Measurement-Based Admission Control: A Revisit

Yuming Jiang, Peder J. Emstad, Victor Nicola [*], Anne Nevin
Center for Quantifiable Quality of Service in Communication Systems [†]
Department of Telematics
Norwegian University of Science and Technology
ymjiang@ieee.org, {peder, nicola, anne.nevin}@q2s.ntnu.no

Abstract

Measurement-based admission control (MBAC) is an important technique for providing statistical service guarantees and many MBAC algorithms have been proposed in the literature. It has also been investigated in a previous work that these algorithms achieve nearly the same utilization for a given packet loss rate or in short the same loss-load curve. Based on this investigation, it has been argued that further research on better MBAC equations will likely be a fruitless endeavor since there is little room to improve existing algorithms in terms of utilization. In this paper, we investigate why the MBAC algorithms investigated in the previous work have the same loss-load curve. We find this not surprising. We hence argue that it should be other issues than the loss-load curve that should be taken into account in designing an MBAC algorithm. Based on this argument, we revisit the assumptions made by these MBAC algorithms and examine the rationales of these assumptions when these MBAC algorithms are used in different network scenarios. Finally, we discuss some remaining challenges for MBAC research.

1 Introduction

Real-time multimedia applications are increasingly becoming an indispensable part of Internet traffic. These applications include voice over IP, streaming audio and video, Internet gaming and real-time video conferencing. Such applications, on the one hand, are both delay and loss sensitive and, on the other hand, can tolerate some loss and delay. As a result, statistical service guarantees have attracted a lot of research interest in the past decade. Statistical service guarantees do not provide hard loss rate or delay bound guarantees as deterministic service guarantees do. Nevertheless, for both types of service guarantees, admission control is typically required.

Usually, admission control for deterministic service guarantees uses worst-case analytical bounds as its basis. Such admission control algorithms typically result in low

[*] On leave from University of Twente, The Netherlands.

[†] "Centre for Quantifiable Quality of Service in Communication Systems, Centre of Excellence", appointed by The Research Council of Norway, funded by the Research Council, NTNU and UNINETT. http://www.ntnu.no/Q2S/

network utilization because of the bursty nature of network traffic. For statistical service guarantees, however, it has long been recognized that *measurement-based admission control* (MBAC) can achieve (possibly much) higher network utilization, since worst cases usually happen rarely in real networks.

In measurement-based admission control, an MBAC algorithm uses the *a priori* source characterizations only for incoming flows and for existing flows that have been in the system, it uses measurements to characterize them. While several MBAC algorithms have been proposed to measure the traffic characteristics of each flow in the system, due to their scalability and feasibility problems, they will not be discussed in this paper. We shall only consider MBAC algorithms that measure the traffic characteristics of the aggregate of flows in the system. In a previous work by Breslau, Jamin and Shenker [1], extensive simulations have been conducted to compare the performance of several representative MBAC algorithms. They found that all MBAC algorithms achieve nearly the same utilization for a given packet loss rate or in short the same loss-load curve. Based on this, they made a comment on MBAC research, which is "further research on better MBAC equations will likely be a fruitless endeavor since "there is little room to improve existing algorithms"[2]. Because of their pioneer position in MBAC research, this comment must be critically examined.

The main purpose of this paper is to investigate analytically why different MBAC algorithms investigated in [1][2] are shown to have the same loss-load curve. The analysis shows that the same loss-load curve is indeed to be expected and not surprising. We hence argue that it should be other issues than the loss-load curve that should be considered when designing or comparing an MBAC algorithm. Based on this argument, the second purpose of the paper is to revisit the assumptions made by existing MBAC algorithms, examine the rationales of these assumptions when these MBAC algorithms are used in different network scenarios, and try to explain from the perspective of their analytical bases why none of the MBAC algorithms can meet their targeted service guarantees. Based on this revisit, the third purpose of the paper is to discuss some remaining challenges which need to be addressed for MBAC research.

2 Network Model and Some Preliminaries

We consider a single link network and investigate MBAC on this link. For ease of exposition, we also assume that there is only one traffic class, all considered flows belong to the same traffic class and the total link bandwidth is allocated to this traffic class. In addition, all flows if admitted share a FIFO buffer with finite buffer size and when the buffer is full or the remaining capacity is not enough to accommodate a new incoming packet, the new incoming packet is simply dropped. While this single-link single-class configuration as adopted in [1] is rather simple, the basic performance aspects of the MBAC algorithms to be investigated are most easily revealed. Unless otherwise specified, we further make the following assumptions as implied in [1]:

(A.1) *Each flow is stationary and independent from other flows.*

(A.2) *Each flow uses a signaling protocol such as RSVP or other mechanism to make its request for service to the network, which contains a traffic descriptor describing the* *behavior of the corresponding traffic of the flow.*

(A.3) *All flows require the same statistical service guarantees such as statistical loss guarantee.*

Suppose for the considered single-link single-class network, the buffer size is B (in bits) and the link capacity is C (in bits/second). For each flow i, we model the arrival of its packets as a point process with stationary increments $A_i(t)$ that denotes the amount of traffic (in bits) generated by flow i in the interval $[s-t, s)$. We assume that $A_i(t)$ has mean $m_i t$ and variance $t^2 v_i(t)$. Let a_i^j be the arrival time of the jth packet of flow i to the network and l_i^j be its packet length (in bits). Here for a_i^j, we assume that multiple packets from the same flow i do not arrive simultaneously.

For the aggregate of n such flows, we define $A(t) \equiv \sum_{i=1}^n A_i(t)$ that represents the total amount of traffic generated by the aggregate in the interval $[s-t, s)$. Clearly $A(t)$ has mean $mt \equiv \sum_{i=1}^n m_i t$. In addition, we denote by a^k the arrival time of the kth packet in the aggregate to the network and l^k its packet length (in bits). We assume that the network has a limit on the maximum packet length, which is denoted by L.

2.1 Poisson Limit Theorem

Owing to the Poisson limit theorem, the arrivals a^k of the aggregated flow tend toward Poisson [3][4][5]. In paticular, as n becomes large, their inter-arrival times τ^k tend toward independent and their marginal distribution tends toward exponential [3][4], for which we denote the mean inter-arrival time by $1/\lambda$. Then, based on results for $M/G/1$ systems (e.g. [6] [7]), we have packet loss probability P_{loss} for the aggregate:

$$P_{loss}(b) = 1 - \frac{1}{\rho + 1/\sum_{l=0}^{b} \beta_l}, \tag{1}$$

where b denotes the buffer size (in packets), ρ denotes the traffic load $\lambda E[S]$ in which $E[S]$ is the mean packet service time, and

$$\beta_l = \begin{cases} 1, & l = 0, \\ (\sum_{k=1}^{l-1} \beta_k \alpha_{l+1-k} + \alpha_l)/(1-\alpha_1), & l = 1, 2, \ldots \end{cases} \tag{2}$$

with,

$$\alpha_k = \int_0^\infty [1 - F_S(t)] \frac{e^{-\lambda t} \lambda^k t^{k-1}}{(k-1)!} dt, \qquad k = 1, 2, \ldots. \tag{3}$$

In (3), $F_S(t)$ denotes the probability distribution function of service demand. If all packets in the network have the same length L(in bits), we have for (1), $b = \frac{B}{L}$, $\rho = \frac{\lambda L}{C}$, and $\alpha_l = 1 - e^{-\rho} \sum_{k=0}^{l-1} \rho^k / k!$.

2.2 Central Limit Theorem

For $A(t)$, owing to the central limit theorem, when n becomes large, it tends toward a Gaussian process with mean $mt = \sum_{i=1}^n m_i t$ and variance $t^2 v(t) = t^2 \sum_{i=1}^n v_i(t)$. In other words, when n becomes large, we have for any $r > m$ and any time $t \geq 0$,

$$Pr\{A(t) > rt + \sigma\} \quad = \quad Pr\{\frac{A(t) - mt}{\sqrt{t^2 v(t)}} > \frac{(r-m)t + \sigma}{\sqrt{t^2 v(t)}}\} = \Psi\left(\frac{(r-m)t+\sigma}{\sqrt{t^2 v(t)}}\right) \quad (4)$$

where $\Psi(x) \equiv \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{u^2}{2}} du$.

With (4), the following results have been used to approximate the tail of the steady-state queue length distribution, $Pr(Q > \sigma)$ (e.g. see [8]). First, it is known that the steady-state queue length of an infinite-buffer queue is given by

$$Q \equiv sup_{t \geq 0}(A(t) - Ct)$$

and hence

$$Pr\{Q > \sigma\} = Pr\{sup_{t\geq 0}(A(t) - Ct) > \sigma\}. \quad (5)$$

With (5) and the fact that $Pr\{sup_{t\geq 0}(A(t) - Ct) > \sigma\} \geq Pr\{(A(t) - Ct) > \sigma\}$, we have a lower bound on $Pr(Q > \sigma)$, which has also been used as an approximation of $Pr(Q > \sigma)$ [8], as follows:

$$Pr(Q > \sigma) \geq \Psi\left(\frac{(C-m)t+\sigma}{\sqrt{t^2 v(t)}}\right). \quad (6)$$

Second, under general condition, large deviations theory states that the following approximation is remarkably accurate:

$$Pr\{sup_{t\geq 0}(A(t) - Ct) > B\} \approx sup_{t\geq 0} Pr\{(A(t) - Ct) > B\}. \quad (7)$$

In addition, since $\Psi(x)$ has an upper bound $\Psi(x) \leq \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} < e^{-\frac{x^2}{2}}$, we have the following approximation for $Pr\{Q > \sigma\}$:

$$Pr(Q > \sigma) \lessapprox e^{-\frac{1}{2} inf_{t\geq 0} \frac{[(C-m)t+\sigma]^2}{t^2 v(t)}}. \quad (8)$$

With $Pr\{Q > \sigma\}$ approximated from (6) or (8), a simple approach for estimating loss probability has often been adopted, which is

$$P_{loss}(B) \approx Pr\{Q > B\}. \quad (9)$$

In addition, another approximation has been used in [9]. This approximation is based on the result that $\frac{Pr\{Q>x\}}{P_{loss}(x)}$ converges to a constant as $x \to \infty$, under network assumptions (A.1)-(A.3) and the assumption that packet service times are independent and identically distributed. Then,

$$P_{loss}(B) \quad = \quad \frac{P_{loss}(a)}{Pr\{Q > a\}} Pr\{Q > B\} \approx \alpha Pr\{Q > B\}, \quad (10)$$

where α is determined from $P_{loss}(0)$ computed exactly from a bufferless system and $Pr\{Q > 0\}$ approximated from (8): $\alpha = e^{\frac{1}{2}\frac{[(C-m)t]^2}{t^2 v(t)}} \Psi\left(\frac{C-m}{\sqrt{v(t)}}\right)$.

3 MBAC Algorithms and Comments: A Review

Fig. 1 depicts the structure of MBAC. It shows that an MBAC algorithm for a network system typically includes three elements: (1) admission decision algorithm; (2) traffic estimator; (3) resource estimator. The MBAC algorithm keeps measuring traffic in the system and/or remaining system resources such as available bandwidth and buffer size. Based on the measurement, the traffic estimator estimates how much traffic is in the system, what its characteristics are, and possibly how many flows there are in the system; the resource estimator estimates how much resource remains in the system. When a new flow requests admission to the system, the MBAC algorithm uses the admission control algorithm to decide if this flow can be admitted. This decision is based on the inputs from the traffic estimator and the resource estimator. In addition, the decision also relies on some input from the requesting flow, which typically includes its quality of service requirement and its traffic description.

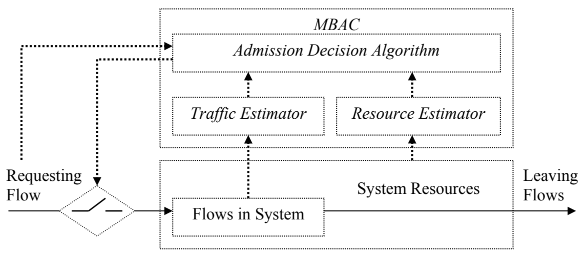

Figure 1: Structure of MBAC

For a specific MBAC implementation, the system in Fig. 1 can be a single node, a domain, or an end-to-end network. If the system is a network domain, the measurement points for the traffic estimator and the resource estimator as well as the admission decision algorithm can be implemented at each node, which results in the *node-by-node MBAC* or *hop-by-hop MBAC*[10], or at the ingress, which results in *ingress MBAC* [11][12], or at the egress, which results in *egress MBAC* [13], or at a central controller such as bandwidth broker in DiffServ [14], which results in *centralized MBAC.* If the system is the whole end-to-end network, the three elements for MBAC may be implemented at end-systems/applications, which results in *endpoint MBAC* [15] [11]. This paper focuses on single-node MBAC algorithms in which the system in Fig. 1 is the single-link. Nevertheless, we believe most discussions in this paper can be extended to the network domain case and the end-to-end network case by viewing the network domain or the end-to-end network as a blackbox or single node.

In the last decade, many MBAC algorithms have been proposed, which explictly or implicitly aim to achieve highest network utilization for a required level of statistical service guarantees. This is done by estimating a priori through measurements the level of such service guarantees that will result. Corresponding to the three elements for MBAC, these algorithms focus on new approaches for better estimating (in terms of accuracy) the characteristics of system traffic and/or that of remaining system resources, and/or on new admission decision algorithms for higher network utilization. While the resource estimator in Fig. 1 has been used in MBAC algorithms for the network domain and end-to-end network cases, it is rarely used in single-node MBAC algorithms since for the single-node case, the system capacity is usually known a priori.

For the traffic estimator, several approaches have been adopted in the literature to achieve better estimation of traffic characteristics of the system. These approaches include Time Window, Exponential Averaging, Point Sample, Adaptive Sampling, and Kalman filter [1] [10] [16]. As shown in Figure 1, the traffic estimator is decoupled from the admission decision algorithm for an MBAC algorithm.

For admission decision shown in Figure 1, many algorithms have been proposed. Based on the assumptions and analyses that they are built upon, these algorithms can be broadly put into the following categories:

Measured Sum (MS): This is possibly the simplest admission decision algorithm for MBAC. Its idea is to measure the load of existing traffic in the system. A requesting flow i is admitted if the following test succeeds [10]:

$$m_i + m < C' \tag{11}$$

where m denotes the measured existing traffic load or rate, m_i denotes the requested rate by flow i, $C' = \alpha C$, C is the total link bandwidth as defined earlier, and α here is a targeted utilization.

Effective Bandwidth (EB): Many MBAC algorithms have been developed based on the concept of effective bandwidth (EB) or equivalent capacity or estimated bandwidth. In these algorithms, the effective bandwidth of each flow and/or the existing aggregate flow in the system is calculated based on Gaussian distribution [18], Hoeffding bounds (HB) [19], measured bandwidth requirement in Measure CAC (MC) [20], or other distributions. With these algorithms, the admission condition is simply

$$m_i^e + m^e < C \tag{12}$$

where m^e denotes the effective bandwidth of the existing aggregate traffic and m_i^e that of the requesting flow. m_i^e and m^e are calculated based on the target loss probability (assuming it is the same for all flows) and the distribution function used for approximating traffic. A class of admission decision algorithms that indirectly use effective bandwidth in admission decision have also been investigated. Particularly, several admission decision algorithms, which are motivated by special choices of Chernoff bound, are proposed in [21]. These choices of Chernoff bound correspond to different tangents to the effective bandwidth function, which include tangent at peak (TP), tangent at arbitrary location, tangent of slope one, and tangent at origin.

Gaussian Approximation (GA): Motivated by the central limit theorem, Gaussian characterization of traffic has been directly used for admission control. Particularly, given a target loss probability or a buffer overflow probability, it is compared with that calculated from (9) [8] [22] or (10) [9][23]. If the target loss is higher than that from (9) or (10), the requesting flow is admitted; otherwise, it is rejected.

A related approach for MBAC is aggregate traffic envelope (TE) [24] [25]. In this approach, Gaussian distribution is used to approximate the stochastic rate envelope of the aggregate traffic, with which admission test is performed. This test can be considered as a variation of the measured sum (MS) algorithm. Specifically, it relies on the mean rate $m(t)$ and variance $v(t)$ of the maximal rate envelopes [24] [25] of the aggregate flow. It admits a requesting flow i if the following test is successful [25]:

$$m_i(t) + m(t) + \alpha v(t) < C, \tag{13}$$

where $m_i(t)$ denotes the rate requested by the incoming flow, α is a controlling parameter set to adjust the performance of this MBAC algorithm, and given a target loss probability $P_{loss} \leq \epsilon$, α is a solution to $sup_t \frac{v(t)G(\alpha)}{m(t)} = \epsilon$,where based on Gaussian approximation, $G(\alpha) = \frac{1}{\sqrt{2\pi}} e^{-\alpha^2/2} - \alpha\Psi(\alpha)$.

In [1], extensive investigation has been conducted through simulation to compare several representative measurement-based admission control algorithms. These MBAC algorithms include MS (Measured Sum [10]), HB (Hoeffding Bounds [19]), TP (Tangent at Peak [21]), MC (Measure CAC [20]) and TE (Traffic Envelope [25]).

Two criteria are used for the investigation and comparison. One is the loss-load behavior and the other is the ability to meet a service guarantee target. The first criterion is intended to show how well an algorithm is able to balance the conflicting goals of providing service guarantees and achieving high network utilization. The second criterion is used to evaluate how well an algorithm can meet its target service guarantee based on its associated tuning parameter.

In [1], it is found that each algorithm has nearly the same loss-load curve and this holds across different traffic models. In addition, the converged loss-load curve does not depend on any particular coupling between estimation and decision processes and the slight variations between the loss-load curves for different MBAC algorithms are caused by the estimation process while not the decision algorithm. In other words, the measurement estimators can be decoupled from the admission control algorithm as shown in Fig. 1. Based on these findings, a critical comment has been made in [2] as:

(C.1) *Further research on better admission decision algorithms with regard to loss-load performance will likely be a fruitless endeavor. This is because results in [1] suggest that there is little room to improve existing algorithms.*

For the second criterion, the performance metric investigated in [1] is loss probability. It is found in [1] that all algorithms are unable to achieve performance close to the targeted one in a consistent manner. In other words, none of the algorithms provides tuning parameters that are useful as performance targets. At best, these parameters can be seen as largely uncalibrated knobs for increasing or decreasing network utilization and the actual performance. With regard to the second criterion, the following comment is drawn in [1]:

(C.2) *None of the investigated MABC algorithms in [1] can reliably match actual performance to targeted service guarantees. The ability of future algorithms to improve in this regard is an open question.*

4 Loss-Load Curve Analysis

In this section, we investigate the *actually achieved* loss-load curve performance of various MBAC algorithms studied in [1], under the same settings as used in [1]. Particularly, it is assumed that the number of flows in the system is large (in the order of 100), all flows are independent, and all packets have the same size. Then, from the Poisson limit theorem, it is known that packet arrivals in the system tend toward a Poisson process regardless which MBAC algorithm is adopted by the system. Motivated by this, we also consider an ideal reference system, which has the same settings as in [1] but with one difference. This difference is that in the ideal system, the flow number is ∞, no matter how much the traffic load ρ is. Note here that in the system investigated in [1], the number of flows in the system varies when the network utilization changes: the less the traffic load, the smaller the flow number. According to the Poisson limit theorem, the arrival process of the ideal system matches with a Poisson process regardless of the type of each flow in the system. In addition, from results presented earlier, its packet loss probability can be calculated from (1).

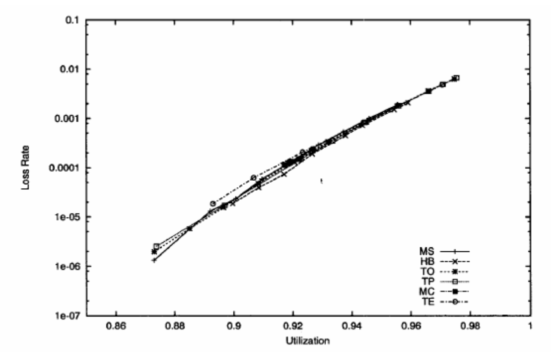

Figure 2: Loss-load curves for MBAC

Fig. 2, reproduced from [2], depicts the loss-load curves for the various MBAC algorithms investigated in [1]. Fig. 2 shows that as commented in [1][2], all MBAC algorithms have the same loss-load performance. Notice that while not shown in [1], the authors have mentioned that the same loss-load curve is achieved when various source models are used, including long range dependent traffic, traffic from a video source and heterogeneous traffic. We believe such matching is owed to the fact that when a large number of flows, as about 300 for Fig. 2 [1], are aggregated, the aggregate flow will have (approximately but very accurately) a Poisson arrival process. Since for this Poisson arrival process, the loss-load curve is fixed given the same buffer size and the same service rate, this matching is not surprising. We believe that the loss-load curve obtained analytically for the ideal system will match with the curves shown in Fig. 2 regardless of the source models used for the investigation. In this regard, we agree on comment (C.1) by providing analytical support for it. Due to the same reason, for delay-load curve, we can have a similar result. To summarize, we have the following:

(B.1) *Loss-load curve should not be used as a performance metric for comparing different MBAC algorithms, since the same loss-load curve can generally be expected for each of them.*

(B.2) *Delay-load curve should not be used as a performance metric either, since the same delay-load curve can also generally be expected for each of them.*

5 Assumptions for MBAC: A Revisit

Existing MBAC algorithms are often based on particular assumptions and requirements. Besides (A.1)-(A.3), the following are some assumptions and requirements that are (fully or partially) needed:

(A.4) *Each incoming flow has a descriptor with* sufficiently detailed *traffic information such as effective bandwidth or mean rate and/or rate variance.* Most existing MBAC algorithms are based on the knowledge of traffic descriptors provided with admission requests. A traffic descriptor, along with measured (or estimated) information of existing traffic and network resources, are used in a mathematical model for predicting the QoS (throughput, loss, or delay) that can be provided. If the computed QoS meets some minimum requirement, the flow is admitted; otherwise it is rejected.

Many MBAC algorithms assume the traffic information carried by the traffic descriptor can be directly used for admission test without further complex calculation. For example, it is often assumed that mean rate or respectively the effective bandwidth of an incoming flow is provided by its descriptor so that (11) or respectively (12) can be immediately used for admission test. Because of this need, a lot of work has been conducted in the literature to calculate the effective bandwidth of a traffic source.

Some other algorithms assume that the traffic descriptor includes sufficient traffic information that can be easily used with the adopted mathematical model. For example, most existing Guassian approximation-based MBAC algorithms are built upon the assumption that the mean rate and rate variance of an incoming flow are either explicitly provided by its descriptor or implicitly obtained based on other assumptions particularly the next assumption (A.5).

(A.5) *Flows are homogeneous.* Many MBAC algorithms (e.g., those based on Gaussian approximations [23][26]) assume homogeneity of flows. This assumption leads to considerable simplification of mathematical models used for QoS predictions of both the incoming flow as well as the existing (aggregate) traffic. For example, as discussed for (A.4), most existing GA-based MBAC algorithms need to know the mean rate and rate variance of the incoming flow so as to apply Gaussian approximation results presented in Section 2.2. By assuming homogeneity of flows, if the number of flows in the system is known as assumed in the next assumption (A.6), the mean rate and rate variance of each flow (as well as the incoming flow) can be estimated based on the measurements of the existing aggregate of flows in system. Admission test can be further conducted using results presented in Section 2.2, since the required mean rate and rate variance for the new aggregate (to be formed if the incoming flow is admitted) can be alculated simply from the measured mean rate and rate variance of the existing aggregate and the estimated mean rate and rate variance of the incoming flow (e.g. [23] [26]).

(A.6) *The number of flows in system and/or other flow state information are known.* Very often, this assumption is used together with assumption (A.5). As discussed for (A.5), only by knowing the number of flows in the system can the admission test be performed by the related MBAC algorithms. This assumption also implies that the system not only knows when a new flow comes but also needs to be informed when an existing flow leaves the system. While the former is usually implied by admission control, the latter needs other supporting mechanisms (e.g. timeout [17]).

(A.7) *When flows are heterogeneous, sufficiently detailed traffic information of an incoming flow is explicitly provided by its descriptor or predicted/estimated based on flow states in system.* In the presence of heterogeneous flows, traffic descriptors containing sufficiently detailed traffic information of incoming flows are usually explictly assumed as in (A.4). Or, to avoid this requirement, some approaches can be used to predict or estimate the traffic information or resource requirement of an incoming flow. One approache uses an empirical distribution of traffic information or resource requirement for a typical flow to make admission decision [27]. Because of this, (A.6) or the number of flows in system is needed [27].

(A.8) *The same traffic descriptor is used at all nodes.* To date, all MBAC algorithms that rely on traffic descriptors use the same descriptor to examine the admissibility of an incoming flow at various nodes along its path.

(A.9) *Each flow in the aggregate experiences the same service guarantees as the aggregate.* The main objective of MBAC algorithms is to guarantee the QoS requirement, not only for the incoming flow examined for admission, but also for all existing flows. A common, yet gross, assumption as given in (A.3) is that all flows require the same service guarantees. In addition, it is implicitly yet commonly assumed that each flow in the aggregate experiences the same service guarantees as the aggregate.

(A.10) *Besides rate guarantee, either statistical delay guarantee or statistical loss guarantee is provided.* Most existing MBAC algorithms are designed to provide either statistical loss guarantee or statistical delay guarantee in addition to rate guarantee. Few consider both of them.

In the following, we briefly discusse the rationables of assumptions (A.1)-(A.10).

Because of the random nature of network flows, (A.1) generally is a reasonable assumption. Particularly, at the first node where some flows are aggregated, the independence of these flows at this node holds very well. However, in the downstream nodes, if some of these flows are present concurrently, the independence assumption among these flows needs further examination. In core networks, since at each node a large number of flows may join and leave, the independence assumption may hence still be a reasonable assumption for obtaining tractable network performance.

Assumption (A.2) can be easily satisfied under IntServ, where RSVP is used as the signaling protocol to communicate both traffic specification and service specification of an incoming flow to all routers along its path. For DiffServ, a separate control plane or bandwidth brokers may be used to for this purpose. In a recently proposed Flow-Aware

Networking architecture, the signaling protocol is not needed, since each node uses a pre-defined descriptor for each flow [17].

Assumptions (A.3) and (A.9) are implicit yet fundamental assumptions for most (if not all) existing algorithms. For scalability considerations, they have made MBAC algorithms simpler. However, while (A.3) is reasonable particularly for flows within the same traffic class, (A.9) is clearly questionable and needs further examination.

Although assumptions (A.4) and (A.7) make MBAC much easier, they may in general not be practical in real networks. Because unless an incoming flow has been studied for its whole duration, its complete traffic information cannot be provided a priori. Particularly, if the flow is real-time, these two assumptions can hardly be met. Recently, less stringent assumptions have been considered. For this, peak rate and token-bucket are two commonly used descriptors. These two descriptors are generally reasonable, since in both IntServ and DiffServ, they are usually explictly given or enforced at the ingress edge. Flow-Aware networks go one-step further: these two assumptions are not needed since a pre-defined descriptor is used for MBAC.

Assumption (A.5) has serious limitation for use in real networks, since naturally flows are heterogeneous in such networks. While assumption (A.6) has been the basis for many algorithms, the knowledge of flow number and states in the system is not available in DiffServ networks which distinguish traffic classes rather than individual flows. In addition to well recognized scalability concerns, associated computational and storage complexities for extracting or estimating such knowledge make (A.6) infeasible in core-stateless networks.

For assumption (A.8), a traffic descriptor of a given flow at the ingress of the network is no longer valid characterization of the same at internal network nodes. While the impact of this misrepresentation may be insignificant at nodes switching large volumes of traffic, this may not be the case at (or close to) the ingress of the network. Fortunately, most MBAC algorithms rely on active on-line measurements (rather than traffic descriptors) to characterize the aggregate traffic which helps limit the impact of such misrepresentation. Clearly, if available, accurate characterization of an individual flow as it traverses network nodes may help to enhance the efficiency and robustness of MBAC algorithms.

While assumption (A.9) is critical to most (if not all) existing MBAC algorithms, it is worth highlighting that the statistical service guarantee to the aggregate is not the same as to an individual flow in the aggregate. However, due to feasibility consideration, (A.9) is reasonable and perhaps necessary (although surprisingly no validation has been provided in the literature.) The final point (A.10) states a fact, which is most existing MBAC algorithms only deal with either statistical delay guarantee or statistical loss guarantee but not both. With rapidly increasing multimedia applications that are both delay-sensitive and loss-sensitive, there is a need for MBAC algorithms designed to provide both statistical delay and statistical loss guarantees.

6 MBAC: Remaining Challenges

An essential requirement of MBAC algorithms is to provide service guarantees to all active flows (existing and to be admitted). Upon the arrival of a new request, not only that request, but all existing flows must be examined for QoS violations. In the presence of heterogeneous flows, this requirement poses difficult challenges and many of them remain open which are briefly presented below. First, based on the revisit of current assumptions for MBAC, we believe that more work is needed to enhance the feasibility and robustness of MBAC algorithms, despite the comment (C.1) in [1][2]. For this, the following are some general challenges that remain:

(R.1) *Relax some of the underlying and restrictive assumptions.*

(R.2 *Reduce dependence on the requirement for the declaration and estimation of traffic parameters, and/or limit the impact of uncertainties associated with these parameters.*

(R.3) *The failure of existing MBAC algorithms to meet their QoS targets is a serious problem that needs to be addressed.*

(R.4) *With QoS targets met, improve network utilization.*

Second, all discussed algorithms adopt a common approach to perform MBAC, which is to treat the aggregate of existing flows as one flow. This is based on a widely adopted assumption that all flows in the same aggregate share the same QoS requirements and an individual flow in the aggregate experiences the same statistical delay and/or loss guarantees as the aggregate (as stated in (A.3) and (A.9)). While this treatment helps achieve scalability in MBAC, it is difficult to provide or maintain service guarantees to each individual flow without over-provisioning of network resources. More adequate treatment of this problem would require maintaining information on the flow number and states (as in IntServ networks), which is known to raise serious scalability concerns. In consequence, the trade off between the salability of MBAC and the granularity of service guarantees poses the following challenge:

(R.5)] *Provide statistical service guarantees to individual flows with a scalable MBAC.*

Third, most MBAC results relate to nodal statistical service guarantees. These results can be easily extended to the end-to-end (e2e) case by assuming independence between service guarantees in different nodes. However, how well this extension works remain un-investigated. Particularly, in some access networks where the admissible number of flows is not large and in cases where most part of an aggregate shares the same path (as in MPLS or VPN networks), the independence assumption may not be valid. This results in the following challenge:

(R.6) *Derive e2e statistical service guarantees from nodal statistical service guarantees.*

Fourth, related to extending nodal statistical service guarantees to the end-to-end case, a commonly used approach is to use the same traffic descriptor of an incoming flow for MBAC at every node along its path. While this clearly reduces the complexity of admission control, it is well understood that due to multiplexing, the initial traffic descriptor at ingress edge is no longer accurate for describing the traffic of this flow in the network core. Hence, we have:

(R.7) *Investigate the validity or impact of using the same traffic descriptor of an incoming flow for admission control at every node along its path.*

The next challenge relates to more accurate analysis of statistical service guarantees with MBAC. While each proposed MBAC is based on its own analysis of statistical service guarantees, (C.2) has shown that none of such MBAC algorithms can meet their QoS targets. One reason for this is that the analyses used by these MBAC algorithms are built upon various approximations, for example as given in Sec. 2.2. It would hence be expected, as another challenge, that better approximations or analyses will allow an MBAC algorithm to have improved performance in meeting its QoS targets

(R.8) *Provide more and/or better methods for analyzing statistical service guarantees from measurements.*

While not explictly stated, up to this point, we have assumed that: *(A.11) flows after admitted are aggregated in the FIFO manner and (A.12) tail-drop is used for buffer management.* By relaxing these assumptions, the following challenge arises:

(R.9) *Investigate statistical service guarantees under non-FIFO and/or non-tail-drop for MBAC.*

The final two remaining challenges discussed in this paper relate to the two estimators shown in Fig. 1. Particularly, the investigation in [1] has shown that different traffic estimators can affect the performance of an MBAC algorithm. Although such effect does not show much impact on its general loss-load curve, it is responsible for the (slight) variations in the loss-load performance. For resource estimation, we notice that very few existing MBAC algorithms have addressed it and more effort is needed.

(R.10) *Estimate and predict traffic information for MBAC based on accurate measurements.*

(R.11) *Estimate and predict available network resources for MBAC in more realistic network environments including wireless and/or mobile networks.*

7 Concluding Remarks

We have revisited several issues pertaining to MBAC for providing statistical service guarantees to real-time multimedia applications. In particular, we reviewed the Poisson Limit Theorem and Central Limit Theorem and their results on which many existing MBAC algorithms are based. Also, we reviewed several existing MBAC algorithms and briefly summarized two critical comments on MBAC research. In addition, we investigated why different MBAC algorithms have the same loss-load performance. We found that rather than a surprise, this same loss-load curve should be expected for different MBAC algorithms. Based on this finding, we hence argued that it should be other issues than the loss-load curve that should be taken into account in designing an MBAC algorithm. Based on this argument, we further revisited the assumptions made by existing MBAC algorithms and examined the rationales of these assumptions when they are used in different network scenarios. Finally, we discussed some remaining challenges for MBAC research. Of these challenges, one is on designing MBAC algorithms that can be used in (core-)stateless networks where per-flow states are not available. Another is on designing MBAC algorithms that can provide per-flow statistical service guarantees. This challenge arises from the fact that most (if not all) existing MBAC algorithms do not provide such guarantees although they are aimed for this. Our future work will manage to address these two challenges.

References

[1] L. Breslau, S. Jamin, and S. Shenker. Comments on the performance of measurement-based admission control algorithms. In *IEEE INFOCOM*, 2000.

[2] L. Breslau, S. Jamin, and S. Shenker. Measurement-based admission control: What is the research agenda? In *International Workshop on Quality of Service (IWQoS)*, 1999.

[3] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. Internet traffic tends toward Poisson and independent as the load increases. In D. D. Denison, M. H. Hansen, C. C. Holmes, B. Mallick, and B. Yu, editors, *Nonlinear Estimation and Classification*. Springer, 2003.

[4] J. Cao and K. Ramanan. A Poisson limit for buffer overflow probabilities. *IEEE INFOCOM'02*.

[5] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido. A nonstationary Poisson view of Internet traffic. In *IEEE INFOCOM*, 2004.

[6] R. B. Cooper. *Introduction to queueing theory*. North Holland, 1981.

[7] S. S. Lavenberg, editor. *Computer Performance Modeling Handbook*. Academic Press, 1983.

[8] J. Choe and N. B. Shroff. A central-limit-theorem-based approach for analyzing queue behavior in high-speed networks. *IEEE/ACM Trans. Networking*, 6(5):659–671, Oct. 1998.

[9] H. S. Kim and N. B. Shroff. Loss probability calculations and asymptotic analysis for finite buffer multiplexers. *IEEE/ACM ToN*, 9(6):755–768, Dec. 2001.

[10] S. Jamin, P. B. Danzig, S. J. Shenker, and L. Zhang. A measurement-based admission control algorithm for integrated services packet networks. *IEEE/ACM ToN*, 5(1):56–70, Feb. 1997.

[11] F. P. Kelly, P. B. Key, and S. Zachary. Distributed admission control. *IEEE JSAC*, Dec. 2000.

[12] B.-K. Choi and R. Bettati. Endpoint admission control: Network-based approach. In *IEEE International Conference on Distributed Computing Systems*, 2001.

[13] C. Cetinkay, V. Kanodia, and E. W. Knightly. Scalable services via egress admission control. *IEEE Trans. Multimedia*, 3(1):69–81, Mar. 2001.

[14] S. Blake and et al. An architecture for Differentiated Services. *IETF RFC 2475*, Dec. 1998.

[15] L. Breslau, E. W. Knightly, S. Shenker, I. Stoica, and H. Zhang. Endpoint admission control: Architectural issues and performance. In *ACM SIGCOMM*, 2000.

[16] Z. Dziong, M. Juda, and L. G. Mason. A framework for bandwidth management in ATM networks - aggregate equivalent bandwidth estimation approach. *IEEE/ACM ToN*, 5:134–147, 1997.

[17] J. W. Roberts. Internet traffic, QoS and pricing. *Proceedings of the IEEE*, to appear.

[18] R. Guerin, H. Ahmadi, and M. Naghshineh. Equivalent capacity and its application to bandwidth allocation in high-speed networks. *IEEE JSAC*, 9(7):968–981, Sept. 1991.

[19] S. Floyd. Comments on measurement-based admission control for controlled-load services. Technical report, Lawrence Berkeley National Laboratory, 1996. Technical Report.

[20] 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 *IEEE ATM'97*.

[21] R. J. Gibbens and F. P. Kelly. Measurement-based connection admission control. In *Proc. the 15th International Teletraffic Congress (ITC)*, 1997.

[22] D. Y. Eun and N. B. Shroff. A measurement-analytic approach for QoS estimation in a network based on the dominant time scale. *IEEE/ACM ToN*, 11(2):222–235, Apr. 2003.

[23] M. Grossglauser and D. N.C. Tse. A framework for robust measurement-based admission control. *IEEE/ACM ToN*, 7(3):293–309, June 1999.

[24] E. W. Knightly. Second moment resource allocation in multi-service networks. In *ACM Sigmetrics*, 1997.

[25] J. Qiu and E. W. Knightly. Measurement-based admission control with aggregate traffic envelopes. *IEEE/ACM ToN*, 9(2):199–210, April 2001.

[26] M. Grossglauser and D. N.C. Tse. A time-scale decomposition approach to measurement-based admission control. *IEEE/ACM ToN*, 11(4):550–563, Aug. 2003.

[27] D. Tse and M. Grossglauser. Measurement-based call admission control: Analysis and simulation. In *IEEE INFOCOM*, 1997.