A Framework for Robust Measurement-Based Admission Control

Matthias Grossglauser, *Member, IEEE*, David N. C. Tse, *Member, IEEE*

Abstract—**Measurement-based Admission Control (MBAC) is an attractive mechanism to concurrently offer Quality of Service (QoS) to users, without requiring a-priori traffic specification and on-line policing. However, several aspects of such a system need to be clearly understood in order to devise robust MBAC schemes, i.e., schemes that can match a given QoS target despite the inherent measurement uncertainty, and without the tuning of external system parameters.**

We study the impact of measurement uncertainty, of flow arrival and departure dynamics, and of estimation memory on the performance of a generic MBAC system in a common analytical framework. We show that a *certainty equivalence* assumption, i.e., assuming that the measured parameters are the real ones, can grossly compromise the target performance of the system. We quantify the improvement in performance as a function of the length of the estimation window and an adjustment of the target QoS. We demonstrate the existence of a *critical time-scale* over which the impact of admissin decisions persists. Our results yield new insights into the performance of MBAC schemes, and represent quantitative and qualitative guidelines for the design of robust schemes.

I. Introduction

The traditional approach to admission control requires an *a priori* traffic descriptor in terms of the parameters of a deterministic or stochastic model. However, it is generally hard or even impossible for the user or the application to come up with a tight traffic descriptor *before* establishing a flow. *Measurement-based admission control* (MBAC) avoids this problem by shifting the task of traffic characterization from the user to the network, so that admission decisions are based on traffic measurements instead of an explicit specification (cf. Fig. 1). This approach has several important advantages. First, the user-specified traffic descriptor can be trivially simple (e.g. peak rate). Second, an overly conservative specification does not result in an overallocation of resources for the entire duration of the session. Third, when traffic from different flows are multiplexed, the QoS experienced depends often on their *aggregate* behavior, the statistics of which are easier to estimate than those of an individual flow. This is a consequence of the law of the large numbers. It is thus easier to predict aggregate behavior rather than the behavior of an individual flow.

Relying on measured quantities for admission control raises a number of issues that have to be understood in order to develop robust schemes.

- **Estimation error**. There is the possibility of making errors associated with any estimation procedure. In the context of MBAC, the estimation errors can translate into erroneous flow admission decisions. The effect of these decision errors has to be carefully studied, because they add another level of uncertainty to the system, the first level being the stochastic nature of the traffic itself. Assuming *certainty equivalence* up-front, i.e. assuming that the estimated parameters are the real parameters, is dangerous, as we simply ignore its impact on the quality of service.
- **Dynamics and separation of time-scale**. A MBAC is a dynamical system, with flow arrivals and departures, and parameter estimates that vary with time. Since the estimation process measures the in-flow burst statistics, while the admission decisions are made for each arriving flow, MBAC inherently links the flow and burst time-scale dynamics. Thus, the question of impact of flow arrivals and departures on QoS arises. Intuitively, each flow arrival carries the potential of making a wrong decision. We therefore expect a high flow arrival rate to have a negative effect on performance. On the other hand, the impact of a wrong flow admission decision on performance also depends on how long it takes until this error can be corrected - that is, on flow departure dynamics.
- **Memory**. The quality of the estimators can be improved by using more past information about the flows present in the system. However, memory in the estimation process adds another component to the dynamics of a MBAC. For example, it introduces more correlation between successive flow admission decisions. Moreover, using too large a memory window will reduce the adaptability of MBAC to non-stationarities in the statistics. A key issue is therefore to determine an appropriate memory window size to use. For this, a clear understanding of the impact of memory on both estimation errors and flow dynamics is necessary.

Matthias Grossglauser is with the Network Mathematics Research Department, AT&T Labs - Research, Florham Park NJ 07932, USA (e-mail: mgross@research.att.com).

David N. C. Tse is with the Dept. of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, USA (e-mail: dtse@eecs.berke ley.edu).

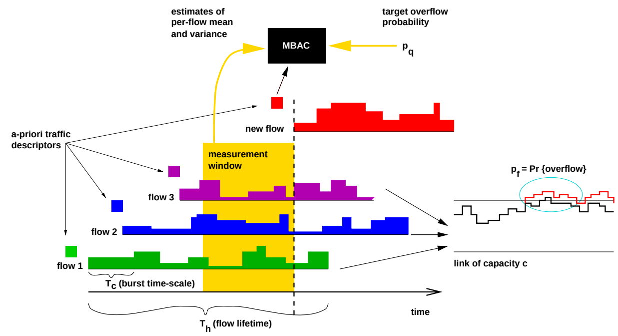

Fig. 1: Traditional admission control makes decisions based on the a-priori traffic descriptors of the existing and the new flow. Measurement-based admission control (MBAC) only uses the new flow's traffic descriptor, but estimates the behavior of the existing flows.

Because of the complex interplay of all these aspects of the MBAC problem, most of the past work has either been analytical but focused on only one of the aspects, or relied primarily on simulations to evaluate MBAC algorithms. In this work, we take a different approach. Using a simple model, we study all of the above issues in a unified analytical framework. The goal is to shed insight on the design of robust MBAC schemes which can make QoS guarantees in the presence of measurement uncertainty, without requiring the tuning of external system parameters.

Due to the complexity of the problem, approximations are made in the performance analysis of the MBAC schemes. These approximations are justified by limit theorems in the *heavy traffic regime*, where the system size is large and when scaling up the size of the system, we exploit the additional statistical regularity by increasing the system utilization while keeping the QoS constant. This is in contrast to the *large deviations regime*, where the system utilization is asymptotically constant, but where the QoS-requirement is scaled with the system size[1]. The heavy traffic regime allows us to use Gaussian approximations and to compute the quantities of interest in terms of first and second-order statistics of the traffic processes.

The rest of the paper is organized as follows. We analyze the *impulsive load model* in Section II and the *continuous load model* in Section III. In Section IV, we apply the insights gained in the analysis to study the problem of choosing the appropriate memory window size of the estimators. In Section V, we comment on some of the assumptions made in our analysis. In Section VI, we discuss how our results relate to previous work in measurement-based admission control. We conclude the paper in Section VII.

II. Impulsive Load Model

The network resource considered is a bufferless single link with capacity c. Flows arrive over time and, if admitted, stay for a random holding time (cf. Fig. 1). The bandwidth requirements of a flow fluctuate over time while in the system. We assume that the statistics of the bandwidth fluctuations of each flow are identical, stationary and independent of each other, with a mean bandwidth requirement of μ and variance σ^2. An important system parameter is the normalized capacity $n := \frac{c}{\mu}$, which measures the system size in terms of the mean bandwidth of the flows. Resource overload occurs when the instantaneous aggregate bandwidth demand exceeds the link capacity, and the quality of service is measured by the steady-state overflow probability p_f. See Fig. 1.

To study the various issues outlined in the introduction, we will first analyze a simpler variation of this model, in which an infinite burst of flows arrives at time 0 and admission control decisions are made then, based on the initial bandwidths of the flows. After time 0, no more flows are accepted, and existing flows stay in the system forever. We call this the *impulsive load* model. This model permits us to study the impact of the measurement errors on the number of admitted flows and on the overflow probability, without the need to worry about flow dynamics. In the next section, we will extend our analysis to the fully dynamical model, where new flows arrive continuously.

The number of admissible flows m^* is the largest integer m such that

$$\Pr\left\{\sum_{i=1}^{m} X_i(t) > c\right\} \le p_q. \qquad (1)$$

where $X_i(t)$ is the bandwidth of the ith flow at time t. (Recall that $c := n\mu$ is the total capacity of the link.) For large system size n, the number of admissible calls will be large, and by the Central Limit Theorem,

$$\frac{1}{\sigma\sqrt{m}}\left[\sum_{i=1}^{m} X_i(t) - m\mu\right] \sim N(0,1)$$

Thus, if the parameters μ and σ^2 are known *a priori*, then the number of flows m^* to accept should satisfy:

$$Q\left[\frac{n\mu - m^*\mu}{\sigma\sqrt{m^*}}\right] = p_q. \qquad (2)$$

where $Q(\cdot)$ is the ccdf of a $N(0,1)$ Gaussian random variable[2]. Because the AC has perfect knowledge of the statistics, the actual steady state overflow probability

$$p_f := \Pr\left\{\sum_{i=1}^{m^*} X_i(t) > c\right\}$$

satisfies the QoS requirement. For reasonably large capacities, it follows from solving (2) that m^* is well approximated by:

$$m^* = n - \frac{\sigma\alpha_q}{\mu}\sqrt{n} + o(\sqrt{n}) \qquad (3)$$

where $\alpha_q := Q^{-1}(p_q)$ and $o(\sqrt{n})$ denotes a term which grows slower than $\sqrt{n}$. Note that n is the number of flows that can be carried on the link if each has constant bandwidth μ. Thus, the term $\frac{\sigma\alpha_q}{\mu}\sqrt{n}$ in the above expression can be interpreted as the safety margin left to cater for the (known) burstiness of the traffic.

Now, consider the situation when a MBAC does not know μ and σ *a priori*, but relies on an estimation of these parameters from the initial bandwidth of the flows and uses the estimates in a *certainty equivalent* fashion. Invoking again the central limit approximation for large systems, the number of flows M_0 the MBAC admits should satisfy:

$$Q\left[\frac{n\mu - M_0\hat{\mu}}{\hat{\sigma}\sqrt{M_0}}\right] = p_q, \qquad (4)$$

where

$$\hat{\mu} = \frac{1}{n}\sum_{i=1}^{n} X_i(0), \qquad \hat{\sigma} = \left[\frac{1}{n-1}\sum_{i=1}^{n}(X_i(0) - \hat{\mu})^2\right]^{\frac{1}{2}} \qquad (5)$$

The criterion (4) is the same as (2), but with the true mean μ and standard deviation σ replaced by the *estimated* mean $\hat{\mu}$ and standard deviation $\hat{\sigma}$ respectively.[3] Note that the number of flows M_0 admitted under the MBAC is now random, depending on the random bandwidths of the flows at time 0. This is a consequence of the fact that the admission control decisions are made based on measurements rather than known parameters. Also, the scheme considered here is an example of a *memoryless* MBAC, since the admission control decisions are made based on the current bandwidths only.

We now want to approximate the average overflow probability

$$p_f := \Pr\left\{\sum_{i=1}^{M_0} X_i(t) > c\right\}$$

in steady state (i.e. for t large) and compare it to the target p_q. To do this, we first find an approximation for the distribution of M_0, the number of flows admitted.

[1] A large deviations analysis of a related measurement-based admission control problem can be found in [21]. In general, one would expect the large deviations approximations to be more accurate if the QoS target is very stringent (say 10^{-6} to 10^{-9}) and the utilization low, and the heavy traffic regime to be reasonable when the QoS target is larger (say 10^{-3} to 10^{-4}) and the utilization high. It is particularly appropriate when the number of multiplexing flows is large.

[2] Note that here, as in the sequel, we are ignoring the fact that m^* is an integer and therefore eqn. (2) cannot be satisfied exactly in general. In the regime of large capacities, however, the approximation is good and the discrepancy can be ignored.

[3] Observe here that the estimation is based on n flows. In a more precise model, the estimation should be based on M_0 flows, the number to be admitted. However, in a large system, M_0 will be close to n and the discrepancy in replacing M_0 by n in the estimators are small.

For large capacities, by the law of large numbers, the estimated mean $\hat{\mu}$ will be close to the true mean μ, and the estimated variance $\hat{\sigma}^2$ will be close to the true variance σ^2. A more precise approximation of the deviation of these estimated quantities from the true values is given by the Central Limit Theorem:

$$\begin{aligned}\hat{\mu} &= \frac{1}{n}\sum_{i=1}^{n} X_i(0) = \mu + \frac{1}{\sqrt{n}}\left\{\frac{1}{\sqrt{n}}\left[\sum_{i=1}^{n} X_i(0) - n\mu\right]\right\} \\ &= \mu + \frac{\sigma Y_0}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right) \qquad (6)\end{aligned}$$

for large n. Here, $Y_0 \sim N(0,1)$, and can be interpreted as the scaled aggregate bandwidth fluctuation at time 0 around the mean. Similarly, the estimated standard deviation can be written as:

$$\hat{\sigma} = \sigma + \frac{Z_0}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right) \qquad (7)$$

where Z_0 is Gaussian. These two approximations imply that the deviation of the estimates from the respective true quantities is of order $1/\sqrt{n}$. Now, as mentioned earlier, if the estimates were *exactly* equal to their true values, then the number of flows admitted M_0 would be precisely m^*. This suggests that we can approximate the distribution of M_0 by a *linearization* of the relationship (4) around a nominal operating point (m^*, μ, σ) (i.e. the operating point under perfect knowledge):

$$\frac{n\mu - (m^* + \Delta_M)(\mu + \frac{\sigma Y_0}{\sqrt{n}})}{(\sigma + \frac{Z_0}{\sqrt{n}})\sqrt{m^* + \Delta_M}} = \alpha_q$$

Expanding the left hand side, using eqn. (2), we get

$$\frac{\Delta_M}{\sqrt{n}} + \frac{\sigma}{\mu} Y_0 = o(1)$$

and hence

$$M_0 = m^* - \frac{\sigma}{\mu} Y_0 \sqrt{n} + o(\sqrt{n}). \qquad (8)$$

Thus, we see that the effect of estimation error is an order $\sqrt{n}$ Gaussian fluctuation around m^*, the number of sources admitted under perfect knowledge (cf. top part of Fig 2). Note also that the randomness in the number of flows admitted is due mainly to the error in estimating the mean (Y_0) rather than the error in estimating the standard deviation (Z_0).

Substituting eqn. (3) into (8), we get M_0 in terms of the system size n:

$$M_0 = n - \frac{\sigma}{\mu}(Y_0 + \alpha_q)\sqrt{n} + o(\sqrt{n}) \qquad (9)$$

A precise statement of the result is as follows. The proof can be found in Appendix B (proof of Theorem III.1).

Proposition II.1: Let $M_0^{(n)}$ be the random number of flows admitted under the MBAC when the capacity is $n\mu$. Then the sequence of random variables $\{\frac{M_0^{(n)} - n}{\sqrt{n}}\}$ converges *in distribution* to the random variable $-\frac{\sigma}{\mu}(Y_0 + \alpha_q)$.

We now proceed with an explicit approximation of the overflow probability. The randomness in the aggregate traffic load at some future time is due both to the randomness in the number of flows admitted as well as the randomness in the bandwidth demands of those flows. This can be approximated with the help of the following lemma, which is an extension of the Central Limit Theorem for a sum of a random number of random variables:

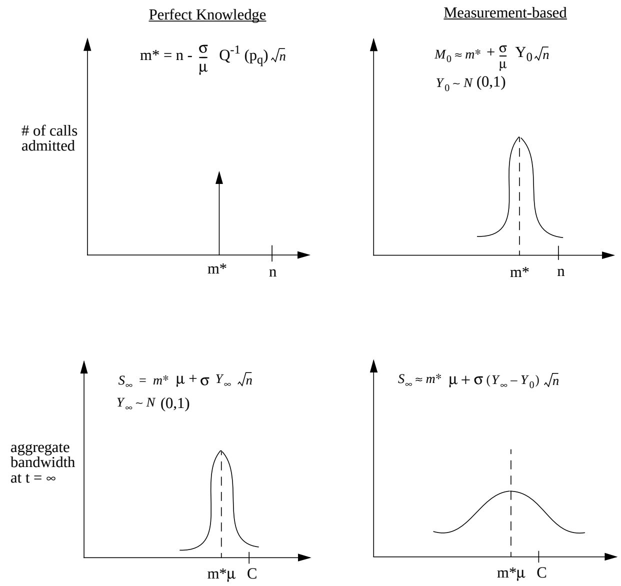

Fig. 2: Uncertainty due to fluctuation in the number of flows (top) and in the aggregate bandwidth (bottom), for an admission controller with perfect knowledge (left) and for an MBAC (right).

Lemma II.2: [3, p. 369, problem 27.14] Let $X_1, X_2, \ldots$ be independent, identically distributed random variables with mean μ and variance σ^2, and for each positive n, let V_n be a random variable assuming positive integers as values; it need not be independent of the X_m's. Let $W_n = \sum_{i=1}^{V_n} X_i$. Suppose as $n \to \infty$, $\frac{V_n}{n}$ converges to 1 almost surely. Then as $n \to \infty$,

$$\frac{W_n - V_n\mu}{\sigma\sqrt{n}}$$

converges in distribution to a $N(0,1)$ random variable.

Applying this lemma, the aggregate load at time t can be approximated by:

$$S_t := \sum_{i=1}^{M_0} X_i(t) = M_0\mu + \sigma Y_t\sqrt{n} + o(\sqrt{n}) \qquad (10)$$

Here $Y_t \sim N(0,1)$ and can be interpreted as an approximation for the scaled aggregate bandwidth fluctuation at time t:

$$\frac{1}{\sigma\sqrt{n}}\left[\sum_{i=1}^{n} X_i(t) - n\mu\right] \qquad (11)$$

Intuitively, eqn. (10) means that the fluctuation of the aggregate load is approximately the linear superposition of two effects: the random number of flows together with the random bandwidth fluctuation around the mean. Substituting eqn. (9) into (10), we get

$$S_t = n\mu + \sigma(Y_t - Y_0 - \alpha_q)\sqrt{n} + o(\sqrt{n})$$

Thus, for large n, the overflow probability at time t is:

$$\Pr\{S_t > n\mu\} \approx \Pr\{Y_t - Y_0 > \alpha_q\}$$

This expression gives us an interpretation of how overflow occurs in a MBAC system: it is a combination of an aggregate bandwidth estimation error at admission time (Y_0) and a fluctuation of the aggregate bandwidth (Y_t) at time t after the flows have been accepted. Contrast this with the case with perfect knowledge, where the overflow probability at time t is simply $\Pr\{Y_t > \alpha_q\}$, due to bandwidth fluctuation at time t.

To get the overflow probability in steady state, we set $t = \infty$, in which case Y_∞ is independent of Y_0. Therefore, the difference $Y_\infty - Y_0$ is a Gaussian random variable with mean 0 and variance $2\sigma^2$. The overflow probability is therefore

$$p_f \approx Q\left(\frac{\alpha_q}{\sqrt{2}}\right). \tag{12}$$

We summarize this result more formally in the following proposition:

Proposition II.3: Suppose the target overflow probability QoS is p_q. Let $p_f^{(n)}$ be the actual average steady state overflow probability using the certainty equivalent MBAC for capacity $n\mu$. Then as the system size grows:

$$\lim_{n\to\infty} p_f^{(n)} = Q\left(\frac{Q^{-1}(p_q)}{\sqrt{2}}\right)$$

Note that for the AC with perfect knowledge, the overflow probability is exactly p_q. This is because the aggregate bandwidth fluctuation stems only from the fluctuation of the individual flow bandwidths (cf. lower left part of Fig. 2). On the other hand, in the measurement-based case, the variance of the aggregate bandwidth is doubled because the number of flows also fluctuates due to measurement error (cf. lower right part of Fig. 2). The $\sqrt{2}$ factor is therefore the effect of measurement error, and has quite a tremendous impact on the overflow probability p_f. For example, if $p_q = 1.0e-5$, then the actual performance in the MBAC system would be $p_f \approx 1.3e-3$, a difference of two orders of magnitude. In other words, if we want to achieve $p_f = p_q$ using a MBAC in this impulsive load model, then we have to adjust the target overflow probability under certainty equivalence.

$$p'_q = Q\left(\sqrt{2}\alpha_q\right) \qquad \text{or} \qquad \alpha'_q := Q^{-1}(p'_q) = \sqrt{2}\alpha_q. \tag{13}$$

Using the approximation $Q(x) \approx \frac{\phi(x)}{x}$ for small $Q(x)$, where ϕ is the pdf of $N(0,1)$, we see that

$$p'_q \approx \frac{\alpha_q}{2\sqrt{\pi}} p_q^2$$

Thus, we see that to achieve a target p_q in this setting, we need to set p'_q roughly to be the square of the target probability. This conservatism leads to a loss in system *utilization* compared to the scheme with perfect knowledge of the statistics. The average utilization (in terms of bandwidth) for the certainty equivalent scheme using parameter p'_q instead of p_q is given by $E(M_0)\mu$, or $c - \sigma\alpha'_q\sqrt{n}$, as implied by eqn. (8). The average utilization for the perfect knowledge scheme, on the other hand, is given by $m^*\mu$, or $c - \sigma\alpha_q\sqrt{n}$, as inferred from (3). Thus, if we pick α'_q to be $\sqrt{2}\alpha_q$, this translates to a loss of utilization of $(\sqrt{2}-1)\sigma\alpha_q\sqrt{n}$.

Proposition II.3 has several surprising aspects. First, it is a *universal* result in the sense that the performance of the certainty equivalent scheme does not depend on the stationary distribution of the flow nor its mean and variance. Second, although the estimators are unbiased, the net impact on the performance of the system is negative. Thus there is an inherent asymmetry between the effects of over-estimation and under-estimation. Third, the impact of the estimation error does not vanish as the system size becomes large, even though the estimates become more and more accurate. Fourth, for a large system, the degradation in performance of the certainty equivalent scheme is due mainly to the estimation error in the *mean* μ of the bandwidth distribution and not to that in the *standard deviation* σ.

To get more insights into the last two phenomena, let us perform the following deterministic sensitivity analysis. Define the following function:

$$p_f(\mu, \sigma, m) := Q\left[\frac{c - m\mu}{\sigma\sqrt{m}}\right]$$

which is the overflow probability when there are m flows in the system each with mean rate μ and variance σ^2. Suppose first that we measure only μ, but that σ is known exactly. The number of flows admitted $m(\widehat{\mu})$ depends on the measured value $\widehat{\mu}$ and is given by the certainty-equivalent admission criterion (compare with (4)):

$$p_f(\widehat{\mu}, \sigma, m(\widehat{\mu})) = p_q. \tag{14}$$

Note that the *actual* overflow probability p_f for a given $m(\widehat{\mu})$ is $p_f(\mu, \sigma, m(\widehat{\mu}))$. The *sensitivity* of the overflow probability with respect to the measured $\widehat{\mu}$ is the deviation of p_f from its target value p_q if $\widehat{\mu}$ deviates slightly from its target value μ. For small deviations, we can simply use the derivative of p_f with respect to $\widehat{\mu}$.

$$s_\mu := \frac{\partial}{\partial\widehat{\mu}} p_f(\mu, \sigma, m(\widehat{\mu}))\bigg|_{\widehat{\mu}=\mu}.$$

Using (14), this derivative can be computed as:

$$s_\mu = -\frac{\phi(\alpha_q)\mu}{\sigma}\sqrt{m^*}.$$

Similarly, the sensitivity with respect to measured $\widehat{\sigma}$, assuming μ known, is given by:

$$s_\sigma = -\frac{\alpha_q\phi(\alpha_q)}{\sigma}$$

Now observe that the sensitivity of the system performance on the knowledge of the standard deviation, s_σ, does not depend on the system size. Therefore, increasing the system size - and therefore improving the quality of the estimator $\widehat{\sigma}$ - results in a *diminishing* net impact on the overflow probability. On the other hand, the sensitivity s_μ *increases* with the system size, approximately as $\sqrt{n}$, while the variance of the estimator $\widehat{\mu}$ decreases approximately as $1/\sqrt{n}$. This suggests that the net impact of the uncertainty in the mean bandwidth estimate does not diminish as the system size grows, and also explains why the deviation from p_f from the target overflow probability p_q is asymptotically independent of n: both effects cancel out. The increased sensitivity to the mean estimate arises because when there are more flows in the system, and therefore more statistical regularity in the aggregate bandwidth, the system is driven closer to full utilization, which makes it more susceptible to admission mistakes.

III. The Continuous Load Model

In the impulsive load model, arrivals only occur at time 0 and admitted flows do not depart from the system. We shall now consider a dynamical model, where flows arrive *continuously* over time and stay for an exponentially distributed holding time with mean T_h. We assume a worst-case scenario, where the effective arrival rate is infinite, i.e. there are always flows waiting to be admitted into the network. Thus, admission control decisions are made continuously at all times. Clearly, the performance of any admission control algorithm under finite arrival rate will be no worse than its performance in this model. Another advantage of this model is that we need not worry about the specific flow arrival process which may be hard to model in practice. Furthermore, we let $\rho(t)$ denote a flow's auto-correlation function, where

$$\rho(t) := \frac{E[(X_i(0)-\mu)(X_i(t)-\mu)]}{\sigma^2}.$$

A. Memoryless MBAC

We first look at a memoryless scheme that only bases admission decisions on estimates of the mean and variance based on the current bandwidths of the flows. Assume that the system starts at time 0. Our goal is to find the overflow probability at an arbitrary time t, particularly at $t = \infty$ which yields the steady-state overflow probability. We do this by first analyzing the dynamics of the number of flows in the system.

At any time t, the MBAC estimates the admissible number of flows M_t. As in (4), M_t is given by:

$$Q\left[\frac{n\mu - M_t\hat{\mu}(t)}{\hat{\sigma}(t)\sqrt{M_t}}\right] = p_q, \tag{15}$$

where

$$\widehat{\mu}(t) = \frac{1}{n}\sum_{i=1}^{n} X_i(t), \qquad \widehat{\sigma}(t) = \left[\frac{1}{n-1}\sum_{i=1}^{n}(X_i(t)-\widehat{\mu}(t))^2\right]^{\frac{1}{2}} \tag{16}$$

Observe that M_t is random and depends only on the current bandwidths $X_i(t)$'s of the flows. An approximation is given by eqn. (8):

$$M_s = n - \frac{\sigma}{\mu}(Y_s + \alpha_q)\sqrt{n} + o(\sqrt{n}) \tag{17}$$

where $\{Y_t\}$ is a stationary zero-mean Gaussian process with unit variance and auto-correlation function $\rho(t)$ (that of an individual flow), and can be interpreted as the scaled aggregate bandwidth fluctuation of the flows around the mean.

The *actual* number of flows N_t in the system at time t is no less than M_t since there are always flows waiting to be admitted. On the other hand, N_t can be strictly greater than M_t, as flows that were admitted earlier stay for a certain duration and thus N_t cannot perfectly track the fluctuations of M_t (see Fig. 3). To compute N_t, first observe that if s^* is the last time at or before time t that flows were admitted, then the number of flows in the system at time s^* is precisely the same as number of flows admissible at time s^*, i.e. $N_{s^*} = M_{s^*}$. In between time s^* and time t, no new flows were admitted. Hence, if we let $D[s,t]$ be the number of flows departed in time interval $[s,t]$, then

$$N_t = N_{s^*} - D[s^*,t] = M_{s^*} - D[s^*,t] \tag{18}$$

On the other hand, for *any* $s \le t$,

$$N_t = N_s + A[s,t] - D[s,t] \ge N_s - D[s,t] \ge M_s - D[s,t] \tag{19}$$

where $A[s,t]$ is the number of flows *admitted* during $[s,t]$. Thus we conclude from (18) and (19) that

$$N_t = \sup_{0\le s\le t}\{M_s - D[s,t]\} \tag{20}$$

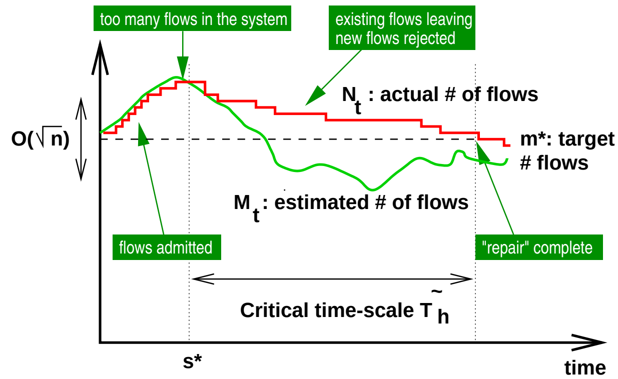

Fig. 3: The relationship between the current estimate of admissible number of flows M_t and the actual number of flows N_t. The time-scale $\widetilde{T_h}$ is the typical time for the system to recover from admission errors.

It is clear from Fig. 3 and eqn. (20) that flow departures have a *repair* effect to past mistakes made by the MBAC. Eqn. (17) tells us that the the fluctuations of the estimated number of admissible flows M_s around the perfect knowledge operating point m^* is of the order of $\sqrt{n}$. Thus, it takes of the order of $\sqrt{n}$ flows to depart to rectify past errors in accepting too many flows. How much time on the average is needed for this to occur? Since the flow departure rate is of the order of n/T_h, this "repair time" is on the order of $\sqrt{n}/(n/T_h) = T_h/\sqrt{n}$. We call this the *critical time-scale* $\widetilde{T_h}$ of the dynamical system: admission errors at time s has little influence on the future much beyond $s + \widetilde{T_h}$, as many flows would have been departed by then to repair the errors. Thus, $\widetilde{T_h}$ is the natural time-scale to analyse the full dynamics of the system. To make such analysis more convienent, let us scale the flow holding time $T_h = \widetilde{T_h}\sqrt{n}$, so that the critical time-scale $\widetilde{T_h}$ can be viewed as fixed as n grows large.[4] Under this scaling, the number of flows departed in time $D[s,t]$ can be approximated as:

$$D[s,t] = \frac{t-s}{\widetilde{T_h}}\sqrt{n} + o(\sqrt{n}). \tag{21}$$

Such linearization is valid for large n such that we are focusing on the time-scale during which only a small fraction (order $1/\sqrt{n}$)of the flows depart from the system.

Combining (17), (20) and (21) yields an approximation for the number of flows N_t in the system at time t. The following limit theorem makes such an approximation precise.

Theorem III.1: Let $\{N_t^{(n)}\}$ be the process describing the evolution of the number of flows in the system. Assume

[4] It should be emphasized that in reality, the system size n and the average flow holding time T_h are independent system parameters. The scaling of T_h as $\sqrt{n}$ is done solely to enable us to study *both* the effect of traffic fluctuation and flow departures in the asymptotic analysis. For, if the holding time T_h were fixed as n grows, the critical time-scale would approach zero, leading to an asymptotic model where any admission errors can be immediately restored by flow departures.

condition B.6 is satisfied. As $n \to \infty$, for each t, $\frac{N_t^{(n)}-n}{\sqrt{n}}$ converges in distribution to

$$\frac{\sigma}{\mu} \sup_{0\le s\le t} \left\{ -Y_s - \frac{\mu(t-s)}{\sigma \widetilde{T_h}} - \alpha_q \right\}$$

where $\{Y_t\}$ is defined as above.

The proof uses the machinery of *weak convergence* and is given in Appendix B. Condition B.6 contains mild technical assumptions on the individual flow processes; these are also stated in the appendix. These assumptions hold for a very broad class of models. For example, they hold if each individual flow is a Markov modulated fluid

Once we obtained an approximation for N_t, we can immediately deduce an approximation for the aggregate load S_t at time t and hence the steady-state overflow probability p_f, using the same argument as for the impulsive load model.

Proposition III.2: Let $S_t^{(n)}$ be the aggregate load at time t and $p_f^{(n)}(t)$ be the overflow probability at time t. As $n \to \infty$, $\frac{S_t^{(n)}-n\mu}{\sigma\sqrt{n}}$ converges in distribution to

$$\sup_{0\le s\le t} \left\{ Y_t - Y_s - \frac{\mu}{\sigma \widetilde{T_h}}(t-s) - \alpha_q \right\}$$

and the overflow probability $p_f^{(n)}(t)$ converges to:

$$\Pr\left\{ \sup_{0\le s\le t} \left\{ Y_t - Y_s - \frac{\mu}{\sigma \widetilde{T_h}}(t-s) \right\} > \alpha_q \right\}.$$

For brevity, we will define the important parameter:

$$\beta := \frac{\mu}{\sigma \widetilde{T_h}}. \tag{22}$$

The steady-state overflow probability can then be approximated by taking $t = \infty$ in Prop. III.2 and using stationarity of $\{Y_t\}$ to get:

$$p_f \approx \Pr\left\{ \sup_{s\le 0} \{Y_0 - Y_s + \beta s\} > \alpha_q \right\} \tag{23}$$

Interestingly, one can interpret the limiting overflow probability at time t as that of the length of a certain *queue* at time t exceeding α_q. The queue is one which has a constant service rate of β, with the amount of work arriving in time interval $[s, t]$ given by $Y_t - Y_s$.

B. Analysis of Overflow Probability

Our next step is to analyze the approximation to the overflow probability given by eqn. (23). Since the process $\{Y_t\}$ is stationary and symmetrically distributed around 0, we can rewrite (23) as

$$p_f \approx \Pr\left\{ \sup_{t\ge 0} \{Y_{-t} - Y_0 - \beta t\} > \alpha_q \right\}.$$

This can be interpreted as the *hitting* probability of a Gaussian process $\{Y_{-\tau} - Y_0\}$ on a moving boundary $y = \beta\tau + \alpha_q$. While there is no known closed-form solution to this problem, an approximation can be obtained by applying results due to Bräker [13], [14] on hitting probabilities of locally stationary Gaussian processes, extending the results by [7] for stationary processes. Define

$$\sigma^2(t) := E[(Y_{-t} - Y_0)^2] = 2[1 - \rho(t)]$$

to be the variance of $Y_{-t} - Y_0$ (recall that Y_t has zero mean and unit variance). Assume the single-sided derivatives of $\rho(t)$ at $t = 0$ exist and are finite, let $v^+(0)$ be the right derivative of the function $\sigma^2(t)$ at $t = 0$. [5] Then an approximation to the hitting probability is given by:

$$\Pr\left\{ \sup_{t\ge 0} \{Y_{-t} - Y_0 - \beta t\} > \alpha_q \right\}$$
$$\approx \frac{1}{2} \int_0^\infty v^+(0) \frac{\alpha_q + \beta t}{\sigma^3(t)} \phi\left(\frac{\alpha_q + \beta t}{\sigma(t)} \right) dt \tag{24}$$

where $\phi(x)$ is the $N(0, 1)$ probability density function. The integrand above can be viewed as an approximation to the first hitting time density at time t; integrating over all t yields the probability that hitting occurs at all. This is an approximation in the sense that as $\alpha_q \to \infty$, the ratio of the left-hand and the right-hand sides approaches 1. Hence this approximation is good when p_q is small.

While this yields an approximation that can be computed numerically for general auto-correlation functions, we would like to get more analytical insights. To that end, consider the specific auto-correlation function:

$$\rho(t) = \exp\left(-\frac{|t|}{T_c} \right). \tag{25}$$

With this choice of the auto-correlation function, $\{Y_t\}$ is the well-known Ornstein-Uhlenbeck process. The parameter T_c governs the exponential drop-off rate of the correlation function, and is a natural *correlation time-scale* for the burst dynamics of the traffic. Substituting this into the approximation (24) and rescaling the time variable, we get:

$$p_f \approx \gamma \int_0^\infty \frac{(\alpha_q + t)}{[2(1 - \exp(-\gamma t))]^{\frac{3}{2}}} \phi\left(\frac{\alpha_q + t}{\sqrt{2(1 - \exp(-\gamma t))}} \right) dt \tag{26}$$

where

$$\gamma := \frac{1}{\beta T_c} = \frac{\widetilde{T_h}}{T_c} \cdot \frac{\sigma}{\mu}.$$

One can think of γ as the separation between the flow and burst scales, although note that $\widetilde{T_h}$ is the scaled holding time. If we make a time-scale separation assumption, i.e. $\gamma \gg 1$, then

$$p_f \approx \gamma \int_0^\infty \frac{(\alpha_q + t)}{2^{\frac{3}{2}}} \phi\left(\frac{\alpha_q + t}{\sqrt{2}} \right) dt = \frac{\gamma}{2\sqrt{\pi}} \exp\left(-\frac{1}{4}\alpha_q^2 \right) \tag{27}$$

Note that the first approximation is via $\exp(-\gamma t) \approx 0$ for $\gamma \gg 1$.

It is interesting to compare this overflow probability for the continuous-load model with the corresponding result for the impulsive load model under long flow durations, given in Proposition II.3. To do this, we first use the approximation $\frac{\phi(x)}{x} \approx Q(x)$ and rewrite (27) in terms of the flow parameters as

$$p_f \approx \frac{\widetilde{T_h}}{2T_c} \cdot \frac{\sigma\alpha_q}{\mu} Q\left(\frac{\alpha_q}{\sqrt{2}} \right) \tag{28}$$

For the impulsive load model, the overflow probability is approximately $Q(\frac{\alpha_q}{\sqrt{2}})$. Eqn. (28) tells us that in the

[5] i.e. $v^+(0) := \lim_{t\to 0+} \frac{\sigma^2(t) - \sigma^2(0)}{t}$.

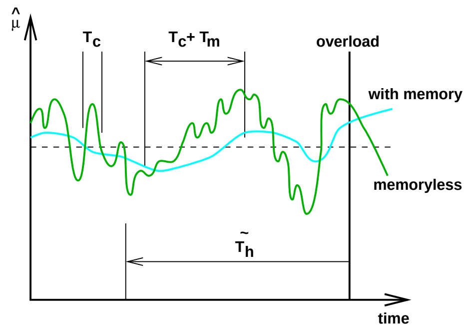

Fig. 4: For the memoryless estimator, the overflow probability depends on the ratio of the correlation time-scale T_c and of the critical time-scale $\widetilde{T}_h$. An estimation memory window of length T_m reduces the variance of the bandwidth estimator, and also smoothes its fluctuation to a time-scale of roughly $T_c + T_m$.

regime of separation of time-scales, the corresponding overflow probability can be much worse in the continuous-load model. This is because while in the impulsive load model estimation errors can occur only at a *single* point in time (time 0), whereas in the continuous-load model, estimation errors can occur up to roughly $\widetilde{T}_h$ before time t to have a significant impact on the number of flows at time t. The shorter the traffic correlation time-scale T_c, the faster the memoryless mean bandwidth estimates fluctuates, and the larger the probability of having an under-estimation at some time in the interval. Hence, the overflow probability in the continuous-load model increases with the separation of time-scale $\frac{\widetilde{T}_h}{T_c}$. For example, note the multiple peaks (underestimations of μ) within the interval of length $\widetilde{T}_h$ in Fig. 4: each of these peaks could potentially cause overload within the critical time-scale $\widetilde{T}_h$. The lesson is that it's not only important to consider the estimation error at a single time-instant, but also the chance of making error any time in the interval defined by the effective flow holding time-scale $\widetilde{T}_h$. Note also that since $\widetilde{T}_h$ decreases as $\frac{T_h}{\sqrt{n}}$, where T_h is the actual mean holding time, the overflow probability decreases roughly as $\frac{1}{\sqrt{n}}$.

We can also write the above approximation as (using again $\frac{\phi(x)}{x} \approx Q(x)$),

$$p_f \approx \frac{\widetilde{T}_h}{\sqrt{2}T_c} \frac{\sigma}{\sqrt{2\pi}\mu} \left(\sqrt{2\pi}\alpha_q p_q\right)^{\frac{1}{2}} \tag{29}$$

C. MBAC with Estimation Memory

We see that the memoryless scheme suffers from two problems . First, the estimation error at a specific admission time instant is large, and in fact has impact which is of the same order of magnitude as that due to the statistical fluctuations of the bandwidths when the correct number of flows are admitted. Second, the correlation time-scale of the estimation errors is the same as that of the traffic itself; thus, in the regime when the flow holding time is much larger than the traffic correlation time-scale ($\widetilde{T}_h \gg T_c$), the probability of having a large under-estimation of mean bandwidth at *some time* during the time-scale $\widetilde{T}_h$ is high. A strategy which, as we will see, counters both these difficulties is to use more memory in the mean and variance estimators.

To be more concrete, let us consider using the first-order auto-regressive filter with impulse response:

$$h(t) := \frac{1}{T_m} \exp\left(-\frac{t}{T_m}\right) u(t)$$

to estimate both the mean and the variances. (Here, $u(t)$ is the unit step function.) Thus, in place of the memoryless estimators in eqn. (16), the MBAC would use:

$$\begin{aligned} \widehat{\mu}_m(t) &= \int_0^\infty \left[\frac{1}{n}\sum_{i=1}^n X_i(t-\tau)\right] h(\tau)d\tau \\ \widehat{\sigma}_m^2(t) &= \int_0^\infty \left[\frac{1}{n-1}\sum_{i=1}^n (X_i(t-\tau) - \widehat{\mu}_m(t))^2\right] h(\tau)d\tau \end{aligned}$$

Note that the estimates are obtained by an exponential weighting of the past bandwidths of the flows. The parameter T_m governs how the past bandwidths are weighted; it can be thought of as a measure of the *estimation window length*. The relationship between $\widehat{\mu}_m(t)$ and the memoryless estimator $\widehat{\mu}(t)$ is simply $\widehat{\mu}_m = \widehat{\mu} * h$, where $*$ is the convolution operation.

Corresponding to Theorem III.1 and Prop. III.2 in the memoryless case, we can show:

Theorem III.3: For the system of size n, let $\{N_t^{(n)}\}$ be the process describing the evolution of the number of flows in the system. Assume condition B.6 is satisfied. If we scale the flow holding time as $T_h^{(n)} = \widetilde{T}_h\sqrt{n}$, where $\widetilde{T}_h$ is a fixed constant, then as $n \to \infty$, for each t, $\frac{N_t^{(n)}-n}{\sqrt{n}}$ converges in distribution to

$$\frac{\sigma}{\mu} \sup_{0\le s\le t} \left\{-Z_s - \frac{\mu(t-s)}{\sigma\widetilde{T}_h} - \alpha_q\right\} \tag{30}$$

where $Z_t = (h * Y)_t$ and $\{Y_t\}$ is a zero-mean, unit-variance stationary Gaussian process with autocorrelation function identical to that of an individual flow. The overflow probability $p_f^{(n)}(t)$ at time t converges to:

$$\Pr\left\{\sup_{0\le s\le t} \left\{Y_t - Z_s - \frac{\mu}{\sigma\widetilde{T}_h}(t-s)\right\} > \alpha_q\right\}.$$

One can interpret Z_t as the error in the *filtered* estimate of the mean bandwidth of a flow at time t. The steady-state overflow probability under the MBAC with memory can therefore be approximated by:

$$p_f \approx \Pr\left\{\sup_{t\ge 0}(Z_{-t} - Y_0 - \beta t) > \alpha_q\right\}$$

This is again a hitting probability of a Gaussian process ($\{Z_{-t} - Y_0\}$) on a moving boundary, and an approximation of such a probability is given by [13], [14]:

$$p_f \approx \frac{\gamma T_c}{T_c + T_m} \int_0^\infty \frac{(\alpha_q + t)}{\left[\sigma_m(\frac{t}{\beta})\right]^3} \phi\left(\frac{\alpha_q + t}{\sigma_m(\frac{t}{\beta})}\right) dt + Q\left(\alpha_q\sqrt{1 + \frac{T_c}{T_m}}\right) \tag{31}$$

where

$$\sigma_m^2\left(\frac{t}{\beta}\right) := E[(Z_{-\frac{t}{\beta}} - Y_0)^2] = \frac{2T_c + T_m}{T_c + T_m} - \frac{2T_c}{T_c + T_m}\exp(-\gamma t)$$

Now, under separation of time-scales, $\gamma \gg 1$, we have the approximation that

$$\sigma_m^2\left(\frac{t}{\beta}\right) \approx \frac{2T_c+T_m}{T_c+T_m}$$

in which case the above integral can be explicitly computed as:

$$p_f \approx \frac{\gamma T_c}{\sqrt{(T_c+T_m)(2T_c+T_m)}} \cdot \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{T_c+T_m}{2(2T_c+T_m)}\alpha_q^2\right) + Q\left(\alpha_q\sqrt{1+\frac{T_c}{T_m}}\right) \quad (32)$$

To compare this result to the memoryless case, let us first use the approximation $Q(x) \approx \frac{\phi(x)}{x}$ to rewrite (32) in terms of p_q and also the flow parameters:

$$p_f \approx \frac{\widetilde{T_h}}{\sqrt{(T_c+T_m)(2T_c+T_m)}} \cdot \frac{\sigma}{\sqrt{2\pi}\mu}\left(\sqrt{2\pi}\alpha_q p_q\right)^{\frac{T_c+T_m}{2T_c+T_m}} + Q\left(\alpha_q\sqrt{1+\frac{T_c}{T_m}}\right) \quad (33)$$

Comparing eqn. (32) to eqn. (27), we can see explicitly the effect of memory. Let us look at the first term in (32), which corresponds to (29). The exponent is $\frac{T_c+T_m}{(2T_c+T_m)}$ which is $\frac{1}{2}$ when there is no memory (as we had in the memoryless scheme), monotonically increases with T_m, and reaching a value of 1 for infinite memory. This effect can be explained by the fact that the variance of the mean bandwidth estimate, $E[Z_t^2]$, is $\frac{T_c}{T_c+T_m}$ and decreases monotonically to zero with more memory. Thus the inaccuracy in the estimates and hence the inaccuracy in the number of flows accepted decreases (cf. Fig. 4). Furthermore, increasing the amount of memory has an additional effect of *smoothing* the mean bandwidth estimates; thus, not only are the *individual* bandwidth estimates more accurate, they also fluctuate less so that the probability of having an under-estimation at *some time* over an interval of length $\widetilde{T_h}$ is reduced. This is reflected in the smaller pre-factor $\frac{\widetilde{T_h}}{\sqrt{(T_c+T_m)(2T_c+T_m)}}$ in the first term of (33) replacing the factor $\frac{\widetilde{T_h}}{\sqrt{2}T_c}$ in the memoryless case. This can be interpreted as increasing the correlation time-scale by T_m, the estimation window length.

In the limit for large T_m, we always have exactly the right number of flows in the system and the overflow occurs due only to the fluctuation of bandwidth requirements of flows in the system, and not to the fluctuation of the number of flows in the system. This is now given by the second term in (33).

IV. Robust Measurement-Based Admission Control

In this section, we discuss how our results can be used to make MBAC robust. We use simulations with synthetic and actual traffic sources to verify these insights. The details of the simulation setup are described in Appendix A

A. Robust MBAC with Known T_c

In this section, we assume that the correlation time-scale T_c is known to the MBAC. Our goal is to verify the validity of the formulas presented in the previous section. Formula (33) can be used to choose the memory size and to adjust the target overflow probability p_q' in the MBAC such that the overflow probability meets the QoS requirement, i,e, choose T_m and p_q' such that $p_f(T_m, T_c, p_q') = p_q$. The shorter T_m, the more conservative the choice of p_q' has to be, resulting in a loss of bandwidth. This loss of utilization can be quantified. The average utilization (in terms of bandwidth) of the system is given by $\mu E[N_t]$, where N_t is the (stationary) number of flows in the system at time t. Eqn. (30) allows us to approximate this when p_q' is used as the target overflow probability:

$$\mu E[N_t] \approx n\mu + \sigma\sqrt{n} E\left[\sup_{s\le t}\left\{-Z_s - \frac{\mu(t-s)}{\sigma\widetilde{T_h}}\right\}\right] - \sigma Q^{-1}(p_q')\sqrt{n}$$

Since the other terms do not depend on p_q', we see that the difference in utilization in using p_q' and p_q is simply

$$\sigma\sqrt{n}\left[Q^{-1}(p_q') - Q^{-1}(p_q)\right] \quad (34)$$

This allows us to quantify the impact on the utilization on using a more conservative target overflow probability.

We now describe the simulations we have performed to verify that our formulas can be used to perform robust measurement-based admission control. We proceed in two steps. First, we compare the overflow probability p_f obtained through simulation to the value predicted by theory. Second, we invert (32) to obtain an adjusted target overflow probability p_q' such that $p_f(T_m, T_c, p_q') = p_q$. We then simulate the system with this adjusted target overflow probability in order to check if the overflow probability p_f really is close to the target overflow probability p_q regardless of the other parameters.

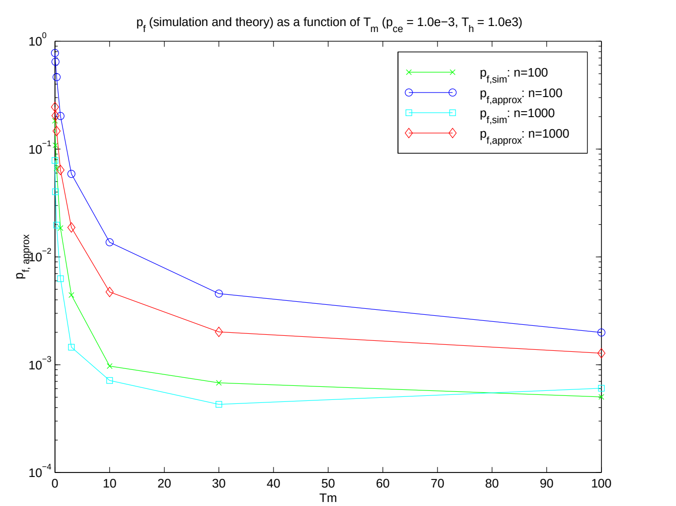

Fig. 5: The overflow probability p_f as predicated by theory (eqn. (32)) and obtained by simulation ($T_h = 1000$, $T_c = 1.0$, $p_q' = 1.0e-3$).

Figure 5 shows the overflow probability p_f as a function of the memory window size T_m. The most striking aspect of this figure is that for small memory window size T_m, the overflow probability p_f can be orders of magnitude larger than the target overflow probability p_q. This confirms that the memory window size T_m is a crucial parameter in achieving a desired QoS target. If it is chosen too small, then the performance of the MBAC can degrade dramatically. We

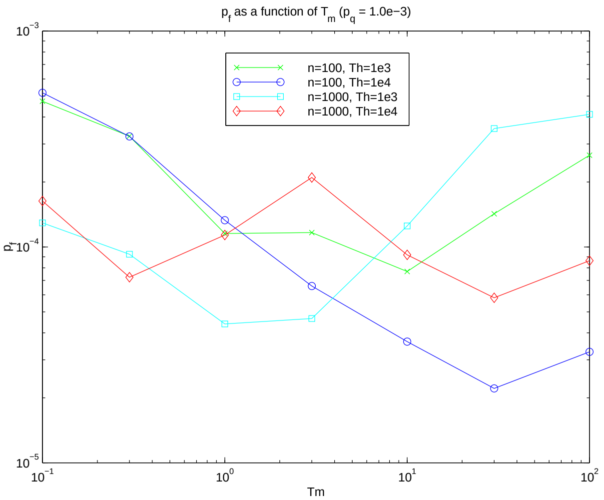

Fig. 6: The simulated overflow probability p_f for the synthetic traffic model using the adjusted target overflow probability p'_q.

observe that the overflow probability predicted by theory is slightly conservative with respect to the simulated value. We attribute this offset to the assumptions in our model, such as ignoring the discreteness of the number of flows. However, the shape of the graphs correspond very well; in particular, the knee, corresponding to the value of T_m beyond which using a longer memory window size has little additional benefit, is well matched. Figure 6 demonstrates that our formulas can be used to perform robust measurement-based admission control. We see that with an adjusted overflow probability target, the actual overflow probability is slightly smaller than p_q over the whole range of parameters (cf. Fig. 6). It is important to note that for small T_m, the adjusted target overflow probability p'_q can be very small ($< 1e-10$) with respect to the target overflow probability p_q of $1e-3$.

B. Robust MBAC with Unknown T_c

So far, we have assumed that the correlation time-scale parameter T_c and the flow holding time T_h are known to the MBAC algorithm so that an adjusted QoS parameter p'_q can be computed. In practice, it is not difficult to obtain a good estimate of the average holding time T_h of flows. On the other hand, the correlation time-scale T_c and more generally the correlation structure of the traffic is hard to estimate reliably. Therefore, we would like to design the MBAC such that its performance is good over a wide range of values for T_c. We claim that this can be accomplished by choosing the memory window length T_m on the order of the critical time-scale $\widetilde{T_h}$. For concreteness, let us pick the window size T_m to be $\widetilde{T_h}$ and examine the performance of the system for a range of T_c.

First, assume T_c is small with respect to $\widetilde{T_h}$. This is the separation of time-scale regime and formula (33) applies and holds for all T_m. Using the fact that $T_m = \widetilde{T_h} \gg T_c$, we get the further approximation:

$$p_f \approx \left(\frac{\sigma \alpha_q}{\mu} + 1\right) p_q \qquad (35)$$

which is of the order of p_q. In this regime, the effect of the estimator memory effectively smoothes the fluctuations of the traffic and obtain a reliable estimate of the mean traffic rate. Although this result is derived using the simple exponential auto-correlation function (25), it can be easily shown that in this regime, the detailed correlation structure is not relevant and a similar approximation holds for other auto-correlation functions. We call this the *masking regime* because the memory window size masks the impact of the parameter T_c on the overflow probability p_f; the fluctuation time-scale of the mean estimator is determined by T_m alone (cf. Fig. 4).

Next, let us consider the other extreme, when T_c is much longer than $\widetilde{T_h}$. In this case, $\gamma = \frac{\widetilde{T_h}}{T_c}\frac{\sigma}{\mu} \ll 1$, and we have the approximation:

$$\sigma_m^2(t) \approx \frac{T_m}{T_c + T_m}.$$

Substituting this into the general formula (31) and evaluating the integral, we get:

$$p_f \approx \frac{1}{\sqrt{2\pi}} \frac{T_c}{\widetilde{T_h}} \frac{\sigma}{\mu} \exp\left[-\left(\frac{T_c}{\widetilde{T_h}}\right)^2 \alpha_q^2\right]$$

which definitely meets the target QoS since $T_c \gg \widetilde{T_h}$ in this regime. In contrast to the masking regime, the time-scale of the estimator fluctuation is dominated by T_c. The memory window is effectively useless in this regime, as it does not reduce estimation errors. However, the fluctuation of the estimators around their mean is at a time-scale longer than the critical time-scale. This is precisely the regime where the repair effect makes overflow unlikely. Therefore, we call this the *repair regime*.

For T_c in between the two extremes, there is no closed-form expression for the overflow probability, and we resort to a numerical integration of the formula (31) to study the performance of the MBAC. This is shown in Fig. 7, where we plot the overflow probability as a function of $T_m/\widetilde{T_h}$ and T_c. We see that while for small $T_m/\widetilde{T_h}$ the performance is not robust, the QoS is satisfied over a wide range of T_c once the memory window size is chosen to be a significant fraction of $\widetilde{T_h}$. This is further corroborated by simulation results shown in Fig. 8.

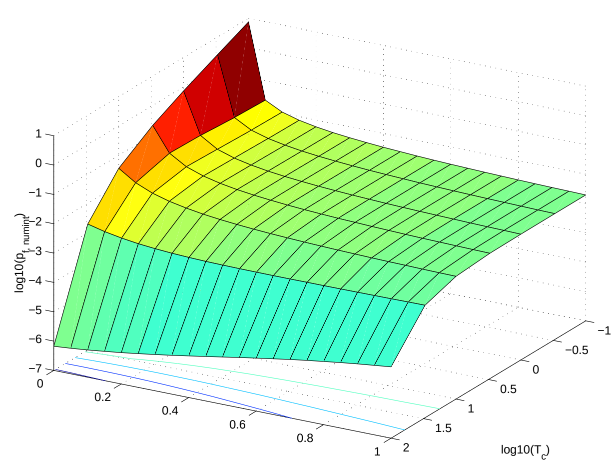

Fig. 7: The overflow probability p_f obtained by numerical integration of (31), as a function of the normalized memory window size $T_m/\widetilde{T_h}$ and of the correlation time-scale T_c.

The above analysis and simulations are based on a traffic model with correlation at a single time-scale. In practice, traffic fluctuations may occur at multiple time-scales. In

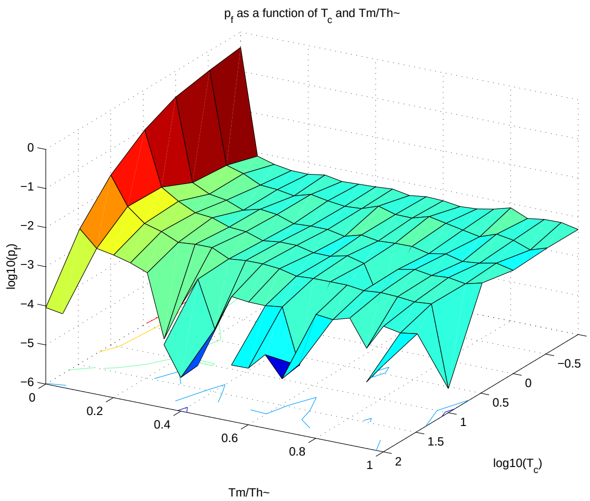

Fig. 8: The simulated p_f over the same parameter range as in Fig. 7.

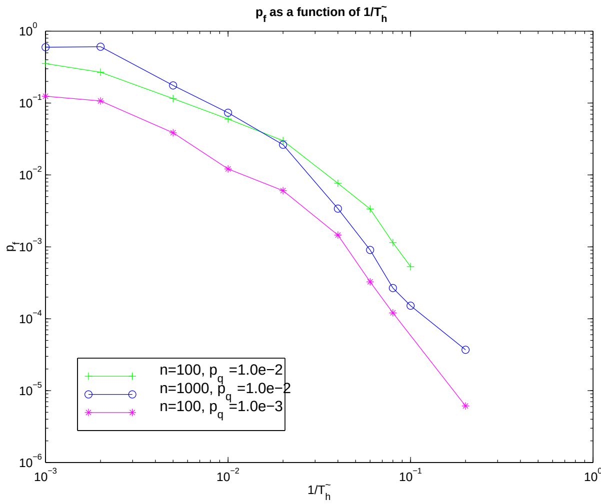

Fig. 9: The overflow probability for Starwars sources with memoryless estimation ($T_m = 0$).

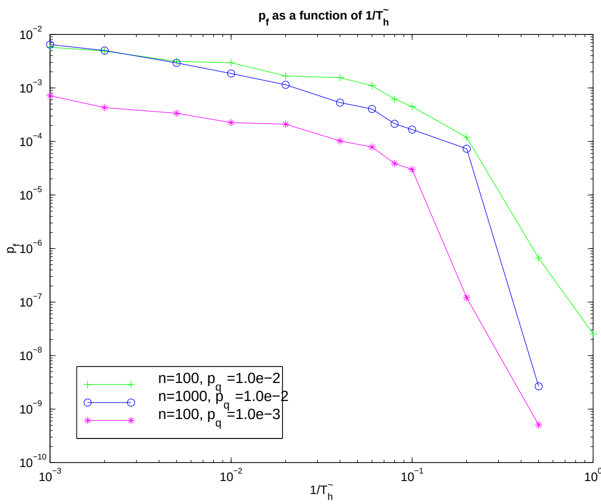

Fig. 10: The overflow probability for Starwars sources with $T_m = \widetilde{T_h}$.

particular, several studies of various types of network traffic have found phenomenon of long-range dependence (LRD) [18], [8], [1], [6]. However, based on the intuition gained from the single time-scale model, we expect that a memory window size on the order of $\widetilde{T_h}$ is again appropriate here. As before, flow departures dictate a critical time-scale $\widetilde{T_h}$ over which the statistics of the future behavior of the traffic has to be predicted. A memory window of $\widetilde{T_h}$ allows the simultaneous *smoothing* of the fluctuations faster than $\widetilde{T_h}$ for reliable estimation and the *tracking* of fluctuations at a time-scale larger than $\widetilde{T_h}$. The statistics of the long-term fluctuations of long-range dependence is therefore irrelevant.

To provide some support for this hypothesis, we present simulation results on an actual traffic trace. Figure 9 and 10 show the overflow probability when the flow is a piecewise CBR version of the MPEG-1 encoded Starwars movie [10]. This particular trace has been shown to exhibit long-range dependence [8]. We vary the average flow holding time and plot the overflow probability as a function of $1/\widetilde{T_h}$. As with the synthetic traffic above, we see that the performance is not robust under memoryless estimation. When $\widetilde{T_h}$ is large (corresponding to small T_c in Fig. 7), the performance misses the target by 1 or 2 orders of magnitude. On the other hand, we note that with the choice of memory window size $T_m = \widetilde{T_h}$, the MBAC is robust (cf. Fig. 10). Apparently, the strong long-term fluctuations of this traffic do not degrade the performance of the MBAC.

V. Discussion

A. Critical Time-Scale

Our analysis has demonstrated the fundamental importance of the critical time-scale $\widetilde{T_h}$ as *the time-scale over which the effect of admission decisions persists.* This insight leads to two important principles for the design of robust and efficient MBAC schemes. First, traffic fluctuations on a time-scale longer than the critical time-scale fall into the *repair regime*; these fluctuations should be tracked by the MBAC so that they can be compensated for by flow admissions and rejections. Second, spare link bandwidth should be set aside to absorb fluctuations at a time-scale shorter than $\widetilde{T_h}$, as these fluctuations are too fast to be compensated for by the repair effect. A consequence is that a robust MBAC should *predict* the fluctuation statistics over a time-scale of $\widetilde{T_h}$, rather than *estimate* the long-term statistics of the traffic. In this context, it does not matter whether or not the traffic is stationary or not over a time-scale much longer than $\widetilde{T_h}$, or if the traffic exhibits long-range dependence (LRD).

By setting the measurement window size to be $\widetilde{T_h}$, our scheme implicitly embodies the first principle: effective tracking of traffic fluctuations slower than $\widetilde{T_h}$. On the other hand, the scheme sets aside spare bandwidth of the order $\sigma\sqrt{n}$. Since σ^2 is the long-term variance of a flow, this leads to an over-conservative spare bandwidth allocation when much of the fluctuation is actually slower than $\widetilde{T_h}$. (This can be seen in Fig. 7, where the actual overflow probability p_f drops rapidly with increasing traffic correlation time-scale T_c.) In a sequel to this paper [12], we propose a novel MBAC design which goes one step further. By appropriate filtering of the traffic measurements, the MBAC scheme simultaneously tracks slow fluctuations and estimates the variance of fast fluctuations, so that the appropriate amount of spare bandwidth can be set aside.

B. Heterogeneous Flows

Although the results in this paper are derived under the ideal assumption of identical traffic statistics across flows, many of the ideas are in fact extensible to a heterogeneous environment. The key concept behind our approach is the existence of an appropriate operating point about which the load fluctuates. We express this operating point in terms of the number of admissible flows m^*. However, when flow statistics are heterogeneous, the operating point should be thought of as a *target utilization level* of $m^*\mu$. This level depends on the statistics of the individual flows *only* through the statistics of the aggregate traffic (mean and variance in the central-limit framework). Moreover, as long as there are many independent flows in the system and no single flow dominate the entire link, the traffic fluctuations can be well-approximated as Gaussian even in the heterogeneous case. In [12], we extend the framework developed here to analyze the performance of the MBAC design under heterogeneous flows.

VI. Related Work

Past work on measurement-based admission control [5], [19], [15] have either ignored measurement errors or assumed a static situation where calls do not arrive or depart the system and there is arbitrarily long time to make accurate measurements. Here we discuss three more recent papers which are closer in spirit to our work.

Jamin *et al.*, in [16], presented a specific algorithm for measurement-based admission control of predictive traffic, and evaluated its performance through simulation. The algorithm relies on measurements of the maximum delay and maximum bandwidth over a measurement interval. There are several tuning parameters in the algorithm (sampling window size S, measurement window size T, utilization target, back-off factor λ) that are found to have a significant impact on performance. We believe that our work offers some insight into the impact of these system parameters. In particular, the measurement window size T is very similar to our measurement time-scale T_m. Also, λ is a parameter that controls an *overestimation* of the actual measured delay - in other words, it controls conservativeness, which in our work is represented through the parameter p'_q.

The MBAC algorithm proposed by Casetti, Kurose and Towsley [4] recognizes the importance of the measurement window size as a system parameter. The authors propose an adaptive algorithm to determine an appropriate window size. While this is an improvement over the fixed window length parameter in [16], the adaptive algorithm itself has external tuning parameters. It is not clear if the overall system is easier to tune.

Gibbens *et al.* [9] studied *memoryless* measurement-based admission control in a decision-theoretic framework. Their work takes into account the impact of measurement errors on performance and also considers the call dynamics. However, there are some significant differences between their and our work. First, a perfect time-scale separation is explicitly built into their model by assuming that the network states seen by successive call arrivals are independent. This makes it difficult to evaluate the performance of MBAC schemes with memory and also the effect of traffic correlation on a system with very high call arrival rates. Indeed they only focused on *memoryless* schemes. Moreover, our results show that the condition for time-scale separation is rather subtle, as it depends, among others parameters, on the system size. Second, while they also observed that a memoryless certainty equivalent scheme can perform poorly, their remedy is quite different. They relied on essentially two mechanisms: the use of a Bayesian prior on the call statistics and network state-independent call rejection. The first mechanism serves to smooth out the fluctuation in successive memoryless estimates, as the observations are weighted by a fixed prior. The second mechanism counters very high call arrival rates, by not accepting calls until one has left the system. In contrast, we propose the use of an appropriate amount of memory in the estimator, which as we have seen deals with both these problems. Our framework, without *a priori* assuming time-scale separation, allows us to evaluate the performance as a function of the amount of memory used. We believe the appropriate use of memory is a natural and effective strategy, particularly when no reliable prior exists.

Our approach of abstracting away enough details of the measurement and admission decision process in order to focus on the fundamental issues of measurement uncertainty and system dynamics is corroborated by recent work by Jamin and Shenker [17]. They simulate several specific MBAC algorithms that have been proposed in the literature, and find them to be essentially equivalent with appropriate tuning of system parameters. In our work, we attempt to study, in a sense, the common denominator of these proposed schemes, and focus on how the "performance tuning knobs", such as memory window size and degree of conservativeness, should be set to achieve robustness.

VII. Conclusion

In this paper, we have presented a framework for studying the performance of admission control schemes under measurement uncertainty and flow dynamics. Using heavy-traffic approximations, the analysis of the resulting dynamical system is simplified via linearization around a nominal operating point and by Gaussian approximations of the statistics via central limit theorems. The insights gained include:

- quantification of the impact of estimation errors on the QoS performance of MBAC schemes;
- identification of a *critical time-scale* for which the effect of admission decisions persist;
- demonstration of precisely how the memory time-scale of the estimators affects performance and what the appropriate choice of memory time-scale is to achieve robust performance.

These insights are directly applicable to the design of robust MBAC schemes. Such schemes do not have to rely on external tuning parameters to achieve the desired performance despite the inherent measurement uncertainty and the complicated system dynamics.

References

[1] J. Beran, R. Sherman, and W. Willinger. Long Range Dependence in Variable Bit Rate Video Traffic. *IEEE Trans. on Communications*, 43(3):1566–1579, February 1995.

[2] P. Billingsley. *Convergence of Probability Measures.* Wiley,New York, 1968.

[3] P. Billingsley. *Probability and Measure (3rd Ed.).* Wiley, 1995.

[4] C. Casetti, J. Kurose, and D. Towsley. An Adaptive Algorithm for Measurement-based Admission Control Integrated Services Packet Networks. In *Proc. Workshop on Protocols for High Speed Networks*, Sophia Antipolis, France, October 1996.

[5] Costas Courcoubetis et al. Admission Control and Routing in ATM Networks using Inferences from Measured Buffer Occupancy. In *ORSA/TIMS special interest meeting*, Monterey, CA, January 1991.

[6] M. Crovella and A. Bestavros. Self-Similarity in World Wide Web

Traffic: Evidence and Possible Causes. In *Proc. ACM Sigmetrics '96*, pages 160–169, Philadelphia, PA, May 1996.
[7] J. Cuzick. Boundary Crossing Probabilities for Stationary Gaussian Processes and Brownian Motion. *Transactions of the American Mathematical Society*, pages 469–492, February 1981.
[8] M. W. Garrett and Walter Willinger. Analysis, Modeling and Generation of Self-Similar VBR Video Traffic. In *Proc. ACM SIGCOMM '94*, pages 269–280, London, UK, August 1994.
[9] R.J. Gibbens, F.P. Kelly, and P.B. Key. A Decision-theoretic Approach to Call Admission Control in ATM Networks. *IEEE JSAC, Special issue on Advances in the Fundamentals of Networking*, August 1995.
[10] M. Grossglauser, S. Keshav, and D. Tse. RCBR: A Simple and Efficient Service for Multiple Time-Scale Traffic. *IEEE/ACM Transactions on Networking*, December 1997.
[11] M. Grossglauser and D. Tse. A Framework for Robust Measurement-Based Admission Control. In *Proc. ACM SIGCOMM '97*, Cannes, France, September 1997.
[12] M. Grossglauser and D. N. C. Tse. A Time-Scale Decomposition Approach to Measurement-Based Admission Control. In *Proc. IEEE INFOCOM '99*, New York, March 1999.
[13] H. U. Bräker. High boundary excursions of locally stationary Gaussian processes. In *Proc. of the Conference on Extreme Value Theory and Applications*, Gaithersburg, Maryland, USA, May 1993.
[14] H. U. Bräker. *High boundary excursions of locally stationary Gaussian processes.* PhD thesis, Universität Bern, Switzerland, 1993.
[15] I. Hsu and J. Walrand. Dynamic Bandwidth Allocation for ATM Switches. *Journal of Applied Probability*, September 1996.
[16] S. Jamin, P. B. Danzig, S. Shenker, and L. Zhang. A Measurement-Based Admission Control Algorithm for Integrated Services Packet Networks. *IEEE/ACM Transactions on Networking*, 5(1), February 1997.
[17] Sugih Jamin and Scott Shenker. Measurement-Based Admission Control Algorithms for Controlled-Load Service: A Structural Examination. CSE-TR-333-97, University of Michigan, Ann Arbor, April 1997.
[18] Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson. On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Trans. on Networking*, 2(1):1–15, February 1994.
[19] H. Saito and K. Shiomoto. Dynamic Call Admission Control in ATM Networks. *IEEE Journal on Selected Areas of Communications*, 9:982–989, 1991.
[20] D.W. Stroock and S.R.S. Varadhan. *Multidimensional Diffusion Proceeses.* Springer Verlag, 1979.
[21] D. Tse and M. Grossglauser. Measurement-Based Call Admission Control: Analysis and Simulation. In *Proc. IEEE INFOCOM '97*, Kobe, Japan, April 1997.

Appendix

I. Simulation Setup

We simulate the admission controller under infinite load and we measure the resulting overflow probability p_f on a bufferless link of capacity c. We describe the details of the simulation setup.

We model traffic flows as fluid flows, i.e., we do not simulate individual packets. In particular, we use a piecewise constant-rate traffic model, where the fluid rate only changes at certain points in time, and remains constant between these points [10]. The advantage of this traffic model is that it lends itself to very efficient simulation.

We use two types of flows. The first type is based on a stochastic model. Each flow is the realization of an independent, identically distributed, stationary fluid process. This fluid process is modulated by an underlying renewal process; the fluid rate is constant on the time interval between two consecutive renewals. The fluid rate in each interval is chosen independently according to the marginal rate distribution, which in this case is Gaussian with $\sigma/\mu = 0.3$. The renewal time distribution is exponential with mean T_c, which implies that the autocorrelation function of the traffic rate process is precisely as in (25).

The second flow type is based on an actual video trace, namely a two-hour MPEG-1 encoded version of the Starwars movie [8]. We use the smoothing algorithm described in [10] to transform this trace into a piecewise constant-rate flow. For this flow model, $\sigma/\mu = 0.33$. Each flow is a subsequence of the trace (with wrap-around at the end of the trace) with a randomly chosen starting point. Note that this trace exhibits the LRD property, i.e., its autocorrelation function decreases subexponentially [8]; also, its empirical marginal rate distribution is not Gaussian. For both traffic models, the flow holding time is exponentially distributed with mean T_h.

We periodically sample the empirical overflow probability p_f over intervals of length $2 \cdot \max(\widetilde{T_h}, T_m, T_c)$. This sample period is long enough to give approximately independent samples, as the only "memory" in the system is due to flow dynamics, estimation memory, and traffic correlation. By chosing a multiple of the maximum of their respective time-scales, we are assured that consecutive samples represent essentially independent observations of the system.

After each sample period, we compute the (two-sided) 95% confidence interval around the estimated value of p_f. We terminate simulations when (a) the 95% confidence interval is less than 20% of the estimated value (i.e., we are confident that the estimated p_f is close enough to the actual p_f), or (b) the estimated value plus the (one-sided) confidence interval is at least two orders of magnitude below the target overflow probability (i.e., we are confident that the estimated p_f is considerably lower than the target overflow probability p_q). We also discard all initial samples until the simulated MBAC has reached steady state.

II. Weak Convergence Results for Heavy-Traffic Approximation

In this appendix, we will prove Theorem III.1, giving a rigorous justification of the heavy traffic approximations we used.

Definition B.1: The space $\mathcal{D}[0,\infty]$ is the space of all real-valued functions on $[0,\infty)$ that are continuous from the right and have limits from the left. There is a metric (Skorohod metric) defined on $\mathcal{D}[0,\infty]$ such that it is complete and separable.

Definition B.2: Let $\{Z_t^{(n)}\}$ be a sequence of processes whose sample paths are in $\mathcal{D}[0,\infty]$. $\{Z_t^{(n)}\}$ is said to *converges weakly* to $\{Z_t\}$ if for every continuous function $f : \mathcal{D}[0,\infty] \to \Re$, $E[f(\{Z_t^{(n)}\})] \to E[f(\{Z_t\})]$.

With a slight abuse of notation, we will use $\xrightarrow{\mathcal{D}}$ to denote weak convergence of processes as well as convergence in distribution for random variables. We shall use the following theorem to verify weak convergence.

Theorem B.3: A sequence of processes $\{Z_t^{(n)}\}$ converges weakly to $\{Z_t\}$ if all finite-dimensional distributions converge and $\{Z_t^{(n)}\}$ is *tight*, i.e.

1) For every $\eta > 0$, there exists an $a > 0$ such that

$$\Pr\left\{|Z_0^{(n)}| > a\right\} \le \eta \qquad \forall n.$$

2) For every $T > 0$, $\epsilon, \eta > 0$, there exists a $\delta \in (0,1)$ and an integer n_0 such that

$$\Pr\left\{\sup_{|t_1-t_2|\le\delta, 0\le t_1,t_2\le T} |Z_{t_1}^{(n)} - Z_{t_2}^{(n)}| > \epsilon\right\} \le \eta \qquad \forall n \ge n_0.$$

We will use the following theorems [2].

Theorem B.4: (Continuous-Mapping Theorem for Processes) Let $\{Z_t^{(n)}\}$ be a sequence of processes whose sample paths are in $\mathcal{D}[0,\infty]$. If $h : \mathcal{D}[0,\infty] \to \mathcal{D}[0,\infty]$ is continuous and $\{Z_t^{(n)}\} \xrightarrow{\mathcal{D}} \{Z_t\}$, then $g(\{Z_t^{(n)}\}) \xrightarrow{\mathcal{D}} g(\{Z_t\})$.

Theorem B.5: Let $\{W_t^{(n)}\}$ and $\{Z_t^{(n)}\}$'s be processes defined on the same probability space, and $g : \mathcal{D}[0,\infty] \times \mathcal{D}[0,\infty] \to \mathcal{D}[0,\infty]$ is continuous. If $\{W_t^{(n)}\} \xrightarrow{\mathcal{D}} \{W_t\}$ and $\{Z_t^{(n)}\}$ converges weakly to a deterministic process $\{Z_t\}$, then $g(\{W_t^{(n)}\},\{Z_t^{(n)}\}) \xrightarrow{\mathcal{D}} g(\{W_t\},\{Z_t\})$.

We need the following technical conditions on the flow processes.

Assumptions B.6: 1) The sample paths of the individual flow processes $\{X_i(t)\}$ are in $\mathcal{D}[0,\infty]$.

2) The mean bandwidth estimates $\{\hat{\mu}_s^{(n)}\}$ converges weakly to the constant process μ.

3) The standard deviation estimates $\{\hat{\sigma}_s^{(n)}\}$ converges weakly to the constant process σ.

4) If we define

$$Y_t^{(n)} := \frac{1}{\sigma\sqrt{n}}\sum_{i=1}^{n}[X_i(t)-\mu]$$

to be the scaled and centered sum of the individual flows, then as $n \to \infty$, $\{Y_t^{(n)}\}$ converges weakly to $\{Y_t\}$, which is a stationary zero-mean Gaussian process with unit variance and auto-correlation function $\rho(t)$ (that of an individual flow).

The fourth condition says that the aggregation of the individual flows satisfies a functional central limit theorem. It holds for a very broad class of models for the individual sources. For example, it can be shown [20] that the condition holds if $\{X_i(t)\}$ is a K-state continuous-time Markov fluid, in which case the limiting process $\{Y_t\}$ is a linear functional of a $K-1$-dimensional diffusion process.

To prove the main theorem, we need the following lemma, which can be viewed as a functional law of large number for the process describing the evolution of the number of flows in the system.

Lemma B.7: The process $\{\frac{N_t^{(n)}}{n}\}$ converges weakly to the deterministic process taking on a constant value of 1 for all t.

Proof: By solving (15), we get for each s,

$$M_0^{(n)} = \frac{1}{4(\hat{\mu}^{(n)})^2}\left(\sqrt{(\hat{\sigma}^{(n)})^2\alpha_q^2 + 4n\mu\hat{\mu}^{(n)}} - \hat{\sigma}^{(n)}\alpha_q\right)^2 \quad (36)$$

Using assumptions (2) and (3) in B.6, together with Theorem B.5, we can see that the process $\{\frac{M_s^{(n)}}{n}\}$ converges to the process taking on a constant value. Now, for all $t \geq 0$,

$$\frac{M_t^{(n)}}{n} \leq \frac{N_t^{(n)}}{n} \leq \sup_{0\leq s\leq t}\frac{M_s^{(n)}}{n}$$

Since $\{\frac{M_s^{(n)}}{n}\}$ converges weakly to the constant process 1, so does the process $\{\sup_{0\leq s\leq t}\frac{M_s^{(n)}}{n}\}$, by the continuous mapping theorem. Hence $\{\frac{N_t^{(n)}}{n}\}$ must converge weakly to the constant process 1. ■

Proof: Proof of Theorem III.1

Using (36), we get for each s,

$$\frac{M_s^{(n)}-n}{\sqrt{n}} = \frac{\sqrt{n}(\mu-\hat{\mu}_s^{(n)})}{\hat{\mu}_s^{(n)}} + \frac{(\hat{\sigma}_s^{(n)})^2\alpha_q^2}{2(\hat{\mu}_s^{(n)})^2\sqrt{n}} - \frac{\hat{\sigma}_s^{(n)}\alpha_q}{2(\hat{\mu}_s^{(n)})^2}\sqrt{\frac{(\hat{\sigma}_s^{(n)})^2\alpha_q^2}{n} + 4\mu\hat{\mu}_s^{(n)}} \quad (37)$$

By assumption B.6, we know that $\{\sqrt{n}(\mu-\hat{\mu}_s^{(n)})\} \xrightarrow{\mathcal{D}} -\sigma Y_s$, where $\{Y_s\}$ is a zero mean Gaussian process with auto-correlation function ρ. Also, $\{\hat{\mu}_s^{(n)}\}$ converges weakly to the constant process μ and $\{\hat{\sigma}_s^{(n)}\}$ converges weakly to the constant process σ. By Theorem B.5,

$$\{\frac{M_s^{(n)}-n}{\sqrt{n}}\} \xrightarrow{\mathcal{D}} \{-\frac{\sigma}{\mu}(Y_s+\alpha_q)\} \quad (38)$$

Next, we will show that for fixed $t > 0$, $\{D[s,t]\}$ as a process in s converges weakly to the deterministic process $\{\frac{s-t}{\widetilde{T_h}}\}$ on $[0,t]$. First, let us fix an $s < t$. Define now two random variables $D^u[s,t]$ and $D^l[s,t]$. $D^l[s,t]$ is the number of flows departing from the system when there are $N(s)$ flows in the system at time s and no more flows enter the system in $[s,t]$; $D^u[s,t]$ is the number of flows departing from the system when there are $W := \sup_{\tau\in[s,t]} N(\tau)$ flows at time s and no more flows enter the system in $[s,t]$. It can be seen that for every x,

$$\Pr\left\{D^l[s,t] \geq x\right\} \leq \Pr\left\{D[s,t] \geq x\right\} \leq \Pr\left\{D^u[s,t] \geq x\right\} \quad (39)$$

Using Chebyshev's bound, we have for every $\epsilon > 0$,

$$\Pr\left\{\left|\frac{D^u[s,t]}{\sqrt{n}} - \frac{s-t}{\widetilde{T_h}}\right| > \epsilon\right\} \leq \frac{\mathrm{E}\left[\left(\frac{D^u[s,t]}{\sqrt{n}} - \frac{s-t}{\widetilde{T_h}}\right)^2\right]}{\epsilon^2}$$

The expectation can be computed using the fact that the flows have exponential holding time and depart from the system independently:

$$\mathrm{E}\left[\left(\frac{D^u[s,t]}{\sqrt{n}} - \frac{s-t}{\widetilde{T_h}}\right)^2\right] = \frac{\mathrm{E}[W]}{n}q + \frac{\mathrm{E}[W^2]-\mathrm{E}[W]}{n}q^2 - \frac{(s-t)^2}{\widetilde{T_h}^2} \quad (40)$$

where q is the probability that a given flow leaves the system some time in $[s,t]$, and is given by

$$q := 1 - \exp\left(\frac{-t}{\widetilde{T_h}\sqrt{n}}\right) \quad (41)$$

By Lemma B.7 and the continuous mapping theorem, as $n \to \infty$,

$$\mathrm{E}\left[\frac{W}{n}\right] \to 1 \qquad \mathrm{E}\left[\frac{W^2}{n^2}\right] \to 1$$

Substituting this into (42) shows that

$$\lim_{n\to\infty}\mathrm{E}\left[\left(\frac{D^u[s,t]}{\sqrt{n}} - \frac{s-t}{\widetilde{T_h}}\right)^2\right] = 0$$

Hence $\frac{D^u[s,t]}{\sqrt{n}}$ converges in probability, and hence in distribution, to $\frac{s-t}{\widetilde{T_h}}$. Using a similar argument, one can show the same thing for $D^l[s,t]$. By (40), this imply that for fixed s and t, $\frac{D[s,t]}{\sqrt{n}} \xrightarrow{\mathcal{D}} \frac{s-t}{\widetilde{T_h}}$. A standard argument in the theory of convergence in distribution implies that for all k and $s_1, \ldots, s_k \in [0,t]$, $(\frac{D[s_1,t]}{\sqrt{n}}, \ldots \frac{D[s_k,t]}{\sqrt{n}}) \xrightarrow{\mathcal{D}} (\frac{s_1-t}{\widetilde{T_h}}, \ldots, \frac{s_k-t}{\widetilde{T_h}})$, i.e. finite-dimensional distributions converge. To show weak convergence as a process, we need to verify tightness, according to Theorem B.3. The first condition is trivially satisfied. For the second condition,

$$\begin{aligned}
&\Pr\left\{\sup_{|s_1-s_2|\le\delta, 0\le s_1,s_2\le t} \frac{D[s_1,s_2]}{\sqrt{n}} > \epsilon\right\} \\
\le\ &\Pr\left\{\sup_{0\le k\le \frac{t}{\delta}} \frac{D[k\delta,(k+1)\delta]}{\sqrt{n}} > \epsilon\right\} \\
\le\ &\left(\frac{t}{\delta}+1\right)\sup_k \Pr\left\{\frac{D[k\delta,(k+1)\delta]}{\sqrt{n}} > \epsilon\right\} \\
\le\ &\left(\frac{t}{\delta}+1\right)\frac{1}{\epsilon^2}\sup_k \mathrm{E}\left[\frac{1}{n}(D[k\delta,(k+1)\delta])^2\right] \\
\le\ &\left(\frac{t}{\delta}+1\right)\frac{1}{\epsilon^2}\left(\frac{\mathrm{E}[U]}{n}p + \frac{\mathrm{E}[U^2]-\mathrm{E}[U]}{n}p^2\right) \qquad (42)
\end{aligned}$$

where $U := \sup_{\tau\in[0,t]} N_\tau^{(n)}$ and

$$p = \Pr\{\text{a flow departs in time } [k\delta,(k+1)\delta]\} = 1-\exp\left(\frac{\delta}{\widetilde{T_h}\sqrt{n}}\right).$$

By direct calculation, (45) is in turn equal to

$$\frac{1}{\epsilon^2}\left(\frac{t}{\delta}+1\right)\left(\frac{\delta^2}{\widetilde{T_h}^2}+o(1)\right)$$

where the $o(1)$ term goes to zero as $n \to \infty$. Thus, by appropriate choice of n and δ, (44) can be made arbitrarily small. this verifies the tightness of $\{\frac{D[s,t]}{\sqrt{n}}\}$ and hence its weak convergence.

Combining the weak convergence of $\{\frac{D[s,t]}{\sqrt{n}}\}$ and $\{\frac{M_t^{(n)}-n}{\sqrt{n}}\}$, it follows that

$$N_t^{(n)} \xrightarrow{\mathcal{D}} \sup_{0\le s\le t}\left\{-\frac{\sigma}{\mu}\left(Y_s - \frac{\mu(t-s)}{\sigma\widetilde{T_h}} + \alpha_q\right)\right\}$$

■

Matthias Grossglauser (S '92, M '99/ACM S '96, M '99) received the Diplôme d'Ingénieur en Systèmes de Communication from the Swiss Federal Institute of Technology (EPFL), the M.Sc. from the Georgia Institute of Technology, both in 1994, and his Ph.D. from the University of Paris 6 in 1998. He did most of his thesis work at INRIA Sophia Antipolis, France. Since November 1998, he has been a member of the Networking and Distributed Systems Laboratory of AT&T Labs - Research. He received the 1998 Cor Baayen Award from the European Research Consortium for Informatics and Mathematics (ERCIM). His research interests are in network traffic analysis and modeling, resource allocation, and network management.

David N. C. Tse (S '91, M '96) received the B.A.Sc. degree in systems design engineering from University of Waterloo, Canada, in 1989, and the M.S. and Ph.D. degrees in electrical engineering from Massachusetts Institute of Technology in 1991 and 1994 respectively. From 1994 to 1995, He was a postdoctoral member of technical staff at AT& T Bell Laboratories. Since 1995, he has been an assistant professor in the Department of Electrical Engineering and Computer Sciences in the University of California at Berkeley. He received a 1967 NSERC 4-year graduate fellowship from the government of Canada in 1989, a NSF CAREER award in 1998, and the Best Paper Award at the Infocom 1998 conference for his work with Stephen Hanly. His research interests are in the areas of communication networks, wireless communications and information theory.