

Dynamic Resizing of Utilization Target in Measurement-Based Admission Control

Chih-Chiang Chuang¹

Yeali S. Sun²

Ying-Dar Lin¹

Department of Computer and Information Science¹
National Chiao-Tung University
Hsinchu, Taiwan
ydlin@cis.nctu.edu.tw
TEL: +886-3-5731899
FAX: +886-3-5721490

Department of Information Management²
National Taiwan University
Taipei, Taiwan
sunny@im.ntu.edu.tw
TEL:886-02-23630231 ext.2870
FAX:886-02-23621327

Abstract

A major issue in the design of the MBAC algorithms relies on how to maximize link utilization while meeting applications' QoS constraints. Studies have claimed that many proposed MBAC algorithms would give the same performance if the parameters were properly tuned. In this paper, we study the effects of tuning two of the most important parameters in measurement-based admission control algorithms: *Measurement Window* and *Utilization Target*. We show that the measurement window is not a good tuning knob for enforcing rigid loss probability constraint. We propose a new parameter called Soft Utilization Target combined with a simple flow admittance policy called the *Latch* algorithm. The simulation results show that this scheme can achieve better enforcement of loss probability constraint while maintaining higher link utilization.

Keyword: admission control, measurement-based, Integrated Services, controlled-load services

1. Introduction

The approach of Measurement-Based Admission Control (MBAC) [1] is preferred in integrated services networks [2], [3], [4], [5] in order to achieve better link utilization. In measurement-based approach, a priori source characterization is used only for incoming flows. Measurements will be used to characterize aggregated behavior of flows that have been in place. The goal is to exploit the potential multiplexing gain. This is based on the observations that the flow description given at the session setup time or resource reservation phase is usually a loose

bound of the traffic that the flow actually generates. The aggregated link utilization would be lower than the expectation evaluated simply by flow specifications. MBAC is proposed to overcome the problem where admission decisions and control are based on on-line measurement.

Several issues are raised in the use of measured-based admission control. First is the feasibility of implementing real-time on-line measurement methods for high-speed networks in approximating traffic loads. The second issue is the possible violation of QoS constraints such as the utilization target [1] and loss probability [6], [7], caused by the measurement errors. The third issue again is to maximize link utilization under the QoS constraints.

Jamin et al., in [1], presented an algorithm in which three performance-related measurement parameters are defined: Measurement Interval size S , Measurement Window T , and Utilization Target v . They showed that how these parameters are determined may have significant impact on the performance. In [8], they contended that the MBAC algorithms developed so far could deliver the same performance if the parameters were properly tuned. Instead of developing new algorithms, properly tuning the measurement parameters of the algorithm in [1] would be a good choice. However, they did not provide solutions on how to tune these parameters. Furthermore, the algorithm in [1] can satisfy the utilization target constraint but with occasional violation of the loss probability constraint. It is appropriate for predictive services. However, it would not be suitable for applications with rigid QoS requirements.

In this paper, we adopt the evaluation method for

aggregated traffic by the Normal distribution approximation [6] to examine how can we achieve the loss probability constraint by resizing the Measurement Window and Utilization Target in [1]. Here we define a new parameter called *Soft Utilization Target*, denoted by v' , which is dynamically adjusted to enforce packet loss constraint in contrast to the fixed utilization target used in [1]. We will show that Soft Utilization Target is a finer performance knob than the Measurement Window. One can easily tune the value of the Soft Utilization Target to deliver a certain loss probability constraint, while the appropriate Measurement Window may not exist at all. The method for tuning Soft Utilization Target depends on the traffic which behavior is difficult to obtain using on-line measurements. We propose a *Latch Algorithm* to address the issue. In the algorithm, no new flows are admitted if the system admission controller forecasts a possible failure on loss probability constraint may occur until any flow departs.

The method proposed in the paper can also be used to support Controlled-Load Services [4] to satisfy both utilization target and loss probability constraints. Although Controlled-Load Services do not have explicit QoS constraints, the system manager can thus specify the QoS constraints such as utilization target and loss probability.

The rest of our paper is organized as follows. In Section 2, we first review the MBAC parameters in [1] and introduce the basic model we use. Then we show the weakness of Measurement Window in enforcing loss probability constraint and present our solution idea, which resizes Soft Utilization Target. In Section 3, we derive the method to resize Soft Utilization Target for loss probability constraint and introduce the Latch Algorithm. We also arranged the pseudo codes of the admission control in this section. The simulation results are given in Section 4, we will discuss the result of loss probability constraint and performance of the admission control. Section 5 concludes our paper.

2. Measurement Window and Loss Probability

Constraint

2.1 Measurements in a MBAC

In [1], the admission control algorithm works as follows. Measurement Window T specifies the time span of each measurement; Measurement Interval of size S specifies the time interval between two measurements. First, the MBAC computes the average load Y of a measurement interval. For a measurement window containing k measurement intervals, the *highest* average load of the k intervals is taken as the estimated load of the next measurement window. For example, in Figure 2.1, the size of the measurement window is three. There are three measurement intervals - $S11$, $S12$ and $S13$ - in measurement window $T1$. At the end of the period, the highest average load $Y12$ is obtained and used as the estimated load of the link in the admission control decisions during the next measurement period. Note that in this method, the load estimation is not dynamically updated during the period. When there are sudden load changes due to the arrival of traffic bursts, without immediate update on the network state, the admission decisions may cause violation of QoS constraints.

A new flow requesting the rate r is admitted if

$$Z < v \cdot C - r \quad \text{EQ(1)}$$

where Z is the load estimate obtained from the previous measurement window, v is the link utilization target, and C is the link capacity. Here we notice two important relationships between Z , v and the performance of measurement-based admission control. First, because the load estimate used is based on the measurement of the last period, the longer the duration of the measurement window, the less responsiveness of the admission control, thus the higher the possibility of violating the QoS constraints. Second, with a smaller value of the utilization target, the admission control decisions would be more conservative in admitting new flows. Here, we propose to dynamically resize measurement window or utilization target to adjust

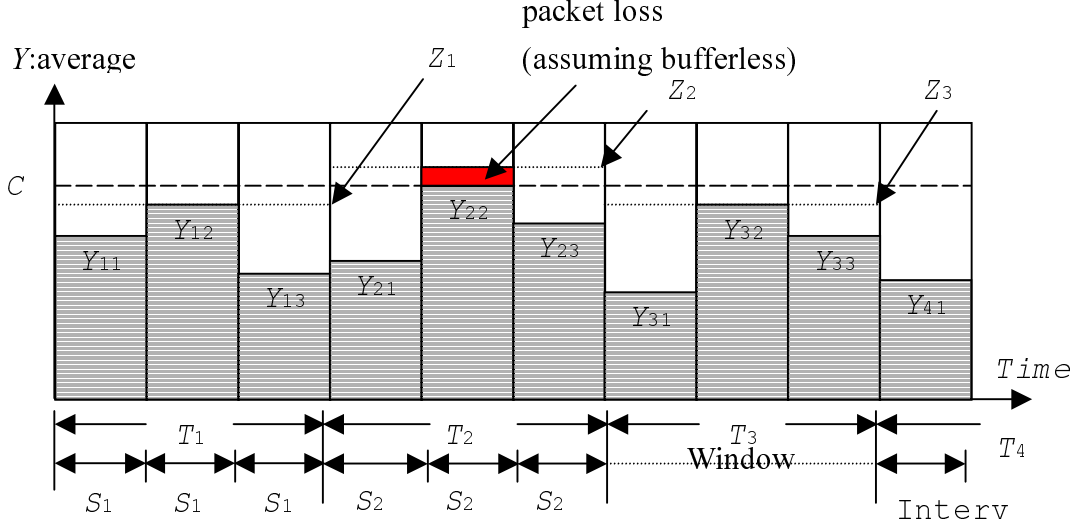


Figure 2.1 Measurements in a Measurement-Based Admission Control Algorithm

the resulting loss probability.

In this paper, we consider a high-speed link with a large number of flows passing through it. It is therefore assumed that the averaged loads of the measurement intervals are in steady state and independently and identically distributed with a normal distribution. The probability density function of the standard normal random variable, x , is

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right),$$

and the standard normal cumulative distribution function is

$$N(y) = \int_{-\infty}^y \phi(x) dx.$$

For the sake of simplicity, we assume links have no buffers.

An arriving packet is dropped if it finds the outgoing link is busy. The dark shadow area shown in Figure 2.1 corresponds to the amount of packets dropped during T_2 . Let the loss probability constraint be denoted by P_{qos} . The goal of the system is to ensure

$$P(\tilde{Y} > C) \leq P_{qos}$$

Let the mean and standard deviation of \tilde{Y} be denoted by $\mu_{\tilde{Y}}$ and $\sigma_{\tilde{Y}}$, respectively. We have

$$1 - N\left(\frac{C - \mu_{\tilde{Y}}}{\sigma_{\tilde{Y}}}\right) \leq P_{qos}. \quad \text{EQ(2)}$$

We will use this equation to evaluate the performance of the tuning knobs: Measurement Window and Utilization Target.

2.2 Measurement Window

To evaluate how the size of the measurement window affects the accuracy of the load estimation, one would need to find the relationship between the expected value of \tilde{Z} and the Measurement Window size k , i.e. the function F that renders $E[\tilde{Z}] = F(k)$. Since the load estimate of a measurement window is set to the maximum average load of the k measurement intervals, we

have $\tilde{Z} = \max\{\tilde{Y}_1, \tilde{Y}_2, \tilde{Y}_3, \dots, \tilde{Y}_k\}$, where

$\tilde{Y}_1, \tilde{Y}_2, \tilde{Y}_3, \dots, \tilde{Y}_k$ are independently identically distributed

as \tilde{Y} . Given $\frac{\tilde{Y} - \mu_{\tilde{Y}}}{\sigma_{\tilde{Y}}} \sim \text{Normal}(0,1)$, we can obtain the

p.d.f. of \tilde{Z} as

$$f_{\tilde{Z}}(z) = kN\left(\frac{z - \mu_{\tilde{Y}}}{\sigma_{\tilde{Y}}}\right)^{k-1} \phi\left(\frac{z - \mu_{\tilde{Y}}}{\sigma_{\tilde{Y}}}\right) \cdot \frac{1}{\sigma_{\tilde{Y}}},$$

by the order statistics [9]. Now, we can evaluate the expected value of \tilde{Z} as

$$E[\tilde{Z}] = \int_{-\infty}^{\infty} z f_{\tilde{Z}}(z) dz = \int_{-\infty}^{\infty} z \cdot kN\left(\frac{z - \mu_{\tilde{Y}}}{\sigma_{\tilde{Y}}}\right)^{k-1} \phi\left(\frac{z - \mu_{\tilde{Y}}}{\sigma_{\tilde{Y}}}\right) \cdot \frac{1}{\sigma_{\tilde{Y}}} dz$$

Let x be equal to $\frac{z - \mu_{\tilde{Y}}}{\sigma_{\tilde{Y}}}$ and rewrite the equation, we have

$$\begin{aligned} E[\tilde{Z}] &= \int_{-\infty}^{\infty} (\sigma_{\tilde{Y}} \cdot x + \mu_{\tilde{Y}}) k N(x)^{k-1} \phi(x) dx \\ &= \sigma_{\tilde{Y}} \cdot f(k) + \mu_{\tilde{Y}} \end{aligned} \quad \text{EQ(3)}$$

where

$$f(k) = \int_{-\infty}^{\infty} x \cdot k \cdot N(x)^{k-1} \phi(x) dx. \quad \text{EQ(4)}$$

From EQ(3), we can see that the expectation of the largest average load during a measurement window period is the mean traffic load plus the standard deviation times a factor which is a function of the measurement window size. Given k , EQ(4) can be evaluated using numerical approximation. Figure 2.2 shows a sketch of $f(k)$. One can see that $f(k)$ grows fast when k is small. After around 25, it starts to saturate. It indicates that indefinitely increasing the duration of the measurement window (observing a longer period of time) only has limited effect on the accuracy of the load estimation.

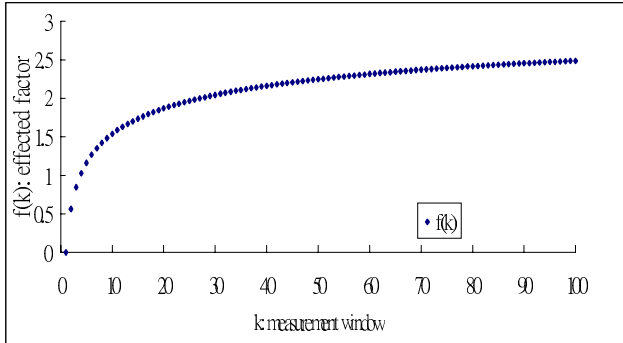


Figure 2.2 A coefficient function of load variance

Next, we will show how to apply this result on run-time measurement. Let μ and σ denote the mean and standard deviation of the measured load \tilde{Y} , respectively. We are interested in the heavy load behavior, namely when the link utilization approaches the utilization target v , i.e.

$$E[\tilde{Z}] \approx v \cdot C \quad \text{EQ(5)}$$

From EQ(4) and EQ(5), we have

$$\sigma \cdot f(k) + \mu \approx v \cdot C. \quad \text{EQ(6)}$$

Let the utilization target v be set to its maximum value 1

and assume the equality of EQ(6) holds. Then we have

$$f(k) = \frac{C - \mu}{\sigma}. \quad \text{With EQ(2), we have } 1 - N(f(k)) \leq P_{qos}.$$

Solving for $f(k)$, as $N(\cdot)$ increases monotonically, we have $f(k) \geq N^{-1}(1 - P_{qos})$. We take

$$f(k) = N^{-1}(1 - P_{qos}), \quad \text{EQ(7)}$$

or

$$k = f^{-1}[N^{-1}(1 - P_{qos})], \quad \text{EQ(8)}$$

if f^{-1} exists.

Table 2.1 presented some numerical results to illustrate the relationship between the size of the measurement window and the loss probability constraint. Note that very large window sizes will be needed to achieve finer degrees of loss probability constraint which may require long periods of measurement. Unfortunately, we couldn't find the inverse function of f . That means given the loss probability constraint, we cannot obtain the appropriate size of the measurement window.

Loss probability(P_{qos})	10^{-1}	10^{-2}	10^{-3}	10^{-4}
$N^{-1}(1 - P_{qos})$	1.28	2.32	3.15	3.96
Window size (k)	7	61	>>500	>>500

Table 2.1 The relationship between measurement window size and loss probability constraint

The measurement window-based scheme is conceptually simple. But it may not be feasible to support links with loss probability constraint greater than 10^{-3} due to the large sampling overhead incurred in the run-time measurement. If the sampling interval is fixed, larger measurement window size means longer observation period. In this approach, the scheme may not be able to quickly respond to load changes, thus compromising the link utilization and its serviceability (A new flow request might be rejected due to the inaccurate load estimation.). On the other hand, if one fixes the time duration of the measurement window, one would need to increase the sampling frequency. In this case, the measurement overhead can become a concern to the system performance.

3. Resizing Soft Utilization Target and Latch Algorithm for Loss Probability Constraint

3.1 Soft Utilization Target

In the previous section, we have analyzed the relationship between different loss probability constraints and the sizes of the measurement window. In this section, we look at another important performance knob in measurement-based admission control schemes, i.e. the utilization target. We will study how to set the utilization target to meet the loss probability constraint. First, a new parameter called *Soft Utilization Target*, denoted by v' , is used to better enforce the loss probability constraint. It is different from the utilization target whose value is typically fixed and set by system administrator according to some bandwidth sharing policy between the QoS-constrained traffic and best-effort traffic.

Since the goal is to ensure the packet loss probability at all times is upper bounded by the loss probability constraint as given in EQ(2), by rewriting it, we have $\mu_r \leq [1 - \frac{\sigma_r}{C} N^{-1} (1 - P_{qos})] \cdot C$. Let the soft utilization target be set as follows:

$$v' = 1 - \frac{\sigma}{C} N^{-1} (1 - P_{qos}), \quad \text{EQ(9)}$$

Note that σ is the standard deviation of the measurements $\{\tilde{Y}\}$. If we replace v as v' in EQ(1), the admission control decision inequality equation becomes as follows:

$$E[Y] < v' \cdot C - r \quad \text{EQ(10)}$$

The σ in EQ(9) only describes traffic variation of the existing flows. To compute the variance of the flow aggregated with the newly-requested flow, we assume the existing flows and new one are independent and the variance of the aggregated traffic is the sum of σ^2 and σ_r^2 if the flow is admitted, where σ_r^2 is the variance of the newly-requested flow. For links with hundreds and thousands of flows, one can assume the value of σ^2 will be much larger than that of σ_r^2 . Thus, we will ignore the variation effect introduced by the new call.

3.2 The Latch Algorithm

One of the major consideration in admission control is not to let the admittance of a new flow affect the promised QoS to existing flows. In measurement-based approach, it relies on accurate estimation of the mean and standard deviation of the existing traffic loads. However, due to measurement uncertainty, especially in heavy load conditions, how to correct or handle estimation error becomes one of the major issues in measurement-based admission control. Assuming the estimated parameters are the real parameters could be dangerous. For example, an underestimation of flow parameters may deteriorate the QoS guarantee to admitted flows. Overestimation of the number of permissible flows may as well result in inefficient use of network resources and unnecessary blocking of call requests.

To address the issue, we propose a method called the Latch algorithm to dynamically adjust the admission control policy. First, if there is a violation of the packet loss constraint during a measurement interval, the system will stop admitting new flows until an existing flow departs. This is to prevent further violation of the packet loss constraint. Moreover, we consider that an occurrence of loss constraint violation may imply further occurrence of the event. It certainly shows that the total traffic load was greater than the expected. Therefore, even though the current load estimate is low, one would like to temporarily close flow admittance to reduce the uncertainties introduced by admitting more flows into the system. Generally, such violation is owing to the fluctuations or burstiness of the aggregated traffic. In practice, it is still very difficult to quantize such effect at the flow admission time based on the flow priori traffic descriptions. By doing so, it will certainly help to compensate the estimation error.

Second, if the system rejects a flow, the flow admission will be temporarily closed until a flow departs. The goal is to compensate possible underestimation of the aggregated traffic load.

The two policies proposed are simple and effective in enforcing loss probability constraints. The simulation results shown in the next section will show the performance of maximizing link utilization without loss probability violation by employing this algorithm.

The pseudo codes of the proposed admission control policies are shown in Figure 3.1.

```
extern float Accumulate; // Accumulated packet size
void Packet_Measurements(float Packet_size)
{
    Accumulate+=Packet_size;
}
```

Figure 3.1 (a) Packet measurement procedure

```
extern float Accumulate; // Accumulated packet size
extern float Traffic_Data[]; // Traffic data storage
extern int Traffic_Data_Index;
extern int Latch; // Latch flag
void On_line_Measurements()
{
    float Bufferless_Loss=
        (Accumulate-LINK_BANDWIDTH)/Accumulate;
    if (Bufferless_Loss>P_QoS)
        Latch=1; // Latch policy 1
    Traffic_Data_Index=
        (Traffic_Data_Index+1)%ROTATE_STORAGE;
    Traffic_Data[Traffic_Data_Index]=Accumulate;
    Accumulate=0; // Reset the accumulation
}
```

Figure 3.1(b) On-Line Measurements procedure

```
extern float Traffic_Data[]; // Traffic data storage
extern int Traffic_Data_Index;
extern int Latch; // Latch flag
int New_Flow_Arrival(float Rate)
{
    float Mean, Std_Dev, Soft_UT;
    if (Latch)
        return(REJECT); // Reject if being latched
    Std_Dev=
        Moving_Std_Dev(); // Get the stdev
    // Calculate the Soft Utilization Target
    Soft_UT=
        1-(Std_Dev/LINK_CAPACITY)*Un_Normal(1-P_QoS);
    // Make the admission decision
    if (Traffic_Data[Traffic_Data_Index]
        <Soft_UT*LINK_CAPACITY-Rate)
        return(ACCEPT); // Accept a new flow
    else {
        Latch=1; return(REJECT);
    } // Latch policy 2
}
```

Figure 3.1(c) New Flow Arrival procedure

```
extern int Latch;
void Flow_Departure()
{
    Latch=0;
}
```

Figure 3.1(d) Flow Departure procedure

In Figure 3.1(a), it is the code for packet measurement triggered by the event of packet arrival. In Figure 3.1(b), it shows for the on-line traffic load

estimation. This function is triggered by clock alarm, once per Measurement Interval. It also detects the violation on the loss probability constraint. Figure 3.1(c) is to handle new flow arrivals. Figure 3.1(d) shows the function for handling flow departures.

In this function, when a new flow request is received, the admission control unit will first check if the Latch flag is off. If yes, it computes the current estimate of the standard deviation of the traffic load to obtain a new value of Soft Utilization Target. Then the admission control process is executed based on the new utilization target. If the flow is rejected, the latch flag is turned on and the flow admittance is closed. The function of “Flow Departure” is invoked by a flow departure event. It will turn off the Latch flag and resume the admission control process.

4. Simulation Results

In this section, we study the performance of the proposed method to dynamically resize the parameter of the target utilization in measurement-based admission control on enforcing loss probability constraints. In the simulation, links with two different bandwidths are considered: 1.544Mbps (T1) and 10Mbps (Ethernet). Three types of traffic sources are studied. The first one is the *ON-OFF* traffic source. In the *ON-OFF* model, a flow alternates between the ON and OFF state; the duration of both states are assumed to be exponentially distributed with the mean of 2 seconds. The flow sends 8Kbytes/sec constant bit rate data in the ON state and is idle when in the OFF state. The second model is the Poisson process with the mean of 8Kbytes/sec. Third, we study the effect under heterogeneous traffic. Here we have traffic mixed of the above two types of flows with sending rate varying from 8Kbyte/sec to 16Kbyte/sec. The Measurement Interval is set to one second. We use *Moving-8* method for the on-line calculation of the mean and standard deviation of the traffic load. Namely, the memory of the on-line measurement is eight measurement intervals. The mean μ and standard

deviation σ of \tilde{Y} are computed as follows: $\mu = \frac{\sum_{n=1}^8 Y_n}{8}$

and $\sigma = \left(\frac{\sum_{n=1}^8 (Y_n - \mu)^2}{8-1} \right)^{\frac{1}{2}}$. The loss rate, $\frac{E[(\tilde{Y} - C)^+]}{E[\tilde{Y}]}$ is

evaluated every 8000 measurement intervals, where $(\cdot)^+$ evaluates the positive part of the result. It is assumed that the transmission link is buffer-less.

A. Enforcing Loss Constraints

Figure 4.1 and 4.2 show the resulting packet loss probability vs. the target loss constraint for measurement-based admission control with or without using the Latch algorithm for links with 1.544Mbps and 10Mbps capacity, respectively. A reference line is shown to indicate whether the loss probability constraint is satisfied or not. The data points that lie below the reference line correspond to those experiments where loss constraints are not satisfied.

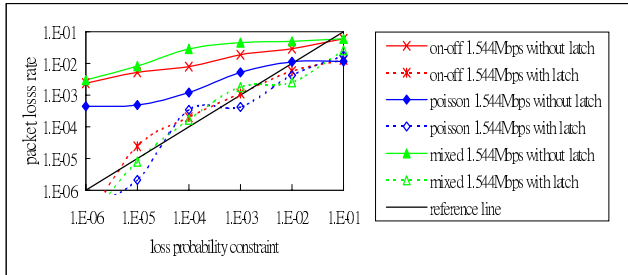


Figure 4.1 Loss probability constraint and simulated loss rate in the 1.544Mbps link bandwidth

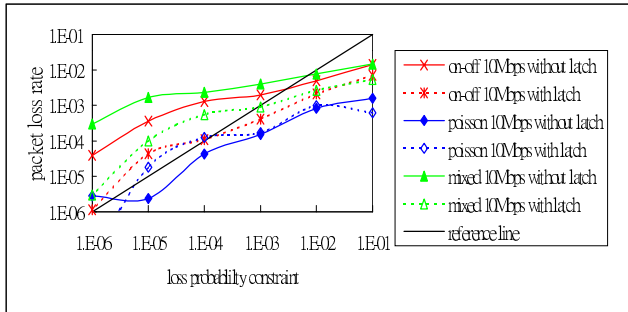


Figure 4.2 Loss probability constraint and simulated loss rate in the 10Mbps link bandwidth

From the figures, one can see that for the transmissions

under the same type of traffic sources, those using admission control with the Latch algorithm achieve better performance in enforcing packet loss constraint than the ones without employing the Latch algorithm. This is mainly due to the blocking of new flows when the link load approaches the target utilization. Next, we compare the performance of the transmissions using the proposed method under different types of traffic sources. From the simulation results, one can see because that the Poisson traffic sources generate the least degree of traffic burstiness, the chances of having loss constraint violation are lowered, thus having the best performance, especially for cases with more rigid loss constraint. When we compare the performances for the same transmissions but under different link capacities, it is hard to tell.

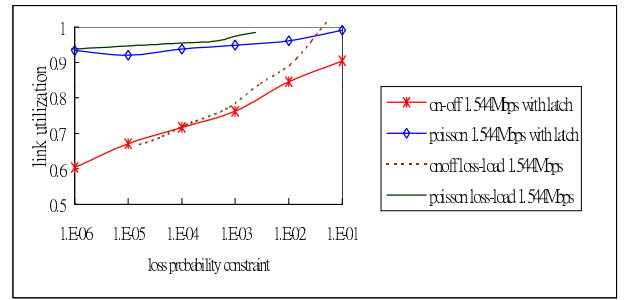


Figure 4.3 Performance of the admission control in the 1.544Mbps link bandwidth

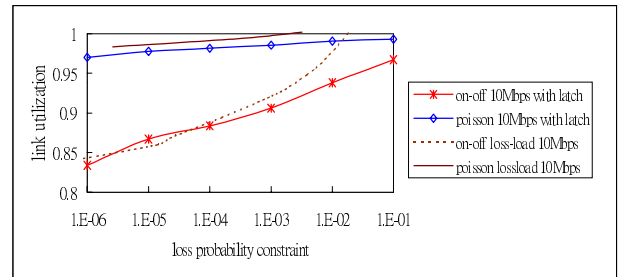


Figure 4.4 Performance of the admission control in the 10Mbps link bandwidth

B. Link Utilization

Next, we would examine the link utilization performance using the proposed method. In Figure 4.3 and Figure 4.4, we compare the link utilization for transmissions under different link capacities. In addition,

the *Loss-Load Curves* [10] are shown in each figure for comparison. The Loss-Load Curve gives the maximum utilization obtained to satisfy the loss probability constraint. One can see that the performance of admission control with the Latch algorithm is very close to the Loss-Load curves.

In figure 4.4, the distance between the Poisson performance curve and the corresponding Loss-Load Curve is about 100Kbps. Given that each Poisson flow is of 64Kbits/sec average rate, the difference is about a flow. But one can achieve better loss constraint when using the Latch algorithm. In the case of the ON-OFF sources, the proposed method appears to be conservative in admitting new flows. It performs better link utilization in the range of tighter loss constraints.

5. Conclusions

It has been shown that tuning certain performance knobs such as the measurement window size is important to the performance of the measurement-based admission control. In this paper, we first analyze how the size of the measurement window affects the accuracy of the load estimation assuming the load samples follow normal distribution. We show that by increasing the number measurement points of a measurement window period, namely a longer period of observation has only limited effect on the accuracy of the load estimation. Next, we studied the tuning performance of another performance knob – utilization target in the guarantee of the loss probability. We propose a new parameter called *Soft Utilization Target* whose value is computed based on the mean of the load measures. Instead of using a fixed target utilization, the soft utilization target is used when the link load is heavy. The Latch algorithm is proposed to throttle flow admission when the link load is heavy and the aggregated traffic is bursty. Generally, the occurrence of the violation of loss probability is owing to the fluctuations or burstiness of the aggregated traffic. It makes the load estimation and admission control even more difficult when

the load is heavy and resources are scarce. The two policies proposed are simple and effective in enforcing loss probability constraints. The simulation results show that for transmissions using admission control with the Latch algorithm achieve better performance in enforcing packet loss constraint. They also achieve the link utilization very close to the maximum utilization obtained to satisfy the loss probability constraint.

References

- [1] S. Jamin, P. B. Danzig, S. J. Shenker, and L. Zhang, "A Measurement-based Admission Control Algorithm for Integrated Service Packet Networks," *IEEE/ACM Transactions on Networking*, vol.5, issue1, Pages 56-70, Feb. 1997.
- [2] R. Braden, Ed., L. Zhang, S. Berson, S. Herzog, S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification," RFC2205, September 1997.
- [3] S. Shenker, C. Partridge and R. Guerin, "Specification of Guaranteed Quality of Service," RFC2212, September 1997.
- [4] J. Wroclawski, "Specification of the Controlled-Load Network Element Service," RFC2211, September 1997.
- [5] J. Wroclawski, "The Use of RSVP with IETF Integrated Services," RFC2210, September 1997.
- [6] M. Grossglauser and D. Tse, "A Framework for Robust Measurement-Based Admission Control", In *Proc. ACM SIGCOMM '97*, pages 237-248, 1997.
- [7] D. Tse and M. Grossglauser, "Measurement-based Call Admission Control: Analysis and Simulation", In *Proc. IEEE INFOCOM '97*, Vol.3, pages 981-989, Kobe, Japan, April 1997.
- [8] S. Jamin and S. Shenker, "Measurement-based Admission Control Algorithms for Controlled-load Service: A Structural Examination", Univ. of Michigan, CSE-TR-333-97, Apr. 1997. Available at <http://irl.eecs.umich.edu/jamin/papers/mbac/clmbac.ps.gz>.
- [9] S. Ross, "A First Course in Probability," Prentice Hall, 4th Ed., Pages 279-282, 1994.
- [10] C. L. Williamson and D. R. Cheriton, "Loss-Load Curves: Support for Rate-Based Congestion Control in High-Speed Datagram Networks," *SIGCOMM '91*, pages 17-28, 1991.