. [9, § 3.1 Local Gaussian Predictor], a measurement-based estimator used in the allocation of resources for new VPNs.

The other statistical estimator-based MBAC algorithm is AC-KQ [20]. This algorithm introduces traffic envelopes (descriptions of the mean and variance of traffic over multiple time-scales) and a loss-boundary mechanism that allows computation of the *effective bandwidth* from the traffic envelope. Computation of the *effective bandwidth* requirements allows construction of an AC algorithm.

AC-LBE, a loss-based estimator is the only MBAC algorithm to use the measured loss-ratio as an admission criterion. This algorithm admits new flow-arrivals based on whether the current measured loss is at or below the target. A moving-average filter is used to reduce variance in the acceptance process, however this algorithm is immature. It is of interest as it makes direct use of measurements of the desired outcome (a target loss-ratio) to drive the AC algorithm – this approach is the only MBAC algorithm presented here that does not use utilisation measurements as input.

The final AC algorithm, AC-T, is not measurement-based. AC-T, the "target" algorithm is provided to compute results based upon an admission process whereupon a given number of flows are present in the network at any time. Such an algorithm allows computation of the performance-frontier illustrated in Section 4 and described more fully therein.

The above AC algorithms are summarised in Table 1 along with the key idea behind each algorithm and several comparison criteria.

Table 1

Summary of admission control algorithms compared. Rate Envelope Multiplexing (REM) describes the approach where the effect of buffering is not taken into account, while Rate-Sharing Multiplexing (RSM) takes into account gains made through buffering. Certainty Equivalence describes MBAC algorithms that are not robust to the random properties of measurements

AC algorithm	Key idea	Buffering effect	Certainty equivalent?
AC-ST	Simple Threshold	RSM	Yes
AC-AR	Acceptance Region	REM	No
AC-CB	*Chernoff Bounds*	REM	Yes
AC-MS	Measured Sum	RSM	Yes
AC-MPFE	Large-Deviation Theory	RSM	Yes
AC-MAE	Large-Deviation Theory	RSM	Yes
AC-MVE	Mean-Variance Estimator	RSM	No
AC-KQ	Traffic Envelope	RSM	No
AC-LBE	Loss-ratio	RSM	Yes
AC-T	Target	–	–

Whether the gains made though buffering are taken into account by an AC is indicated by the column "Buffer Effect" of Table 1. The Rate Envelope Multiplexing (REM) approach does not take the effect of buffering into account, while Rate-Sharing Multiplexing (RSM) takes into account gains made through buffering. For RSM the combined rate at which data enters the buffer may exceed the buffer service-rate for small periods before packet-loss will occur. Because of this, the burst rate and burst duration of sources play an important contribution in the computation of the *effective bandwidth* of sources. For REM the burst duration and burst rate need not be considered. As a result, the *effective bandwidth* of the loss is computed from only the sustained rate and the peak rate. For algorithms that are measurement-based the incorporation of an RSM approach may be implicit, such as that of AC-MS or explicit such as AC-KQ or AC-MPFE.

Table 1 also lists which of the MBAC algorithm are based upon Certainty Equivalence.

Certainty Equivalence in AC algorithms is the use of a static AC algorithm but with the insertion of measurement derived estimations rather than those computed from *a priori* traffic descriptors [28]. The attraction in this approach is the ability to reuse existing ACs. However, CE methods may give too optimistic results due to the reliance upon measured quantities [15]. The random nature of measured quantities must be incorporated into the algorithm. However, several MBAC algorithms document specific solutions to this measurement problem [13,15,19,20]. The approaches of [15] and [20] incorporate computation of the measurement variance directly while in [13,19] a solution is provided using an estimator based upon a Bayesian model that incorporates the error as the prior is developed.

Comparisons among MBAC algorithm are conducted in the Section 4 on the basis of experimental results.

3. Method

This section outlines the method adopted in the examination and comparison of AC algorithms presented in this work. Throughout this study it is assumed that network users will make requests for new flows be admitted using a protocol such as RSVP [5] or ATM Forum's signaling specification version 4.0 [2]. Under these protocols, each service request contains a traffic descriptor of the worst-case behaviour of the traffic requesting admission.

This section also outlines the evaluation environment in Section 3.2, including broad assumptions common to each experiment scenario, then each experiment configuration is outlined in Section 3.3.

3.1. Criteria

A variety of criteria for identifying the *best* MBAC algorithm have been put forward by previous authors. Suggestions for comparison criteria have included packet-loss versus utilisation and flow-acceptance rates versus utilisation [6,7,16,17]. This paper presents results of comparisons made using packet-loss versus utilisation, a common method for illustrating AC algorithm performance.

Packet-loss versus utilisation results are presented for several experiments, each of which uses different combinations of flow characteristics (arrivals and holding times) carrying a variety of traffic types. The experiment configurations are outlined in Section 3.3.

Each MBAC algorithm studied here incorporates parameters to control the performance of the algorithm. The exact control exercised both for algorithms that purport to have a relationship between control parameter and performance, and for those algorithms that do not have such a relationship provides a useful comparison. The performance profile of each MBAC algorithm is studied for a range of control-values using each of the experiment configurations of Sections 3.3 to ensure coverage across a range of different flow and traffic conditions. As these control parameters affect the desired performance criteria, comparing the performance allows an insight into algorithm behaviour.

3.2. Experimental environment

The unique characteristic of this investigation has been that the results are gained using implementations of the algorithms not in a simulator but in an experimental network. An implementation of each algorithm has given access to performance and behaviour aspects of each algorithm not available to a simulation. Relevant aspects of the implementation environment, along with the experimental approach are discussed in this section.

The experimental environment used in the evaluation of MBAC algorithms allows for the implementation of a AC algorithm in a pre-existing framework of connection generation, traffic generation, ATM network and measurement systems [22]. This modularised environment allows us to make direct comparisons of one AC against another when placed under identical connection loads and traffic types. Additionally, we can compare consecutive experiments using one AC algorithm where tuning parameters for that algorithm are adjusted for consecutive runs. In operation, the test environment allows extraction of variables indicating the performance and behaviour of the AC algorithm under test.

For each experiment in this paper the buffer offers an undifferentiated FIFO tail-drop service; the packet-loss due to buffer overflow will be borne by one or more flows currently in progress without any differentiation in the buffer between the flows themselves.

Aside from providing a realistic work load, the mean rate of flow-admission attempts is selected to ensure the AC algorithm under test is placed under a high load. This is needed to ensure that each test contains enough attempts for a meaningful comparison to be made.

3.3. Traffic configuration

The traffic used in this study consisted both of sources generated from deterministic models and traffic created from actual systems. This allows MBAC algorithms to be tested against traditional, Poisson-model traffic, through to sources of traffic currently in evidence in modern networks.

The eight traffic sources used in this study include those based upon Markovian models (TP10S1, VP64S23), Pareto models (PP10S1) and constant-rate models (V64S64) [23]. Additionally, traffic streams VP25S4, RP10S1 and EP6S480k each represent traffic-loads derived not from deterministic models but from actual real traffic loads.

VP25S4 is a controlled-load source representing the carriage of video stream data in packets within a stream with a pre-defined peak-rate, sustained-rate and both maximum and mean burst-sizes. RP10S1 represents IP traffic recorded from a LAN, while EP6S480k is used to represent a stream of IP traffic as would be found connecting sites across a wide-area network.

Using the traffic summarised in Table 2, eight experiments, given in Table 3, were constructed. A summary of the experimental configuration is provided here [23].

Table 2
Summary of traffic sources

Sources	Summary	Parameters
TP10S1	2-state ON-OFF Markov	10 Mbps Peak, 1 Mbps Mean, 1325 octets Mean Burst Size (MBS)
PP10S1	2-state ON-OFF Pareto	10 Mbps Peak, 1 Mbps Mean, 1325 octets MBS
VP64S64	Voice channel uncompressed	64 kbps Peak, 48 octets BS
VP64S23	Voice channel with compression	64 kbps Peak, 22.5 kbps Mean, 2880 octets MBS
VP25S4	Video data stream	25 Mbps Peak, 4 Mbps Mean, 75024 octets MBS
RP10S1	Internet LAN traffic	10 Mbps Peak, 1 Mbps Mean
EP6S480k	Internet WAN traffic	6 Mbps Peak, 480 kbbs Mean

Table 3
Summary of experiments

Experiment	Sources	Summary	Distribution	Mean
Exp1	TP10S1	Flow arrival rate	Markovian	10 fps
		Flow holding time	negative exponential	10 s
Exp2	PP10S1	Flow arrival rate	Markovian	10 fps
		Flow holding time	negative exponential	10 s
Exp3	RP10S1	Flow arrival rate	Pareto	12.5 fps
		Flow holding time	log-normal	160 s
Exp4	VP25S4	Flow arrival rate	Pareto	12.5 fps
		Flow holding time	log-normal	300 s
Exp5	RP10S1	Flow arrival rate	Pareto	12.5 fps
		Flow holding time	log-normal	160 s
	VP25S4	Flow arrival rate	Pareto	12.5 fps
		Flow holding time	log-normal	300 s
Exp6	RP10S1	Flow arrival rate	Markovian	10 fps
		Flow holding time	negative exponential	10 s
	VP25S4	Flow arrival rate	Markovian	10 fps
		Flow holding time	negative exponential	10 s
Exp7	EP6S480k	Flow arrival rate	Markovian	5 fps
		Flow holding time	log-normal	300 s
	VP64S64	Flow arrival rate	Markovian	2.5 fps
		Flow holding time	log-normal	300 s
	VP64S23	Flow arrival rate	Markovian	2.5 fps
		Flow holding time	log-normal	300 s
Exp8	VP64S64	Flow arrival rate	Markovian	5 fps
		Flow holding time	log-normal	300 s
	VP64S23	Flow arrival rate	Markovian	5 fps
		Flow holding time	log-normal	300 s
	EP6S480k	Background flows	constant	8
		Flow holding time	constant	∞

Exp1 through Exp6 are based upon a test network consisting of a simple *dumb-bell* topology between sources and sinks, the networks active component consists of a buffer at the bottleneck with a capacity of 512 packets (27136 octets) and a service-rate of 100 Mbps. 512-packets is a common buffer size found in the line-card interfaces of commercial switch equipment used in the test-environment [11].

Exp1 and Exp2 provide the MBAC algorithm under test with a flow-lifetime and arrivals process that is Poisson, differing only in the type of traffic carried. Exp3 and Exp4 each carry real-world traffic, IP LAN traffic and video traffic respectively. Exp5 combines IP LAN traffic and video traffic to emulate the situation faced at a border switch, which must multiplex a number of heterogeneous lower rate sources onto a higher capacity link. This situation could reasonably have been expected to exist when common desktop bandwidth (e.g., 10 Mbps Ethernet) was substantially lower than backbone or intra-office capacity (e.g., 100 Mbps Ethernet or 155 Mbps OC-3 Sonet).

Using a traffic configuration similar to Exp5, Exp6 uses Markovian distributions, thus providing an interesting comparison between two systems, one with LRD flow properties and one with Markovian properties.

The design of Exp7 and Exp8 was to emulate issues faced in admission control in an ADSL facility. The configuration consisted of a single, uni-directional, bottleneck. The buffer at the bottleneck has a capacity of 512 packets (27136 octets) and a service-rate of 6 Mbps.

Exp7 gave the MBAC algorithm control of constant-rate voice traffic, compressed voice traffic and wide-area network IP traffic. In contrast, Exp8 presented the MBAC algorithm under test with a constant background of IP traffic, providing it with admission control over only the compressed and uncompressed voice traffic.

4. Results

The results of this section present comparison of MBAC algorithm using a number of different criteria. In Section 4.1, results are presented using comparison criteria that have been commonly used by previous authors (e.g., line utilisation versus packet loss, packet loss versus flow acceptance rate). Section 4.2 investigates the mechanism by which a user specifies the performance-objective and how this varies among different MBAC algorithms. By contrasting MBAC algorithm behaviour, this approach allows insight into the independence of such controls for different traffic types.

4.1. Traditional performance criteria

4.1.1. Line-utilisation versus packet-loss

The relationship between the data-loss and line utilisation, herein referred to as the loss-load curve, has been used by a number of previous papers [12,16,17].

Comparisons are intended to show how well each MBAC algorithm performs relative to each other or to an "optimum" loss-load curve. However, such results are shown to give similar relationships between loss and load for a particular configuration of traffic and buffer characteristics independently of the MBAC in use.

While one conclusion made using loss-load results is that such comparison is flawed and reveals little useful information, the omission of such a commonly used comparison would be difficult to justify. In [7] an algorithm, (referred to therein as *Quota*), was used to derive the performance-frontier values for MBAC algorithm behaviour when faced with particular traffic. The performance-frontier may be considered the best possible performance for any given criteria. Such an algorithm allows the computation of optimal results for a given level of flow-arrival activity and traffic by allowing a fixed number of flows to be active at any time. As noted in Section 2, the AC-T algorithm implemented in this work serves an identical purpose.

Figure 1 presents the loss-load results of a number of AC algorithms conducted against a representative sample of the experiments proposed in Section 3. The trend towards a utilisation boundary for a given loss ratio is evident. The variance will be higher as the loss ratio is decreased as a consequence of the experiment run-length noted in [22]. For each graph in Fig. 1, this trend towards increasing variance as the loss ratio is decreased is clear.

Figure 1 illustrates how graphs of data-loss versus line utilisation return a similar function regardless of the AC in use. The small deviations from the loss-load achieved by the target algorithm (AC-T) may be because of MBAC algorithm behaviour [7]; the errors resulting from sampling error are undocumented. It appears that an implicit assumption of correct or near-correct results is made. However, by examining the results gained here it is clear that part of the behaviour may be apportioned to the algorithm but part will also be due to the error arising from the environment.

In [22], an error margin due to sampling error of $\approx 0.6\%$ with 95% confidence is noted, for a loss ratio of 1×10^{-3}. However, when testing MBAC algorithms repeatability, [22] noted an error margin of $\pm 5\%$ with 95% confidence for the same loss-rate. This implies, with 95% confidence that an error margin of up to 5% is present in the results for 1×10^{-3}. For smaller loss-ratio values this error will be increased: with 95% confidence, an error margin of 12% is present in the results for 1×10^{-5}, and the margin due to sampling-error will have increased to $\approx$6% resulting in a total potential error-margin in results of 18%. Such a large error-margin is clear in the results of Fig. 1(a) so its contribution should not be ignored. Alongside the experimental data, Fig. 1 illustrates the results of a fluid flow approximation for the TP10S1 traffic source. While the experimental and theoretical results have similar slope, differences are clear [1]. This may be speculated as being a result of the fluid flow approximation and perhaps further serve to justify experimental evaluations alongside theoretical or simulation studies.

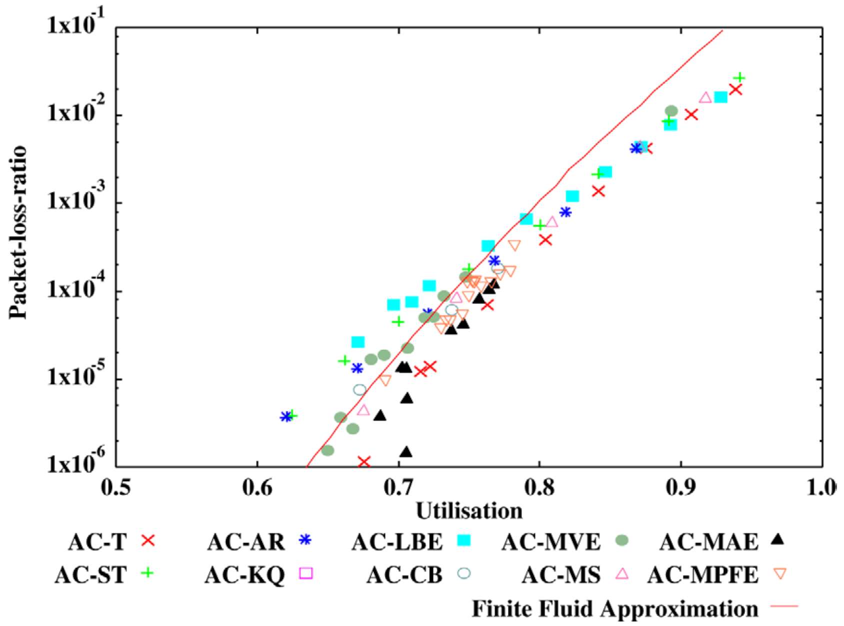

(a)

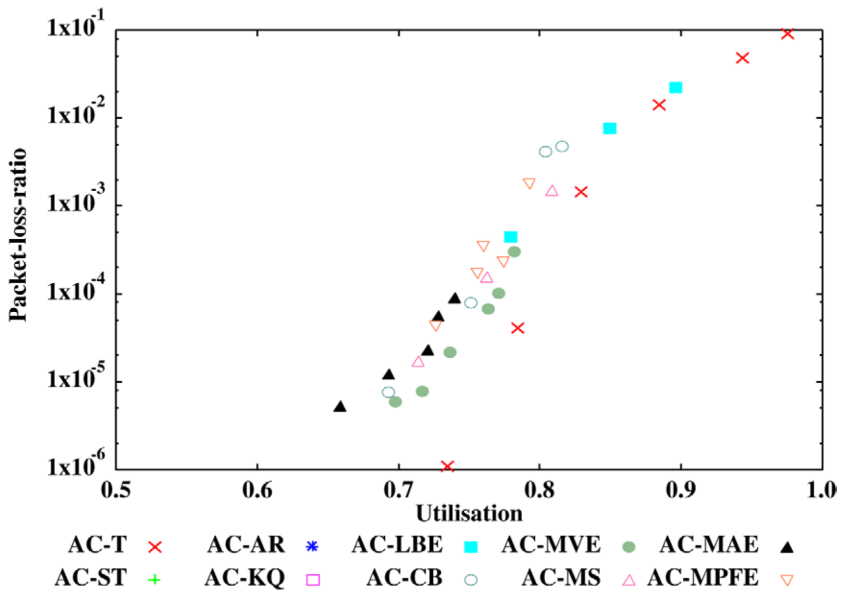

(b)

Fig. 1. Packet-loss-ratio versus line utilisation. (a) Exp1: 2-state ON-OFF Markovian. (b) Exp3: Internet traffic.

Each AC algorithm was operated over a reasonable range of parameters. However in several cases (e.g., AC-KQ for Fig. 1(a)) the reasonable range of operating parameters, e.g., a target loss rate between 1×10^{-1} and 1×10^{-6}, did not generate results that fell within the graphed area. It may be assumed that if an algorithm did not create results to be plotted, the experiment did not generate a sufficient level of loss.

While at first this may seem to be a sure indication of the limitations and the effectiveness of an algorithm, it results from a quite different cause. Algorithms such as AC-KQ maintain an upper boundary on the loss-ratio, the algorithm is based upon worst-case behaviour of the measured traffic. Thus a lack of loss is an indication that to get the algorithm to achieve arbitrary levels of loss would require operation outside what was a reasonable range for the parameters (e.g., an artificially high loss boundary).

Grossglauser in [15,28] noted that MBAC algorithms commonly compensate for errors introduced as part of the measurement process by using a conservative handling of the measurements themselves. A common tuning variable for an MBAC algorithm is how conservatively the measurements will be handled. Subsequently, an MBAC algorithm can achieve a given loss-load curve target. This has been done, however, by removing any safety margin to account for poor measurements and it increases an algorithm's reliance on the measurement characteristics – such as variance due to measurement period.

Therefore, algorithms able to achieve a particular point on the load-loss curve have done so by sacrificing any safety-margin (maintaining the QoS guarantee) for all flows in exchange for higher utilisation. If a QoS constraint, such as a loss ratio, was a bounded agreement across all flows, then such algorithms operating outside the safety margin could not make such a QoS guarantee to all flows in the system.

4.1.2. Flow-blocking versus packet-loss

An alternative comparison criterion was to consider the curve relating flow-blocking and packet-loss [16]. If the packet-loss requirements are strict, a high flow-blocking probability will occur; similarly high packet loss will go with a low blocking probability.

The results presented in Fig. 2 and Fig. 3 allow this idea to be explored further. In the experiments conducted for these results, the relationship between acceptance rate and utilisation is near identical for each MBAC; additionally, as would be predicted by these results, the acceptance rate and packet-loss-ratio also express a clear relationship. This can be seen in Fig. 2(a) and 2(b), particularly when compared with the counterparts expressing the relationship between loss and utilisation: Fig. 1(a) and 1(b). However, in each case the flows carry traffic that is homogeneous and the processes describing flow-arrivals and flow-lifetimes are statistically stable: each using fixed mean values with a distribution based upon an exponential decay. This is of greater interest when the process controlling flow-attempts are varied, and the traffic itself is heterogeneous.

For Fig. 3(a) the results of a heterogeneous experiment are illustrated: unlike the results of earlier Fig. 2(a) and 2(b), different flow-acceptance behaviour is occurring dependent upon the particular AC algorithm. Algorithms such as AC-AR, AC-ST and AC-MVE that do not implement an admission decision dependent upon the declared parameters of new flows have results that are clustered in the top-left of the figure. In contrast, algorithms that use a pessimistic admission process, give results that are clustered in the lower right of the figure. A pessimistic admission process is where an admission decision is only taken if the declared-rate of the new flow may be admitted into the multiplex without impacting current flows in progress. Such a pessimistic admission process is employed by algorithms such as AC-MPFE, AC-MS, or AC-MAE.

The clustering is related to the admission decision used by each particular AC algorithm. Those algorithms that use a pessimistic admission decision will be biased towards flows with low declared peak-rate values. As a result, a global acceptance ratio may serve little value, giving information about an AC algorithm only when it is under one particular flow-load. Emphasising this point, Fig. 3(a) and 3(b) illustrate results of Exp5 and Exp6. Both experiments use the same heterogeneous traffic mix made up of flows with either a 10 Mbps peak-rate or a 25 Mbps peak-rate. However, the flow arrival and departure characteristics are different between the two experiments.

Acceptance ratios, intrinsically tied as they are to the utilisation process, may reveal little that allows comparison of MBAC algorithms. However, an area in which MBAC algorithms do differ from one another is the difference between acceptance rates for a mixture of flows arriving with different traffic types. Figure 4 plots the acceptance ratio for the two different traffic types used in Exp6: a mix of traffic with peak-rates of 10 Mbps and 25 Mbps.

The AC-MS algorithm implements a declared-rate decision: one that compares the declared-rate of new flows to see if this declared rate may be added to the current flow, and it is characterised by a discrimination towards the class of flows that declare a lower peak-rate. Interesting effects of this admission-process are evident as the acceptance rate drops. For AC-MS, as the acceptance rate of all flows falls below 0.45, (marked **A**), the flows declaring a large peak-rate are accepted at ever-lower rates while the flows declaring a small peak-rate maintain

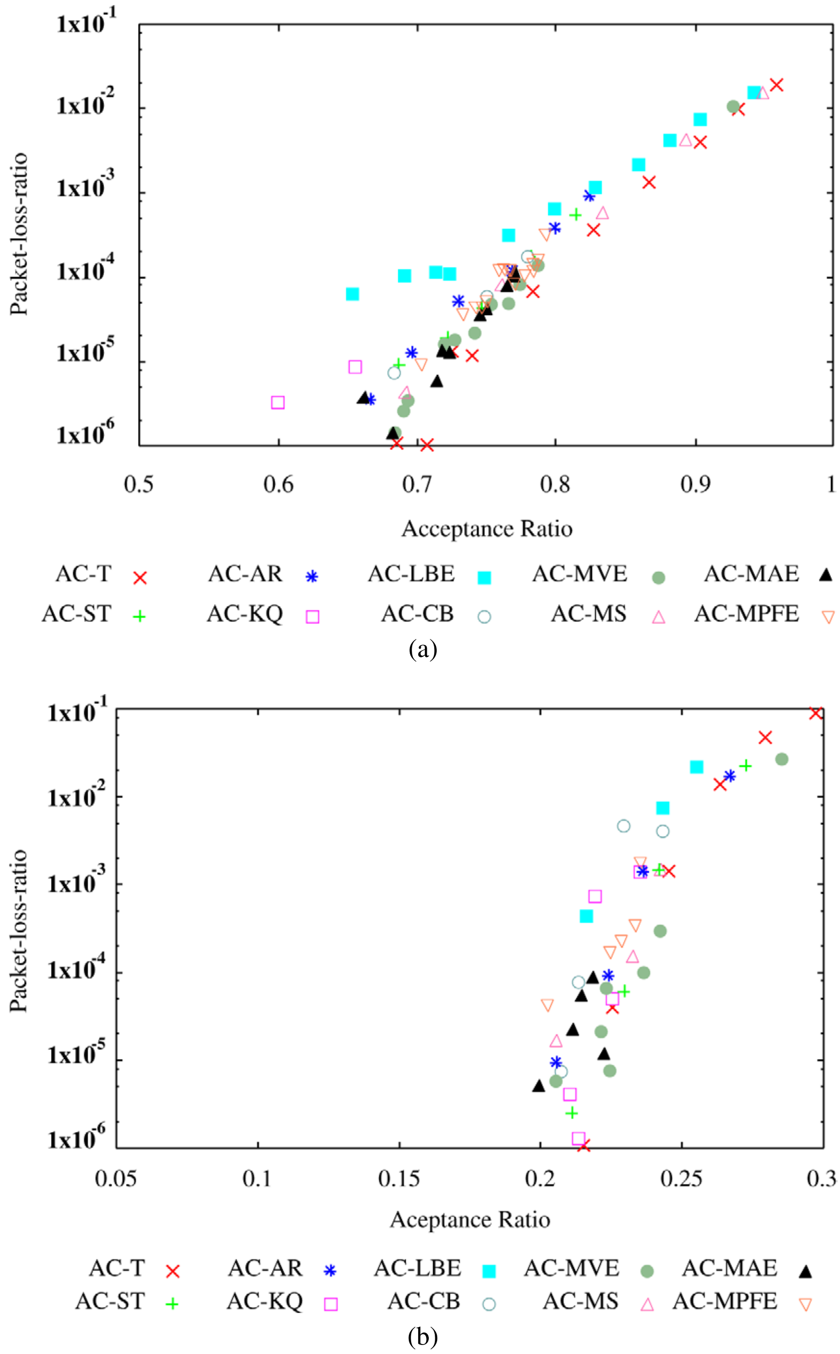

Fig. 2. Packet-loss-ratio versus flow acceptance ratio. (a) Exp1: 2-state ON-OFF Markovian. (b) Exp3: Internet traffic.

a constant level of acceptance. This characteristic is caused by the discrimination of an MBAC algorithm against larger flows, smaller flows are able to be accepted where larger flows are not.

This leads to the conclusion that the admission-process alone will dictate the differences between the admission characteristics of many AC algorithms. Algorithms that differ in admission decision, particularly those using a

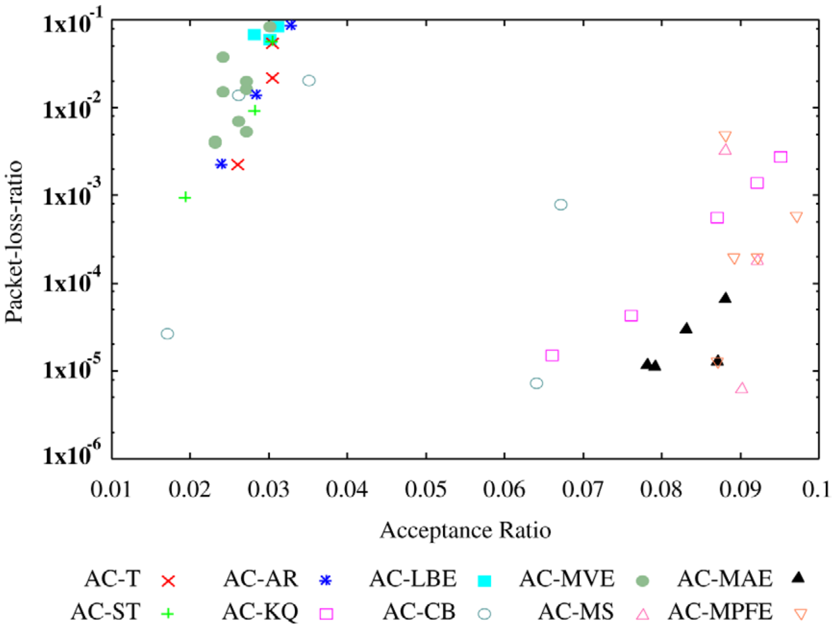

(a)

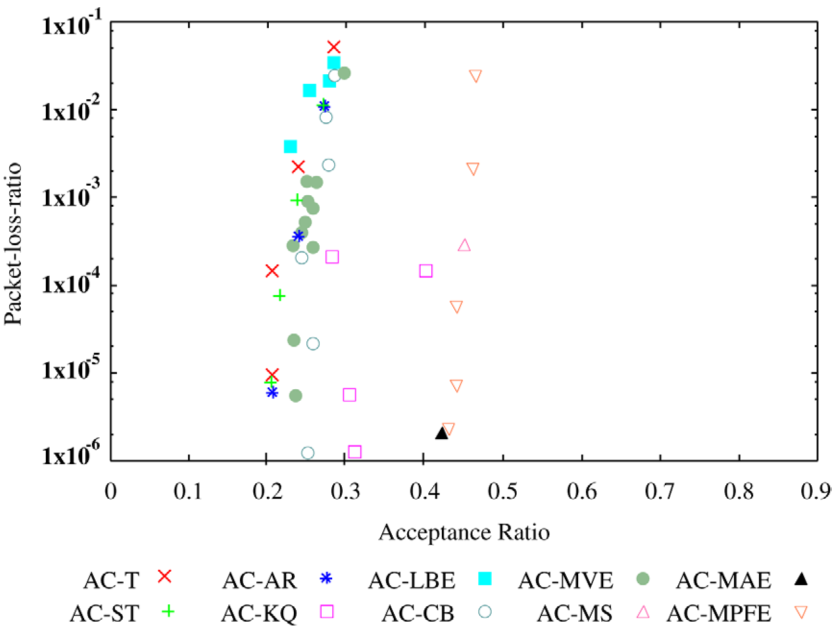

(b)

Fig. 3. Packet-loss-ratio versus flow acceptance ratio. (a) Exp5: Internet traffic/Video Streams. (b) Exp6: Internet traffic/Video Streams.

simple admission decision versus those with a pessimistic admission process will give significantly differing results in the admission for each traffic class. Additionally, algorithms such as AC-CB that incorporate the declared parameters of classes directly into the MBAC algorithm will also generate results different from the generic MBAC algorithm.

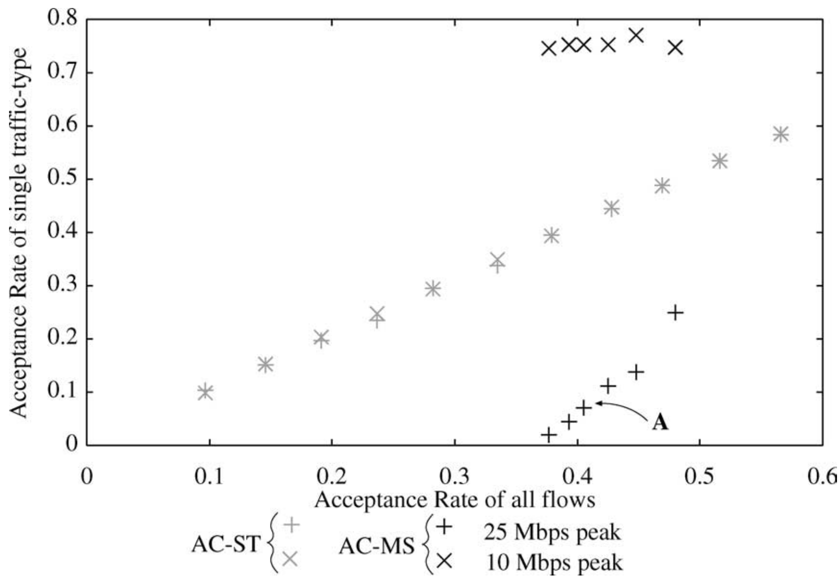

Fig. 4. Flow admittance ratio for traffic type.

4.2. *MBAC algorithm parameters*

The previous section revealed how criteria for performance comparison among AC algorithms commonly used by previous authors (loss-load and loss-acceptance results), do not present a complete picture of an AC algorithm's performance. This section presents results illustrating the control of algorithm by their respective parameters.

Figure 5 gives results in which the measured loss ratio is plotted against a control parameter of each AC algorithm. In each of these figures a data-point represents a single experiment, additionally, lines had been added to assist visualising the relationship between the control parameter and loss ratio. In each set of results it is clear that no universal algorithm exists that is able to control the variety of traffic presented by the experiments of Section 3. Of particular interest in Fig. 5 is the independence the algorithm has from the traffic type. Such independence will relate to the desired outcome (the desired outcome for these examples is a QoS guarantee of packet-loss).

While every algorithm has at least one specific control parameter, many such as AC-MAE or AC-KQ have a variety of other additional parameters. Also, the measurement period is a basic parameter to all MBAC algorithms, although other parameters such as the number of samples also become parameters. This section restricts itself to a study of the principle control parameter as proposed with each algorithm. Readers should be mindful, however, of the existence of many additional parameters incorporated into each algorithm.

Figure 5(a) presents results for experiments using the AC-ST algorithm using a common measurement interval of 10 ms. In contrast, AC-AR, uses a similar threshold-based technique but proposes an integrated adaptation to the acceptance region (from which the threshold value is derived). This implies that while the threshold versus loss-ratio curves would resemble those of AC-ST, the algorithm would adapt the threshold to the conditions of traffic, the QoS guarantee to be met, as well as buffer size and link capacity.

Using a control that stipulates a level of utilisation, AC-MS gives the results of Fig. 5(b). Specifying the utilisation results in different traffic types being treated in a uniform manner. Utilisation shown in Fig. 1 does not express a clear, traffic-independent relation to the loss ratio. However, if the QoS were based upon a guaranteed level of link utilisation, this mechanism may be appropriate.

The AC-CB algorithm (Fig. 5(c)) uses an un-calibrated control as a pre-multiplier on a factor derived from the flow's peak-rate. Thus, this value controls the degree of robustness the algorithm will have to flows. A small scalar allows more admissions and, thus, more (potential) packet-loss by reducing the impact the declared peak-rate has

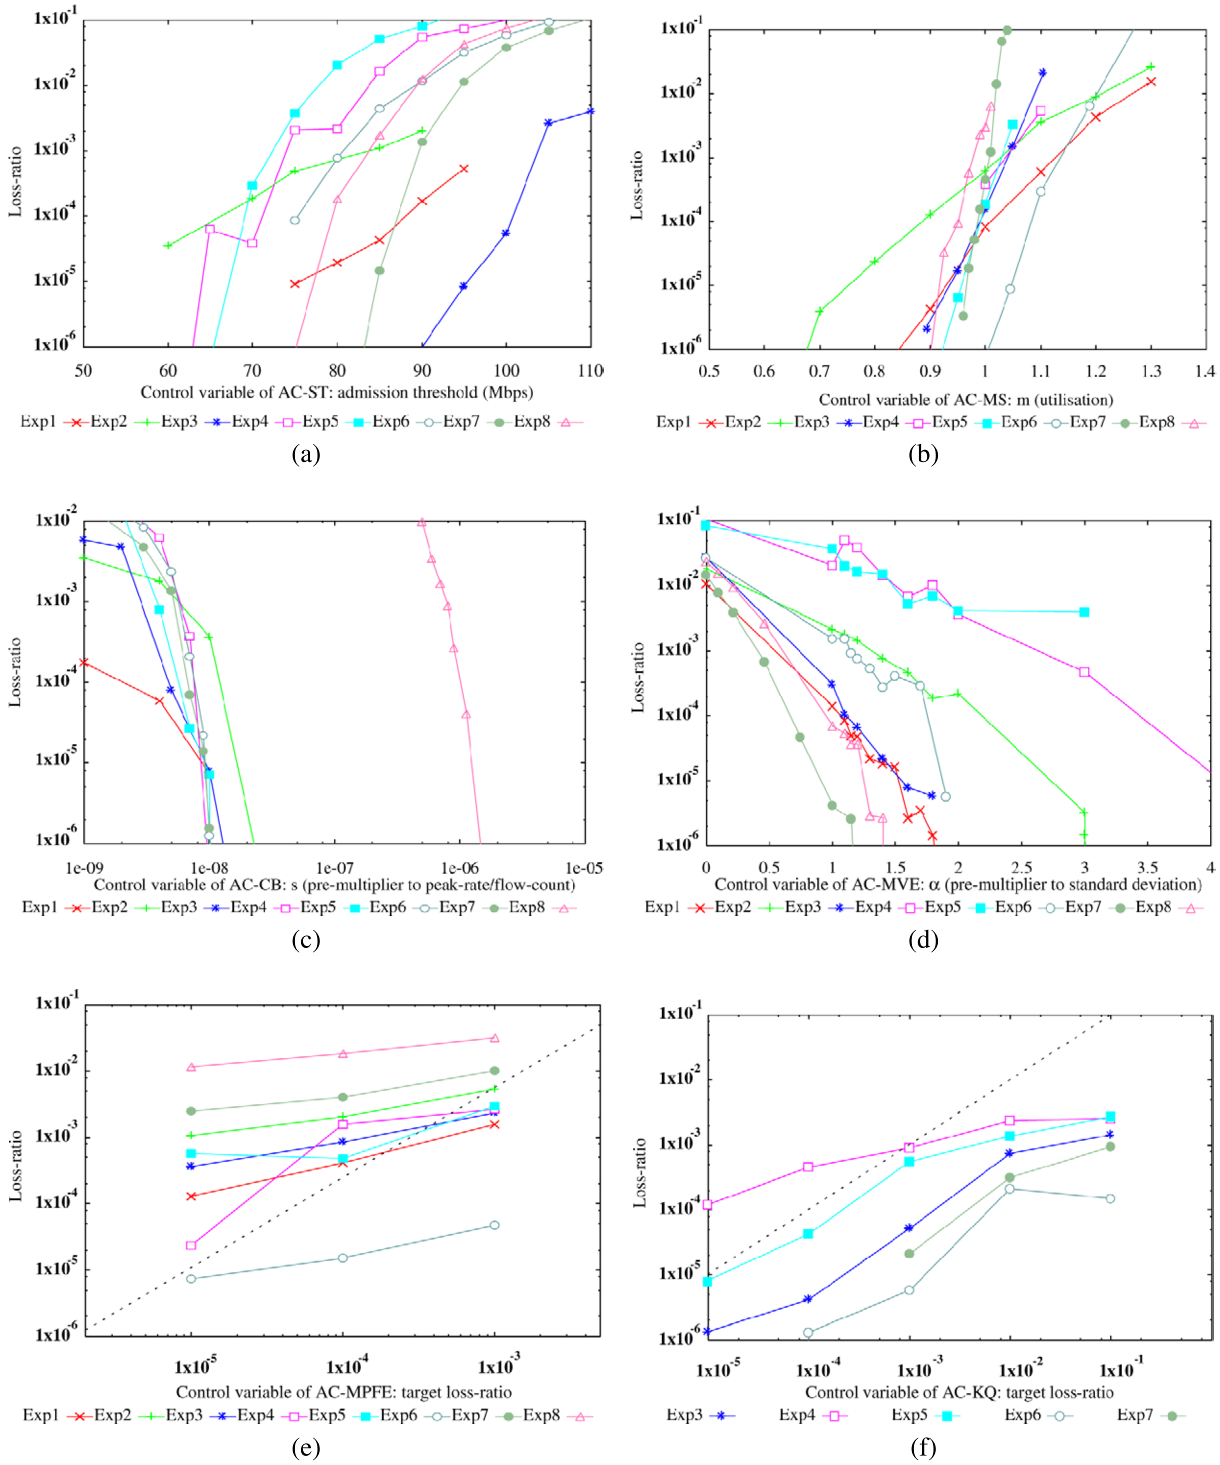

Fig. 5. Packet-loss-ratio versus control parameter. (a) AC-ST. (b) AC-MS. (c) AC-CB. (d) AC-MVE. (e) AC-MPFE. (f) AC-KQ.

upon the system, while a large scalar value has the opposite effect of reducing potential for packet-loss. Clearly, this value is of limited use when selecting an objective as a particular loss ratio.

Using AC-MVE, an algorithm based upon mean and variance measurements, returns the results shown in Fig. 5(d). These results illustrate no clear relationship between the variance of measurements made over one timescale and the measured packet-loss when this estimator is incorporated into an AC algorithm. The control parameter affects the contribution the measured variance will make to the estimate.

Figures 5(e) and 5(f) plot results of algorithms that allow the nomination of a target loss ratio. In each of these figures, a dotted line is added to aid in comparing target with measured loss ratio.

The AC-MPFE algorithms, (Fig. 5(e)), has multiple tuning factors. Aside from a target loss ratio similar to other measurement-based estimators, this MBE will characterise traffic over one period. The disadvantage in using only one characterisation period, (a common block length was used for all the results of Fig. 5(e) is that only one period may not characterise traffic correctly for all desired loss-ratios.

The differences each traffic type may have as well as the different behaviour that may arise for measurements made for the same traffic at different measurement-lengths leads to the need to characterise traffic over several measurement periods, such is the approach of AC-KQ.

The results of AC-KQ also indicate considerable variation for different traffic types, Exp1, Exp2 and Exp8 are not shown, as no significant loss ratio was recorded in these experiments. The objective of AC-KQ's estimator is to limit the loss of flows. Of all the algorithms that allow specification of target loss ratio AC-KQ is the only one where the majority of loss results are maintained below the target loss ratio.

The control parameters of a range of AC algorithms have been compared; several algorithms use parameters related to utilisation, others use parameters intended to control the loss-ratio while those such as AC-CB use a control parameter that may best be described as uncalibrated. Each type of control has particular application: utilisation controls are necessary should it be important to elect the level of network utilisation, while bounding the loss-ratio is critical for a system that attempts to maintain such a QoS guarantee to flows in the network. In contrast, the uncalibrated control of AC-CB may prove useful for traffic that adapts its packet-loss to current conditions, e.g., elastic traffic, although AC-CB is still measurement based and may serve as a unique approach in this respect.

5. Conclusions

We have discussed and compared a subset of MBAC algorithms. However, unlike previous comparisons [7,16,20], the approach here has been to implement the MBAC algorithms in a purpose-built test environment that allows a modular substitution of one MBAC for another between consecutive test runs. The result has allowed a high-fidelity comparison of MBAC algorithms, testing each against real-world traffic sources.

- **Comparison of loss-load curve results provide little useful information.** Our results, derived from an implementation-based comparison using both model-based and real traffic sources support the conclusions of Breslau et al. [6]. Results illustrating the relationship between system loss and load are defined by the traffic and of the resources of the network, not by the admission control process in use.
- **The admission control decision process may be characterised using loss-acceptance curves.** The decision process followed for a new admission attempt, may be usefully characterised from results illustrating the relationship between system loss and the acceptance of flows. Acceptance ratios also provide a useful characterisation of the decision process when faced with heterogeneous admission attempts.
- **Limited relationships exist between MBAC algorithm control parameters and achieved system loss.** Few algorithms provided any sort of calibrated control to elect a particular loss ratio and only one algorithm (AC-KQ) returned acceptable results.
- **Simulator-based comparisons of MBAC algorithms do not identify fundamental flaws that implementation-based comparisons reveal.** The majority of MBAC algorithms have no allowance for the statistical nature of measurements, however this aspect is only revealed when these algorithms are implemented and performance is tested against a wide-variety of traffic conditions.

By providing an implementation-based comparison of MBAC algorithms, this work supports some of the findings of previous authors while simultaneously offering new insight into the Admission Control problem and the use of MBAC algorithms in its solution.

5.1. Future work

Differences between MBAC algorithms lie in the manner in which each approach the desired point on the utilisation curve. Thus, any future comparison of MBAC algorithms must encompass predictability (the ability to achieve a given point on the loss-load curve), stability (the speed at which recovery from changing circumstances can be effected), and fairness (how an MBAC treats flows of different characteristics). These issues are in addition to those of measurement, computation and memory overhead, the relationship between the MBAC decision-process and flow characteristics, or the association between estimator and measurement characteristics. Any future work investigating MBAC algorithm and Measurement-based Estimation in general needs to attend these areas to achieve complete coverage of the topic of measurement-based management.

Acknowledgements

The author extends his thanks to members of the Cambridge University Computer Laboratory who assisted him through the course of this work. Thanks also to Zhong Fan and David K. Hunter for their feedback on earlier drafts of this paper.

References

[1] D. Anick, D. Mitra and M.M. Sondhi, Stochastic theory of a data-handling system with multiple sources, *Bell System Technical Journal* **61**(8) (1982), 1871–1894.

[2] ATM Forum, *Traffic Management Specification, Version 4*, 95-0013R8, October 1995.

[3] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang and W. Weiss, An Architecture for Differentiated Service, RFC 2475, IETF, December 1998.

[4] R. Braden, D. Clark and S. Shenker, Integrated Services in the Internet Architecture: an Overview, RFC 1633, IETF, June 1994.

[5] R. Braden, L. Zhang, S. Berson, S. Herzog and S. Jamin, Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification, RFC 2205, IETF, September 1997.

[6] L. Breslau, S. Jamin and S. Shenker, Measurement-based admission control: What is the research agenda? in: *Proceedings of ACM SIGCOMM'99*, Cambridge, MA, 1999.

[7] L. Breslau, S. Jamin and S. Shenker, Comments on the performance of measurement-based admission control algorithms, in: *Proceedings of IEEE INFOCOM 2000*, Tel Aviv, Israel, 2000.

[8] C. Cetinkaya and E.W. Knightly, Egress admission control, in: *Proceedings of IEEE INFOCOM 2000*, Tel Aviv, Israel, 2000.

[9] N.G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K.K. Ramakrishnan and J.E. van der Merwe, A flexible model for resource management in virtual private networks, in: *Proceedings of ACM SIGCOMM'99*, Cambridge, MA, 1999.

[10] N.G. Duffield, J.T. Lewis, N. O'Connell, R. Russell and F. Toomey, Entropy of ATM traffic streams, *IEEE Journal on Selected Areas in Communications* **13**(6) (1995), 981–990.

[11] FORE Systems Inc., *Fore – ATM Switch Network Modules Product Overview*, October 1999.

[12] R.J. Gibbens and F.P. Kelly, Measurement-based connection admission control, in: *Proceedings of 15th International Teletraffic Congress (ITC15)*, 1997.

[13] R.J. Gibbens, F.P. Kelly and P.B. Key, A decision-theoretic approach to call admission control in ATM networks, *IEEE Journal on Selected Areas in Communications* **13**(6) (1995), 1101–1114.

[14] R.J. Gibbens, S.K. Sargood, C. Van Eijl, F.P. Kelly, H. Azmoodeh, R.N. Macfadyen and N.W. Macfadyen, Fixed-point models for the end-to-end performance analysis of IP networks, in: *Proceedings of the ITC Specialist Seminar on IP Traffic Modeling, Measurement and Management*, Monterey, CA, 2000.

[15] M. Grossglauser and D. Tse, A framework for robust Measurement-Based Admission Control, *ACM Computer Communication Review* **27**(4) (1997), 237–248.

[16] S. Jamin and S.J. Shenker, Measurement-based admission control algorithms for controlled-load service: A structural examination, Technical Report CSE-TR-333-97, University of Michigan, April 1997.

[17] S. Jamin, S.J. Shenker and P.B. Danzig, Comparison of measurement-based admission control algorithms for controlled-load service, in: *Proceedings of IEEE INFOCOM'97*, Kobe, Japan, 1997.

[18] F.P. Kelly, Notes on effective bandwidths, in: *Stochastic Networks: Theory and Applications*, Volume 4 of *Royal Statistical Society Lecture Note Series*, F.P. Kelly, S. Zachary and I.B. Ziedins, eds, Oxford University Press, 1996, pp. 141–168.

[19] P.B. Key, Connection admission control in ATM networks, *BT Technology Journal* **13**(3) (1995).

[20] E.W. Knightly and J. Qiu, Measurement-based admission control with aggregate traffic envelopes, in: *Proceedings of 10th Tyrrhenian International Workshop on Digital Communications*, Ischia, Italy, 1998.

[21] M. Kodialam and T.V. Lakshman, Minimum interference routing with applications to MPLS traffic engineering, in: *Proceedings of IEEE INFOCOM 2000*, Tel Aviv, Israel, 2000.

[22] A. Moore and S. Crosby, An experimental configuration for the evaluation of CAC algorithms, *Performance Evaluation Review* **27**(3) (1999), 43–54.

[23] A.W. Moore, Measurement-based management of network resources, Technical Report 528, Cambridge University Computer Laboratory, April 2002, PhD dissertation.

[24] K. Nichols, S. Blake, F. Baker and D. Black, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers, RFC 2474, IETF, December 1998.

[25] K. Nichols, V. Jacobson and L. Zhang, A Two-bit Differentiated Services Architecture for the Internet, RFC 2638, IETF, July 1999.

[26] J. Schlembach, A. Skoe, P. Yuan and E.W. Knightly, Design and implementation of a scalable admission control, in: *Proceedings of International Workshop on QoS in Multiservice IP Networks*, Rome, Italy, January 2001.

[27] I. Stoica and H. Zhang, Providing guaranteed services without per flow management, in: *Proceedings of ACM SIGCOMM'99*, Cambridge, MA, 1999.

[28] D. Tse and M. Grossglauser, A time-scale decomposition approach to Measurement-Based Admission Control, in: *Proceedings of IEEE INFOCOM'99*, New York, NY, 1999.

[29] J. Wroclawski, Specification of the Controlled-Load Network Element Service, RFC 2211, IETF, September 1997.