Providing QoS with Measurement–Based Admission Control

Martin Reisslein
GMD FOKUS
Kaiserin–Augusta–Allee 31
10589 Berlin, Germany
reisslein@fokus.gmd.de
http://www.fokus.gmd.de/usr/reisslein

February 1999

Abstract

Integrated services networks perform call admission control before accepting a new connection to ensure that the QoS requirements of all connections are met. Rather than relying on *a priori* traffic descriptors (such as leaky buckets), which often poorly characterize the actual traffic, measurement–based admission control bases admission decisions on measurements of the actual traffic. In this paper we first develop a novel Large Deviations (LD) approach to measurement–based admission control. We then provide an extensive review of the existing literature on measurement–based admission control. We conduct simulation studies with traces of MPEG 1 encoded movies to compare the performance of the admission rules in the literature with that of our Large Deviations approach. We demonstrate that our LD approach achieves higher link utilizations. Finally, we compare the performance of measurement–based admission control with that of traditional admission control, which relies on *a priori* traffic descriptors. Our numerical work indicates that measurement–based admission control achieves significant gains in link utilizations over traditional admission control.

1 Introduction

Call admission control is performed in Integrated Services networks to ensure that the connections' Quality of Service (QoS) requirements are met. A call admission test is performed before a new connection is accepted. The new connection is accepted if and only if the network is able to meet the QoS requirements of all already existing connections as well as the new connection.

Traditional call admission tests are based on *a priori* characterizations (e.g. leaky bucket characterizations) of the connections' traffic [1]. Oftentimes, however, it is difficult, if not impossible, to provide an accurate *a priori* characterization of a connection's traffic. This is especially true for traffic emanating from live sources, such as the video traffic from the live coverage of a sporting event. Even if accurate *a priori* characterizations are available, however, traditional call admission tests typically over–provision networking resources. This is because traditional call admission tests

usually assume that the connections are adversarial to the extent permitted by the *a priori* characterizations and transmit worst–case traffic patterns [2, 3, 4, 5, 6]. This assumption, however, is often overly conservative, as in most practical circumstances connections do not transmit worst–case traffic patterns. As a consequence traditional call admission tests typically over–provision networking resources and thus underutilize the network.

Measurement–based admission control is a promising alternative to admission control based on *a priori* traffic descriptors. Instead of relying on *a priori* traffic characterizations, measurement–based admission control bases admission decisions primarily on traffic measurements. Admission decisions are based on measurements of the actual traffic from the already existing connections and an *a priori* characterization of the connection requesting establishment. The *a priori* characterization of the connection requesting establishment can be very simple, such as a peak rate specification. An overly conservative *a priori* characterization does not result in an over–provisioning of resources for the entire lifetime of the new connection, as the new connection — once admitted — is included in the measurements and is no longer characterized by its *a priori* specification. Thus measurement–based admission control is able to exploit the statistical multiplexing effect and achieves high network utilizations.

In this paper we discuss measurement–based admission rules that base admission decisions on measurements of the aggregate traffic from the already existing connections. We do not consider admission rules that require per–flow traffic measurements, as it is difficult to conduct per–flow measurements accurately and cost–efficiently in practice. Admission tests based on aggregate measurements can not enforce per–flow QoS; these would require per–flow measurements which are not practicable. Therefore, in this paper we focus on measurement–based admission rules that provide aggregate QoS. We study the measurement-based admission rules within the smoothing/bufferless multiplexing framework [5, 6]. The key aspects of the smoothing/bufferless multiplexing framework are to (1) pass each connection's traffic through a *buffered* smoother (peak rate limiter) at the connection's input to the network, and (2) use *bufferless* statistical multiplexing inside the network. The bufferless multiplexing inside the network has the advantage that a new connection's *a priori* characterization (e.g. peak rate) does not change as it passes through a bufferless node. Thus the same *a priori* characterization can be used for the admission test at each node traversed by the new connection. However, to simplify the discussion and highlight the measurement aspect of the admission rules we focus on a single bufferless node in this paper.

The contributions of this paper are threefold. First, we develop and evaluate a novel Large Deviations (LD) approach to measurement–based admission control. In this LD approach traffic measurements are used to estimate the logarithmic moment generating function of the aggregate arrival stream. From this estimate of the logarithmic moment generating function we compute an estimate of the loss probability at the node using the LD approximation. A new connection requesting establishment is accepted if the estimated loss probability is less than some miniscule QoS parameter ϵ, say $\epsilon = 10^{-6}$, and rejected otherwise. Secondly, we provide an extensive review of the existing literature on measurement–based admission control. We compare the performance

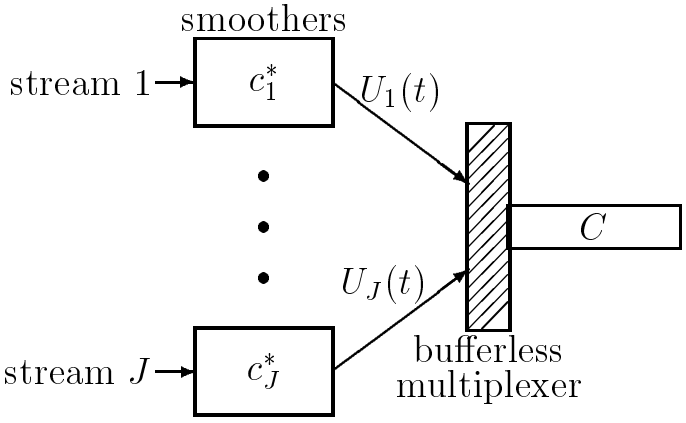

Figure 1: The traffic of connection j is passed through a smoother that limits the peak rate to c_j^*. The smoothed traffic is then multiplexed onto a bufferless link with capacity C.

of the measurement–based admission rules in the literature with that of our LD approach through simulations with traces of MPEG 1 encoded movies. The simulations demonstrate that our LD approach achieves both higher link utilizations and smaller loss probabilities than the existing admission rules. Lastly, we compare the performance of measurement–based admission control with that of traditional admission control which relies exclusively on *a priori* traffic characterizations. We demonstrate that measurement–based admission control achieves significantly higher link utilizations than traditional admission control that relies on leaky bucket characterizations and assumes worst–case on–off traffic patterns.

2 A Large Deviations Approach to Measurement–Based Admission Control

In this section we develop our Large Deviations (LD) approach to measurement–based admission control. We develop a basic admission rule first and study then some important refinements. We focus throughout this paper on a single node consisting of a bufferless multiplexer that feeds into a link of capacity C. We view traffic as fluid. The fluid model, which closely approximates a packetized model with small packets, permits us to focus on the central issues and significantly simplifies notation. Consider a set of J connections. In the smoothing/bufferless multiplexing framework each connection j, $j = 1, \ldots, J$, is passed through a buffered smoother before it is multiplexed onto the bufferless link. The smoother limits the peak rate of connection–j traffic entering the bufferless multiplexer to c_j^* (see Figure 1). Let $U_j(t)$, $j = 1, \ldots, J$, denote the rate at which connection–j traffic arrives to the bufferless multiplexer at time t. The smoother ensures that $U_j(t) \leq c_j^*$ $\forall t \geq 0$. Now regard the jth arrival process as a stochastic process. Let $(U_j(t),\ t \geq 0)$ denote the jth arrival process. Let $X(t)$ denote the aggregate arrival rate at time t:

$$X(t) = \sum_{j=1}^{J} U_j(t),$$

and let $(X(t),\ t \geq 0)$ denote the aggregate arrival process. The expected long–run fraction of traffic lost due to link overflow is

$$P_{\text{loss}} = E\left[\lim_{\Delta\to\infty} \frac{\int_0^\Delta (X(t)-C)^+ dt}{\int_0^\Delta X(t)dt}\right], \tag{1}$$

where the expectation is over all arrival processes and $(x)^+ := \max(0,\ x)$.

Our goal is to develop a measurement–based call admission rule that ensures that P_{loss} is less than some minute ϵ, such as $\epsilon = 10^{-4}$ or 10^{-6}. The call admission decisions are based on measurements of the aggregate arrival rate. In practical systems, however, it is impossible to measure the instantaneous arrival rate $X(t)$. For this reason, we divide time into slots of length T and measure the amount of traffic arriving in an interval of length T. Let X_n denote the amount of traffic arriving in the interval $[nT,\ (n+1)T]$, i.e.,

$$X_n = \int_{nT}^{(n+1)T} X(t)dt.$$

For small T we can reasonably approximate:

$$\int_{nT}^{(n+1)T} (X(t)-C)^+ dt \approx (X_n - CT)^+. \tag{2}$$

This approximation is particularly good when the fluctuation of the aggregate arrival process $(X(t),\ t \geq 0)$ is on a time scale larger than the slot length T. The slot length should therefore be set to the smallest value that allows for meaningful traffic measurements. In practical systems we suggest to set T to a few packet transmission times. Throughout this paper we shall assume that the approximation (2) is exact. Substituting (2) into (1) we obtain:

$$P_{\text{loss}} = E\left[\lim_{N\to\infty} \frac{\sum_{n=0}^N (X_n - CT)^+}{\sum_{n=0}^N X_n}\right]. \tag{3}$$

A practical measurement–based call admission rule can not rely on measurements over an infinite time horizon, but instead must base its decisions on some finite portion of the history of the aggregate streams behavior. We propose to base admissions decisions on the measured aggregate arrivals in the past M slots, i.e., $M \geq 1$ is the measurement window. Before we describe our admission rule in detail we need to introduce some notation. Let k denote the slot in which a new stream with smoother rate c_k^* requests connection establishment. Our admission rule relies on the measured aggregate arrivals in slots $k-M,\ldots,k-1$. Let $x_i,\ i=1,\ldots,M$, denote the measured aggregate arrivals in slot $k-i$.

Now consider the random variable X_k denoting the (not yet measured) aggregate arrivals in slot k. Define the estimated loss probability $P_{\text{loss}}^{\text{est}}$ as follows:

$$P_{\text{loss}}^{\text{est}} = \frac{E[(X_k - (C-c_k^*)T)^+]}{E[X_k]}. \tag{4}$$

$P_{\text{loss}}^{\text{est}}$ is the expected fraction of traffic lost by the already established connections at a bufferless link of capacity $C - c_k^*$ during slot k. Note that we are conservatively setting aside the peak rate

c_j^*	smoother rate (= peak rate) of connection j in bit/sec
C	service rate of multiplexer in bit/sec
ϵ	QoS parameter (= target loss probability)
k	index of slot in which a new connection with smoother rate c_k^* requests establishment
m_X	estimate of average aggregate arrivals in one slot in bit
$P_{\text{loss}}^{\text{est}}$	estimated loss probability; computed from m_X and $\mu_X(s)$ using Large Deviation approximation
P_{loss}	actual loss probability; obtained through simulation
T	slot length in seconds
x_i	measured aggregate arrivals in slot $k-i$ in bit
$\mu_X(s)$	estimate of logarithmic moment generating function of aggregate arrivals in one slot

Table 1: Summary of notation.

c_k^* for the stream requesting establishment. Our strategy is to base admission decisions on $P_{\text{loss}}^{\text{est}}$. If $P_{\text{loss}}^{\text{est}} \leq \epsilon$ connection k is admitted, otherwise it is rejected.

We evaluate $P_{\text{loss}}^{\text{est}}$ using the Large Deviations (LD) approximation. Toward this end, let m_X denote the estimate of $E[X_k]$, the mean of X_k. We compute the estimate m_X by averaging over the aggregate arrivals in slots $k-M, \ldots, k-1$:

$$m_X = \frac{1}{M} \sum_{i=1}^{M} x_i.$$

Furthermore, let $\mu_X(s)$ denote the estimate of $\ln E[e^{sX_k}]$, the logarithmic moment generating function of X_k. (The notation used in this paper is summarized in Table 1.) Again, we compute $\mu_X(s)$ by averaging over the M latest measurements:

$$\mu_X(s) = \ln \frac{1}{M} \sum_{i=1}^{M} e^{s x_i}. \tag{5}$$

Note that (5) generalizes the notion of the sample mean and allows for the estimation of arbitrary moments of the aggregate arrival stream. For instance, we obtain the logarithm of m_X by evaluating the derivative of $\mu_X(s)$ with respect to s at $s = 0$, i.e., $\mu_X'(0) = \ln m_X$. The LD approximation of (4) is given by [1]

$$P_{\text{loss}}^{\text{est}} \approx \frac{1}{m_X s^{\star 2} \sqrt{2\pi \mu_X''(s^\star)}} e^{-s^\star (C - c_k^*)T + \mu_X(s^\star)}, \tag{6}$$

where $s^\star$ is the unique solution to

$$\mu_X'(s^\star) = (C - c_k^*)T. \tag{7}$$

In summary, our basic measurement–based admission rule works as follows: Suppose that in slot k a connection with peak rate c_k^* requests establishment and the QoS requirement is $P_{\text{loss}} \leq \epsilon$. First, we estimate the logarithmic moment generating function of the aggregate arrival stream based on the measurements in the last M slots using (5). We then estimate $P_{\text{loss}}^{\text{est}}$ using the LD approximation (6) and admit connection k if $P_{\text{loss}}^{\text{est}} \leq \epsilon$, otherwise connection k is rejected.

Simulation Results

We now evaluate the basic measurement–based admission rule using traces from MPEG 1 encoded movies. We obtained the frame size traces, which give the number of bits in each video frame, from the public domain [7]. The movies were compressed with the Group of Pictures (GOP) pattern IBBPBBPBBPBB at a frame rate of $F = 24$ frames/sec. Each of the movie traces available from [7] has $N = 40{,}000$ frames, corresponding to about 28 minutes. Let $f_n(j),\ n = 1, \ldots, N$, denote the size of the nth frame of video j in bits. We convert the discrete frame size trace to a fluid flow by transmitting the nth frame of video j at rate $f_n(j) \cdot F$ over the interval $[n - 1/F, n/F]$. The numerical results reported in this paper were obtained with the *"Silence of the Lambs"* (lambs) trace. The lambs trace is the burstiest trace available from the library of traces [7]. Specifically, the lambs trace has an average frame size of 8,048 bit, which corresponds to an average rate of 193.2 kbit/sec. The trace has a peak–to–mean ratio of 18.4 and is therefore considered extremely bursty. The lambs trace furthermore has a Hurst parameter of 0.92, which indicates a high degree of Long Range Dependence [7]. Because of its burstiness and Long Range Dependence the lambs trace poses a particular challenge for admission control. In the numerical experiments reported in this paper all video streams use the lambs trace but each stream has its own independent random phase.

We evaluate the measurement–based admission rule within the smoothing/bufferless multiplexing framework [5, 6]. Each video stream is passed through a smoother before it enters the bufferless multiplexer. The smoother for connection j consists of a buffer which serves the traffic at rate c_j^*. When the smoother buffer is nonempty, traffic is drained from the smoother at rate c_j^*. When the smoother buffer is empty and connection–j's traffic is arriving at a rate less than c_j^*, traffic leaves the smoother exactly at the rate at which it enters the buffer. The smoother thus limits the peak rate of connection–j traffic entering the multiplexer to c_j^*. The smoother rate c_j^* is set to the smallest value that guarantees that the video traffic is delayed by no more than a connection–specific delay limit in the smoother (see [5] for details). We initially set the delay limit for all connections to 10 frame periods, i.e., 10/24 seconds. The corresponding smoother rate for the lambs trace is 731.6 kbit/sec. Throughout this paper we set the rate of the bufferless multiplexer to $C = 45$ Mbit/sec.

We simulate the system consisting of smoothers and bufferless multiplexer on a per frame period basis. Throughout we set the slot length of the measurement algorithm to the length of one frame period, i.e., $T = 1/F$. In the simulation calls arrive according to a Poisson process. We fix the call arrival rate at 1 call/10 frame period; thus the time between call arrivals is exponentially distributed with a mean of 10 frame periods. For each accepted call we draw a random starting frame. The

M	50	100	200	500	1000
J_{avg}	204	201	198	192	183
P_{loss}	$9.8 \cdot 10^{-4}$	$6.0 \cdot 10^{-4}$	$3.8 \cdot 10^{-4}$	$1.9 \cdot 10^{-4}$	$1.3 \cdot 10^{-4}$
M	2000	4000	6000	9000	12000
J_{avg}	171	147	131	107	93
P_{loss}	$7.6 \cdot 10^{-5}$	$5.3 \cdot 10^{-5}$	$4.5 \cdot 10^{-5}$	$2.8 \cdot 10^{-5}$	$3.2 \cdot 10^{-5}$

Table 2: Evaluation of basic measurement–based admission rule. Average number of admitted streams and loss probability for different measurement window lengths. The QoS parameter is set to $\epsilon = 10^{-6}$.

starting frames are independent and uniformly distributed over $[1, N]$. For each call we also draw a random life time. In our first set of experiments we draw the life times from an exponential distribution with a mean of 6,000 frame periods (= 250 seconds). With these parameters the link operates in constant overload. To see this, note that the stability limit of the 45 Mbps link is 232 lambs streams, each with an independent exponentially distributed lifetime with a mean of 6,000 frame periods. We allow the simulation to warm up for 60,000 frame periods. We determined by visual inspection that 60,000 frame periods is sufficient for the systems to reach steady state.

The goal of the simulations is to obtain estimates for the average number of admitted streams, denoted by J_{avg}, and the loss probability P_{loss} (3). We estimate J_{avg} and P_{loss} using the method of batch means [8]. We use a batch size of 6,000 frame periods. In order to ensure the independence of the batches, we separate successive batches by 12,000 frame periods, twice the average life time of a stream. We run each simulation until the width of the 90 % confidence interval of the loss probability is less than 20 % of the corresponding point estimate. We observed that the estimate for the average number of admitted streams converges much faster than the estimate for the loss probability. By the time the confidence interval for the loss probability has converged to less than 20 % of the point estimate, the width of the 90 % confidence interval for the number of admitted streams is typically less than 1 % of the point estimate.

In the first set of simulation experiments we evaluate the basic measurement–based admission rule. We set the QoS parameter to $\epsilon = 10^{-6}$ and run simulations for different values of M, the length of the measurement window. The results are reported in Table 2; in order to avoid visual clutter only point estimates are reported. We observe from the table that the loss probabilities are one to two orders of magnitude larger than the target loss probability $\epsilon = 10^{-6}$.

Figure 2 shows typical sample path plots from the simulation of the basic measurement–based admission rule with M = 6,000. The time axis is in units of slots of length $1/F$ seconds (= 1/24 seconds). Figure 2(a) shows the aggregate arrivals in each slot. Figure 2(b) shows $P_{\text{loss}}^{\text{est}}$ versus time. $P_{\text{loss}}^{\text{est}}$ is only computed (and plotted) when a call requests establishment. Recall that the call inter arrival times are exponentially distributed with a mean of 10 slots. Figure 2(c) shows the number

of admitted connections, J, versus time. We observe that the plots show strong oscillations. The reason for these oscillations lies in the way the logarithmic moment generating function $\mu_X(s)$ (5) (from which the loss probability $P_{\text{loss}}^{\text{est}}$ is computed) is estimated. Recall that $\mu_X(s)$ is estimated by averaging over the measured aggregate arrivals in the past M slots (M = 6,000 slots in the example shown). Consider the period from slot 64,800 to slot 71,000. $P_{\text{loss}}^{\text{est}}$ is larger than $\epsilon = 10^{-6}$ throughout this period and no new connections are admitted. The number of streams and the aggregate arrivals are dropping as connections terminate. $P_{\text{loss}}^{\text{est}}$, however, is slow to reflect this drop in utilization, as it is computed by equally weighing the aggregate arrivals in the past 6,000 slots. $P_{\text{loss}}^{\text{est}}$ remains above the threshold ϵ until around slot 71,000 when the peak in the aggregate arrivals around slot 65,000 leaves the measurement window. After $P_{\text{loss}}^{\text{est}}$ drops below ϵ many new connections are admitted. $P_{\text{loss}}^{\text{est}}$, however, continues to drop and reaches its minimum around slot 72,300 since the newly admitted streams have initially little impact on the estimate of the logarithmic moment generating function $\mu_X(s)$. The number of stream climbs quickly to 221 and the aggregate arrivals exceed the link capacity of $1.875 \cdot 10^6$ bit/slot (=45 Mbps $\cdot$ 1/24 sec) before $P_{\text{loss}}^{\text{est}}$ jumps to $1.4 \cdot 10^{-5}$ around slot 73,000 and stops the acceptance of new streams.

These typical sample plots illustrate the shortcomings of the basic measurement–based admission rule. Weighing the measured aggregate arrivals in the measurement window uniformly results in strong oscillations and poor system performance. The link is underutilized for extended periods of time; hence J_{avg} is relatively small. At the same time, however, the loss probability is fairly large (and off the target loss probability ϵ) as the surges in the number of admitted streams lead periodically to losses. We next try to improve the measurement–based admission rule by weighing the more recent measurement more heavily when estimating the logarithmic moment generating function.

2.1 Non–Uniform Weight Refinement

In order to improve the performance of the Large Deviations admission rule, we now introduce an important refinement. The basic idea of this refinement is to give the recent measurements more weight when estimating the logarithmic moment generating function. In (5) all measurements are weighed uniformly with $1/M$. Our goal is to improve the measurement–based admission rule by assigning weights larger than $1/M$ to recent measurements and weights smaller than $1/M$ to older measurements. Toward this end, let p_i, $i = 1, \ldots, M$, denote weights with $0 \leq p_i \leq 1$ and $\sum_{i=1}^{M} p_i = 1$. Throughout this paper we use exponentially decaying weights:

$$p_i = \frac{e^{-i/\tau_p}}{\sum_{l=1}^{M} p_l}, \quad i = 1, \ldots, M,$$

where τ_p is a tuning parameter. Setting $\tau_p = \infty$ gives uniform weights $p_i = 1/M$ for $i = 1, \ldots, M$, which were used in the basic measurement–based admission rule. For finite τ_p the weights for the recent measurements are larger than the weights for older measurements. With the non–uniform

τ_p	∞	6000	3000	1200	600	300	120	60
J_{avg}	131	132	134	137	157	173	191	198
P_{loss}	$4.5 \cdot 10^{-5}$	$4.4 \cdot 10^{-5}$	$4.0 \cdot 10^{-5}$	$3.1 \cdot 10^{-5}$	$6.2 \cdot 10^{-5}$	$8.1 \cdot 10^{-5}$	$1.6 \cdot 10^{-4}$	$2.1 \cdot 10^{-4}$
M_{eff}	6,000	6,000	6,000	6,000	6,000	4,506	1,912	998

Table 3: Evaluation of measurement–based admission rule with non–uniform weights.

weights the estimates m_X and $\mu_X(s)$ are computed as:

$$m_X = \sum_{i=1}^{M} p_i x_i \quad \text{and } \mu_X(s) = \ln \sum_{i=1}^{M} p_i e^{s x_i}. \tag{8}$$

As before, these estimates are used to compute $P_{\text{loss}}^{\text{est}}$ (4) and the connection requesting establishment is accepted if $P_{\text{loss}}^{\text{est}} \leq \epsilon$ and rejected otherwise. We refer to this call admission rule as *measurement–based admission rule with non–uniform weights.*

Simulation Results

For the evaluation of the measurement–based admission rule with non–uniform weights we use the same simulation approach as in Section 2. The QoS requirement is again set to $\epsilon = 10^{-6}$. We set the length of the measurement window to $M = 6{,}000$. This choice is intuitively appealing as the average lifetime of a stream is 6,000 frame periods. In order to avoid unnecessary computation we ignore measurements that are assigned weights less than 10^{-9}. Thus, for small τ_p the number of measurements considered when estimating the logarithmic moment generating function is less than M; we denote M_{eff} for the number of samples actually used in the estimation. We note that the computational complexity of the LD admission test is $O(M_{\text{eff}})$. We found in our numerical experiments that it takes typically $M_{\text{eff}} \cdot 0.13$ msec to perform one admission test on a SUN ULTRA 10 workstation. Finding s^* (7) is computationally the most expensive step in the LD admission test. We use Newtons method [9] for computing s^* and start Newtons method with the s^* of the last admission test. Provided the call arrivals are not too far apart, s^* changes only little from one admission test to the next and Newtons method finds the new s^* with high accuracy after only a few iterations.

Table 3 gives the average number of admitted streams, the loss probability and M_{eff} for different values of τ_p. We see from the table that the average number of connections increases as τ_p decreases. It is interesting to note that for decreasing τ_p the loss probability first decreases slightly and then increases. However, the loss probability is generally over one order of magnitude larger than the imposed QoS requirement. In Figure 3 we show typical sample path plots from the simulation for $\tau_p = 300$. It is instructive to compare these sample path plots with the plots for the admission rule with uniform weights in Figure 2. When comparing the plots take note of the different scales on the axes; Figure 2 shows the plots over 20,000 time slots while Figure 3 shows the plots over

10,000 time slots and the scales on the y–axes are also different. We observe from the figures that the non–uniform weights have reduced the oscillations significantly. In the next section we try to improve the admission rule by adding the so–called peak rate reservation refinement.

2.2 Peak Rate Reservation Refinement

To motivate the peak rate reservation refinement consider a scenario where a stream, say stream u, is admitted and a few slots later another stream, say stream v, requests establishment. When conducting the admission test for stream v only a few aggregate arrival measurements that include stream–u traffic are available. These few measurements that include stream–u traffic have little impact on the estimated logarithmic moment generating function $\mu_X(s)$. Especially when the measurement window is long and older measurements are assigned relatively large weights, the few samples including stream u have very little influence on $\mu_X(s)$. The new stream u is therefore underrepresented in $\mu_X(s)$ and the aggregate bandwidth demand is underestimated. As a result the estimated loss probability $P_{\text{loss}}^{\text{est}}$ is too small and too many connections are admitted. In summary, the problem with the measurement–based admission rules studied so far is that they "forget" the peak rates of recently admitted streams even though the new stream's traffic is not yet fully reflected in the measurements.

To fix this shortcoming we add a refinement to the measurement–based admission rule with non–uniform weights. This refinement works roughly as follows. We keep a record of peak rates of the recently admitted streams. When conducting an admission test this record is used to compute a reserved peak rate denoted by c^*. The reserved peak rate c^* is computed by assigning weights to the recorded peak rates. Peak rates of relatively new streams are assigned weights close to one, while peak rates of relatively old streams are assigned weights close to zero. Thus, streams that are relatively new are mostly accounted for by the reserved peak rate. On the other hand, stream that have been established for a while are mostly accounted for by the traffic measurements. The reserved peak rate c^* is then subtracted from the link capacity C when computing the estimated loss probability $P_{\text{loss}}^{\text{est}}$.

To make these ideas a little more precise, suppose that in slot k a stream with peak rate c_k^* requests establishment. Let $y_i,\ i = 1, \ldots, M$, denote the peak rates of the admitted streams in slots $k - i,\ i = 1, \ldots, M$. y_i is set to zero if no new stream was admitted in slot $k - i$. Let $q_i,\ i = 1, \ldots, M$, denote weights with $0 \leq q_i \leq 1$. Throughout this paper we use exponentially decaying weights

$$q_i = e^{-i/\tau_q}, \qquad i = 1, \ldots, M,$$

where $\tau_q \geq 0$ is a tuning parameter. The reserved peak rate is computed as

$$c^* = c_k^* + \sum_{i=1}^{M} q_i y_i.$$

Note that the peak rate c_k^* of the stream requesting establishment is assigned a weight of one. We now define the estimated loss probability $P_{\text{loss}}^{\text{est}}$ as the expected fraction of traffic lost by the

		τ_p						
		1200	600	300	120	60	40	30
τ_q	0	137 $3.1 \cdot 10^{-5}$	157 $6.2 \cdot 10^{-5}$	173 $8.1 \cdot 10^{-5}$	191 $1.6 \cdot 10^{-4}$	198 $2.1 \cdot 10^{-4}$	197 $3.2 \cdot 10^{-4}$	198 $3.8 \cdot 10^{-4}$
	50	147 $4.3 \cdot 10^{-6}$	161 $5.0 \cdot 10^{-6}$	174 $8.9 \cdot 10^{-6}$	185 $1.5 \cdot 10^{-5}$	188 $2.7 \cdot 10^{-5}$	190 $3.3 \cdot 10^{-5}$	192 $4.6 \cdot 10^{-5}$
	100	150 $8.0 \cdot 10^{-7}$	165 $9.0 \cdot 10^{-7}$	173 $1.1 \cdot 10^{-6}$	180 $2.9 \cdot 10^{-6}$	183 $4.1 \cdot 10^{-6}$	185 $6.6 \cdot 10^{-6}$	186 $9.0 \cdot 10^{-6}$
	125	154 $4.8 \cdot 10^{-7}$	163 $4.9 \cdot 10^{-7}$	172 $6.4 \cdot 10^{-7}$	178 $1.1 \cdot 10^{-6}$	181 $2.0 \cdot 10^{-6}$	183 $2.9 \cdot 10^{-6}$	184 $4.3 \cdot 10^{-6}$
	200	156 $2.4 \cdot 10^{-9}$	161 $6.9 \cdot 10^{-9}$	166 $3.2 \cdot 10^{-8}$	172 $7.1 \cdot 10^{-8}$	174 $2.3 \cdot 10^{-7}$	175 $1.3 \cdot 10^{-7}$	177 $4.7 \cdot 10^{-7}$

Table 4: Evaluation of measurement–based admission rule with peak rate reservation. Each table entry gives the average number of streams, J_{avg}, and the loss probability, P_{loss}, for a specific combination of the tuning parameters τ_p and τ_q.

established connections at a bufferless link of capacity $C - c^*$, formally:

$$P_{\text{loss}}^{\text{est}} := \frac{E[(X_k - (C - c^*)T)^+]}{E[X_k]}.$$

As before, the estimated loss probability is computed using the LD approximation; the expression for the LD approximation of $P_{\text{loss}}^{\text{est}}$ (6) is modified in the obvious way. The logarithmic moment generating function $\mu_X(s)$ is evaluated using the non–uniform weight refinement (8). Connection k is admitted if $P_{\text{loss}}^{\text{est}} \leq \epsilon$ and rejected otherwise. We refer to this admission rule as the *measurement–based admission rule with peak rate reservation.*

The parameter τ_q is used to tune the peak rate reservation. For $\tau_q = 0$ all the weights are zero and the measurement–based admission rule with peak rate reservation reduces to the admission rule with non–uniform weights. For strictly positive τ_q the weights q_i decay exponentially. The larger τ_q, the larger the peak rate reservation.

Simulation Results

We now evaluate the measurement–based admission rule with peak rate reservation through simulation. We use the same simulation set–up as in Sections 2 and 2.1. The results are reported in Table 4. The table gives the average number of streams, J_{avg}, and the loss probability, P_{loss}, for different combinations of the tuning parameters τ_p and τ_q. Several points are noteworthy here. First, consider the column $\tau_p = 600$. We see that as τ_q increases from zero (i.e., no peak rate reservation) to 100 the average number of streams increases while the loss probability decreases. Loosely speaking the admission rule makes "smarter" admission decisions by reserving more peak

rate; it achieves both higher link utilizations and smaller losses. As τ_q increases further, however, both J_{avg} and P_{loss} drop. Reading along any row of the table we see that for fixed τ_q both J_{avg} and P_{loss} increase with decreasing τ_p.

The goal of this simulation experiment is to find the combination of tuning parameters that gives good on–target performance, i.e., a loss probability nearly equal to ϵ, as well as high link utilizations. We see from the table that the combination $\tau_p = 120$ and $\tau_q = 125$ gives the highest J_{avg} among the combinations with P_{loss} nearly equal to $\epsilon = 10^{-6}$. Unless stated otherwise these tuning parameters are used for all numerical experiments in the remainder of this paper.

Figure 4 shows typical sample path plots from the simulation for $\tau_p = 120$ and $\tau_q = 125$. Comparing these sample path plots with the sample path plots in Figures 2 and 3 we see that the peak rate reservation refinement has eliminated the oscillations that were present in the earlier plots. The plots illustrate the responsiveness of the admission rule with peak rate reservation. $P_{\text{loss}}^{\text{est}}$ remains typically close to the threshold $\epsilon = 10^{-6}$. Moreover, when $P_{\text{loss}}^{\text{est}}$ deviates from ϵ by more than one order of magnitude it very quickly returns to ϵ.

3 Comparison with other Measurement–based Admission Rules

In this section we review the measurement–based admission rules in the existing literature and compare the performance of our Large Deviations based admission rule with that of the admission rules in the literature.

Measured Sum Approach

Jamin *et al.* in their seminal work on measurement–based admission control [10, 11] develop a measurement–based admission rule for Integrated Services networks. They assume a network consisting of buffered multiplexers. Their admission rule consists of a delay criterion and a rate criterion. The delay criterion is designed to keep the delay in the network below a prespecified delay bound. The delay criterion takes the measured delay in the network and a leaky bucket characterization (i.e., average rate r and bucket depth b) of the connection requesting establishment into consideration. The rate criterion strives to keep the link utilizations below prespecified utilization targets. The rate criterion relies on the measured link utilizations and the leaky bucket rate r of the new stream. The new stream is admitted if it passes both the delay criterion and the rate criterion. In order to compare the performance of the admission rule of Jamin *et al.* with that of our Large Deviations based admission rule, we apply the admission rule of Jamin *et al.* to the smoothing/bufferless multiplexing networking architecture [5, 6]. In the smoothing/bufferless multiplexing networking architecture the traffic is delayed by no more than the delay bound in the buffered smoother at the network edge. The bufferless multiplexers inside the network add no further delay. Therefore there is no need to check the delay criterion. This simplifies call admission tremendously since delay measurements, which are difficult to conduct in practice, are no longer required.

We now briefly review the rate criterion; see [10, 11] for more details. First, we review the necessary notation. As before, T denotes the slot length. Note however, that the slots lengths of our Large Deviations based admission rule and the admission rule of Jamin *et al.* are fundamentally different. We base admission decisions on the estimate of the moment generating function of the aggregate arrival stream. To ensure that the estimate of the moment generating function correctly reflects the variability of the aggregate arrival stream we use a slot length short enough to capture individual bursts. Recall that for the numerical work in this paper we use a slot length of $1/F \approx 0.042$ seconds. Jamin *et al.* base admission decisions on the estimate of the average aggregate arrival rate. To obtain a good and stable estimate of the average aggregate arrival rate, they average the aggregate arrivals over a longer slot length and thus avoid capturing individual bursts. This is also reflected in the terminology as they refer to the slot length as *averaging period*. They suggest to set the averaging period to 0.5 seconds. We therefore set the averaging period to $12/F$ for the numerical evaluation of the Jamin *et al.* approach in this paper.

Let $\hat{x}$ denote the estimate of the aggregate arrivals in one averaging period. To estimate $\hat{x}$ Jamin *et al.* employ a time–window measurement mechanism. Let W denote the length of the measurement window in multiples of the averaging period T. The time window measurement mechanism works roughly as follows. The amount of traffic arriving in each averaging period is measured. At the end of the measurement window, that is, after obtaining W measurements, $\hat{x}$ is set to the largest of the W samples. When a new stream is admitted $\hat{x}$ is increased by the arrivals of the new stream (computed from the leaky bucket rate r), i.e., we set $\hat{x} \leftarrow \hat{x} + rT$. Furthermore, when a particular sample value is larger than $\hat{x}$ we do not wait until the end of the measurement window to update $\hat{x}$, but instead set $\hat{x}$ to this larger sample value right away. Finally, let v denote a prespecified utilization target. Jamin *et al.* suggest to set $v = 0.9$. With $v = 0.9$ the admission rule strives to keep the link utilization below 90 %. The rate criterion of Jamin *et al.* is to verify whether

$$\hat{x} + rT \leq vCT. \tag{9}$$

The new stream with leaky bucket rate r is accepted if (9) holds, otherwise it is rejected.

We now compare the performance of the admission rule of Jamin *et al.* with that of our Large Deviations based admission rule. We use the load–loss curve [12] for the performance comparison. The load–loss curve is a plot of the loss probability P_{loss} versus the average number of admitted streams J_{avg}. Both P_{loss} and J_{avg} are obtained through simulation. We employ the simulation approach outlined in Section 2. For all simulations in this section we set the link capacity to C = 45 Mbps. All the traffic streams are *"Silence of the Lambs"* video streams, each with its own independent random phase. We consider two scenarios. Figure 5(a) shows the load–loss curves for the case where the video streams are passed through smoothers (see Figure 1) with a maximum smoothing delay of 10 frame periods before they enter the bufferless multiplexer. Figure 5(b) gives the load–loss curves for the case where the unsmoothed lambs video streams are multiplexed onto the bufferless link. The plots give the load–loss curves of the admission rule of Jamin *et al.* for

different measurement windows W. Specifically, we show the load–loss curves for $W = 100T$, $W = 10T$, $W = 5T$ and $W = T$; recall that the averaging period is set to $T = 0.5$ seconds as suggested by Jamin *et al.* [10, 11]. The curves are obtained by varying the utilization target v. The curve for $W = 100T$, for instance, was obtained by running simulations for $v = 0.9$, 0.925, 0.95, 0.975 and 1.0. Two observations are in order. First, we observe that the actual link utilization differs significantly from the target utilization. For $W = 100T$ and $v = 0.9$ (see Figure 5(a)), for instance, the admission rule admits on average 162.4 unsmoothed lambs streams, which corresponds to an average link utilization of 70 %. The corresponding loss probability is $1.2 \cdot 10^{-6}$. Note that the admission rule of Jamin *et al.* is not designed to take a target loss probability as input. The second noteworthy observation is that for smaller measurement windows W the load–loss curves move towards the lower right corner of the plots. This means that for smaller W the admission control rule performs better; it achieves higher link utilizations and smaller loss probabilities. Jamin and Shenker [12] define the load–loss frontier as the load–loss curve that gives the smallest loss probabilities for the range of link utilizations. We see from the plots that the load–loss frontiers are composed of the load–loss curves for $W = 5T$ and $W = T$. The plots in Figure 5 give also the load–loss points of our Large Deviations based admission rule (LD–MBAC). These points were obtained by setting $\tau_p = 120$ and $\tau_q = 125$ and running simulations for the target loss probabilities $\epsilon = 10^{-6}$ and $\epsilon = 10^{-4}$. As in Section 2 we run the simulations for the LD admission rule until the width of the 90 % confidence interval of the loss probability is less than 20 % of the point estimate. The simulations for the admission rule of Jamin *et al.*, which is computationally less demanding than the LD admission rule, are terminated when the 90 % confidence interval of the loss probability is less than 10 % of the point estimate. The 90 % confidence intervals for the loss probability P_{loss} are plotted in Figure 5. We do not plot the confidence intervals for the average number of connections since these confidence intervals are much tighter. In fact, the widths of the 90 % confidence intervals for J_{avg} are less than 0.1 connection and do not show up on the plots. We observe that the load–loss points of our Large Deviations based admission rule are below the load–loss frontiers of the admission rule of Jamin *et al.* Considering Figure 5(a) we see that for $\epsilon = 10^{-6}$ our admission rule admits on average 178 connections and the loss probability is $1.1 \cdot 10^{-6}$. For the same average link utilization, i.e., for 178 connections, the admission rule of Jamin *et al.* gives a loss probability of roughly $6 \cdot 10^{-6}$. Comparing the Plots 5(a) and 5(b) we observe that the gap in performance is wider when the burstier, unsmoothed video streams are multiplexed. We see from Figure 5(b) that for a given QoS requirement the LD admission rule admits on average 8 unsmoothed lambs video streams more. These numerical results indicate that by measuring individual bursts and capturing the variability of the arrival process in the moment generating function, the LD admission rule utilizes the available link capacity more efficiently.

We conclude the review of the measurement–based admission rule of Jamin *et al.* by briefly discussing an enhancement of this admission rule that is due to Casetti *et al.* [13]. As we have seen in Figure 5 the performance of the admission rule of Jamin *et al.* depends greatly on the tuning of its parameters — the measurement window W and the target utilization v. The appropriate tuning

of these parameters is still an area of ongoing research. Casetti *et al.* propose a feedback control mechanism that automatically tunes the measurement window W. Instead of the target utilization v their adaptive measurement–based admission rule takes a target loss probability as input. The measurement window W is dynamically adjusted based on measurements of the aggregate arrivals and the losses at the multiplexer. Roughly speaking the adjustment works as follows; see [13] for more details. If the aggregate arrivals are above a trigger value, which is computed internally by the feedback algorithm, and the measured loss probability exceeds the target loss probability the measurement window W is increased. This makes the admission rule more conservative and results in the acceptance of fewer new connections. Conversely, if the aggregate arrivals are below the trigger value and the measured loss probability is smaller than the target loss probability the measurement window W is decreased. This results in a less conservative admission rule. Casetti *et al.* demonstrate that their adaptive a measurement–based admission rule works reasonably well for large loss probabilities on the order of 10^{-3} or larger. However, their adaptive admission rule fails for smaller loss probabilities as it is difficult to measure smaller loss probabilities with sufficient statistical significance.

Chernoff Bound Approach

We next review the work on measurement–based admission control by Gibbens *et al.* [14, 15]. Gibbens *et al.* study different variations of Chernoff bound based admission rules. They assume that the multiplexed traffic streams are on–off streams. Recall from Section 2 that $U_j(t),\ j = 1, \ldots, J$, denote the rate of connection–j traffic entering the multiplexer at time t. Let $U_j,\ j = 1, \ldots, J$, denote the associated steady state random variables. Recall that c_j^* denotes the peak rate of connection–j traffic entering the multiplexer, and let r_j denote the average rate of connection j (which is measured). Gibbens *et al.* assume that connection j transmits at rate c_j^* with probability r_j/c_j^* and is silent with probability $1 - r_j/c_j^*$, i.e., $P(U_j = c_j^*) = r_j/c_j^*$ and $P(U_j = 0) = 1 - r_j/c_j^*$. Let $\mu_{U_j}(s)$ denote the logarithmic moment generating function of U_j defined as $\mu_{U_j}(s) := \ln E[e^{sU_j}]$. Clearly,

$$\mu_{U_j}(s) = \ln[1 - \frac{r_j}{c_j^*} + \frac{r_j}{c_j^*} e^{sc_j^*}]. \tag{10}$$

The Chernoff bound [16] gives an upper bound on the probability that the sum of the independent random variables U_j exceeds the link capacity C:

$$P(\sum_{j=1}^{J} U_j > C) \leq e^{s(\sum_{j=1}^{J} \frac{1}{s}\mu_{U_j}(s) - C)}. \tag{11}$$

Notice from the exponent of the Chernoff bound that $\frac{1}{s}\mu_{U_j}(s)$ can be ascribed the meaning of bandwidth; in fact the term $\frac{1}{s}\mu_{U_j}(s)$ is commonly referred to as effective bandwidth of connection j [17]. Gibbens *et al.* view the effective bandwidth as function of the average rate r_j and bound this concave function of r_j by a tangent of the function. Depending on the point at which the tangent of

the effective bandwidth function is constructed different admission rules are derived. Constructing a tangent at $r_j = 0$, for instance, gives the bound

$$\frac{1}{s}\mu_{U_j}(s) \leq \frac{e^{sc_j^*} - 1}{sc_j^*} r_j,$$

which is easily verified by calculating the derivative of μ_{U_j} (10) with respect to r_j at $r_j = 0$. Substituting this bound into the exponent of the Chernoff bound (11) and requiring that the Chernoff bound be less than some miniscule ϵ gives the admission control condition

$$\sum_{j=1}^{J} r_j e^{sc_j^*} \leq C;$$

see Appendix A of [15] for details. Interestingly, the QoS parameter ϵ does not appear in the admission control condition. The condition is tuned with the space parameter s of the moment generating function. The appropriate setting of the tuning parameter s is still the subject of ongoing research. Also, note that this admission control condition requires measurements of the average arrival rates r_j, $j = 1, \ldots, J$, of each individual connection j. In practice, however, it is very difficult to measure per–connection rates accurately and cost efficiently. Gibbens *et al.* therefore propose to assign individual traffic streams based on their peak rate c_j^* to a small number of traffic classes and measure the aggregate arrivals for each class. Gibbens *et al.* employ a simple point sample measurement mechanism. They measure the aggregate arrivals per class over an averaging period T and base admission decisions on the most recent measurement only; older measurements are not considered. (Note that this is equivalent to setting $W = T$ in the measurement window scheme of Jamin *et al.*) Thus, in the simplest case of only one class the admission control condition is

$$(rT + x_1)e^{sc^*T} \leq CT,$$

where r is the leaky bucket rate of the connection requesting establishment, x_1 denotes the measured aggregate arrivals in the most recent averaging period of length T, and c^* is the peak rate of the class. Gibbens *et al.* study this admission rule, which is derived from the tangent of the effective bandwidth function at $r_j = 0$, in the context of a decision theoretic framework for admission control in [14]. In [15] Gibbens and Kelly derive a number of variations of this admission rule by considering tangents at different points of the effective bandwidth function.

Jamin and Shenker [12] compare the performance of their approach to measurement–based admission control with the different variation of the approach of Gibbens *et al.* They find in their extensive simulations that the load–loss frontiers of both approaches coincide. Hence both approaches have the same performance. This surprising result is due to structural similarities of the two approaches, which are studied in detail by Jamin and Shenker [12]. Roughly speaking, both approaches rely on the mean of the measured arrivals (higher moments are not considered). A constant, which represents the new connection, is added to the mean and the new connection is accepted if this sum is less than the link capacity multiplied by a constant.

Our Large Deviations based approach to measurement–based admission control is fundamentally different from the approach of Gibbens *et al.* in that we do not assume on–off traffic. Moreover, we do not bound the effective bandwidth function, which is derived from the logarithmic moment generating function, but use the actual logarithmic moment generating function in our admission rule. The LD admission rule thus takes the mean as well as the higher moments of the measured aggregate arrivals into account.

Hoeffding Bound Approach

Floyd [18] studies measurement–based admission control based on the Hoeffding bound [19, 20]. The Hoeffding bound is a Chernoff–style bound for sums of bounded, independent random variables. Recall that $U_j,\ j = 1, \ldots, J$, are steady state random variables denoting the rate at which connection–j traffic enters the multiplexer. Furthermore, recall that c_j^* and r_j denote the peak rate and average rate of connection j. The Hoeffding bound gives an upper bound on the probability that the sum of the random variables U_j exceeds the sum of the average rates r_j by some positive constant δ:

$$P(\sum_{j=1}^{J} U_j > \sum_{j=1}^{J} r_j + \delta) \leq e^{-2\delta^2 / \sum_{j=1}^{J} c_j^{*2}}.$$

Floyd derives an admission control condition from the Hoeffding bound by requiring that the bound be less than some miniscule QoS parameter ϵ. Clearly,

$$e^{-2\delta^2 / \sum_{j=1}^{J} c_j^{*2}} \leq \epsilon$$

for

$$\delta \leq \sqrt{\frac{\ln(1/\epsilon)}{2} \sum_{j=1}^{J} c_j^{*2}}.$$

Noting furthermore that

$$P(\sum_{j=1}^{J} U_j > C) \leq P(\sum_{j=1}^{J} U_j > \sum_{j=1}^{J} r_j + \delta)$$

for $\sum_{j=1}^{J} r_j + \delta \leq C$, Floyd arrives at the admission control condition

$$\sum_{j=1}^{J} r_j + \sqrt{\frac{\ln(1/\epsilon)}{2} \sum_{j=1}^{J} c_j^{*2}} \leq C. \tag{12}$$

We note that Gibbens and Kelly derive this condition as a special case of their Chernoff bound based admission control conditions. They obtain (12) by bounding the effective bandwidth of an on–off stream $\frac{1}{s}\mu_{U_j}(s)$ (where $\mu_{U_j}(s)$ is given by (10)) by a tangent of slope one at $r_j = 1/s - c_j^*/(e^{sc_j^*} - 1)$; see Appendix A3 of [15] for details. Measurement–based admission control based on the Hoeffding

bound is also studied by Brichet and Simonian [21]. They derive a tighter bound on the effective bandwidth of an on–off stream $\frac{1}{s}\mu_{U_j}(s)$ by considering a series expansion of $\mu_{U_j}(s)$ for small s.

Floyd [18] as well as Brichet and Simonian [21] employ an exponential weighted moving average measurement mechanism. Let $\hat{x}$ denote the estimate of the aggregate arrivals in an averaging period of length T, i.e., $\hat{x}$ denotes the estimate of $T\sum_{j=1}^{J} r_j$. The estimate $\hat{x}$ is updated using the recursion

$$\hat{x} = (1-\omega)\hat{x} + \omega x_i,$$

where ω is the weight used to tune the measurement mechanism and x_i denotes the aggregate arrivals in the just expired averaging period of length T. Now suppose that a new connection $J+1$ with peak rate c^*_{J+1} requests establishment. The new connection is accepted if

$$\hat{x} + T\sqrt{\frac{\ln(1/\epsilon)}{2}\sum_{j=1}^{J} c_j^{*2}} + c^*_{J+1}T \leq CT,$$

and rejected otherwise.

Jamin and Shenker [12] consider the measurement–based admission control approach employing the Hoeffding bound and the exponential weighted moving average measurement mechanism in their simulation studies. They find that the load–loss frontier of the Hoeffding bound approach coincides with the load–loss frontiers of the Jamin *et al.* [11] and Gibbens *et al.* [14, 15] approaches. Hence, these three approaches have the same performance; this is due to their structural similarities [12]. Jamin and Shenker also note that the Hoeffding bound admission rule performs far off target; setting the QoS parameter to $\epsilon = 0.99$, for instance, resulted in an actual loss probability of 10^{-5} in their experiments.

Time–Scale Decomposition Approach

Finally, we provide a brief review of the work on measurement–based admission control by Grossglauser and Tse [22, 23]. Roughly speaking, their approach is to estimate mean and variance of the arrivals from the measurements and estimate the loss probability at the node using the Normal approximation [16]. In [22] they conduct an extensive analysis of this Normal approximation approach for the case when per–flow measurements are available. In [23] they extend the analysis to the more practicable case when only aggregate traffic measurements are available. We focus on this latter case, which is also referred to as the time–scale decomposition approach, in this review. In their analysis Grossglauser and Tse identify a critical time–scale $\tilde{T}_h = T_h/\sqrt{J_{\text{cap}}}$, where T_h denotes the average holding time (lifetime) of a connection and J_{cap} denotes the capacity of the multiplexer (which is the multiplexer rate divided by the average rate of a typical connection). The key idea of the time–scale decomposition approach is to decompose the aggregate arrival process (i.e., the measurements x_i over time) into a low–frequency component and a high–frequency component. The low frequency component is obtained by passing the measurements through a low–pass filter with a cutoff frequency of $1/\tilde{T}_h$, while the high frequency component is obtained through a high–pass filter with a cutoff frequency of $1/\tilde{T}_h$. The low frequency component, which tracks the slow

time–scale fluctuations of the arrival process, is used to estimate the mean of the arrivals, while the high frequency component, which tracks the fast time–scale fluctuations, is used to estimate the variance of the arrivals. Formally, let $g_i,\ i \geq 1$, denote the impulse response of the low–pass filter. Grossglauser and Tse propose to use a geometrically decaying impulse response

$$g_i = \frac{1}{\tilde{T}_h}\left(1 - \frac{1}{\tilde{T}_h}\right)^{i-1}.$$

Suppose that in slot k a new connection with smoother rate (peak rate) c_k^* requests establishment. Recall that $x_i,\ i = 1, \ldots, M_{\text{eff}}$, denote the measured aggregate arrivals in slot $k - i$. Let $J_i,\ i = 1, \ldots, M_{\text{eff}}$, denote the number of ongoing streams in slot $k - i$. Let m_U denote the estimate of the average arrivals per connection in a slot. This estimate is obtained by averaging the measured aggregate arrivals over the ongoing connections and convolving the measurements with the impulse response of the low–pass filter:

$$m_U = \sum_{i=1}^{M_{\text{eff}}} \frac{x_i}{J_i} g_i.$$

Furthermore, let σ_U^2 denote the estimate of the variance of the arrivals per connection in a slot. For details on how this estimate is obtained we refer the reader to [23]. The new connection with peak rate c_k^* is accepted if

$$Q\left(\frac{(C - c_k^*)T - J_1 m_U}{\sqrt{J_1 \sigma_U^2}}\right) \leq \epsilon,$$

where $Q(\cdot)$ denotes the complementary cumulative distribution function of a standard normal random variable, which is given by

$$Q(a) = \frac{1}{\sqrt{2\pi}} \int_a^\infty e^{-y^2/2} dy.$$

Grossglauser and Tse define the loss probability as the long–run fraction of time during which loss occurs, that is, the long–run fraction of slots with aggregate arrivals exceeding CT. We denote this loss measure by $P_{\text{loss}}^{\text{time}}$ to distinguish it from the loss measure P_{loss} (3), which is the long–run fraction of information (bits) lost. Formally, $P_{\text{loss}}^{\text{time}}$ is defined as

$$P_{\text{loss}}^{\text{time}} = E[\lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{N} 1(X_n > CT)].$$

The time–scale decomposition approach is designed to ensure that $P_{\text{loss}}^{\text{time}} \leq \epsilon$. We evaluate the time–scale decomposition approach using the simulation set–up of Section 2. We run simulations for different ϵ and record both loss measures, $P_{\text{loss}}^{\text{time}}$ and P_{loss}. The results are reported in Table 5. Part (a) of the table gives the results for the case where the video streams are passed through smoothers with a maximum smoothing delay of 10 frame periods before being multiplexed. Part (b) gives the results for the case where the unsmoothed video steams are multiplexed. We observe

ϵ	10^{-1}	10^{-2}	10^{-3}	10^{-4}
J_{avg}	211	197	185	180
$P_{\text{loss}}^{\text{time}}$	$1.0 \cdot 10^{-1}$	$1.2 \cdot 10^{-2}$	$1.5 \cdot 10^{-3}$	$1.6 \cdot 10^{-4}$
P_{loss}	$3.8 \cdot 10^{-3}$	$3.2 \cdot 10^{-4}$	$2.8 \cdot 10^{-5}$	$3.6 \cdot 10^{-6}$

(a) maximum smoother delay = 10 frame periods

ϵ	10^{-1}	10^{-2}	10^{-3}	10^{-4}
J_{avg}	205	187	174	163
$P_{\text{loss}}^{\text{time}}$	$1.0 \cdot 10^{-1}$	$1.3 \cdot 10^{-2}$	$1.7 \cdot 10^{-3}$	$2.7 \cdot 10^{-4}$
P_{loss}	$5.2 \cdot 10^{-3}$	$4.8 \cdot 10^{-4}$	$5.4 \cdot 10^{-5}$	$7.1 \cdot 10^{-6}$

(b) no smoothing

Table 5: Evaluation of time–scale decomposition approach of Grossglauser and Tse.

from the table that the time–scale decomposition approach gives good on–target performance. We also observe that P_{loss} is typically almost two orders of magnitude smaller than $P_{\text{loss}}^{\text{time}}$. This is because only the fraction $(X_n - CT)/CT$ of bits is lost in slots with aggregate arrivals exceeding the link capacity CT. The load–loss points (J_{avg}, P_{loss}) (with 90 % confidence intervals for P_{loss}) are plotted in Figure 5. We observe from the plots that the load–loss frontier of the time–scale decomposition approach lies between the load–loss frontiers of the approach of Jamin *et al.* and our LD approach. This indicates that by taking the first two moments of the arrival process into consideration the time–scale decomposition approach can accommodate more connections (on average) than the other reviewed approaches, which take only the first moment into consideration.

We conclude the discussion of the time–scale decomposition approach with two remarks. First, note that the time–scale decomposition approach has no explicit tuning parameters. However, it requires knowledge of the average lifetime of the connections, T_h, and the average rate of a typical connection (to determine the capacity J_{cap}). The time–scale decomposition approach is therefore an attractive admission control approach if these two parameters are fairly well known. If these parameters are unknown or change over time, however, the time–scale decomposition approach faces tuning problems similar to the other approaches.

Secondly, note that the time–scale decomposition approach requires the multiplexer to keep track of the number of ongoing connections in each slot. The multiplexer has to maintain a counter that is incremented when a new connection is accepted and decremented when a connection terminates. The other reviewed approaches as well as the LD approach do not require the multiplexer to maintain this counter (and a history of the counter values). If this information is readily available it can be incorporated in our LD admission rule in a straightforward manner. The logarithmic moment generating function of the aggregate arrivals $\mu_X(s)$ is not estimated using (8), but instead as $\mu_X(s) = J_1 \mu_U(s)$. Here, J_1 denotes the most recent connection count and $\mu_U(s)$ denotes the

estimate of the logarithmic moment generating function of the arrivals per connection, which is obtained as

$$\mu_U(s) = \ln \sum_{i=1}^{M_{\text{eff}}} p_i e^{s x_i / J_i}.$$

We plan to study this version of the LD admission rule in detail in future work.

4 Comparison with traditional Admission Control

In this section we compare measurement–based admission control with traditional admission control that bases admission decisions on *a priori* traffic characterizations.

Adversarial Admission Control

First, we consider an admission rule that takes leaky bucket traffic characterizations as input and assumes that the connections are adversarial to the extent permitted by the leaky bucket characterizations, i.e., transmit worst–case on–off traffic. Suppose that the connection–j traffic at the smoother output is characterized by the traffic constraint function $\mathcal{E}_j(t)$, that is, the amount of traffic leaving smoother j over an interval of length t is less than $\mathcal{E}_j(t)$. Specifically, suppose that $\mathcal{E}_j(t) = \min(c_j^* t,\ \sigma_j + \rho_j t)$, that is, the output of smoother j is constrained by the smoother rate (peak rate) c_j^* and a single leaky bucket $(\sigma_j,\ \rho_j)$, where σ_j is the maximum burst size and ρ_j bounds the long–term average rate of connection j. The adversarial admission rule assumes that each connection transmits worst–case on–off traffic [2, 5], that is, connection j transmits at rate c_j^* with probability ρ_j/c_j^* and is silent with probability $1 - \rho_j/c_j^*$. Recall that $U_j,\ j = 1, \ldots, J$, are steady state random variables denoting the rate at which connection–j traffic enters the multiplexer. The adversarial assumption is that $P(U_j = c_j^*) = \rho_j/c_j^*$ and $P(U_j = 0) = 1 - \rho_j/c_j^*$. The logarithmic moment generating function of U_j, defined as $\mu_{U_j}(s) := \ln E[e^{sU_j}]$, is clearly

$$\mu_{U_j}(s) = \ln[1 - \frac{\rho_j}{c_j^*} + \frac{\rho_j}{c_j^*} e^{s c_j^*}].$$

Let X denote the sum of the random variables $U_j,\ j = 1, \ldots, J$, i.e., $X = \sum_{j=1}^{J} U_j$. Furthermore, let $\mu_X(s)$ and m_X denote the logarithmic moment generating function and mean of X. Assuming that the connections generate traffic independently, i.e., $U_j,\ j = 1, \ldots, J$, are mutually independent, we obtain $\mu_X(s) = \sum_{j=1}^{J} \mu_{U_j}(s)$. [1] The expected fraction of traffic lost at the multiplexer is

$$P_{\text{loss}} = \frac{E[(X - C)^+]}{E[X]}. \tag{13}$$

The Large Deviation approximation of (13) is given by [1]

$$P_{\text{loss}} \approx \frac{1}{m_X s^{\star 2} \sqrt{2\pi \mu_X''(s^\star)}} e^{-s^\star C + \mu_X(s^\star)}, \tag{14}$$

[1] Note that we have redefined $\mu_X(s)$ in this section. In Section 2 $\mu_X(s)$ is defined as the estimate of the logarithmic moment generating function of X_k, the *amount* of traffic (in bit) arriving in slot k. In this section $\mu_X(s)$ is defined as the logarithmic moment generating function of X, the *rate* of arriving traffic (in bit/sec).

where s^* satisfies

$$\mu'_X(s^*) = C.$$

A set of J connections is permissible on a link with capacity C if the loss probability P_{loss} (14) is less than some miniscule QoS parameter ϵ.

For the numerical work in this section we use again the *"Silence of the Lambs"* (lambs) trace. We obtain the traffic constraint function of the lambs trace $\mathcal{E}_{\text{lambs}}(t)$ by following the procedure described in [5]. We give here only a brief outline of this procedure and refer the reader to [5] for details. The first step is to compute the empirical envelope, which is the tightest traffic constraint function of the video sequence. The empirical envelope is however not necessarily concave. It is therefore bounded by a piecewise linear function so that the traffic constraint function can be represented by a cascade of leaky buckets. Given this leaky bucket characterization and a maximum delay in the smoother, which we set again to 10 frame periods (= 10/24 seconds), we obtain the smoother rate $c^*_{\text{lambs}} = 731.6$ kbit/sec by applying the theory developed in [5]. The lambs traffic at the smoother output is characterized by this smoother rate (peak rate) and the leaky bucket $(\sigma_{\text{lambs}},\ \rho_{\text{lambs}}) = (3.16$ Mbyte, 193.2 kbit/sec). For the following numerical evaluation we set the multiplexer rate to $C = 45$ Mbps and assume that all traffic streams are independent lambs video streams. We vary the number of connections J and compute the loss probability for each J using the LD approximation (14). The results are plotted as the solid line (labeled "adversarial") in Figure 6. We defer the discussion of this result to the end of this section.

Histogram–based Admission Control

We next consider an admission rule that is specifically designed for prerecorded sources [24]. This admission rule bases admission decisions on the marginal distribution of the sources' traffic. For video traffic the histogram of the frame sizes is used to compute the logarithmic moment generating function of the video stream. Recall from Section 2 that $f_n(j),\ n = 1, \ldots, N$, denotes the frame size of the nth frame of video j in bits. Also recall from Section 2 that the transmission of a frame is spread out over one frame period of length $1/F$, i.e., the frame $f_n(j)$ is transmitted at rate $f_n(j) \cdot F$ over the interval $[(n-1)/F,\ n/F]$. Let $u_n(j),\ n = 1, \ldots, N$, denote the smoothed frame size trace obtained by simulating the transmission of the original frame size trace $f_n(j),\ n = 1, \ldots, N$, through the buffered smoother with peak rate c^*_j (see Figure 1). The logarithmic moment generating function of the smoothed video stream j is calculated directly from the smoothed frame size trace:

$$\mu_{U_j}(s) = \ln \frac{1}{N} \sum_{n=1}^{N} e^{sFu_n(j)}. \tag{15}$$

A set of J connections is admitted if the loss probability (14) is less than some prespecified QoS parameter ϵ. For the numerical evaluation we assume that all multiplexed traffic streams are smoothed lambs video streams. The dashed line (labeled "histogram") in Figure 6 is the load–loss curve of this admission rule, which is obtained by computing the loss probability (14) for a range of link utilizations.

We also verify the accuracy of the LD approximation through simulation. For this purpose we use the simulation program used to evaluate the measurement–based admission rules in Sections 2 and 3. Instead of employing any of the studied measurement–based admission rules, we fix a maximum number of admissible streams, $J_{\max}$. A connection requesting establishment is accepted if there are currently less than $J_{\max}$ connections in progress, and rejected otherwise. As before, all of the streams are generated from the lambs trace. Each stream has its own independent random phase, which is uniformly distributed over $[1,\ N]$. The lifetime of the streams is fixed at $N =$ 40,000, and the lambs trace is wrapped around to generate the streams. We set the connection inter arrival time to zero, thus there are always $J_{\max}$ connections in progress. We run the simulation until the width of the 90 % confidence interval of the loss probability is less than 10 % of the point estimate. The results are plotted in Figure 6 (labeled "fixed adm. region (simul.)").

The figure also shows the load–loss points of our LD approach to measurement–based admission control (LD–MBAC). These points are obtained by running simulations for $\epsilon = 10^{-4}$ and $\epsilon = 10^{-6}$. The parameters of the admission rule are set to $\tau_p = 120$ and $\tau_q = 1,250$ for these simulations. The time between call arrivals is exponentially distributed with a mean of 67 frame periods. Furthermore, the lifetime of each stream is fixed at $N = 40,000$ frame periods to ensure a fair comparison with the other admission rules. (The traffic constraint function obtained through the procedure described in [5] gives the tightest leaky bucket characterization of the full–length video segment; for shorter segments, however, the characterization may be loose.)

Several points are noteworthy about Figure 6. First, note that the $J_{\max}$–simulation ("fixed adm. region (simul.)") verifies the accuracy of the histogram admission rule. We observe that the histogram admission rule is a little too conservative, but generally very accurate. Secondly, we observe that the adversarial admission rule, which assumes worst–case on–off traffic, results in low link utilizations. Note that by following the procedure described in [5] we have obtained the tightest leaky bucket characterization of the prerecorded lambs trace. The difference in link utilization (horizontal distance) between the "adversarial" and "histogram" curves in Figure 6 therefore gives an indication of the conservatism of the assumption of adversarial on–off traffic. With a QoS parameter of $\epsilon = 10^{-6}$, for instance, the adversarial admission rule admits 147 lambs video streams while the histogram rule admits 169 streams and the measurement–based admission rule admits on average 174.5 streams. In the case of live video transmission where one has to resort to loose leaky bucket characterizations the link utilization with adversarial admission control is even lower, while measurement–based admission control still achieves high link utilizations.

The third noteworthy observation is that measurement–based admission control outperforms histogram–based admission control, which has perfect knowledge of the marginal distribution of the streams' traffic. This can be intuitively explained as follows. The histogram–based admission rule bases admission decisions on the connections' logarithmic moment generating functions $\mu_{U_j}(s)$ (15), which characterize the connections' traffic over their entire lifetime. A new connection is accepted if the long–run fraction of traffic lost due to excursions of the aggregate arrival process X above the link capacity CT is less than ϵ. Most of the time, however, the aggregate arrival

process X is below the threshold CT and the slack capacity $CT - X$ is wasted. Measurement–based admission control bases admission decisions on measurements of the aggregate arrival process X. It admits new connections when slack capacity is available. Conversely, measurement–based admission control stops the acceptance of new connections when the aggregate arrivals are close to the link capacity or even exceed the link capacity. It does not accept any new streams until departing connections have created slack capacity. Measurement–based admission control thus utilizes the link capacity efficiently by taking advantage of the connection arrival and departure dynamics. Note however, that measurement–based admission control is bound to fail when the connection arrival and departure times collude, that is, when the connections arrive roughly at the same time and have identical lifetimes. In the worst–case scenario when all connections arrive in the same time slot the LD–MBAC rule bases admission decisions on the connections' peak rate specification, i.e., it admits 61 smoothed lambs streams ($= C/c^*_{\text{lambs}}$). Traditional admission control, on the other hand, achieves the link utilizations shown in Figure 6 irrespective of the connection arrival and departure dynamics.

5 Conclusion

In this paper we have studied measurement–based admission control for unbuffered multiplexers. We have developed a Large Deviations approach to measurement–based admission control. We have provided an extensive review of the existing literature on measurement–based admission control. We found in our simulations with MPEG–1 encoded videos that the LD admission rule compares favorably with the admission rules in the existing literature. Finally, we compared measurement–based admission control with traditional admission control, which relies on *a priori* traffic characterizations. Our numerical work indicates that measurement–based admission control can achieve significantly higher link utilizations.

In our current research we are addressing the parameter tuning problem. We are investigating the use of feedback control to automatically tune the parameters of he LD admission rule. Our numerical work suggests that the LD admission rule performs on–target and achieves high link utilizations when the estimated loss probability $P^{\text{est}}_{\text{loss}}$ oscillates with small amplitudes around the target loss probability ϵ (see for instance Figure 4(b)). We are therefore studying feedback control policies that tune the parameters of the LD admission rule based on $P^{\text{est}}_{\text{loss}}$; in particular integral control policies that strive to minimize the area under the $P^{\text{est}}_{\text{loss}}$ curve. This approach avoids the difficult problem of measuring the actual, miniscule losses at the multiplexer.

Another avenue for future research is to study the combination of measurement–based admission control with traditional admission control. The goal is to develop a hybrid admission rule that strives to achieve the high link utilizations of measurement–based admission control and at the same time guards against gross mistakes by the measurement–based admission rule due to incorrect parameter settings of colluding connection arrivals/departures.

Acknowledgment: I am grateful to Matthias Grossglauser for clarifying the filtering mechanisms used in the time–scale decomposition approach [23].

References

[1] J. Roberts, U. Mocci, and J. Virtamo (Eds.). *Broadband Network Traffic: Performance Evaluation and Design of Broadband Multiservice Networks, Final Report of Action COST 242, (Lecture Notes in Computer Science, Vol. 1155).* Springer Verlag, 1996.

[2] A. Elwalid, D. Mitra, and R. H. Wentworth. A new approach for allocating buffers and bandwidth to heterogeneous regulated traffic in an ATM node. *IEEE Journal on Selected Areas in Communications*, 13(6):1115–1127, August 1995.

[3] F. LoPresti, Z. Zhang, D. Towsley, and J. Kurose. Source time scale and optimal buffer/bandwidth trade–off for regulated traffic in an ATM node. In *Proceedings of IEEE Infocom*, Kobe, Japan, April 1997.

[4] S. Rajagopal, M. Reisslein, and K. W. Ross. Packet multiplexers with adversarial regulated traffic. In *Proceedings of IEEE Infocom '98*, pages 347–355, San Francisco, CA, April 1998.

[5] M. Reisslein, K. W. Ross, and S. Rajagopal. Guaranteeing statistical QoS to regulated traffic: The single node case. In *Proceedings of IEEE Infocom '99*, New York, NY, March 1999. available at `http://www.fokus.gmd.de/usr/reisslein` or `http://www.eurecom.fr/~ross`.

[6] M. Reisslein, K. W. Ross, and S. Rajagopal. Guaranteeing statistical QoS to regulated traffic: The multiple node case. In *Proceedings of 37th IEEE Conference on Decision and Control (CDC)*, pages 531–538, Tampa, FL, December 1998. available at `http://www.fokus.gmd.de/usr/reisslein` or `http://www.eurecom.fr/~ross/`.

[7] O. Rose. Statistical properties of MPEG video traffic and their impact on traffic modelling in ATM systems. Technical Report 101, University of Wuerzburg, Insitute of Computer Science, Am Hubland, 97074 Wuerzburg, Germany, February 1995.
ftp address and directory of the used video traces:
ftp-info3.informatik.uni-wuerzburg.de /pub/MPEG/.

[8] R. Jain. *The Art of Computer Systems Performance Analysis.* Wiley, 1991.

[9] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C, The Art of Scientific Computing.* Cambridge University Press, Cambridge, MA, second edition, 1992.

[10] S. Jamin, P. B. Danzig, S. J. Shenker, and L. Zhang. A measurement–based admission control algorithm for integrated services packet switched networks. In *Proceedings of ACM SIGCOMM '95*, pages 2–13, 1995.

[11] S. Jamin, P. B. Danzig, S. J. Shenker, and L. Zhang. A measurement–based admission control algorithm for integrated services packet switched networks (extended version). *IEEE/ACM Transactions on Networking*, 5(1):56–70, February 1997.

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

[13] C. Casetti, J. Kurose, and D. Towsley. A new algorithm for measurement–based admission control in integrated services packet networks. In *Proceedings of the Protocols for High Speed Networks Workshop*, October 1996. Available at `ftp://gaia.cs.umass.edu/pub/Case96:MB-Adm.ps.gz`.

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

[15] R. J. Gibbens and F. P. Kelly. Measurement–based admission control. In *Proceedings of 15th International Teletraffic Congress (ITC15)*, pages 879–888, Washington, D.C., June 1997.

[16] P. Billingsley. *Probability and Measure*. Wiley, third edition, 1995.

[17] F. P. Kelly. Notes on effective bandwidths. In F. P. Kelly, S. Zachary, and I. B. Ziedins, editors, *Stochastic Networks: Theory and Applications, Royal Statistical Society Lectures Note Series 4*, pages 141–168. Oxford University Press, 1996.

[18] S. Floyd. Comments on measurement–based admission control for controlled–load services. *submitted*, July 1996. available at `http://www-nrg.ee.lbl.gov:80/floyd`.

[19] W. Hoeffding. Probability inequalities for sums of bounded random variables. *American Statistical Association Journal*, 58:13–30, March 1993.

[20] M. Hofri. *Probabilistic Analysis of Algorithms*. Springer Verlag, 1987.

[21] F. Brichet and A. Simonian. Conservative gaussian models applied to measurement–based admission control. In *Proceedings of IWQoS*, Napa, CA, May 1998. extended version available as Technical report of Project COST 257 at `ftp://www-info3.informatik.uni-wuerzburg.de/pub/cost/cost257/september97/257/td97056.ps`.

[22] M. Grossglauser and D. Tse. A framework for robust measurement–based admission control. In *Proceedings of ACM SIGCOMM*, pages 237–248, Cannes, France, September 1997.

[23] M. Grossglauser and D. Tse. A time–scale decomposition approach to measurement–based admission control. In *Proceedings of IEEE Infocom '99*, New York, NY, March 1999.

[24] M. Reisslein and K. W. Ross. Call admission for prerecorded sources with packet loss. *IEEE Journal on Selected Areas in Communications*, 15(6):1167–1180, August 1997.

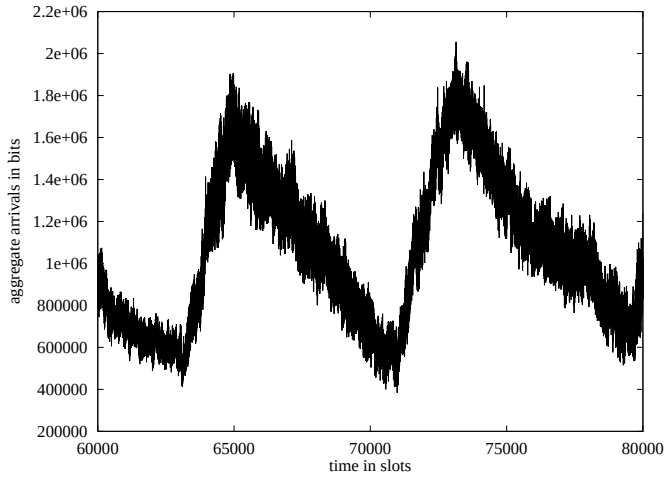

(a) aggregate arrivals versus time

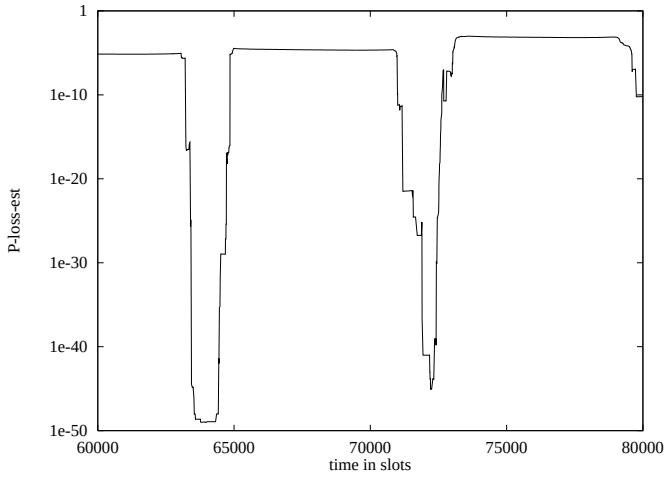

(b) $P_{\text{loss}}^{\text{est}}$ versus time

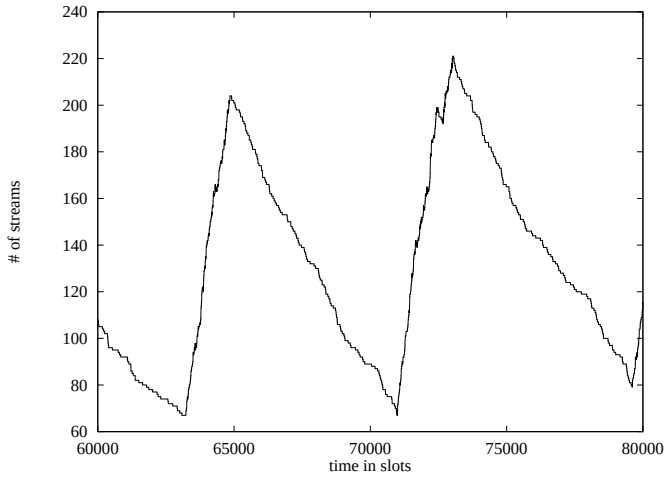

(c) J versus time

Figure 2: Sample path plots from simulation of basic measurement–based admission rule with M = 6,000.

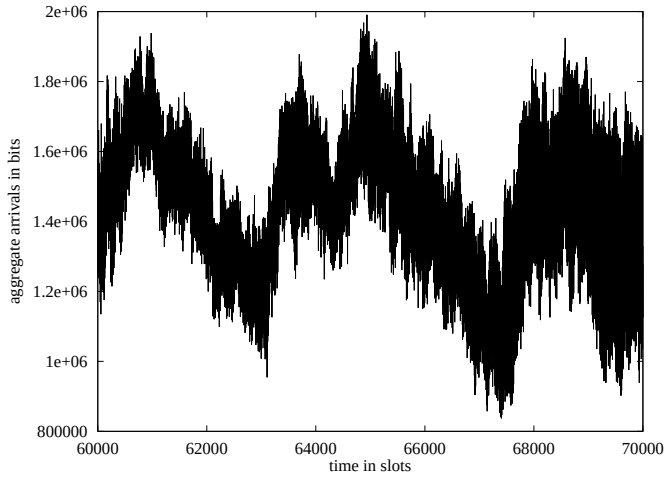

(a) aggregate arrivals versus time

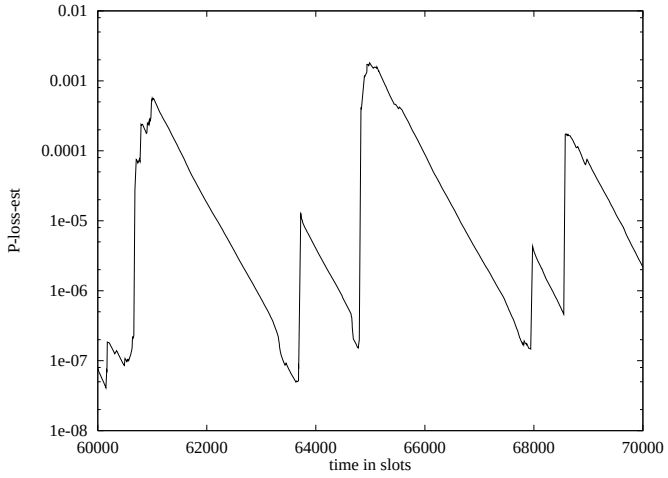

(b) $P_{\text{loss}}^{\text{est}}$ versus time

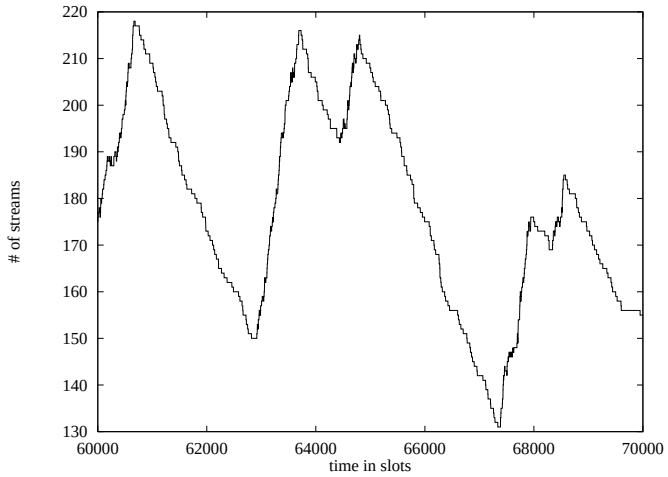

(c) J versus time

Figure 3: Sample path plots from simulation of measurement–based admission rule with non–uniform weights for $\tau_p = 300$.

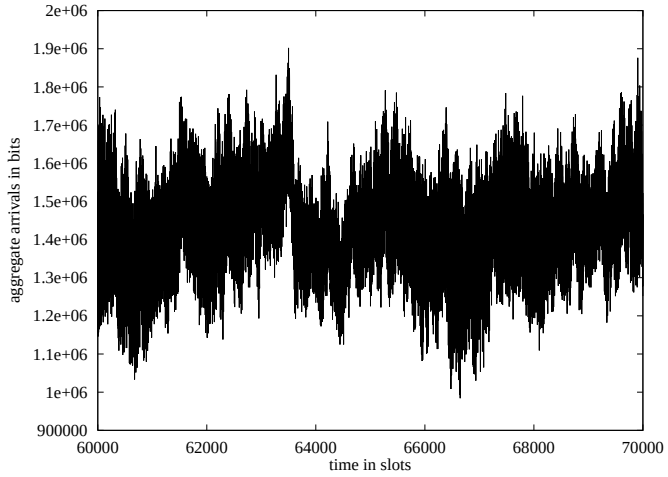

(a) aggregate arrivals versus time

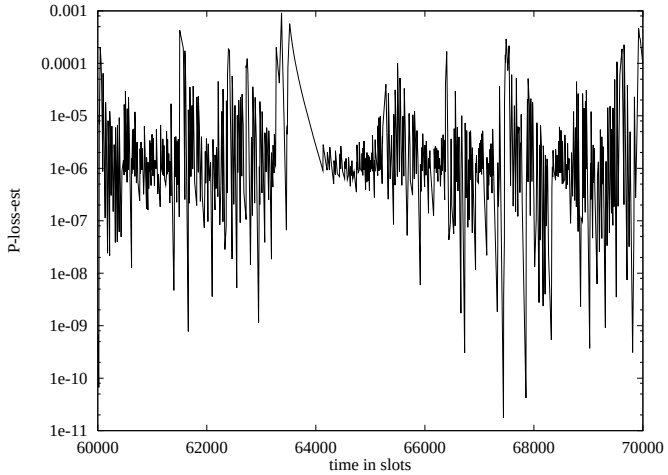

(b) $P_{\text{loss}}^{\text{est}}$ versus time

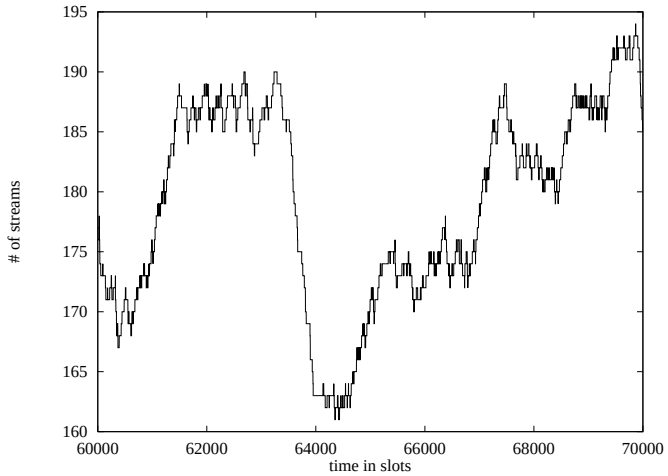

(c) J versus time

Figure 4: Sample path plots from simulation of measurement–based admission rule with peak rate reservation for $\tau_p = 120$ and $\tau_q = 125$.

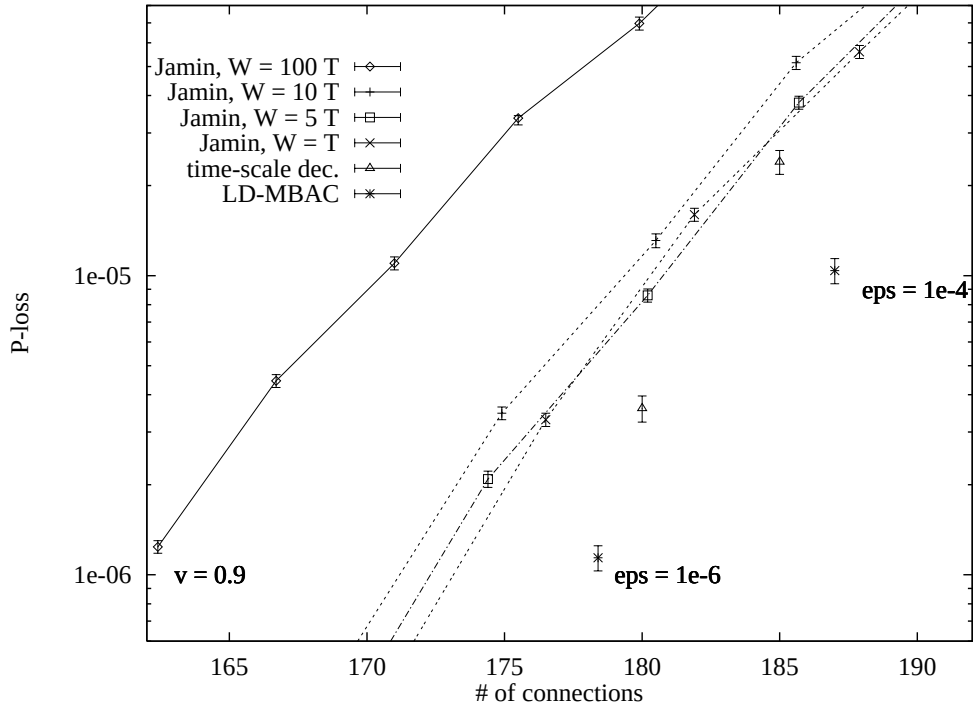

(a) maximum smoother delay = 10 frame periods

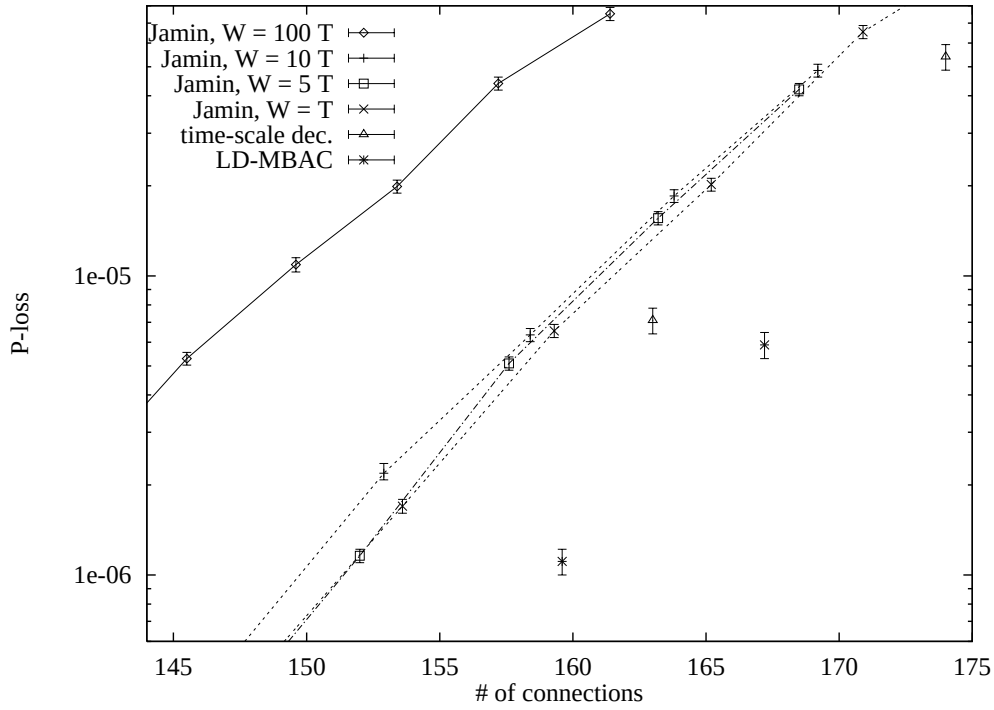

(b) no smoothing

Figure 5: Comparison of admission rules in the literature with our LD based admission rule.

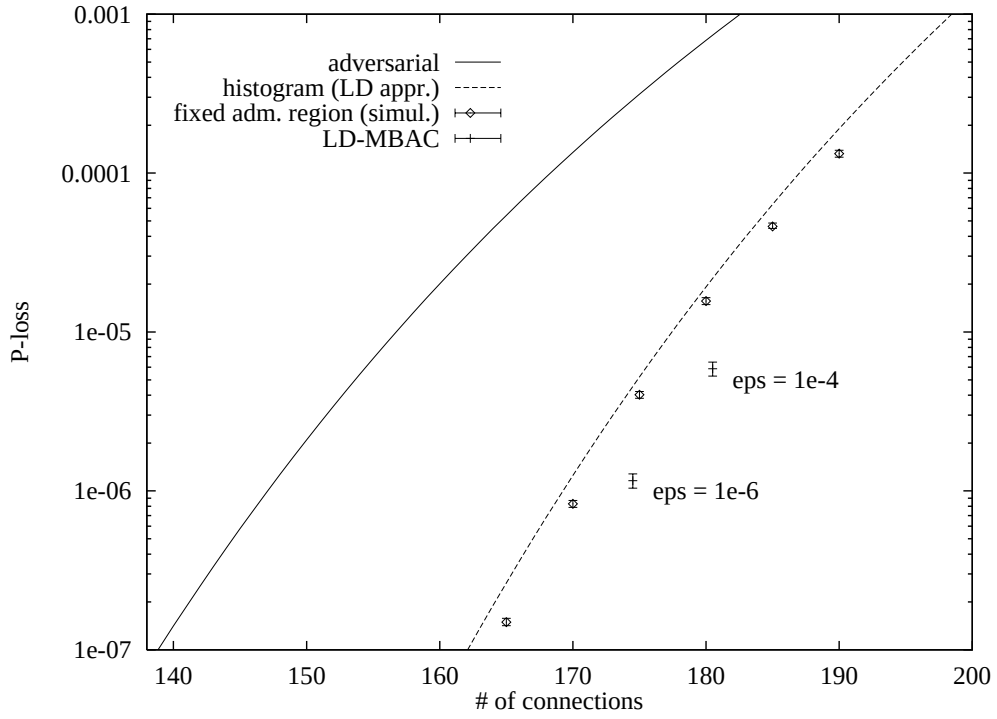

Figure 6: Comparison of measurement–based admission control with traditional admission control.