



HG-RCP: A rate-based flow control mechanism for intranets interconnected by ATM networks

Mahbub Hassan ^{a,*}, James Breen ^{a,1}, Mohammed Atiquzzaman ^{b,2}

^a School of Computer Science and Software Engineering, Monash University, Melbourne, Vic. 3145, Australia

^b Department of Electrical and Computer Engineering, University of Dayton, Dayton, OH 45469-0226, USA

Abstract

A number of distributed corporate intranets can be connected together using the ABR service of an ATM network to form a virtual private network. However, because the available bandwidth of the ABR connection may be occasionally reduced by the ATM network, severe congestion may occur at the gateway connecting an intranet to the ATM network. Such congestion is likely to cause packet loss at the gateway which in turn may significantly reduce the throughput of intranet applications. In order to address the above problem, we propose a novel rate-based flow control mechanism, called Host Gateway Rate Control Protocol (HG-RCP), between the gateway and the intranet hosts. The control mechanism of HG-RCP is modeled as a classical feedback control system and the control parameters are selected to guarantee the stability of the system. The performance of HG-RCP is evaluated using simulation for both limited and unlimited buffers at the gateway. Simulation results confirm that HG-RCP is very stable and fair in allocating the ABR bandwidth among the competing sources in the intranet. It is also demonstrated via simulation that HG-RCP can significantly reduce the amount of buffers required at the gateway and improve the performance of TCP/IP applications in the intranets. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Intranet; TCP/IP; ATM; ABR service; Flow control

1. Introduction

Intranet is a rapidly growing internet-working technology for connecting computers and resources within an organisation using the *standard* Internet

protocols such as Transmission Control Protocol (TCP) and Internet Protocol (IP). Such intranets have enormous benefits over proprietary networking as they provide a consistent and efficient infrastructure to support all communication needs within the organisation. With such intranets in place, organisations can readily connect their networks to the global Internet with minimum cost and effort.

For large corporations with several geographically distributed offices, a corporate wide intranet (sometimes referred to as wide area intranet or virtual private network) is established by interconnecting the individual intranets in remote offices using IP-

* Corresponding author. Tel.: +61 3 9903 2122; Fax: +61 3 9903 1077; e-mail: m.hassan@cssc.monash.edu.au

¹ Tel.: +61 3 9905 3298; e-mail: j.breen@csse.monash.edu.au.

² Tel.: +1 937 229 3183; Fax: +1 937 229 4529; e-mail: atiq@engr.udayton.edu.

based gateways and interoffice links. The most common type of interoffice links are established over a dedicated fixed bandwidth leased line from the public telecommunication service providers as shown by the direct solid lines in Fig. 1. In this figure, there are three remote intranets each with a IP-based gateway for connecting the intranet with a remote intranet. Although the leased lines provide a high speed intranet connectivity, they do not make efficient use of fixed bandwidth for bursty data traffic and hence have diminished cost effectiveness.

The next generation of telecommunications networks, based on the Asynchronous Transfer Mode (ATM) networking technology, are expected to provide a more efficient, flexible and economical way of transmitting bursty, non-real time data traffic through a new mechanism called the Available Bit Rate (ABR) service [4]. The ATM network provides a variable bandwidth to ABR service users by controlling the transmission rates of the users with explicit rate feedback. As ABR service users will use only the hitherto unused bandwidth in the ATM network, they are expected to be charged at a signifi-

cantly lower rate than other service users. The cost of wide area intranet connectivity can be significantly reduced by replacing the leased lines with ABR connections as shown by the dotted lines in Fig. 1.

ABR-based interoffice links, however, may cause potential traffic congestion at the gateways if the ABR bandwidth is suddenly reduced due to congestion in the ATM network. If the traffic from the intranet sources continues to arrive at a higher rate than the reduced bandwidth of the ABR link, the gateway buffers may overflow causing data loss and reduction of *goodput* which is defined as the effective data transmission rate of the intranet sources. It is, therefore, very important to implement a flow control mechanism within the intranet which can effectively control the traffic rates of the intranet sources to avoid buffer overflow at the gateway.

IP is the network layer protocol for the intranets which supports a number of transport layer protocols, among which TCP is the most common. Many intranet applications are supported by TCP/IP, however, there is no flow control mechanism included in

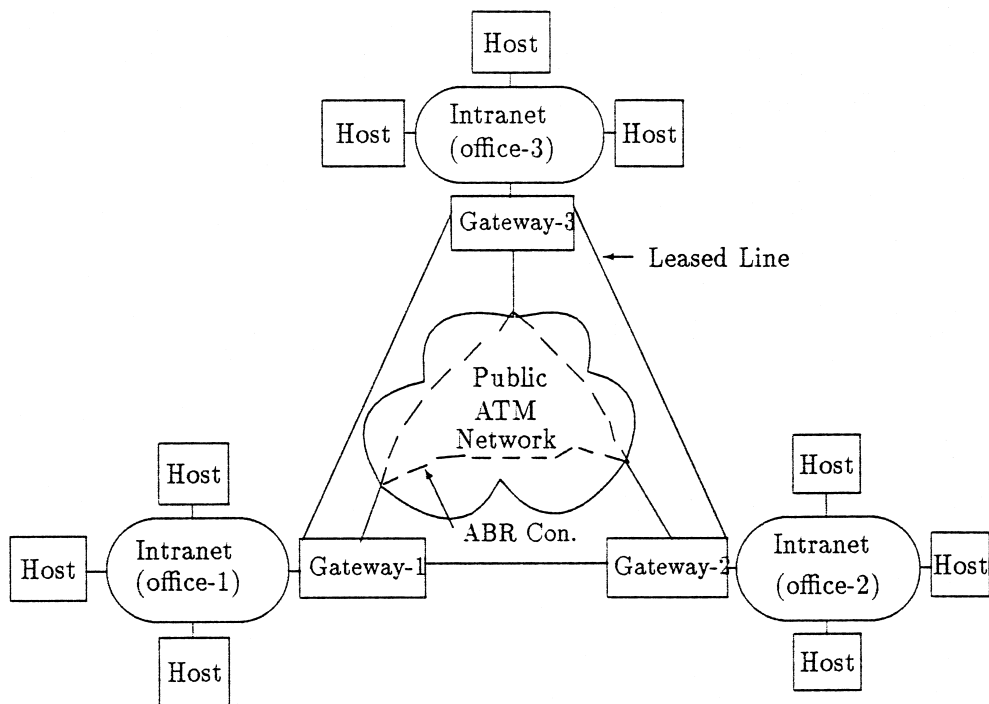


Fig. 1. Interconnection of three remote intranets.

the standard IP; the only flow control available in the standard TCP/IP is the end-to-end window-based flow control within the TCP layer.

Recent studies [10,16,24] have shown that the end-to-end *window-based* flow control algorithm in TCP is not efficient in adjusting the traffic rates of the intranet sources in response to the dynamic ABR bandwidth at the gateway. Window-based flow controls only control the number of packets in the end-to-end path; they do not control the *rate* of transmission (interpacket gap is not controlled). As a result, a window full of packets may still continue to arrive at the gateway at a high rate even though the ABR bandwidth has been reduced. Consequently, large buffers are necessary at the gateway to prevent buffer overflow during congestion caused by the sudden drop of ABR bandwidth.

Implementing large buffers has several disadvantages. Besides increasing the cost, large buffers also increase the queueing delay in the gateway, which may severely degrade the performance of delay-sensitive applications.

In order to reduce the amount of buffer space required at the gateway, it is necessary to effectively control the rate of intranet sources according to the current ABR bandwidth at the gateway. Hence, a mechanism is necessary to explicitly carry the ABR rate from the gateway to the individual sources on the intranet. A number of researchers have previously proposed flow control mechanisms [7–9,17,20,21] to improve the performance of TCP/IP during congestion, however, none of these mechanisms are suitable for intranets with ABR-based interoffice links for reasons explained below.

The mechanisms proposed in [8,9,21] are based on end-to-end window flow control at the transport layer. As mentioned above, such mechanisms do not interact well with the rate-based flow control of the ABR service. As a result, the buffer space required at the gateway to guarantee zero loss will be quite large.

While a few authors proposed some rate-based flow control mechanisms [7,17,20], *they did not consider explicit rate feedback* from the gateway to the sources. Without the explicit rate feedback, the sources implement complex and ad hoc rate increase/decrease algorithms which are likely to cause instability in the system. Stability issue was not

addressed and no stability analyses of the flow control mechanisms were provided in [7,20]. Although the mechanism in [17] was designed using formal methods to guarantee stability, it requires dedicated buffers at the gateway for each competing source, and therefore is not scalable for large intranets with many sources, and cannot be implemented in existing first-in-first-out (FIFO) gateways.

The objective of this paper is to present the design of a flow control mechanism for the intranets which can effectively adjust the rate of the intranet sources according to the current bandwidth of an ABR interoffice link which is dynamically controlled by the ATM network. By adjusting the source traffic rates, the flow control mechanism would greatly reduce the buffer size required to prevent data loss at the gateway and increase goodput of intranet sources. In the design of the flow control mechanism, we wish to avoid (i) modification of TCP or any existing protocols at transport layer or above, (ii) complex rate increase/decrease algorithms at the sources, (iii) instability in the system, and (iv) per source buffering at the gateway.

In this paper, we propose a novel flow control mechanism, called the Host-Gateway Rate Control Protocol (HG-RCP), which achieves the above objectives by enhancing the standard IP to include an *explicit rate feedback* control mechanism. Since all intranet transport layer protocols are based on IP, the proposed mechanism provides an effective flow control for the entire range of intranet applications including those using TCP. It should be noted that *HG-RCP does not modify the TCP or any other transport protocols*. Since HG-RCP uses explicit rate feedback, the sources adjust their rates by simply switching to the new rate as explicitly fed back by the gateway. This eliminates the need for complex rate increase/decrease algorithms at the sources.

The proposed flow control mechanism is modeled as a classical feedback control system, with values of control parameters being obtained through analysis of the feedback control system to ensure the stability of the control mechanism and simultaneously minimise the buffer level fluctuation at the gateway. The stability of HG-RCP is confirmed by the results obtained from a detailed simulation study of a corporate wide intranet established by connecting two remote offices over the ATM ABR service. HG-RCP

does not require per source buffering at the gateway, all competing sources share a common FIFO buffer.

We use the term *conventional TCP/IP* to refer to the standard TCP/IP without the HG-RCP. A comparison of TCP/IP with HG-RCP is made against the best possible implementation (fine TCP clock granularity at the source and packet-based discarding at the gateway) of the conventional TCP/IP and it is shown that (i) TCP/IP with HG-RCP requires fewer buffers at the intranet-ATM gateway to guarantee zero loss, (ii) for the same amount of buffer at the gateway, TCP/IP with HG-RCP achieves higher goodput, and (iii) TCP/IP with HG-RCP is fair in evenly distributing the available bandwidth of the ABR connection among the contending TCP sources on the intranet.

The rest of this paper is structured as follows. A functional description of HG-RCP is provided in Section 2. Section 3 presents analyses on the stability of the control mechanism, derivation of buffer size at the gateway required for zero loss and the amount of control traffic overhead for implementing HG-RCP. The simulation model and the results are provided in Sections 4 and 5. Finally, concluding remarks are given in Section 6.

2. Description of HG-RCP

HG-RCP extends the standard IP to introduce a rate control mechanism which dynamically controls the traffic rates of the intranet sources to maintain the buffer occupancy at the gateway at a desired level. This prevents congestion at the gateway and reduces the amount of buffer required to prevent data loss.

The extension, which only has application within the intranet, comprises two additional Internet Control Message Protocol (ICMP) message types, and the use of one hitherto unused bit in the IP header. HG-RCP requires extension of the standard IP at both the gateway and at the hosts. The behaviour of HG-RCP at the gateway and the hosts are described below.

2.1. Gateway behaviour

For gateways with multiple ABR connections for multiple destinations, HG-RCP considers the buffer

space for each ABR connection as a separate, independent control system. For each buffer, the gateway calculates an aggregate required input rate R once every T seconds interval (T is called the *sampling interval* of the control mechanism) which will eliminate any difference between the desired buffer level Q and the current level q according to the following equation:

$$R = \max(0, \text{ACR} + K(Q - q)) \quad (1)$$

where allowed cell rate (ACR) is the current bandwidth of the ABR connection provided by the ATM network and K is a gain parameter which affects the speed of response and the stability of the control system. It may be worth noting here that if the current buffer level q exceeds Q by a significant margin, R may become zero and the intranet hosts may be asked to stop transmission completely (until q drops to an appropriate lower value) even when the minimum cell rate (MCR) of the ATM ABR connection is greater than zero.

If there are N sources competing for the bandwidth of an ABR connection, the gateway calculates an explicit rate for each competing source as R/N for equal distribution of bandwidth.³ Such explicit rates are calculated for all ABR connections in the gateway. At the end of each sampling interval, the gateway compares the new explicit rates with the ones calculated at the previous intervals and broadcasts the explicit rates which differ from the previous ones to the hosts on the intranet using a new ICMP message, referred to as rate notification (RN) message in this paper.

The gateway identifies each ABR connection and the associated buffer with a unique identifier Q_{id} and associates each explicit rate in the RN message with the corresponding identifier to enable the intranet hosts to control the rates of various sources sending traffic through various ABR connections in the gateway. The rate management within the hosts is described below in more detail.

Since the number of competing sources N for a given ABR connection changes over time, the gate-

³ Non-equal distribution can be easily achieved with HG-RCP by introducing a *rate-multiplier factor* α as shown in [11].

way updates N depending on the number of sources which transmitted one or more IP packets to the ABR connection in the last W seconds.

There are two types of traffic generated by the hosts, one with local destinations (does not travel through the ATM link) and the other with remote destinations (travels through the ATM link). It is the second type of traffic, the remote traffic, that needs to be rate controlled in the hosts. However, the hosts do not know whether a destination is local or remote. Therefore, some mechanism is necessary for the gateway to notify the host of the destinations for which it should control the rate.

With HG-RCP, a new ICMP message, referred to as queue notification (QN) message, is used by the gateway to notify a host of the destination network addresses for which the host should control the rate. In addition to knowing which destination network address is a remote one, the hosts also need to know the ATM link in the gateway (the gateway may be connected to several ATM links) over which the traffic to this remote address will travel. This is because the gateway manages different buffers and different rates for different ATM links. The QN message includes all the necessary information for the hosts as described below.

As mentioned earlier, each ABR connection in the gateway is identified by a unique Q_{id} . When an IP packet is received for which the host did not impose rate control (because the host did not know if the destination was a remote one), the gateway sends a QN message to the host. A QN message includes the destination network address of the IP packet and the Q_{id} of the ABR connection over which the IP packet is transmitted.

It is necessary for the gateway to generate an ICMP QN message when it receives an IP packet which was not rate controlled by the host (because the host was yet to determine if the destination was a remote one). To enable the gateway to determine if an IP packet was rate managed or not, the hosts assign a hitherto unallocated bit in the IP header, referred to as Q-bit in this paper, as described below.

Initially, the host is not able to determine whether the destination of an outgoing IP packet is local or remote and hence the IP packet is placed in the (default) non-rate-managed queue. The Q-bit of an IP packet transmitted from the non-rate-managed

queue is not modified (it remains zero). If the Q-bit of an IP packet transmitted over an ABR connection is not set, the gateway immediately sends a QN message to the originating host with the destination address and identification number of the ABR connection so that the subsequent IP packets to this address will be rate-managed according to the explicit rates for this ABR connection as fed back by the gateway via RN messages.

For IP packets with the Q-bit set, the gateway does not need to generate any QN message. The gateway clears the Q-bit before forwarding the IP packet to an ABR connection to preserve the original header of the IP packet, as the Q-bit only has relevance within the intranet.

2.2. Host behaviour

The rate management within hosts is performed at the IP layer. Existing IP software at the host typically maintains a single non-rate-managed queue where the application or transport processes deposit IP packets for transmission. The IP module waits for an IP packet to be transmitted through the LAN interface and processes the next IP packet from the queue as soon as the previous IP packet is transmitted.

With HG-RCP implemented as an enhancement of the IP layer, the IP layer at the host maintains a set of rate-managed queues (for ABR connections in the gateway) in addition to the default non-rate-managed queue. Transmission of IP packets from a rate-managed queue is controlled by the explicit rate feedback from the gateway via the ICMP RN message. Each rate-managed queue, corresponding to a particular ABR connection, is identified by the unique identifier Q_{id} in the ICMP QN message.

The queue management at the host IP is illustrated in Fig. 2. The host IP software maintains a Queue Mapping Table (QMT) to associate traffic going out of the intranet to the rate-managed queues, according to the destination of the traffic. Each entry in the table stores the ABR connection identifier for a destination network address. Traffic not destined to leave the intranet is handled via the default non-rate managed queue.

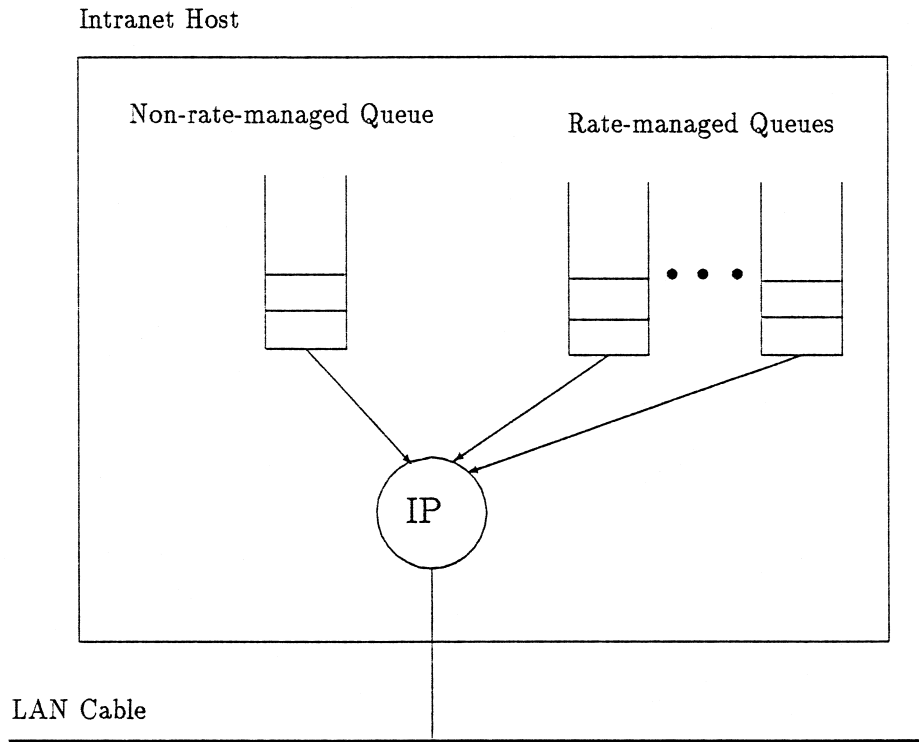


Fig. 2. Queue management in the host IP for HG-RCP.

When the host IP software receives an ICMP QN message, it extracts the destination network address and the ABR connection identifier Q_{id} from the message, updates the QMT with a new entry and moves any IP packets for that destination network address from the non-rate-managed queue to the rate-managed queue identified by the Q_{id} . Initially, the QMT is empty; it gradually builds up as new QN messages are received.

When an application in a host submits an IP packet to the IP handler, the destination network address of the IP packet is checked against the QMT. The IP packet is placed in the non-rate-managed queue if destination network address is not found in the QMT (at this point it is assumed that the IP packet will not leave the intranet), otherwise it is placed in the rate-managed queue identified by the Q_{id} extracted from the QMT and the Q-bit is set in the header.

The rate management for each individual rate-managed queue, which is implementation dependent,

can be achieved in several ways. A possible implementation is described in [11].

3. Analysis of HG-RCP

In this section, we present three different analyses. First, we show how to select the parameters K and T for ensuring stability of the control mechanism and at the same time minimise the buffer level variance around the desired setting Q . Then, we derive the buffer size required at the intranet-ATM gateway to ensure zero packet loss. Finally, we provide an analysis of the control traffic overhead for implementing HG-RCP. Some of these analyses can also be found in [12].

3.1. Stability analysis

In this section, we analyse the stability of the HG-RCP using classical feedback control theory and

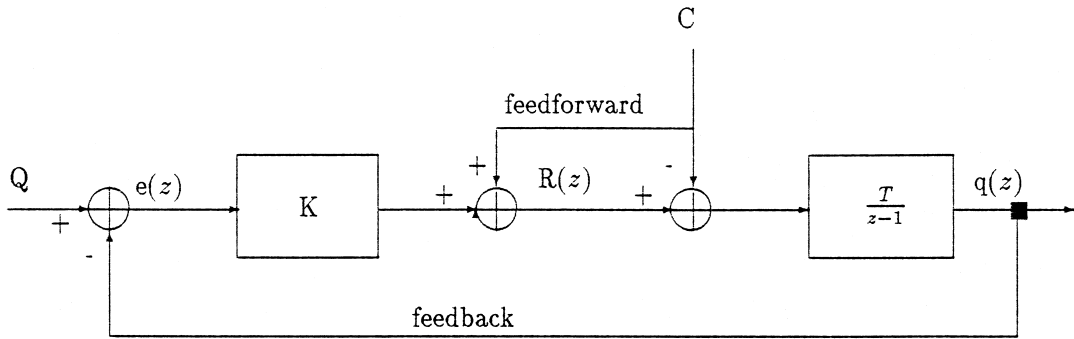


Fig. 3. Block diagram of the proportional feedback controller augmented with the noise feedforward.

derive a relationship between parameters K and T which will ensure stability of the control mechanism. The control mechanism of HG-RCP can be designed as a proportional feedback controller with noise feedforward as shown in Fig. 3. The current bandwidth C of the ABR connection, which is controlled by the ATM network, is modeled as noise in this system as it affects the buffer level and is beyond the control of HG-RCP. The goal of the control system is to maintain the queue length q (the output of the system) at a desired set point Q . The input to the queue is the aggregate traffic rate R from all the sources in the intranet. The input rate R is controlled by the proportional controller with gain K .

The relationship between the input rate R and the queue length q can be expressed as follows:

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

i.e.,

$$q(n+1) = q(n) + TR(n) - TC, \tag{2}$$

where $q(n+1)$ and $q(n)$ are the buffer levels measured at the $(n+1)$ th and n th sampling intervals respectively. Taking the z -transform of Eq. (2), the following transfer function between queue output q and queue input R is obtained as:

$$q(z) = \frac{T}{z-1} (R(z) - C). \tag{3}$$

The input rate R of the feedforward system is obtained as:

$$R(n) = C + K(Q - q(n)) = C + Ke(n), \tag{4}$$

where e is the error in the queue length. It can be seen from Eq. (4) that there is no error ($e = 0$) in the steady state (when $R = C$). Although R can assume negative values in Eq. (4), HG-RCP only considers zero or positive values as suggested in Eq. (1). By taking the z -transform of Eq. (4), the transfer function for the rate controller is obtained as:

$$R(z) = C + Ke(z). \tag{5}$$

We apply the *frequency-response* method [23] for analysing the stability of HG-RCP. With this method, the *phase margin* (PM) of the system determines the stability of the control system. To ensure stability, the phase margin has to be positive.

The phase margin of a control system can be determined by looking at the open-loop transfer function of the system which is obtained by multiplying all the block transfer functions in the loop. Thus the open-loop transfer function of Fig. 3 is obtained as $KT/(z-1)$. From the definition of PM [23],

$$PM = 180^\circ + \angle \frac{KT}{z-1}, \tag{6}$$

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

$$\left| \frac{KT}{z-1} \right| = 1. \tag{7}$$

From Eqs. (6) and (7), the phase margin can be calculated as:

$$PM = \tan^{-1} \frac{\sqrt{4 - (KT)^2}}{KT}. \quad (8)$$

Therefore, for a stable system ($PM \geq 0$), K should be selected smaller than or equal to $2/T$, where T is expressed in seconds. For a good stability, we consider a phase margin of 60 degrees, which gives a $K = 1/T$. Therefore, for a given sampling interval of T seconds, the stability of HG-RCP can be ensured by selecting K as the inverse of T .

3.2. Derivation of buffer level variance

Although the control system of Fig. 3 will keep the buffer level at the desired setting of Q in the steady state, sudden fluctuation in the available bandwidth of the access-link will cause momentary shifts of the buffer level from Q . Therefore, the variance σ_q^2 of the buffer level around the desired setting Q is an important performance factor 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 analyse the buffer level variance σ_q^2 as a function of K . This analysis will allow us to select a K which provides a low buffer level variance.

We need to obtain the transfer function between the noise source and the buffer level in order to calculate the variance in the buffer level. The available bitrate B can be modeled by the following difference equation:

$$B(n+1) = \alpha\beta(n) + (1-\alpha)\beta w(n), \quad (9)$$

where $W(n)$ is a random noise source and β is a constant to indicate the magnitude of the noise.⁴ For $\alpha = 1$, there is no noise and the available bitrate remains constant; for $\alpha = 0$, B is a pure random number. In practice, α will be a fraction between 0 and 1.

⁴ For an ATM-ABR access-link, a possible noise source may be a VBR (Variable Bit Rate) connection multiplexed with the ABR connection within the ATM network. The ABR rate is affected whenever the VBR rate is changed.

Taking z -transform of Eq. (9) we get

$$\frac{B(z)}{W(z)} = \frac{(1-\alpha)\beta}{z-\alpha}. \quad (10)$$

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

$$\begin{aligned} \frac{Q(z)}{B(z)} &= \frac{(z^{-1}-1)T/(z-1)}{1+KT/(z-1)} \\ &= \frac{T(1-z)}{z(z-(1-KT))}. \end{aligned} \quad (11)$$

Multiplying Eqs. (10) and (11) we obtain the transfer function between noise source W and buffer level Q as:

$$\frac{Q(z)}{W(z)} = \frac{(1-\alpha)\beta T(1-z)}{z(z-\alpha)(z-(1-KT))} = \frac{b(z)}{\alpha(z)}. \quad (12)$$

Applying variance calculation theorem [1] and Cauchy's residue theorem [18],

$$\begin{aligned} E[q^2] &= \sigma_q^2 = \frac{1}{2\pi j} \oint \frac{b(z)b(z^{-1})dz}{a(z)a(z^{-1})z} \\ &= \Sigma \left[\text{residues of } \frac{b(z)b(z^{-1})}{a(z)a(z^{-1})z} \right. \\ &\quad \left. \text{at poles inside unit disc} \right] \\ &= \frac{(1-\alpha)\beta^2 T^2 [(\alpha^2 - \alpha + 2)KT - 2(\alpha - 1)^2]}{(1+\alpha)(1-\alpha + \alpha KT)(\alpha - 1 + KT)(2 - KT)}. \end{aligned} \quad (13)$$

Eq. (13) provides the buffer level variance as a function of the controller gain K . The optimum value K is the one for which the expression at the right-hand of Eq. (13) is minimised. In deriving the optimum K , we make the following observations:

1. *Optimum K does not depend on the noise parameter β .* This results is obtained as follows. For a given value of α and T , we equate the first derivative of the right-hand side of Eq. (13) to zero and solve for K . We find that the solution for K does not contain the parameter β .
2. *Optimum K does not significantly depend on the noise parameter α .* For a range of practical values of T , the optimum K as a function of α is shown in Fig. 4. It can be seen, the optimum K




Fig. 4. Optimum K as a function of noise parameter α .

decreases only slightly for increasing α . The amount of decrease is negligible.

3. *Optimum K is inversely proportional to T and is roughly equal to $1/T$.* This result can be obtained from Fig. 4. For $K = 1/T$, the control system has a phase margin of 60° (from Eq. (8)), which guarantees a good stability for the control system.

The third observation above confirms that the proposed control mechanism is stable for the choice of gain K for which the buffer variance is minimised. This is a very significant result as it allows one to design the control system to minimise the buffer level variance without sacrificing the stability of the system or vice versa.

The first two of 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 which is required to minimise the buffer requirement at the access gateway.

From the above analyses, we suggest that K should be selected as $1/T$, where T is expressed in seconds, to guarantee a stable system and simultaneously minimise the buffer level variance around the desired setting Q . In the simulation experiments presented in Section 4, we have, therefore, used $K = 1/T$.

3.3. Buffer size for zero loss

In this section, we obtain the required buffer size at the gateway to guarantee zero loss at the intranet-

ATM gateway. We achieve this by deriving the maximum queue length for the worst possible (pathological) scenario.

Once the gateway broadcasts a zero explicit rate for a Q_{id} corresponding to a particular ABR connection, all the IP software at the hosts will stop transmitting from the rate-managed queues identified by the Q_{id} until the explicit rate for this Q_{id} is increased in a subsequent RN message. The only traffic that might still enter the queue in the gateway are the *first* IP packet of new connections whose remote destination network addresses were not found in the QMT. For a corporate wide intranet with limited number of remote offices, the amount of such traffic will be minimal and can be safely ignored.

The minimum queue length q_{min} that triggers a zero explicit rate in the gateway is obtained from Eq. (1) as:

$$q_{min} = \frac{PCR + K \times Q}{K} \text{ cells,}$$

where peak cell rate (PCR) is the maximum ABR bandwidth that is offered by the ATM network. The reason for considering the maximum bandwidth for the ABR link is to analyse the queue length in the worst possible situation.

If a sampling interval ends when $q \geq q_{min}$, the gateway will broadcast a zero explicit rate for this ABR connection and the queue length will stop increasing beyond q . However, if the sampling interval ends when q is just below q_{min} , a positive explicit rate (R/N) may be broadcast. If many new sources start transmission to this particular link (these sources were not taken into consideration in the calculation of the explicit rate), they will all transmit at the explicit rate broadcast on the intranet and the incoming traffic will be much greater than R . Since these new sources will not be identified as active ones until the end of the sampling interval, traffic may enter the queue in the gateway at the full speed of the LAN (the LAN in the intranet is connected to the gateway) in the worst case.

For a LAN speed of B bps, the maximum amount of traffic that may be received during one sampling interval of T seconds is BT bits or $BT/(8 \times 48)$ cells (an ATM cell has 48-byte payload). Therefore, the maximum size the queue can grow before trig-

gering a zero explicit rate (and hence stop increasing) is:

$$\begin{aligned} q_{\max} &= q_{\min} + \frac{BT}{8 \times 53} \\ &= \frac{PCR + K \times Q}{K} + \frac{BT}{8 \times 48} \text{ cells.} \end{aligned} \quad (14)$$

Eq. (14) suggests that with HG-RCP implemented in the intranet, the maximum queue length is bounded and it *does not depend on the number of sources*. Hence HG-RCP scales well with large intranets having many sources.

It should be emphasised that the maximum queue length q_{\max} derived in Eq. (14) is for pathological situations where the traffic enters the gateway at the maximum LAN speed even when the queue length is already above the desired level Q . With actual intranets, it may not be possible to receive traffic at full LAN speed due to reasons such as consumption of LAN bandwidth by LAN frame header, interframe gap, etc. Therefore, with actual networks, the maximum queue length may be slightly lower than the one predicted by Eq. (14).

3.4. Control traffic overhead

The control traffic overhead for HG-RCP is the extra traffic generated by the intranet-ATM gateway in transmitting ICMP RN and QN messages to the intranet hosts. In this section, we analyse the control traffic and derive an upperbound for the amount of such traffic.

HG-RCP is designed to minimise the amount of control traffic by (i) restricting the generation of RN messages to those intervals where the new explicit rates differ from the previous ones, (ii) generating a single RN message for all hosts instead of separate messages for individual hosts, and (iii) generating a QN message to a host for a particular destination network only once for the first IP packet.

In the worst case, the gateway broadcasts one ICMP RN message at each interval of T seconds. Therefore, the upperbound U_m on the amount of traffic due to RN messages is obtained as:

$$U_m = \frac{\text{length of an RN message in bits}}{T} \text{ bps.} \quad (15)$$

Since an RN message only contains a floating point number (representing the explicit rate for the contending intranet sources) for a given ABR connection in the gateway, the length of an RN message is quite small. A possible format is suggested in [11] where an ICMP RN message has a length of 8 bytes (for one ABR queue at the gateway) resulting in a 28-byte IP packet (the ICMP message is encapsulated in an IP packet which has a 20-byte header). Therefore, for an intranet with a 10-Mbps LAN, the overhead is less than 0.5% of the total LAN bandwidth for a sampling interval of 10 ms. The percentage overhead decreases as the sampling interval or the LAN speed increases.

The percentage overhead in an intranet due to QN messages approaches zero in the long run as the QMT tables in the host IPs are updated with the addresses of the remote intranets. Once the QMT tables in the hosts contain the addresses of all remote intranets of a corporation, the intranet-ATM gateway does not send any QN message. Therefore, the total control traffic overhead of HG-RCP is mainly due to RN messages, which is only a tiny fraction of the total LAN bandwidth as shown above.

4. Simulation

In order to validate the results (stability of the control system and the buffer size required in the gateway for zero packet loss) obtained through analysis in the previous section, we simulate a detailed implementation of HG-RCP in a corporate wide intranet. In addition to the above validation objectives, we also wish to investigate the fairness of HG-RCP in distributing the ABR bandwidth among the competing sources on the intranet, and study the sensitivity of HG-RCP to various control parameters.

The buffer size requirement for zero loss can be obtained by measuring the maximum queue length reached in the simulation, and the stability of the control mechanism can be assessed by observing the queue length as a function of time. The fairness of HG-RCP can be judged by comparing the average goodput of individual sources at the end of the simulation. Sensitivity analysis for a particular control parameter can be obtained by observing the effect of change in a parameter value on the maxi-

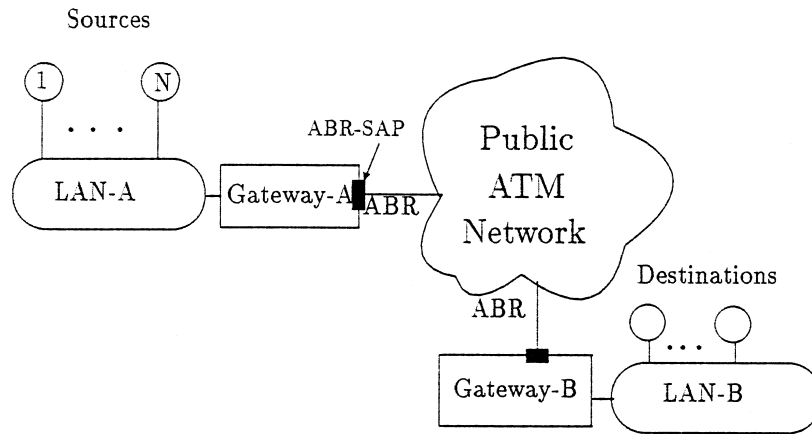


Fig. 5. Network configuration for simulation experiments.

num and average queue lengths and the goodput of the sources.

The simulation model of a corporate wide intranet achieved by interconnecting two remote offices over the ATM ABR service is shown in Fig. 5. There are N hosts attached to LAN-A with a single TCP source in each host; the destinations of all TCP sources are in LAN-B. The hosts are attached to the constant-rate LAN medium while the gateway is connected to both the LAN and the variable-rate ABR service of the public ATM network. The start of each TCP connection is staggered by 50 ms to minimise any possible synchronisation effects. Detailed descriptions for the TCP and ABR implementations in the simulation program are presented in the following sections.

4.1. TCP

TCP sources in LAN-A transfer *large files* (files are large enough to have data ready to be transmitted during the entire simulation period) to TCP destinations in LAN-B. Use of such large files allows us to study the performance of the TCP algorithm in the steady-state and has been used previously by other researchers to evaluate the performance of TCP over ATM networks [5,19,6,22]. It should be noted, however, that although the TCP applications have infinite data to transmit, the intranet sources *may not appear infinite* because of the TCP's flow control.

The values of the TCP parameters used in the simulation are shown in Table 1. Timeout gain fac-

tor, deviation gain factor and timeout deviation coefficient are three parameters used to update the retransmission timeout according to the algorithm described in [2].

Only one way traffic, from sources in LAN-A to destinations in LAN-B, is simulated. The only traffic that travels in the reverse path are the ACKs generated by the destination TCPs. Since a TCP generates very short ACKs, these ACKs are not explicitly simulated (i.e., their processing time and consumption of LAN bandwidth are ignored).

The *slow-start* and *congestion-avoidance* algorithms described in [13] are implemented in the TCP simulation. The ACK-delay timer at the TCP destination is not implemented as there is no data traffic simulated in the return path. An ACK is generated by the destination immediately upon receiving a TCP segment.

Table 1
Values of TCP parameters

Parameter	Value
Maximum segment size	1460 octets
Maximum window size	44 segments
Header size	20 octets
IP header size	20 octets
Minimum retransmit-timeout	0.3 s
Maximum retransmit-timeout	10 s
Maximum retransmission retry	16
Timeout gain factor	0.125
Deviation gain factor	0.25
Timeout deviation coefficient	4.0

In order to make a fair comparison between HG-RCP and conventional TCP/IP, we simulated the best possible implementation of the conventional TCP over ATM (to achieve the highest performance) by, (i) simulating a very fine-grain TCP-clock and, (ii) a packet-based discarding at the gateway. Coarser clock-granularity results in longer timeout interval in the TCP resulting in longer delay in detecting the loss of an IP packet at the gateway, which in turn reduces TCP goodput [22]. With packet-based discarding, the entire IP packet is discarded at the gateway if there is insufficient buffer space to hold all the cells of that packet. This avoids the waste of bandwidth and loss of throughput associated with transmitting cells from *partial* packets [22].

4.2. ABR

Gateway-A is simulated as the source end system (SES) and Gateway-B as the destination end system (DES) of the ABR connection using the algorithms recommended by the ATM Forum [4]. Since only one way traffic (from Gateway-A to Gateway-B) is simulated, the SES does not generate any backward RM cells and the DES merely turns around any RM cell it receives. It is assumed that no cells are lost within the ATM network. This is a valid assumption, as the cell loss for ABR connections are expected to be minimal.

The ABR parameter values used in the simulation are shown in Table 2. A PCR peak cell rate (PCR) of 23 585 cells/s corresponds to a maximum bandwidth of 10 Mbps (equal to the LAN speed). One per cent of the PCR (235 cells/s) is guaranteed by the ATM network as minimum cell rate (MCR).

All the ATM switches along the path of the ABR connection are assumed to be explicit rate (ER)

capable, i.e., the switches can explicitly convey the maximum allowed rate to the SES via the ER field in the RM cells. With ER switches, it is safe for the SES to update the ACR by simply replacing it with the ER value in the RM cell returned by the ATM network (hence $RIF = 1.0$).⁵ This allows the gateways to increase the rate as fast as allowed by the ABR flow-control.

Since HG-RCP is a flow control scheme primarily designed to increase the throughput of intranet applications by controlling the buffer level at the gateway during congestion, it is appropriate that it be put to test against the worst case of congestion that might occur in the gateway. If satisfactory results are achieved for such a worst case, it can be concluded with confidence that the scheme will perform well against other cases. This *worst-case testing* principle has been adopted in the past by other researchers to evaluate the performance of various flow control schemes [3,8,21,22].

The worst case congestion in Gateway-A due to dynamic update of ABR bandwidth occurs when the bandwidth is reduced from one extreme value, PCR, to another, MCR. This worst-case bandwidth variance is simulated by introducing a background VBR source as described below.

The ABR connection is multiplexed with a VBR connection. The VBR traffic that enters the ATM network (and the ATM switches) is modeled using the well-known ON-OFF model with exponential ON (active) and OFF (silent) periods. The VBR source transmits cells at the peak rate during the ON-period and completely stops transmitting during the OFF-period. As VBR has a higher priority than ABR, any bandwidth not consumed by the VBR traffic is made available to the ABR users.

The peak rate of the VBR source is close to the link capacity.⁶ Therefore, when the VBR source is

Table 2
Values of ABR parameters

Parameter	Value
PCR	23585 cells/s
MCR	350 cells/s
ICR	500 cells/s
Nrm	32
RIF	1.0
ADTF	500 ms

⁵ The other type of switches are called EFCI or 'binary' switches. These switches only use a single EFCI bit in the cell header to convey the presence or absence of congestion in the switch. For EFCI switches, the SES should increase the rate more conservatively ($RIF < 1.0$) to avoid unnecessary oscillations [15].

⁶ Other researchers [16] have also simulated such a single VBR source with a peak bandwidth close to the link capacity to introduce extreme bandwidth variability in a congested link shared by ABR connections.

Table 3
Values of some other simulation parameters

Parameter	Value
LAN bandwidth	10 Mbps
LAN frame format	Ethernet
ATM adaptation layer	AAL5
Round trip delay in ATM network	60 ms
mean ON-period for background VBR source	100 ms
mean OFF-period for background VBR source	400 ms

active, there is very little bandwidth available for the ABR connection and the ABR flow-control allows an MCR to the SES (Gateway-A). When the VBR source is silent, there is ample available bandwidth and the ABR connection is allowed to transmit at the PCR. The mean lengths of the ON and OFF periods of the VBR source along with some other simulation parameters are shown in Table 3.

Each host IP maintains a rate-managed queue for the ABR connection at the gateway. It is assumed that the QMTs at the host IPs contain the address of the destination network (LAN-B) from the start of the simulation and hence the gateway does not generate any ICMP QN messages during the simulation.

In each simulation run, the system went through a warm up (initial transient) period during which the observations were discarded (these observations were not used in the calculation of the final output from the simulation). The simulation length was automatically controlled by an on-line monitor which stopped the simulation when 95% confidence level was achieved with 5% relative precision for measuring the aggregate goodput for all the traffic sources on the intranet.

5. Results

In this section we compare the results obtained from simulation experiments for the conventional TCP/IP, against the TCP/IP augmented with HG-RCP. As mentioned in Section 4, for fair comparison, we have simulated a very fine-grain TCP clock at the source, and used the *packet discard* (instead of *cell discard*) strategy at the gateway which gives the best possible performance with conventional TCP/IP.




Fig. 6. Queue length as a function of time for conventional TCP/IP and TCP/IP with HG-RCP for unlimited buffers.

We analyse simulation results for both unlimited and limited buffers at the gateway. Simulations with unlimited buffers allow us to observe the queue length as a function of time and the maximum queue length (as there is no packet loss) which helps verifying the stability of the control mechanism and dimensioning the buffer at the gateway for zero loss guarantee. We also measure the average queueing length and average delay at the gateway which are important performance measures for delay sensitive traffic. In the absence of packet loss at the gateway for unlimited buffers, we also present a sensitivity analysis for some parameters of HG-RCP. Finally, the effect of non-zero intranet propagation delay on the buffer size requirements for zero loss is observed by measuring the maximum queue length.

For limited buffers, which may cause packet loss at the gateway, we compare the goodput and fairness

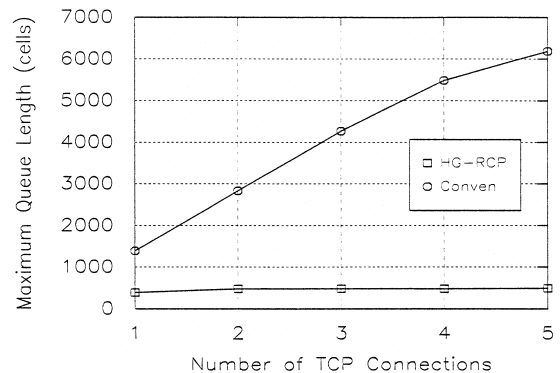


Fig. 7. Maximum queue length for conventional TCP/IP and TCP/IP with HG-RCP ($T = 10$ ms) for unlimited buffers.

Table 4
Maximum buffer size for varying sampling interval T

Sampling period T (ms)	K	Maximum buffer size (cells)	
		Simulation	Theoretical bound
100	10	3896	5062
50	20	1880	2581
20	50	835	1092
10	100	489	596

(in distributing the ABR bandwidth among all the contending sources on the intranet) of conventional TCP/IP and TCP/IP with HG-RCP.

5.1. Unlimited buffers

5.1.1. Stability

For unlimited buffers, we simulated five connections for conventional TCP and TCP with HG-RCP to represent a typical scenario where multiple sources in the intranet are contending for the ABR bandwidth at the gateway. Initially a zero intranet propagation delay, i.e., the propagation delay between the gateway and the hosts, was assumed as all the hosts within an organisation are likely to be located within a relatively short range. The queue length was recorded at the arrival of an IP packet at the ABR-SAP in the gateway. The queue length against time for conventional TCP/IP and TCP/IP with HG-RCP are shown in Fig. 6 for the first 5 s of the simulation.

As can be seen from Fig. 6, the queue length for the conventional TCP/IP increases beyond 3000 cells. These long spikes demonstrate the failure of TCP's window-based flow control in effectively controlling the source rate and hence providing a stable control system for the intranet. With HG-RCP implemented, the spikes in the queue length are reduced by orders of magnitude; the queue length stabilises around 100 cells ($Q = 100$ in the simulation). The small spikes for the HG-RCP are due to a sudden drop of ACR at the gateway and are quickly brought under control by controlling the source rates. Hence, HG-RCP proves to be a stable system.

5.1.2. Buffer size for zero loss

Next, we examine the maximum buffer size required to guarantee a zero loss in the gateway by running many simulations with an increasing number

of simultaneous TCP connections in successive simulations. Fig. 7 plots the maximum queue length for the varying number of TCP connections. It can be clearly seen that without HG-RCP, the buffer size required for zero loss increases almost linearly as a function of the number of contending sources. However, with HG-RCP implemented, the maximum buffer size is independent of the number of TCP connections and remains constant below the worst case value as predicted by our analysis in Eq. (14).

Although the maximum queue length does not depend on the number of simultaneous TCP connections, Eq. (14) suggests that it is a function of T and K . Table 4 compares the maximum queue lengths observed in the simulations with those obtained from Eq. (14) for a varying T .⁷ As predicted by the analysis, the maximum queue length observed in the simulations decreases for decreasing T and remains below the worst case values obtained from Eq. (14).

5.1.3. Delay

The average queue length and the average queuing delay at the gateway are important performance criteria for interactive, delay-sensitive application e.g., TELNET. The average queue lengths and the average delays as a function of the number of simultaneous TCP connections are shown in Figs. 8 and 9. As was the case with the maximum queue lengths, the average queue lengths and the average queuing delays remain constant with HG-RCP implemented whereas they increase almost linearly with the conventional TCP/IP.

⁷ In each simulation, the gain K was selected as $1/T$ (T in seconds) to obtain a 60 degrees phase margin for a good stability of the control system.

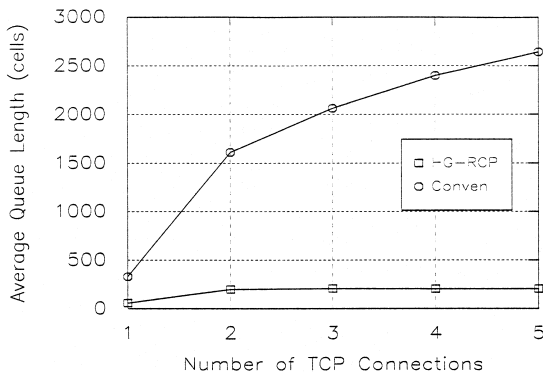


Fig. 8. Average queue length for conventional TCP/IP and TCP/IP with HG-RCP ($T = 10$ ms) for unlimited buffers.

5.1.4. Sensitivity analysis

It is important to know the sensitivity of a control mechanism to the values selected for the various parameters. If the control system is too sensitive to the parameter values, it will be absolutely necessary to select the correct values; the performance will degrade otherwise. It is, therefore, desirable to design a control mechanism which can effectively perform over a wide range of parameter values. In this section, the sensitivity of HG-RCP to various parameters are assessed using the results obtained from the simulations experiments.

As described in Section 2, HG-RCP has four parameters: T , K , Q and W . The effect of the first two parameters are easily obtained through analysis. The selection of T directly affects the buffer size required for zero loss in the gateway as given by Eq. (14). The smaller the T the better the performance (as confirmed by the simulation results in Table 4) at the expense of increased control traffic overhead on the intranet. For a given T , the value of K is selected using Eq. (8) to ensure stability ($K = 1/T$ provides a good stability) of the control mechanism.

It is not easy to analytically obtain the sensitivity of HG-RCP for the parameters Q and W . We, therefore, analyse the sensitivity of Q and W using simulation results as shown below. Table 5 shows the maximum and the average queue lengths and the aggregate TCP goodputs for increasing values of Q . Although the queue length increases for increasing Q (as expected), there is no significant change in the TCP goodputs. This suggests that a small value of Q

can safely be selected without significantly affecting the performance of the TCP connections.

The performance of HG-RCP for different values of W is shown in Table 6. It can be seen that the performance of HG-RCP is not too sensitive to the values selected for W . This is because any error made in calculating the number of active sources (and subsequently the explicit rate to the sources) is corrected by the new calculations in the next interval.

5.1.5. Effect of intranet propagation delay

The intranet propagation delay, i.e., the delay between the gateway and the intranet hosts, was assumed to be zero (negligible) in the analysis in Section 3 and the simulations results presented so far. To investigate the effect of propagation delay on the buffer size requirement for zero loss, in this section, we present results from simulations with non-negligible intranet propagation delays.

For intranets of reasonable sizes, Table 7 compares the maximum queue lengths for a range of reasonable propagation delays with the one obtained for zero propagation delay. It can be seen that the maximum queue length increases for longer propagation delays which is intuitively expected as it takes longer for the sources to receive the rate feedback from the gateway. However, the increase in the maximum queue lengths remains within 10% of the queue length for zero delay.

It is worth emphasising here that an intranet has a limited size (remains within a local site) and hence

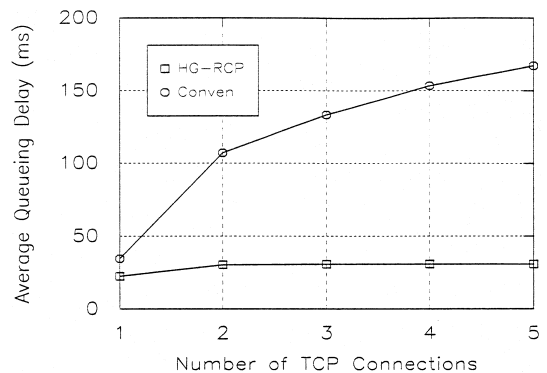


Fig. 9. Average queuing delay in the gateway for conventional TCP/IP and TCP/IP with HG-RCP ($T = 10$ ms) for unlimited buffers.

Table 5

Performance of HG-RCP with different values of the reference queue length Q for 5 TCP connections and unlimited buffers ($T = 10$ ms, $K = 100$, $W = 150$ ms)

Q (cell)	Maximum queue length (cell)	Average queue length (cell)	TCP goodput (Mbps)
20	407	129.64	6.4092
50	437	156.37	6.3958
100	489	205.11	6.4304
200	595	304.68	6.4321
300	676	403.47	6.4315
400	785	504.28	6.4314
500	893	604.01	6.4317

Table 6

Performance of HG-RCP with different values of W for 5 TCP connections and unlimited buffers ($T = 10$ ms, $K = 100$, $Q = 100$)

W (ms)	Maximum queue length (cell)	Average queue length (cell)	TCP goodput (Mbps)
50	486	206.23	6.4139
100	480	205.08	6.4304
150	489	205.11	6.4304
200	489	205.11	6.4304
500	489	204.72	6.4301

the propagation delays between the gateway and the hosts are not expected to be very large. HG-RCP, therefore, is expected to perform well for most intranets.

5.2. Limited buffers

In reality, the allocation of buffer space for each queue (for each ABR connection) in the gateway will be limited. It is, therefore, necessary to evaluate the performance of our proposed flow control scheme when IP packets are discarded in the gateway due to lack of buffer space. Goodput and fairness achieved from simulation for limited buffers are presented below.

Table 7

Maximum queue lengths for increasing intranet propagation delays for 5 TCP connections with unlimited buffers ($T = 10$ ms, $K = 100$, $Q = 100$, $W = 150$ ms)

Propagation delay (ms)	Maximum queue length (cell)
0.0	489
0.1	501
0.2	505
0.5	536

5.2.1. Goodput

For limited buffers, we increase the load at the gateway by increasing the number of simultaneous TCP connections from five (as was used in the previous section) to eight in order to increase the congestion at the gateway which will result in a higher probability of packet loss due to buffer overflow. The buffer size was varied to observe the TCP goodput as a function of the buffer size in the gateway. It can be seen in Fig. 10 that it takes much




Fig. 10. Goodput for 8 TCP connections as a function of buffer size with conventional TCP/IP and TCP/IP with HG-RCP ($T = 10$ ms, $K = 100$, $Q = 100$ cells, $W = 150$ ms).

Table 8
Comparison of fairness between conventional TCP/IP and TCP/IP with HG-RCP

Buffer size (cells)	Fairness index	
	Conventional	HG-RCP
800	0.86	0.87
1200	0.80	0.87
1600	0.77	0.87
2000	0.57	0.87
2400	0.52	0.87

larger buffers for the conventional TCP/IP to achieve TCP goodputs comparable to those obtained from TCP/IP with HG-RCP. Thus for the same goodput, HG-RCP reduces the required amount of buffer at the gateway compared to the conventional TCP/IP and for the same amount of buffer, it increases TCP goodput.

It should be noted here that Fig. 10 shows the best possible TCP goodput achievable with the conventional TCP/IP. In fact, the goodput would be lower than those in Fig. 10 if, (i) coarser clock-granularity was implemented in the TCP sources, and (ii) cell-based discarding, instead of packet-based discarding, was implemented at the gateway.

5.2.2. Fairness

Fairness is considered as one of the important performance measures for a flow control scheme. For fairness, we consider the fairness index proposed by Jain [14] which is calculated as:

$$\text{Fairness} = \frac{(\sum x_i)^2}{n \sum x_i^2}, \quad (16)$$

where n is the number of competing sources and x_i is the ratio of actual throughput to optimal throughput. The closer the ratio is to 1, the fairer it is, with 1 indicating perfectly fair sharing of the bottleneck bandwidth.

Since in our case all the sources have the same bottleneck (the ATM ABR connection at the gateway), equal share of the bottleneck bandwidth is considered the optimum allocation and hence the optimal throughput for a source is obtained as the total throughput achieved by all the sources divided by the number of sources. In this case, therefore, the optimal throughput is the same for all the sources.

Using Eq. (16), we obtain the fairness index for conventional TCP/IP and TCP/IP with HG-RCP. Table 8 shows the fairness achieved for different buffer sizes. It is interesting to note that although the conventional TCP/IP succeeds to increase the total goodput as the buffer size is increased (as shown in Fig. 10), it only does so at the price of decreasing the fairness. The fairness achieved by TCP/IP with HG-RCP is close to 90% which is much better than the ones obtained by conventional TCP/IP.

6. Conclusion

We have proposed a novel rate-based flow control mechanism, HG-RCP, for intranets interconnected by the ATM ABR service. The proposed flow control mechanism enhances the standard IP to regulate the traffic rate of the intranet hosts by providing explicit rate feedback from the intranet-ATM gateway to the intranet hosts. HG-RCP provides an effective flow control for the entire range of intranet applications including the ones based on TCP/IP. We have compared TCP/IP with HG-RCP against the best possible implementation of the conventional TCP/IP. We have carried out analysis and simulation to show that TCP/IP with HG-RCP performs significantly better than the conventional TCP/IP.

Through analysis, we have shown how to select the parameters to ensure stability of the proposed flow control mechanism and minimise the buffer level fluctuation at the intranet-ATM gateway. Simulation results have confirmed that HG-RCP is a stable flow control mechanism.

We have analysed the performance sensitivity of HG-RCP to the values of the various parameters. Results suggest that HG-RCP is not too sensitive to

the parameters which allows the network operators to select the parameter values from a wide range without significantly affecting the performance of HG-RCP.

We have shown that TCP/IP with HG-RCP has a number of advantages as compared to the conventional TCP/IP. Firstly, TCP/IP with HG-RCP requires smaller buffers at the intranet-ATM gateway to guarantee zero loss resulting in more cost-effective gateway design. Secondly, for the same amount of buffer at the gateway, TCP/IP with HG-RCP achieves higher goodput than conventional TCP/IP. Thirdly, it has been shown that TCP/IP with HG-RCP is also fairer than conventional TCP/IP in evenly distributing the available bandwidth of the ABR link among the contending TCP sources on the intranet.

References

- [1] K.J. Astrom, B. Wittenmark, *Computer Controlled Systems*, 2nd ed., Prentice-Hall, Englewood Cliffs, NJ, 1990.
- [2] D.E. Comer, D.L. Stevens, *Internetworking with TCP/IP, Vol. II: Design, Implementation and Internals*, 2nd ed., Prentice-Hall, Englewood Cliffs, NJ, 1994.
- [3] D.E. Comer, R. Yavatkar, A rate-based congestion avoidance and control scheme for packet switched networks, Technical Report 152-89, Department of Computer Science, Purdue University, West Lafayette, IN 47907, October 1989.
- [4] The ATM Forum Technical Committee, *Traffic Management Specification, Version 4.0.*, ATM Forum, April 1996.
- [5] C. Fang, H. Chen, J. Huchins, A simulation study of TCP performance in ATM networks, in: *Proc. GLOBECOM'94*, 1994, vol. 2, pp. 1217–1223.
- [6] C. Fang, A. Lin, On TCP Performance of UBR with EPD and UBR-EPD with a Fair Buffer Allocation Scheme, *ATM Forum Contribution 95-1645*, December 1995.
- [7] G.G. Finn, A connectionless congestion control algorithm, *ACM Comp. Commun. Rev.* 19 (5) (1989) 12–31.
- [8] S. Floyd, TCP and explicit congestion notification, *ACM Comp. Commun. Rev.* 24 (5) (1995) 8–23.
- [9] S. Floyd, V. Jacobson, Random early detection gateways for congestion avoidance, *IEEE/ACM Trans. Network.* 1 (4) (1993) 397–413.
- [10] M. Hassan, Impact of variable bandwidth on the performance of data communications over ATM-ABR based LAN interconnection, in: *Proc. 20th Australasian Computer Science Conference*, Sydney, Australia, February 1997, pp. 422–429.
- [11] M. Hassan, Performance of data communications over asynchronous transfer mode networks, Ph.D. thesis, Monash University, Melbourne, Vic. 3145, Australia, 1997.
- [12] M. Hassan, H. Sirisena, M. Atiquzzaman, A traffic management mechanism for intranets with available bit rate access to the Internet, in: *Proc. SPIE Conference on Performance and Control of Network Systems*, Dallas, TX, November 3–5, 1997, pp. 470–478.
- [13] V. Jacobson, Congestion avoidance and control, in: *Proc. ACM SIGCOMM*, Stanford, August 1988.
- [14] R. Jain, Fairness: How to Measure Quantitatively?, *ATM Forum Contribution 94-881*, September 1994.
- [15] R. Jain, S. Kalyanaraman, S. Fahmy, R. Goyal, S. Kim, Source behaviour for ATM ABR traffic management: an explanation, *IEEE Commun. Mag.* 34 (11) (1996) 50–57.
- [16] 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.
- [17] S. Keshav, A control-theoretic approach to flow control, *ACM Comp. Commun. Rev.* 21 (4) (1991) 3–15.
- [18] E. Kreyszig, *Advanced Engineering Mathematics*, 4th ed., Wiley, New York, 1979.
- [19] M. Perloff, K. Reiss, Improvements to TCP performance in high-speed ATM networks, *Commun. ACM* 38 (2) (1995) 90–100.
- [20] W. Prue, J. Postel, Something a Host Could Do with Source Quench: The Source Quench Introduced Delay (SQUID), *RFC 1016*, July 1987.
- [21] K.K. Ramakrishnan, R. Jain, A binary feedback scheme for congestion avoidance in computer networks, *ACM Trans. Comp. Syst.* 8 (2) (1990) 158–181.
- [22] A. Romanow, S. Floyd, Dynamics of TCP traffic over ATM networks, *IEEE J. Sel. Areas Commun.* 13 (4) (1995) 633–641.
- [23] Stefani, Savant, Shahian and Hostetter, *Design of Feedback Control Systems*, Saunders College Publishing, 1994.
- [24] N. Yin, S. Jagannath, End-to-End Traffic Management in IP/ATM Internetworks, *ATM Forum Contribution 96-1406*, October 1996.



Mahbub Hassan is Senior Lecturer in the School of Computer Science and Software Engineering, Monash University, Melbourne, Australia. He received a Ph.D. from Monash University and M.Sc. in Computer Science from University of Victoria, British Columbia, Canada. He has organized and chaired special sessions on a variety of topics in the area of ATM and IP networks in several international conferences. He serves in the Editorial Board of *IEEE*

Communications Magazine and has guest edited for *Real Time Imaging Journal* and *Journal of Supercomputing*. He has published many technical articles on TCP/IP and ATM networks. He is currently authoring a book titled 'Performance of TCP/IP over ATM Networks'. In addition to conducting teaching and research at Monash University, Dr. Hassan also serves as consultant to external bodies on projects involving ATM and high speed networks.



Jim Breen, who received the B.Sc. and MBA degrees from the University of Melbourne in 1970 and 1980 respectively, is an Associate Professor in the School of Computer Science and Software Engineering at Monash University, where he was formerly chairman of the Digital Systems department, and is now coordinator of graduate programs in digital communications. Prior to joining Monash University in 1985, he worked in the telecommunications and defence

industries as a networking specialist and consultant. Professor Breen's research activities include network protocol design and performance, telecommunications policies, and computer lexicography.



Mohammed Atiquzzaman received his M.Sc. and Ph.D. degrees in Electrical Engineering and Electronics from the University of Manchester, Institute of Science and Technology, UK. Currently he is a faculty member in the Department of Electrical & Computer Engineering at the University of Dayton, Ohio. He is a senior editor of the IEEE Communications Magazine and serves on the editorial boards of the Computer Communications journal and the

Telecommunication Systems journal. He has guest edited many special issues of various journals, including a feature topic on Switching and Traffic Management for Multimedia in the IEEE Communications Magazine, special issues on ATM Switching and ATM Networks of the International Journal of Computer Systems Science & Engineering, Next Generation Internet in the European Transactions on Telecommunications. He has also served in the technical program committee of many national and international conferences, including IEEE Infocom and IEEE Globecom. His current research interests are in Broadband ISDN and ATM networks, Internet Quality of Service, Integrated Services, multimedia over high speed networks, and ATM switching (see <http://www.engr.udayton.edu/faculty/matiquzz/>). He has over 100 refereed publications. He can be contacted at atiqu@ieee.org.