Study of an Adaptive Rate Control Scheme under Unequal Propagation Delays *

Mohamed Abdelaziz† Ioannis Stavrakakis ‡†

†Dept. of Electrical Engineering and Computer Science
University of Vermont
Burlington, VT 05405

‡Dept. of Electrical and Computer Engineering
Northeastern University
Boston, MA 02115

Abstract

In this paper, a queueing model is developed and analyzed to evaluate the performance of a number of remote and local connections subject to a simple two-level (Binary) adaptive rate control scheme. An efficient technique is developed for the computation of the individual cell loss probabilities of the local and remote connections; the dimensionality of the formulated model is independent of the associated propagation delays. Numerical results quantify the bias against remote connections of buffer management schemes which are fair with respect to connections with equal propagation delays.

1 Introduction

Preventive congestion control has been widely considered for the congestion management of Asynchronous Transfer Mode (ATM) networks (see [1] and references therein). Recently, adaptive as well as hybrid rate control schemes have been proposed ([2], [3]) in an attempt to improve network utilization by dynamically allocating unused network bandwidth. The ATM Forum has defined a new service denoted as Available Bit Rate (ABR), which allows the delay tolerant traffic sources to dynamically share the excess network capacity through the use of adaptive rate control schemes.

As opposed to open-loop schemes where no competition for network bandwidth takes place once a call is set-up, a fundamental characteristic of adaptive rate control schemes is that existing connections compete, at the cell level, for the available bandwidth. This competition necessitates the development of adaptive rate control mechanisms and buffer management policies that will guarantee fairness. Fairness is defined in this work in the sense that equal performance measures are delivered to the different connections within the same class. The issue of fair allocation of resources under adaptive rate control schemes is considered in [4] in the presence of negligible propagation delays and in [5] where all controlled sources have the same propagation delays. In ATM networking environments where geographically dispersed areas are spanned by the supported services, it is not an uncommon situation that sources located at significantly different physical distances will share the same network access node. Since the effectiveness of adaptive rate control schemes is known to be affected by the propagation delays, it is expected that some unfairness may arise in such a situation. In order to quantify the resulting unfairness, it is necessary to be able to isolate the performance of individual sessions subject to adaptive rate control. In this paper, a queuing model is developed for the performance evaluation of adaptively controlled connections with significantly different propagation delays. The queueing analysis presented in section 3 provides for the computation of per-session performance measures under the buffer management policies considered in section 4. Finally, numerical results are presented in section 5.

2 Adaptive Rate Control

The model considered in this work is shown in Fig. 1 where a number of sources share a common network access node. The access node is modeled by a finite buffer receiving deterministic service. Sources are divided into two groups: a group of n_l local sources which has a round-trip propagation delay of T_l and a group of n_r remote sources which has a round-trip propagation delay of T_r.

Under an adaptive rate control mechanism, a traffic de-

*Research supported by the Advanced Research Project Agency under Grant F49620-93-1-0564 monitored by the Air Force Office of Scientific Research (AFOSR).

†Ioannis Stavrakakis is on leave from the University of Vermont

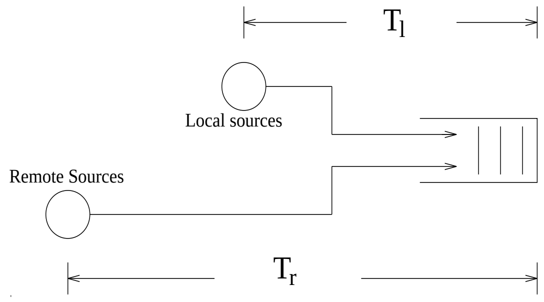

Figure 1: Network access model

scriptor to be modulated is identified (for instance, the arrival rate) and sources adjust its value dynamically based on either explicit or implicit information about the state of the network-premises. The network state is determined in terms of some pre-identified state descriptor such as the node load or buffer occupancy. By comparing the state descriptor to a threshold value, the network state is determined to be in an under-load state or overload state.

The network state is reviewed periodically over intervals of length T_u, referred to as update intervals. Either the same or different update intervals may be used to update the sources sharing the network node. At the end of each update interval, the network state is communicated to the traffic sources. Individual sources receive the network state information after a time interval equal to their respective propagation delays and respond by modulating the pre-identified traffic descriptor. The modulated traffic is received at the network node a round-trip propagation delay later after the network state information is sent.

In this work, sources modulate their (instantaneous) transmission rates based on the state of the network buffer. A two-level (Binary) adaptation scheme is considered, where sources switch to a high arrival rate (λ_h) if the network is under-loaded and switch to a low arrival rate (λ_l) if the network is overloaded. Both equal and unequal update intervals are considered. The two level scheme is chosen for tractability, and to help gain insight into the associated issues of fairness.

3 Performance Analysis

In this section, the queueing analysis is carried out in discrete-time. Time is slotted and a slot duration is equal to the transmission time of one cell. The length of the update interval of the local (remote) sources T_u^l (T_u^r) is taken to be equal to the round-trip propagation delay of the local (remote) sources $T_l(T_r)$. With minor changes, the presented analysis can be used when both the local and remote sources are updated at the same update interval ($T_u = T_r$). The propagation delay of the remote sources T_r is considered to be an integer multiple of that of the local sources, that is, $T_r = M \times T_l, M \in Z^+$, where Z^+ is the set of positive integers. Time is divided into frames each of length T_r slots; each frame is divided into M subframes, each of length T_l slots. Cell arrivals from both source groups arrive at a finite size network buffer of capacity Q and are served according to the FIFO service policy. If the number of cell arrivals over a slot exceeds the available buffer space, the buffer management policy determines which cells to be admitted to the queue. Several buffer management policies are considered in section 4.

Arrivals from the local (remote) sources are represented by a two state aggregate process $\{r^l\}, r^l \in \{0,1\}$ ($\{r^r\}, r^r \in \{0,1\}$), that will be referred to as the local (remote) arrival process. Let $f_i^l(k)$ ($f_i^r(k)$) denote the probability mass function (PMF) for the number of cells generated in each slot when the local (remote) process is in state $i, i \in \{0,1\}$. N_i will denote the maximum number of cell arrivals per slot from state i. State transitions for the local and remote arrival processes are governed by the two-level adaptive rate control mechanism as determined by the state of the network-state descriptor.

Let $\{q_k\}_{k\geq 0}$ be the buffer occupancy process embedded at subframe boundaries with state space $S^Q = \{0,1,\cdots,Q\}$. Let q_{th} be the buffer threshold value indicating the network state as overloaded ($q \geq q_{th}$) or underloaded ($q < q_{th}$). Let $\{r_k^l\}_{k\geq 0}$ ($\{r_k^r\}_{k\geq 0}$) be the local (remote) arrival process embedded at the subframe boundaries with state space $S = \{0,1\}$. Notice that due to the difference in propagation delays the state of process $\{r_k^l\}_{k\geq 0}$ at the beginning of subframe k is determined by the state of the buffer process at the beginning of subframe $k-1$, while the state of process $\{r^r\}_{k\geq 0}$ at the beginning of the same subframe is determined by the state of the buffer process at the beginning of subframe $(k-M)$. In order to obtain a system state description that evolves according to a Markov process embedded at subframe boundaries, the following two indicator processes are introduced to keep track of the state of the remote process at subframe boundaries:

- $\{I_k\}_{k\geq 0}$: it is an indicator process describing the index of subframe k, that is, its position within the current frame; its state space is $S^I = \{0,1,\cdots,M-1\}$.
- $\{C_k\}_{k\geq 0}$: it is an indicator process— defined precisely below— embedded at subframe boundaries, which carries the state of the remote process at the first subframe (subframe 0) to the current subframe. It may change at subframe 0 and remains unchanged for the following $M-1$ subframes; its state space is $S = \{0,1\}$.

Thus the key processes evolve as follows:

$$I_{k+1} = (I_k + 1) \bmod M$$

$$r_{k+1}^l = \begin{cases} 1 & \text{if } q_k < q_{th} \\ 0 & \text{if } q_k \geq q_{th} \end{cases}$$

$$C_{k+1} = \begin{cases} 1 & \text{if } q_k < q_{th} \text{ and } I_k = 0 \\ 0 & \text{if } q_k \geq q_{th} \text{ and } I_k = 0 \\ C_k & \text{if } I_k \neq 0 \end{cases}$$

$$r^r_{k+1} = \begin{cases} r^r_k & \text{if } I_k \neq M-1 \\ C_k & \text{if } I_k = M-1 \end{cases}$$

Given that $r^l_k = i$ and $r^r_k = j$ at the beginning of subframe k, the finite queue will evolve over subframe k under the i.i.d. batch arrival process with PMF obtained through the convolution of the PMFs f^l_i and f^r_j. It is easy to establish that the stochastic process $\{I_k, r^r_k, C_k, r^l_k, q_k\}_{k \geq 0}$, embedded at subframe boundaries, is a finite state Markov chain. Its state space, S_m, and associated transition probabilities can be found in [6]. The multi-dimensional state-space of the process $(I_k, r^r_k, C_k, r^l_k, q_k)$ can be mapped lexicographically into a one dimensional space l_k. Due to the periodic evolution of I_k, the transition matrix of the one dimensional process l_k, $\mathbf{P}$, takes the canonical form of a periodic Markov chain of period M given by:

$$\mathbf{P} = \begin{bmatrix} 0 & \mathbf{P}_0 & 0 & \cdots & 0 \\ 0 & 0 & \mathbf{P}_1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \mathbf{P}_{M-2} \\ \mathbf{P}_{M-1} & 0 & 0 & \cdots & 0 \end{bmatrix}$$

where the matrices $\mathbf{P_i}, \mathbf{i} \in \{\mathbf{0}, \cdots, \mathbf{M-1}\}$ are those corresponding to state transitions from states $(i, \cdot, \cdot, \cdot, \cdot)$ to states $(j, \cdot, \cdot, \cdot, \cdot)$, where $j = (i+1) \bmod M$. Each $\mathbf{P_i}$ is of order $m_i \times m_{i+1}$ with $m_0 = 4(Q+1), m_1 = 4(Q+1), m_n = 8(Q+1), 2 \leq n \leq M-1$. The stationary probabilities of the process l_k, or equivalently $\pi(I, r^r, C, r^l, i)$ can be obtained [7] by solving M systems of m_i linear equations each, $0 \leq i \leq M-1$, the largest of which is composed of $8(Q+1)$ equations. This is instead of solving the larger system of order $8(M-1)(Q-1)$. This effectively eliminates the dependence of the computational storage complexity on the ratio of the propagation delays $M = T_r/T_l$. Notice that since the Markov process is embedded at the subframe boundaries, the complexity is independent of T_l. Thus it is possible to solve for systems with both large local and remote propagation delays.

3.1 Per-session Cell Loss Probability

In this section, the cell loss probability of an arbitrary cell from the local arrival process is derived. The cell loss probability for the remote arrival process and the overall cell loss probability can be derived in a similar fashion.

In the following, let r^l, r^r be states of the local and remote arrival processes at the beginning of the current subframe, respectively. Let i be the buffer state at the beginning of the k^{th} slot in the subframe. Let $w_l(r^l, r^r, i, k), 1 \leq k \leq T_l$ be a random variable describing the number of local cells lost over the period starting from the beginning of the k^{th} slot till the end of the current subframe (over a period of $T_l - k$ slots). Clearly, $w_l(r^l, r^r, i, 1)$ describes the number of losses over the period of the subframe. It can be easily shown that the loss probability of local sources (L_l) can be derived in terms of the expected value of $w_l(.,.,.,.)$, $\overline{w_l}(.,.,.,.)$, as follows:

$$L_l = \frac{\sum_{(n,r^r,C,r^l,i) \in S_m} \pi(n, r^r, C, r^l, i)\overline{w_l}(r^r, r^l, i, 1)}{\sum_{(n,r^r,C,r^l,i) \in S_m} \pi(n, r^r, C, r^l, i)\lambda_{r^l} T_l}$$

Let w_1 be a random variable describing the number of local cells lost during the k^{th} slot. The random variables $w_l(r^l, r^r, i, k)$ and their expected values can be derived recursively by noting that the number of cells lost starting from the k^{th} slot in the current subframe till the end of the subframe is equal to the number of cells lost over the k^{th} slot plus the number of cells lost from the beginning of the $(k+1)^{st}$ slot till the end of the subframe, as follows:

For $i \in \{0, \cdots, Q\}$, $1 \leq k \leq T_l - 1$:

$$w_l(r^l, r^r, i, k) = \begin{cases} 0 + w_l(r^l, r^r, j, k+1) & \text{if } i \xrightarrow{1} j, \\ & \quad j \neq Q \\ m + w_l(r^l, r^r, Q, k+1) & \text{if } w_1 = m, \\ & \quad 0 \leq m \leq N_{r^l} \end{cases}$$

$$w_l(r^l, r^r, i, T_l) = \begin{cases} m & \text{if } w_1 = m, 0 \leq m \leq N_{r^l} \\ 0 & \text{otherwise} \end{cases}$$

Let $p_l(m, i, r^l, r^r)$ denote the probability that $w_1 = m$, the buffer occupancy is i and the arrival processes are in states r^l, r^r, respectively. The probabilities $p_l(m, i, r^l, r^r)$, needed for the evaluation of $\overline{w}(.,.,.,.)$, depend on the buffer management policy and will be derived in section 4.

4 Buffer management Policies

The buffer management policy determines the number of cells from each of the local and the remote arrival processes which are discarded over any slot. It is therefore necessary to consider the impact of the buffer management policy on the per-session performance in conjunction with the rate control scheme.

Let i denote the buffer state at the beginning of the current time slot. It is assumed that the cell receiving service during that time slot leaves the queue immediately at the beginning of the slot and thus the remaining buffer capacity that can be allocated to arriving cells during the current time slot is $Q - i + 1$. Let k_l and k_r be the number of local and remote cells arrived during the current time slot, respectively. The local process cell discarding probabilities $p_l(m, i, r^l, r^r)$, introduced in subsection (3.1), can be evaluated as:

$$p_l(m, i, r^l, r^r) = \\ \sum_{k_l=0}^{N_{r^l}} \sum_{k_r=0}^{N_{r^r}} p_l[(m, i, r^l, r^r)|(k_l, k_r)] f_{r^l}(k_l) f_{r^r}(k_r)$$

where the conditional probabilities $p_l[(m, i, r^l, r^r)|(k_l, k_r)]$ are determined by the adopted buffer management policy. A number of policies are briefly presented here, detailed presentation can be found in [6].

- Equal Remaining Capacity (ERC) Policy
  According to this policy, the remaining number of available buffer positions at any time slot is split equally between the two arrival processes. If some of the buffer positions allocated to one of the two processes are not used, they may be utilized by the other arrival process.

- Random Allocation (RA) Policy
  According to the RA policy, two bounds on the number of local (remote) cells that may be lost during the current time slot are defined: an upper bound m_u and a lower bound m_w identified by the maximum and minimum number of cells that can be lost in any time slot. The number of local (remote) cells lost m^l (m^r) is uniformly distributed between the two bounds.

- Fixed Distribution policy (FD Policy)
  This policy favours certain loss patterns over others. As with the random policy, an upper bound m_u and a lower bound m_w are identified. However, instead of the uniform distribution of losses between the two bounds in the random policy, a specific distribution $g(m), m_w \leq m \leq m_u$ is identified and cell losses are distributed accordingly. A simple $g(m)$ is considered in this work, where losses are split equally, if possible, whenever they occur.

5 Numerical Results

For the results derived in this section, the arrival processes are assumed to follow a binomial distribution with parameters $N_l, \lambda_l/N_l$ when in state 0 and with parameters $N_h, \lambda_h/N_h$ when in state 1.

The effect of the ratio of the propagation delays M on the cell loss probability is shown in Fig. 2 under the ERC policy where different update intervals are used, $\lambda_h = 0.8, \lambda_l = 0.10, N_l = 1, N_h = 20$, $Q = 60, q_{th} = 50$ and $T_l = 2$. The unfairness is also clear from Fig. 3, where the ratio of the remote loss probabilities to the local loss probabilities increase as M increases. The effectiveness of the FD policy and the RA policy over the ERC policy in reducing the differences in cell loss probabilities becomes apparent as M increases.

The throughputs obtained by the FD policy is shown in Fig. 4. Note that as M increases, the ratio between the update intervals ($T_r/T_l = M$) increases and the unfairness as expressed by the differences in throughputs increase. The throughputs obtained by the RA and ERC policies are almost identical to that obtained by the FD policy and are not plotted. This is expected, since the buffer management policies primarily control the losses. For the above parameters, losses are in small orders of magnitude and the throughputs are mainly controlled by the time individual sources operate at the different access rates. This can be seen from Fig. 5, for equal update intervals, where the differences in throughputs are negligible and are mainly due to the losses.

Fig. 6 shows the loss probability for the ERC policy under equal update intervals ($T_u = T_r$) for the same parameters considered before. Although the ratio of remote to local cell losses tend to be smaller than when unequal update intervals is employed (Fig. 2), the overall losses tend to be higher due to the laxed control over the local sources. (local sources are updated at longer intervals (T_r) than before (T_l)).

Fig. 7 shows that the unfairness increases as the difference between λ_h and λ_l increases. This can also be seen from Fig. 8 where the ratio of remote and local loss probabilities are shown. The dependence of the unfairness on the choice of the overall parameters in the system can be seen by contrasting Fig. 9 where $\lambda_h = 0.49, M = 4, N_l = 10, N_h = 10$, $Q = 40, q_{th} = 20$ and $T_l = 10$ with Fig. 7 where $\lambda_h = 0.8, M = 4, N_l = 10, N_h = 20$, $Q = 60, q_{th} = 50$ and $T_l = 2$. In the former figure (Fig. 9), the loss probabilities of the local and the remote processes are almost equal.

References

[1] J. Bae and T. Suda. Survey of traffic control schemes and protocols in atm networks. In *Proceedings of the IEEE*, pages 170–189, 1991.

[2] A. Atai and J. Hui. Rate-based feedback traffic controller for atm networks. In *ICC'94*, pages 1605–1615, New Orleans, 1994.

[3] P. Skelly, G. Pacifici, and M. Schwartz. A cell and burst level control framework for integrated video and image traffic. In *Infocom'94*, pages 346–354, Toronto, Canada, 1994.

[4] D. Chiu and R. Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN Systems*, pages 1–14, 1989.

[5] A. Mukherjee and J. C. Strikwerda. Analysis of dynamic congestion control protocols - a fokker-planck approximation. In *Sigcomm'91*, pages 159–169, Zurich, Switzerland, Sept 1991.

[6] Mohamed Abdelaziz. *Performance Evaluation of Traffic Control Mechanisms in ATM Networks.* Ph.D. thesis, University of Vermont, 1995.

[7] J. Hunter. *Mathematical Techniques of Applied Probability, Discrete Time Models: Techniques and Applications*, volume II. Academic Press, 1983.

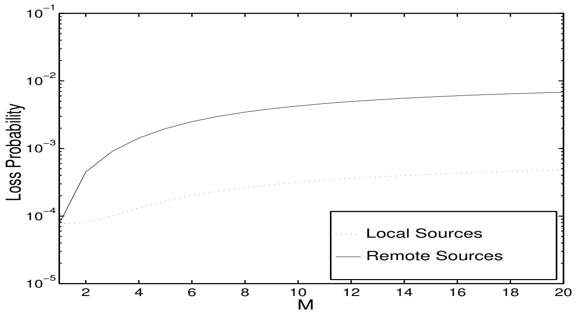

Figure 2: ERC Policy, unequal update intervals

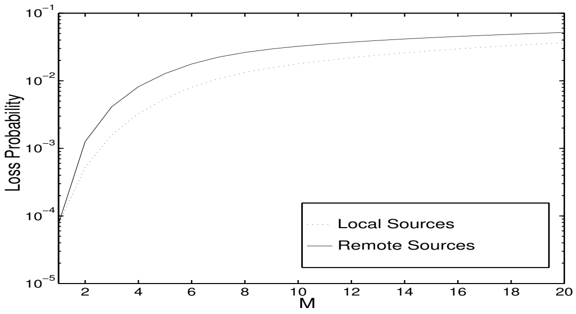

Figure 6: ERC Policy, equal update intervals

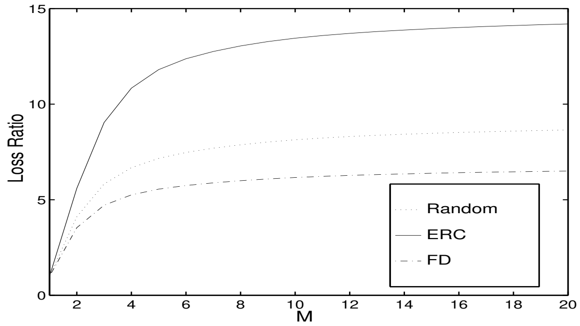

Figure 3: Ratio of remote cell losses to local cell losses

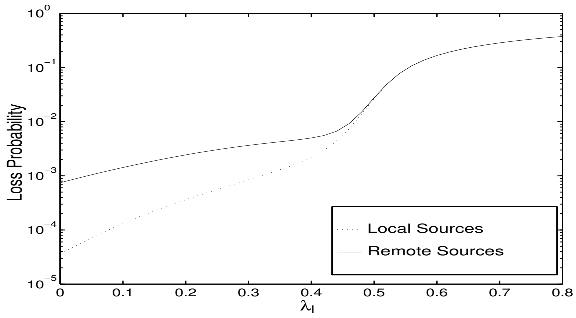

Figure 7: Effect of changing λ_l, unequal update intervals

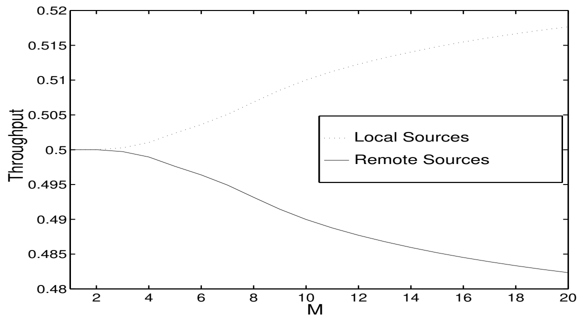

Figure 4: FD Policy, unequal update intervals

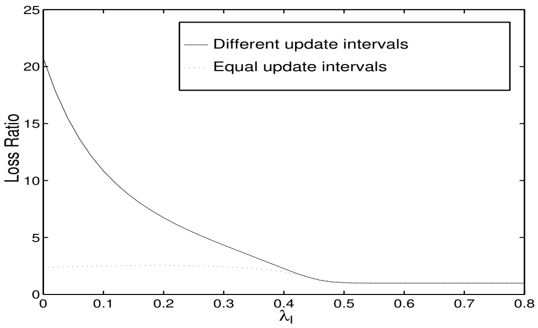

Figure 8: Ratio of remote to local cell losses (ERC policy)

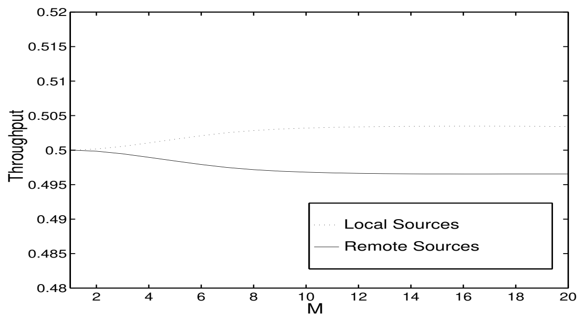

Figure 5: FD Policy, equal update intervals

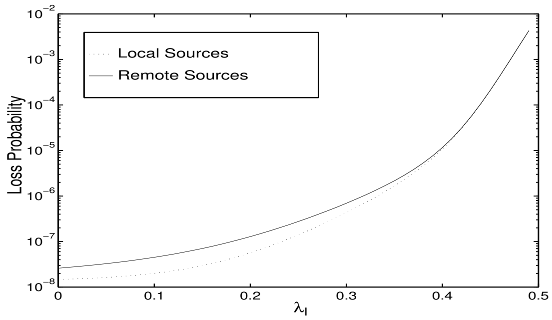

Figure 9: Effect of changing λ_l, unequal update intervals