ELSEVIER

Computer Communications 22 (1999) 1296–1306

computer communications

A congestion control mechanism for enterprise network traffic over asynchronous transfer mode networks

M. Hassan[a,*], H. Sirisena[b,1], M. Atiquzzaman[c,2]

[a]School of Computer Science and Software Engineering, Monash University, Melbourne 3145, Australia
[b]Department of Electrical and Electronic Engineering, University of Canterbury, Christchurch 8020, New Zealand
[c]Department of Electrical and Computer Engineering, University of Dayton, Dayton, OH 45469-0226, USA

Abstract

Available bit rate (ABR) is an emerging asynchronous transfer mode (ATM)-based telecommunication service, which dynamically allocates bandwidth to the users according to the available bandwidth in the network. However, research shows that ABR-based connections suffer from congestion at LAN–ATM gateway when ATM network abruptly reduces the bandwidth available to the gateway. Gateway congestion may result in high packet loss. This paper proposes a gateway congestion control mechanism, which controls the rate of outgoing enterprise traffic according to the available bandwidth at the gateway. Our analysis shows how to achieve a stable system and keep the packet loss below a desired threshold. Analytical results are validated by simulation. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Enterprise networks; Intranet; Traffic management; Asynchronous transfer mode networks

1. Introduction

Intranet is a rapidly growing internetworking technology for connecting the hosts within an enterprise network using the *standard* Internet protocols such as transmission control protocol/internet protocol (TCP/IP). Such intranets have enormous benefits over proprietary networking as it provides ready access to the global Internet which connects millions of hosts all over the world. World-wide corporations are rapidly deploying local intranets for their organizations, and they are also connecting them together in a single enterprise network.

The enterprise network of an organization is connected to the global Internet via an access-gateway and access-link. Traffic that flows between the enterprise network and the rest of the Internet travels through this gateway and access-link. At the access-gateway, traffic from many hosts within the organization is statistically multiplexed onto the access-link. If the traffic rate momentarily exceeds the access-link bandwidth, some packets will be lost. In order to prevent or reduce such a loss, the gateway has some buffer to temporarily accommodate the excess traffic. These data are transmitted later, once the traffic burst is over.

The access-link is usually established over a leased line from a telecommunication service provider. The bandwidth of the leased line is constant and the amount of bandwidth is provisioned according to the traffic profile of the organization, to ensure minimum congestion in the gateway.

For bursty data traffic, the bandwidth of the leased line is not utilized efficiently. The organization has to pay for the bandwidth for the entire duration whether there is traffic or not. The cost of maintaining leased lines is, therefore, very high.

With asynchronous transfer mode (ATM) [1], a high-speed cell-switching network, set to become the world-wide standard [2] for the next generation of telecommunication networks, it is expected that such leased lines will be replaced by ATM virtual circuits (VCs). Interconnection of an enterprise network to the global Internet via such an ATM VC is shown in Fig. 1.

Unlike the constant bandwidth leased lines, ATM networks provide an available bit rate (ABR) service for data communications, where the bandwidth of the VC is controlled by the ATM network, and varies between a specified minimum and maximum negotiated during the VC set-up process. If there are enough resources available in the ATM network, the gateway may be allowed to transmit at the maximum rate. However, if congestion occurs within the ATM network, the gateway may be asked to reduce the transmission rate.

---

* Corresponding author. Tel.: + 61-3-9903-2122; fax: + 61-3-9903-1077.

E-mail addresses: mahbub@insect.sd.monash.edu.au (M. Hassan), h.sirisena@elec.canterbury.ac.nz (H. Sirisena), atiq@engr.udayton.edu (M. Atiquzzaman)

[1] Tel.: + 64-3-366-7001; fax: + 64-3-364-2761.

[2] Tel.: + 1-937-229-3183; fax: + 1-937-229-4529.

PII: S0140-3664(99)00127-9

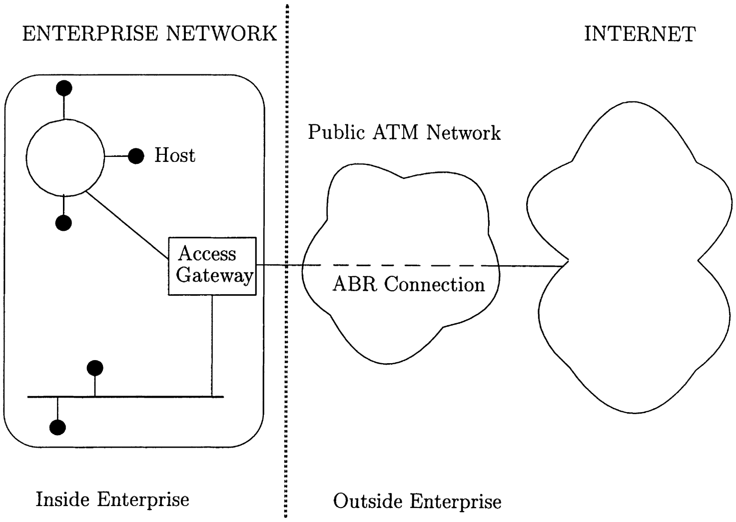

Fig. 1. Interconnection of an enterprise network to the Internet via ATM ABR service.

As an ABR connection allows a gateway to utilize only the available bandwidth within the ATM network, it is expected that the cost of data communications will be significantly reduced with the use of the ABR service. Such cost reduction may become a major driving force in the near future for replacing existing leased lines with ATM ABR connections for connecting enterprise networks to the global Internet.

One major problem with the ABR service is the potential congestion that might be experienced in the access-gateway due to occasional throttling of the access-link bandwidth by the ATM network. The ABR feedback control simply shifts the congestion within the ATM network to the edge of the ATM network (in this case the access-gateway), and hence increased buffering is required in the access-gateway [3–7] to cope with this congestion. If there is not enough buffer in the gateway, some data will be lost which may significantly degrade the performance of enterprise applications. While implementing large buffers in the gateway may prevent data loss, it will increase the delay in the gateway. Such undesirable delay may cause retransmissions from the TCP hosts, which would exacerbate the congestion in the gateway. It is, therefore, desirable to implement a traffic management mechanism in the intranet to control the transmission rates of the intranet hosts according to the current bandwidth of the access-link and keep the buffering requirement in the gateway to a minimum. Reducing the buffering requirement also helps reduce the production cost of access-gateways. The design objectives of such a traffic management mechanism would be to—(i) guarantee the stability of the traffic control system, (ii) minimize the buffer level variance (or fluctuation) at the access-gateway due to fluctuation in the ABR access-link bandwidth, and (iii) minimize the packet loss at the access-gateway due to buffer overflow.

There is no suitable traffic management mechanism available in the standard Internet protocol suite, which can effectively control the transmission rates of intranet hosts connected to the Internet using an ABR access-link. It was shown in Ref. [3] that the implicit flow-control mechanism embedded in the Internet protocol suite is not capable of controlling the traffic rates of the intranet hosts satisfactorily; the system remains unstable susceptible to large buffer level variances, and hence has the potential for high packet loss for small buffers. Due to the lack of stability and large buffer level variance, quite large buffers are required in the gateway to minimize the packet loss and improve the performance of intranet applications.

In this paper, the design of a traffic management mechanism for enterprise networks is presented, taking the varying bandwidth of the access-link into consideration. Implementation of this mechanism in the intranet is expected to significantly reduce the buffering requirement in the access-gateway.

We have formulated the traffic management problem as a classical feedback control system; the current buffer level and the bandwidth of the ABR link in the gateway are periodically fed back to a “controller” which calculates the allowed intranet traffic rate to the gateway to keep the buffer level at a desired set point. Modeling the traffic management mechanism as a feedback control system allows us to study the stability and the buffer level variance using the tools of control theory. The optimization of the control parameters for the stability of the control system and minimum buffer level variance at the gateway is obtained through analysis. An estimate for packet loss probability for finite buffers in the gateway is also presented using the buffer level variance obtained from the analysis of the control system.

The remainder of the paper is organized as follows. A description of the proposed traffic management mechanism is provided in Section 2. The design of the control system and the analyses for optimization of the control parameters to guarantee stability and minimize buffer level variance are presented in Section 3. An estimate for the packet loss

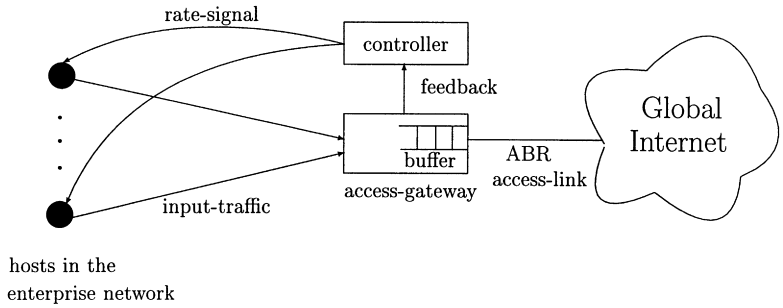

Fig. 2. Schematic diagram of the proposed traffic management mechanism.

probability for finite buffers in the gateway is derived in Section 4. The simulation results are presented in Section 5. Finally, a conclusion is provided in Section 6.

2. Proposed traffic management mechanism

The proposed traffic management mechanism uses feedback rate signals from the access-gateway to the intranet hosts to control the traffic rates of the hosts to keep the buffer level in the gateway at a desired set point q_r (see Fig. 2). The explicit rates allowed to individual hosts (via the rate signals) are controlled by a *controller* in the gateway. The controller calculates the optimum aggregate rate R for the entire intranet according to the current buffer level q and the access-link bandwidth B in the gateway which are periodically fed back to the controller. The aggregate rate R is then distributed to individual hosts (via the rate signals) according to a desired rate allocation policy. A schematic diagram of the proposed traffic management mechanism is shown in Fig. 2.

At the heart of the proposed traffic management mechanism lies the controller in the gateway. If the controller does not calculate the aggregate intranet traffic rate R correctly, the control system may become unstable and the buffer level q in the gateway may far exceed (large variance) the desired set point q_r causing buffer overflow and high packet loss for small buffers. The performance of the traffic management mechanism, therefore, depends significantly on the performance of the controller in the gateway.

In this paper, the controller is modeled as a classical feedback control system with periodic feedback of the current buffer level q to the controller. The design and the performance analysis of the controller is presented in Section 3.

3. Controller design and analyses

The controller and the buffer in the gateway form a feedback control system, where the current buffer level q is periodically fed back to the controller for effective calculation of the optimum intranet traffic rate R. By analyzing the control system, it is possible to study the stability, buffer level variance and packet loss probability of the proposed traffic management mechanism. In this section, the design of the control system is presented and the selection of the control parameters to guarantee stability and minimize buffer level variance is obtained through analyses (a concise version of these analyses is presented in Ref. [8]). Estimation of packet loss probability at the gateway is treated in Section 4.

3.1. Design

In this section, the design of the feedback control system is presented.

3.1.1. Buffer

This is a key component of the control system. The goal of the control system is to maintain the output of a Buffer, the buffer level q, at a desired set point q_r. The input to the Buffer is the aggregate traffic rate R from all the hosts in the local intranet sending traffic to the Internet through the access-gateway. An appropriate R is calculated by the Controller.

In addition to the controlled input R, there is another input to the Buffer—the transmission rate B of the access-link. This is considered a disturbance to the system, and is beyond the control of the Controller. Any fluctuation in B will cause q to deviate from q_r. This effect of disturbance on the output variable q is fed back to the controller, which adjusts R to bring q back to q_r.

The buffer level dynamics can be expressed as an approximation of the following *fluid flow* model:

$$m' = w_{\text{in}} - w_{\text{out}},$$

where m' is the rate of change of mass, w_{in} is the rate of flow in, and w_{out} is the rate of flow out. Therefore, the aggregate rate R will affect the buffer level according to the following equation:

$$\frac{q(n+1) - q(n)}{T} = R(n) - B(n),$$

$$q(n+1) = q(n) + TR(n) - TB(n) \tag{1}$$

where $q(n+1)$ and $q(n)$ are the buffer levels measured at

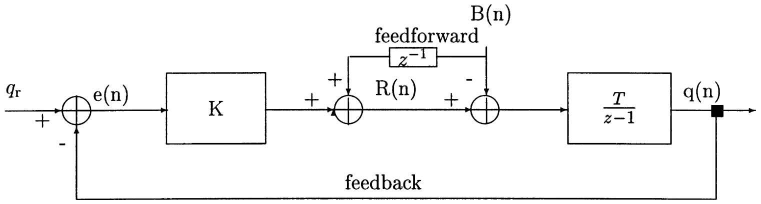

Fig. 3. Block diagram of a proportional controller with disturbance feed forward.

the $(n+1)$th and nth sampling instants, respectively, and T is the sampling period.

Taking the z-transform of Eq. (1), the following transfer function between Buffer output (occupancy level) q and Buffer input R is obtained as:

$$q(z) = \frac{T}{z-1}(R(z) - B(z)). \tag{2}$$

3.1.2. Controller

The main function of this component is to calculate the required input rate R to the access-link as a function of the current buffer level q. R is the output of the Controller and at the same time the input to the Buffer. The input to the Controller is the *error*, the difference between the desired buffer level q_r, and the current observed queue length q.

The transfer function of the Controller depends on the choice of controller. The simplest controller is a *proportional* controller [9] whose output is directly proportional to the input. The adjustment of R according to a proportional controller is simply given by:

$$R(n) = K(q_r - q(n)) = Ke(n) \tag{3}$$

where K is the *gain* of the controller and e is the observed error in the output.

The advantage of using a proportional controller is its simplicity; there is only one parameter (gain K) to adjust, and the stability analysis of the control system is quite simple. However, a major drawback of a proportional controller is the non-zero steady-state error in the output as explained below.

In the steady state, the input rate R is equal to the output rate (transmission rate of the access-link) B. This results in (from Eq. (3) the finite error (B/K) for the proportional feedback controller. In other words, if a simple proportional controller is used, the buffer level will stabilize at $q_r - (B/K)$ instead of at q_r.

The performance of a proportional feedback controller can be significantly improved (by reducing or eliminating the steady state error) by adding *integral*[3] control, which yields a PI (proportional plus integral) controller. However, integral control typically reduces the stability of the system. In order to increase the stability, *derivative*[4] control is usually used in conjunction with a PI controller to form a PID controller [9].

A PID controller provides an acceptable degree of error reduction simultaneously with acceptable degree of stability. However, a PID controller has three parameters to adjust (K, K_1 and K_2), one for each type of controller (proportional, integral and derivative). This increases the complexity of the control system. Moreover, it responds to a disturbance only *after* an error has been caused by a disturbance.

If a PID controller were to be used for the control system, the network administrator would have to adjust three control parameters to achieve the optimum performance. Adjusting such parameters can be quite costly; a wrong selection could result in large error, which in turn could degrade the performance of the control mechanism significantly.

Our objective is to design a control system, which provides an acceptable level of error reduction simultaneously with an acceptable level of stability with a *minimum* number of control parameters. One way to reduce the steady state error is to measure the disturbance and take appropriate actions before the disturbance has any effect (error) on the output. This type of control is called the feed forward control [9].

Disturbance feed forward can be a very effective strategy for reducing error if (i) the effect of the disturbance on the output is known, and (ii) the disturbance can be measured easily and accurately. In addition to reducing the error, disturbance feed forward has the benefit of not affecting the feedback loop (and hence, the stability properties of the controller) present in a feedback control system. And, importantly, disturbance feedforward can counteract the effect of the disturbance *before* significant error is caused.

With the proportional rate control system described above, both the effect of disturbance on the buffer level and the measurement[5] of disturbance can be easily obtained. From Eq. (3), it can be seen that for a given K, the higher the B, the greater the magnitude of the error in the steady state ($R = B$). The steady state error can be eliminated if the disturbance B is fed forward to the controller in the manner shown in Fig. 3.

The input rate R of the feed forward system is now

---

[3] An integral controller uses the equation [9]: $R(n) = R(n-1) + K_1 e(n)$, where R is the input and e the error in the output.

[4] A derivative controller uses the equation [9]: $R(n) = K_2[e(n) - e(n-1)]$, where R is the input and e the error in the output.

[5] The disturbance in this case is simply the transmission rate of the access-link, which can be obtained from the ABR rate feedback.

obtained as:

$$R(n) = B(n-1) + K(q_r - q(n)) = B(n-1) + Ke(n). \quad (4)$$

The reason that $B(n-1)$ is fed forward instead of $B(n)$ is because the controller can only act on the information gathered in the last sampling period. It can be seen from Eq. (4) that there is no error ($e = 0$) in the steady state (when $R = B$). By taking the z-transform of Eq. (4), the transfer function of the the feed forward controller is obtained as:

$$R(z) = z^{-1}B + Ke(z). \quad (5)$$

Feeding forward B eliminates the steady state error experienced under the proportional feedback controller. However, unlike integral control action, feed forward does not reduce the stability of the system; the stability properties remain the same because the feedback loop is unchanged. In addition to this, the number of control parameters (gain parameters) remains one (only K), which significantly simplifies the operation and tuning of the control system.

3.2. A note on linearity

The design presented above is based on a *linear* system: it is assumed that the buffer level can grow without limit in both directions of the target level q_r. In practice, however, only *finite* buffers are used at the gateway, and hence the buffer level may hit a limit at both ends. The linear design also assumes that input rate R can take any value including negative ones. In practice, R cannot be negative and so R must be calculated as:

$$R(n) = \max(0, B(n-1) + Ke(n)). \quad (6)$$

Therefore, the performance of our proposed control mechanism will depend on how frequently the system hits the limits on the buffer level and the input rate R. If the system is controlled to avoid hitting these limits as much as possible, the system will behave essentially as a linear system. In the following section, we analyze the performance of our linear design. The effectiveness of the linear design in modeling practical systems with finite buffers is evaluated via simulation in Section 5.

3.3. Performance analyses

In this section, the performance of the feedback control system is analyzed using tools from classical control theory.

3.3.1. Stability analysis

Stability is a very important requirement for any traffic rate control system. If the system is not stable, the buffer level in the gateway may increase without causing buffer overflow.

A variety of methods exist to assess the stability of a linear control system. In the *frequency-response* method [9], the phase margin (PM) is a measure of the relative stability of the system. A system is unstable for negative phase margin and stable for positive phase margin. Therefore, the gain K has to be selected to achieve a positive phase margin, the larger the phase margin the greater the relative stability as well as its robustness against system parameter changes.

The phase margin of a control system is determined by looking at the open-loop transfer function of the system, which from Fig. 3 is obtained as $(KT/z-1)$. From the definition of PM [9],

$$\mathrm{PM} = 180^\circ + \angle KT/(z-1), \quad (7)$$

where $\angle KT/(z-1)$ is the phase of $KT/(z-1)$ at the *gain crossover frequency*, which is defined as the frequency at which the magnitude of $KT/(z-1)$ is one (0 dB). Thus

$$\left|\frac{KT}{z-1} = 1\right|. \quad (8)$$

From Eqs. (7) and (8), the phase margin is obtained as:

$$\mathrm{PM} = \tan^{-1}\frac{\sqrt{4-(KT)^2}}{KT}. \quad (9)$$

From Eq. (9), it can be deduced that for a stable system ($\mathrm{PM} \geq 0$), K should be less than or equal to $(2/T)$.

3.3.2. Buffer level variance

Although the control system of Fig. 3 will keep the buffer level at the desired set point of q_r in the steady state, sudden fluctuation in the available bandwidth of the access-link will cause momentary shift of the buffer level from q_r. Therefore, the variance σ_q^2 of the buffer level around the desired set point q_r is an important performance metric for the intranet traffic management system. If the variance is large, there may be buffer overflow and data loss in the access gateway. In this section, we analyze the buffer level variance σ_q^2 as a function of K. This analysis will allow us to select a K, which minimizes the buffer level variance.

In order to calculate the variance in the buffer level, we first need to model the changes that occur to the available bandwidth B. One of the main contributors to the change of available bandwidth in the ATM network is the bursty traffic rate of VBR sources. If the ABR connection is multiplexed with several VBR sources, $B(n)$ is the difference between the ABR peak rate C and the aggregated rate of VBR sources. If the VBR sources are modeled as ON–OFF sources (in the literature, bursty traffic expected for ATM networks are usually modeled as ON–OFF sources [10,11]), then (see Appendix) the available bandwidth B can be modeled by the following first order stochastic difference equation

$$B(n+1) = \alpha B(n) + (1-\alpha)\beta W(n), \quad (10)$$

where $W(n)$ is a random (white-noise) signal, α is the correlation coefficient of $B(n)$, and β is a constant indicating the magnitude of the noise. For $\alpha = 1$, there is no noise and the ABR remains constant; for $\alpha = 0$, B is a pure random number. In practice, α will be a fraction between 0 and 1.

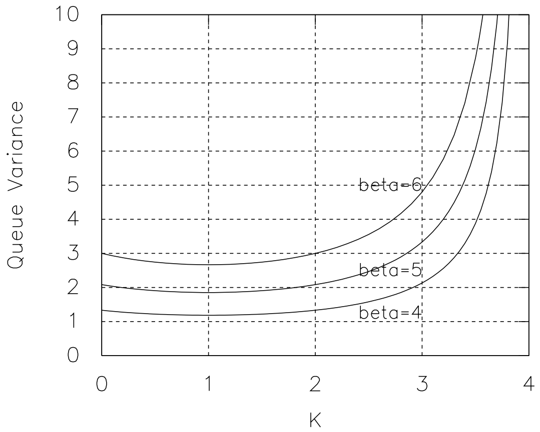

Fig. 4. Queue level variance as a function of $K(\alpha = 0.5,\ T = 0.5)$.

Taking z-transform of Eq. (10) we get,

$$\frac{B(z)}{W(z)} = \frac{(1-\alpha)\beta}{z-\alpha}. \tag{11}$$

From the block diagram in Fig. 3 we get the transfer function between q and B as:

$$\frac{q(z)}{B(z)} = \frac{(z^{-1}-1)T/z-1}{1+\dfrac{KT}{z-1}} = \frac{T(1-z)}{z(z-(1-KT))}. \tag{12}$$

Multiplying Eqs. (11) and (12), we obtain the transfer function between noise source W and buffer level q as:

$$\frac{q(z)}{W(z)} = \frac{(1-\alpha)\beta T(1-z)}{z(z-\alpha)(z-(1-KT))} = \frac{b(z)}{a(z)}. \tag{13}$$

From Eq. (13), the discrete power spectral density S_q of the buffer level q obtained as,

$$S_q = \frac{b(z)b(z^{-1})}{a(z)a(z^{-1})}. \tag{14}$$

From S_q the variance of q can be computed using the well-known *contour integral* [12] and *Cauchy's residue theorem* [13] as,

$$\begin{aligned} \mathrm{var}(q) = \sigma_q^2 &= \frac{1}{2\pi j}\oint \frac{b(z)b(z^{-1})\mathrm{d}z}{a(z)a(z^{-1})z} \\ &= \sum\left[\text{residues of } \frac{b(z)b(z^{-1})}{a(z)a(z^{-1})z} \text{ at poles inside unit disc}\right] \\ &= \frac{2(1-\alpha)^2\beta^2T^2}{(1+\alpha)(1-\alpha+\alpha KT)(2-KT)} \quad \text{for } 0 < K < \frac{2}{T}. \end{aligned} \tag{15}$$

Note that the contour integral defining the variance is valid only as long as the control system is strictly stable, i.e. only if $0 < K < (2/T)$.

For a given ABR connection, the noise parameters α and β may be estimated by analyzing traces of earlier data. If the change in available bandwidth is primarily caused by VBR sources multiplexed with the ABR connection in the ATM network, both α and β can be obtained analytically from the properties of the VBR sources (see Appendix).

3.3.3. Results

Eq. (15) provides the buffer level variance as a function of the controller gain K. It would be an important design criterion to select the K which minimizes the buffer level variance. Let us denote this value of K by $\hat{K}$.

By examining the first derivative of Eq. (15), it can be shown (the proof is omitted due to space limitation) that the buffer level variance is an increasing function of K for $\alpha \le (1/3)$. However, for usual operation of the system, α will be greater than 1/3 (see Appendix). For $\alpha > (1/3)$, it can be shown (the proof is omitted due to space limitation) that the variance function of Eq. (15) is a convex function with a minima occurring in the range $0 < \hat{K} < (1/T)$. By equating the first derivative of Eq. (15) to zero, we get

$$\hat{K} = \frac{3\alpha-1}{2\alpha T} \quad \text{for } \alpha > \frac{1}{3} \tag{16}$$

Fig. 4 shows the queue level variance as a function of K for $\alpha > (1/3)$ and for different values of β. A number of important observations are in order:

1. *$\hat{K}$ does not depend on the noise parameter β.* This result is obtained from Eq. (16) and is verified in Fig. 4, where the minimum variance is obtained at $K = 1$ for all three different values of β. Therefore, in obtaining $\hat{K}$, the only noise parameter, one needs to estimate α.
2. *Variance curve is very flat near the minima.* This allows a "reasonable" error margin for estimating α and hence $\hat{K}$; the variance is not significantly affected by a "reasonable" amount of error in estimating $\hat{K}$.
3. $\hat{K} \le (1/T)$. This result is obtained by replacing $\alpha = 1$ in Eq. (16). From Eq. (9), a minimum of 60° of phase margin is obtained for $K \le (1/T)$, which suggests a very stable system. This observation confirms that the proposed control mechanism has very good relative stability for the choice of gain K for which the buffer variance is minimized.

The above observations suggest that the proposed feedback control mechanism is robust to system variations, and hence can guarantee a low buffer level variance for a wide range of bandwidth fluctuation for the access-link. This is expected to minimize the buffer requirement, and hence the queueing delay and delay variation at the access gateway.

The variance obtained from Eq. (15) may also be used to estimate the probability of buffer overflow (or packet loss) at the access-gateway for a given buffer size. In the following section, a technique is presented to estimate such packet loss probability.

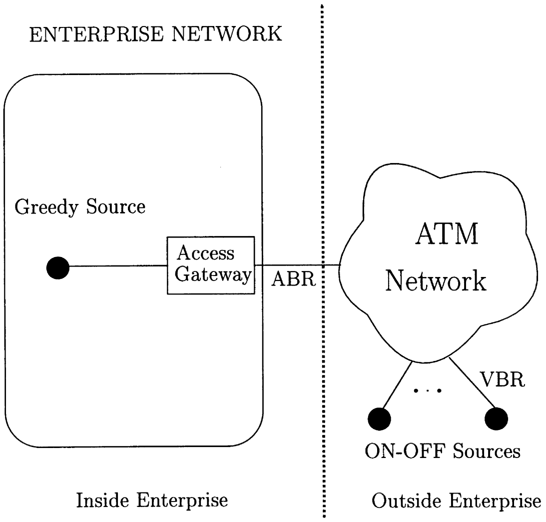

Fig. 5. Network model for the simulation.

4. Packet loss probability

In deriving the packet loss probability for a given buffer size, we first assert that *the buffer level fluctuates around the steady state value* (*desired level* q_r) *with a normal distribution*. A proof substantiating the above assertion is presented next followed by the technique to estimate packet loss probability. The validity of the above assertion and the packet loss probability estimating technique are further verified through simulation in Section 5.

4.1. Normal distribution

The buffer level q is related to the random noise W by the linear transfer function of Eq. (13). Hence, $q(k)$ is the convolution of $W(k)$ with the system's impulse response, i.e.

$$q(k) = h(k)W(0) + h(k-1)W(1) + \cdots + h(0)W(k). \quad (17)$$

Now $W(k)$ is a sequence of independent random numbers having an arbitrary amplitude probability distribution. Thus, Eq. (17) implies that $q(k)$ is a weighted sum of independent random numbers; the effect of the weighting is merely to scale the amplitude probability distribution of each term. Hence, by the central limit theorem [14], the amplitude probability distribution of $q(k)$ tends to normal as k tends to infinity.

4.2. Packet loss probability

In the previous section, we showed that the buffer level distribution can be closely approximated by a normal distribution. For our control system, the mean is q_r and the standard deviation is σ_q as obtained from Eq. (15). The probability of buffer overflow p is calculated as $Pr\{q \geq L\}$, where q is the instantaneous buffer level. Converting the normal distribution with parameters q_r and σ_q into a *standard* or *unit* normal distribution [14], this probability becomes equivalent to

$$P = 1 - \Pr\left\{X < \frac{L - q_r}{\sigma_q}\right\} = 0.5 \times \operatorname{erfc}\left(\frac{L - q_r}{\sqrt{2}\sigma_q}\right), \quad (18)$$

where erfc is the *error function complementary* [15], X is a normal variable with mean 0 and variance 1, and hence can be evaluated using standard function tables.

It should be noted, however, that the probability in Eq. (18) actually defines the probability of the buffer level exceeding L assuming *infinite* buffer size. With *finite* buffers, the loss probability is expected to be slightly lower than the one predicted by Eq. (18). The effectiveness of Eq. (18) in estimating the packet loss probability for finite buffers is evaluated using simulations in Section 5.

5. Simulation

5.1. Simulation model

The simulated network model is shown in Fig. 5. There is a "greedy" source in the enterprise network trying to send traffic to the Internet as fast as possible via the access gateway. The gateway is connected to the ATM network via the ABR service. Hence, the transmission rate of the gateway is dynamically controlled by the ATM network depending on the amount of bandwidth within the ATM network that becomes available.

In order to exercise variability in the available bandwidth in the ATM network, the ABR connection is multiplexed with a number of VBR connections inside the ATM network. Each VBR source can be characterized by a two-phase ON–OFF model with exponential ON and OFF periods. A VBR source transmits traffic at a constant rate during the ON period and does not send any traffic during the OFF period. VBR has higher priority than ABR, any bandwidth not consumed by the VBR sources is made available to the ABR source by updating the transmission rate of the access gateway whenever an ON–OFF source changes its state from ON to OFF or OFF to ON. It is assumed that there is enough ATM bandwidth for the gateway to transmit at the peak rate (PCR of the ABR connection) when all ON–OFF sources are in the OFF-state.

The current queue level q and the average bandwidth (averaged over the last T seconds) B of the gateway is measured every T seconds, and the allowed rate R for the greedy source is calculated as (according to Eq. (6))

$$R = \max(0, B + K(q_r - q)) \quad (19)$$

The greedy source is then immediately notified of the new allowed rate R. The propagation delay between the gateway and the greedy source is ignored in the simulation because an intranet for a single campus is limited within a "small" boundary.

In all simulations, the target queue level, q_r, is set to half of the buffer size. In practice, the "headroom" $L - q_r$ is

Table 1
Values of simulation parameters

Parameter	Value
Number of VBR sources, N	4
Transmission rate during ON period for VBR source A	10 pk/s
Mean length of ON and OFF period for VBR source, t_0	2 s
Peak transmission rate of gateway	50 pk/s
Sampling interval, T	0.5 s
Target queue level, q_r	half of buffer size

chosen to keep the buffer overflow (packet loss) probability less than a specified value. Similarly, the "tailroom" q_r should be chosen so that the probability of buffer underflow (link idling) is less than a certain value. A simple choice is to make both probabilities the same; in such a case $q_r = L/2$.

The length of all simulations was dynamically controlled until the values of the estimated variables (e.g. packet loss rate) reached steady state, and 95% confidence interval was achieved with 5% relative precision. With such adaptive length simulations, the number of packets simulated varied between a few million packets to tens of millions of packets depending on the buffer size simulated at the gateway.

The values of simulation parameters are summarized in Table 1. Results obtained from the simulations are presented and compared with the analytical results in the following section.

5.1.1. Results

In this section, the results obtained from the simulations are presented and compared with the ones obtained from analysis to assess the validity of our analysis. In particular, four aspects of our analysis are assessed in the following subsections—(i) stability of the control system, (ii) normal distribution of queue length, (iii) the existence of an optimum gain K to minimize the packet loss probability at the gateway, and (iv) calculation of packet loss probability as a function of buffer size at the gateway (this assists buffer dimensioning).

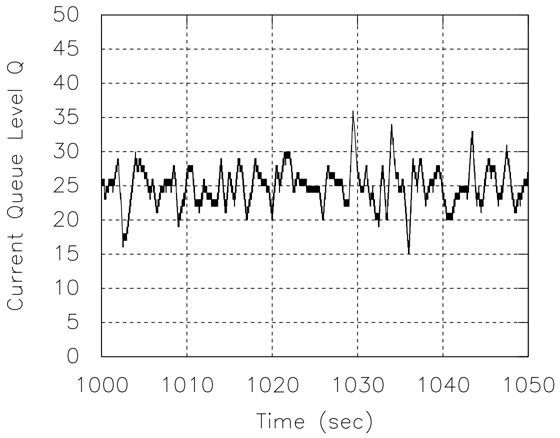

Fig. 6. Current queue level as a function of time.

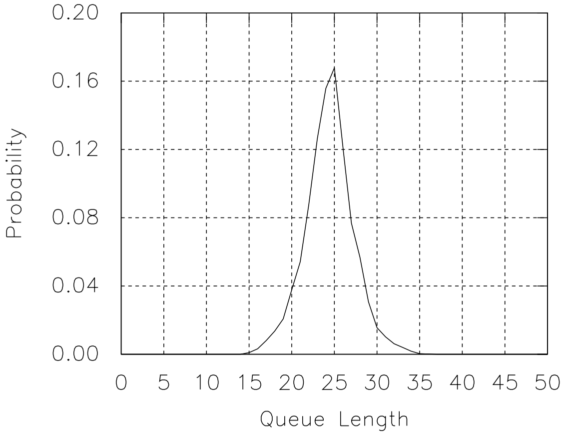

Fig. 7. Queue length distribution obtained from simulation.

5.1.2. Stability

We simulate a buffer size of 50 with target queue level set to $q_r = 25$ and investigate the stability of the proposed control mechanism by recording the current queue level for 50 s (100 samples) of simulated time after 1000 s (after 2000 initial samples) into the simulation. The controller gain K was set to 1.62, which is the optimum K to minimize buffer level variance.

Fig. 6 shows the current queue level as a function of time. It can be seen that the queue level fluctuates around the target queue level 25, which demonstrate that the system is quite stable.

5.1.3. Queue length distribution

Our packet loss probability analysis in Section 4 is based on the assertion that the queue level fluctuates around the target level q_r with a normal distribution. We substantiated our assertion by making use of a central limit theorem. In this section, we investigate the accuracy of our assertion by examining the distribution of the queue levels obtained from the simulation. We simulated a buffer size of 50 with a target queue level of $q_r = 25$.

Fig. 7 shows the queue length distribution obtained from the simulation between 1000 and 1100 s into the simulation (200 sample times for a sample period of $T = 0.5$ s). There were a total of 5989 queue lengths recorded during this time period. The x-axis shows the queue lengths and the y-axis plots the fraction of times a given queue length occurred in the sample of 5989 queue lengths. The curve is roughly "bell-shaped" around 25 confirming that the queue length can be approximated by normal distribution with the mean q_r.

5.1.4. Optimum K

In this section, we compare the packet loss probability P

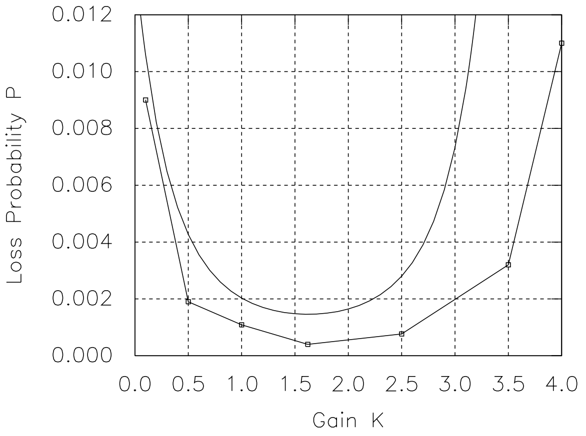

Fig. 8. Packet loss probability as a function of gain *K*.

(as a function of gain *K*) as obtained from simulation of a finite buffer of size $L = 20$ with the one obtained from analysis (see Eq. (18)). For calculating P using Eq. (18), we need to calculate the variance, and hence need to obtain the values of noise parameters α and β. For our simulation model, α and β are obtained as shown in Eqs. (A.7) and (A.5) in the Appendix. For $T = 0.5$ and $t_0 = 2$, we get $\alpha = 0.73$, and an optimum gain $K = 1.62$ (from Eq. (16)). For $N = 4$, $A = 10$, $T = 0.5$ and $t_0 = 2$, we get $\beta = 23.2$.

The loss probability P as a function of K, as obtained from the simulation and analysis (using Eq. (15) for $\alpha = 0.73$ and $\beta = 23.2$.) are shown in Fig. 8. As expected, the loss probabilities obtained from *finite* buffer simulation are lower than the ones obtained from our linear analysis based on *infinite* buffer. However, the most important observation here is that the simulation results confirm the existence of an optimum gain K which minimizes the packet loss probability at the gateway and the minimum K is obtained near 1.62 as obtained from analysis. This confirms that our analysis is very effective in deriving the optimum K.

5.1.5. Packet loss probability

As mentioned in Section 4, Eq. (18) can be used for buffer dimensioning as it allows us to dimension the buffer for a target packet loss probability. In this section, we present simulation results for packet loss rates and compare them with analytical ones for a number of different buffer sizes for target loss rates below 1% to evaluate the effectiveness of Eq. (18). The packet loss rates obtained from simulations are plotted in Fig. 9 along with the loss probabilities obtained from Eq. (18). For both simulation and analysis, we consider the optimum gain $K = 1.62$ to minimize the buffer level variance.

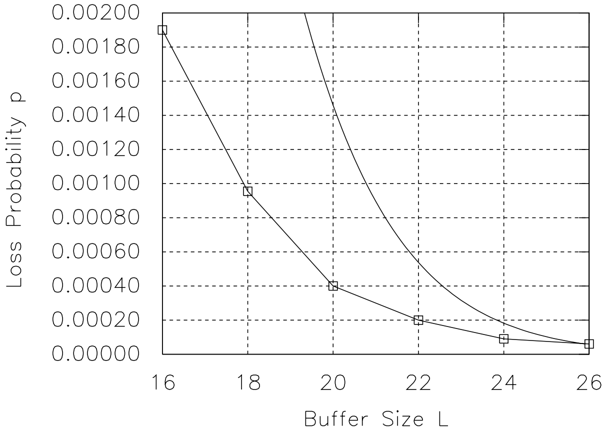

Fig. 9. Packet loss rate as a function of buffer size.

As predicted, the simulation results are lower than the analytical values (see Fig. 9), as we simulated *finite* buffers, whereas the linear analysis is based on *infinite* buffer. However, the discrepancy between simulation and analysis decreases for smaller (and probably for more realistic) loss probabilities. For example, in order to keep the packet loss probability below 0.001, our analysis suggests a buffer size 21 (see Fig. 9), but the simulation shows that a buffer size of only 18 will do the job; this results in over dimensioning the buffer by 3. However, for a target loss probability of 0.0002, the buffer will be over dimensioned by only 2. Hence, our packet loss probability formula presented in Section 4 can be effectively used for buffer dimensioning purposes for low loss probabilities.

6. Conclusion

The need for traffic rate control at the access-gateway of enterprise networks connected to the Internet using the ATM ABR service is introduced in this paper. The design of a rate-based traffic management mechanism between the gateway and the enterprise network has also been presented. The proposed traffic management mechanism was analyzed using a proportional feedback controller which was augmented with a noise feedforward to minimize the effect of bandwidth fluctuation in the ATM ABR access-link on the performance of the control system. Implementation of the proposed traffic management mechanism is expected to reduce the buffer requirement at the access gateway of the intranet resulting in less expensive access gateway design. Alternatively, for a given amount of buffer, the proposed mechanism minimizes or eliminates buffer overflow and packet loss at the access gateway by keeping the buffer level at a desired set point; this in turn improves the performance of applications in the enterprise network.

The variance of the buffer level, at the access gateway, around the desired threshold was studied through analysis. The proportional gain of the controller was optimized to achieve the minimum possible variance of the buffer level. It was demonstrated that the dependence of the optimum gain on the noise parameters is minimal, which makes the traffic management mechanism robust to system variations, such as fluctuations in the ABR access-link bandwidth.

For a given amount of buffer at the access gateway, the

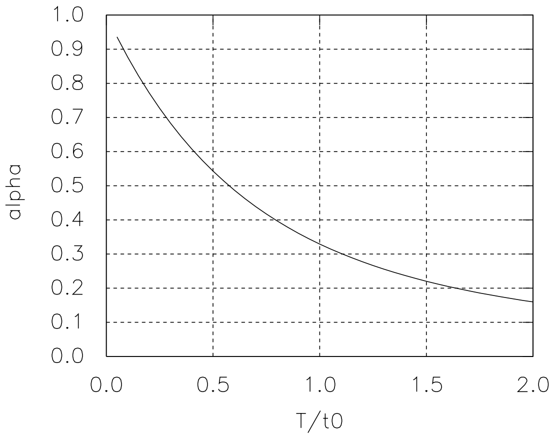

Fig. 10. Value of α as a function of (T/t_0), the ratio of the sampling interval to the mean ON and OFF period of VBR sources.

probability of packet loss at the gateway due to buffer overflow was estimated as a function of the buffer level variance.

As the selection of optimum gain depends on the value of noise parameter α, it is necessary to estimate α for a given ABR access-link. This estimation can be performed off-line from the log data. We are currently working on an adaptive estimation technique, which will estimate the value of α dynamically.

Appendix A. Calculation of α and β

In our model, the ABR rate $B(n)$ during the nth sampling interval is the difference between the peak ABR rate C and the rate $V(n)$ of the ON–OFF VBR sources averaged over the interval. The ON–OFF sources can be modeled by an equation similar to Eq. (10) with the same α and β. Thus, α and β can be calculated from the properties of the ON–OFF sources.

Consider now a single ON–OFF source which has a rate $v(t)$ at time t equal to A in the ON state and zero in the OFF state. The ON and OFF epochs are both exponentially distributed with mean t_0. Then, the mean and autocovariance functions of $v(t)$ are [16]

$$\bar{v} = E\{v(t)\} = \frac{A}{2},$$

$$R_v(\tau) = E\{\tilde{v}(t)\tilde{v}(t+\tau)\} = \frac{A^2 \mathrm{e}^{-2|\tau|/t_0}}{4}, \tag{A.1}$$

where $\tilde{v}(t) = v(t) - \bar{v}$.

Let $V(n)$ represent the average bandwidth during the nth sampling period of duration T, i.e.

$$V(n) = \frac{1}{T}\int_{(n-1)T}^{nT} v(t)\,\mathrm{d}t. \tag{A.2}$$

Then the exponential form of the autocovariance of Eq. (A.1) implies that $V(n)$ satisfies a first-order stochastic difference equation of the form of Eq. (10). Invoking the stationarity of $v(t)$, the mean and variance of $V(n)$ are obtained as

$$E\{V(n)\} = \frac{1}{T}\int_0^T E\{v(t)\}\,\mathrm{d}t = \tilde{v} = \frac{A}{2},$$

$$\mathrm{var}\{V(n)\} = \sigma_{\mathrm{v}}^2 = E\left\{\frac{1}{T}\int_0^T \tilde{v}(t)\,\mathrm{d}t \times \frac{1}{T}\int_0^T \tilde{v}(t')\,\mathrm{d}t'\right\}$$

$$= \frac{A^2 t_0^2[2T/t_0 + \mathrm{e}^{-2T/t_0} - 1]}{8T^2}. \tag{A.3}$$

Applying the variance calculation theorem [12] and *Cauchy's residue theorem* [13] on Eq. (11) we get,

$$\sigma_{\mathrm{v}}^2 = \frac{(1-\alpha)\beta^2}{(1+\alpha)}. \tag{A.4}$$

From the above equation β is obtained as:

$$\beta = \sigma_{\mathrm{v}}\sqrt{\frac{(1+\alpha)}{(1-\alpha)}}. \tag{A.5}$$

In order to calculate α, we also need to find

$$\mathrm{cov}\{V(n)V(n-1)\}$$

$$= E\left\{\frac{1}{T}\int_0^T \tilde{v}(t+T)\,\mathrm{d}t \times \frac{1}{T}\int_0^T \tilde{v}(t')\,\mathrm{d}t'\right\} \tag{A.6}$$

$$= \frac{A^2 t_0^2 (1 - \mathrm{e}^{-2T/t_0})^2}{16T^2}.$$

Then, we have

$$\alpha = \frac{\mathrm{cov}\{V(n)V(n-1)\}}{\mathrm{var}\{V(n)\}} = \frac{(1-\mathrm{e}^{-2T/t_0})^2}{2\mathrm{e}^{-2T/t_0} + 4T/t_0 - 2}. \tag{A.7}$$

If there are N identical, but independent ON–OFF sources, then α is unchanged but β is increased by a factor of $\sqrt{N}$.

In order to observe the effect of the sampling interval T on α, we plot α as a function of (T/t_0) (from Eq. (A.7)), the ratio of sampling interval to the mean ON and OFF period of VBR sources, as shown in Fig. 10. The figure demonstrates that $(T/t_0) \leq 1$, remains above (1/3). For a reasonable performance, the sampling interval is selected such that $(T/t_0) \leq 1$, it is confirmed that for usual operation of the system, there exists an optimum gain K which yields a minimum buffer level variance (see Section 3.3.3).

References

[1] M.D. Prycker, Asynchronous Transfer Mode, 3, Prentice Hall, Englewood Cliffs, NJ, 1995.

[2] ITU, Recommendation I.121: Broadband Aspects of ISDN, ITU, Geneva, 1991.

[3] M. Hassan, Impact of variable bandwidth on the performance of data communications over ATM-ABR based LAN interconnection, in: 20th Australasian Computer Science Conference, Sydney, Australia, pp. 422–429, February 1997.

[4] S. Kalyanaraman, R. Jain, S. Fahmy, R. Goyal, S. Kim, Performance of TCP over ABR on ATM backbone and with various VBR traffic patterns, ATM Forum Contribution: 96-1294, October 1996.

[5] N. Yin, S. Jagannath, End-to-end traffic management in IP/ATM internetworks, ATM Forum Contribution: 96-1406, October 1996.

[6] S. Jagannath, N. Yin, End-to-end TCP performance in IP/ATM internetworks, ATM Forum Contribution: 96-1711, December 1996.

[7] W. Ren, K. Siu, H. Suzuki, G. Ramamurthy, Performance of TCP over ATM with legacy LAN, ATM Forum Contribution: 96-1716, December 1996.

[8] M. Hassan, H. Sirisena, M. Atiquzzaman, A traffic management mechanism for intranets with available bit rate access to the internet, in: SPIE Conference on Performance and Control of Network Systems, Dallas, USA, pp. 470–478, November 1997.

[9] C.C. Bissell, Control Engineering, 2, Chapman & Hall, New York, 1994.

[10] J.W. Roberts, editor. Management Committee of the COST224 Project, Performance Evaluation and Design of Multiservice Network, Commission of the European Communities, 1992.

[11] E.P. Rathgeb, Modeling and performance comparison of policing mechanisms for ATM networks, IEEE Journal on Selected Areas in Communications 9 (3) (1991) 325–334.

[12] K.J. Astrom, B. Wittenmark, Computer Controlled Systems, 2, Prentice Hall, Englewood Cliffs, NJ, 1990.

[13] E. Kreyszig, Advanced Engineering Mathematics, 4, Wiley, New York, 1979.

[14] S.M. Ross, Introduction to Probability and Statistics for Engineers and Scientists, Wiley, New York, 1987.

[15] M. Abramowitz, I.A. Stegun, Handbook of Mathematical Functions with Formulas, Graphs and Mathematical Tables, National Bureau of Standards, Washington, DC, 1970.

[16] A. Papoulis, Probability, Random Variables, and Stochastic Processes, McGraw-Hill, New York, 1965.