Comparative Evaluation of Adaptive Stochastic and ERICA Switch Algorithms for ABR Traffic Management in ATM Networks

Slobodan S. Petrovic (IEEE Student Member)
School of Electrical and Information Engineering J03, University of Sydney, Sydney, NSW 2006, Australia
Email: ssp@ee.usyd.edu.au Fax: +61.2. 9351.3847, Phone: +61.2.9351.2842

Abstract

Performances of two explicit rate-based congestion control schemes for ABR service in ATM networks: Explicit Rate Indication for Congestion Avoidance (ERICA) and proposed Adaptive Stochastic (AS), are compared under heavy load conditions. Neither of the schemes dictates a particular switch architecture. The requirements tested include: utilization, queuing delay, queuing delay variance and queue size. A simple network topology involves an ATM switching node modeled as a single server queuing system, fed by a number of highly bursty traffic sources. Small and large bursts are used in simulations. The controllable traffic sources' allowed cell rates (ACR) are dynamically shaped by the explicit feedback messages from the switch. The results of a simulation study suggest that AS scheme can provide for higher priority traffic in ATM LANs shorter queuing delay, queuing delay variance, maximum and mean queue size, while retaining ERICA scheme's utilization .

1. Introduction

Future high speed networks are expected to use the Asynchronous Transfer Mode (ATM) in which the information is transmitted using short fixed-size cells consisting of 48 bytes of payload and 5 bytes of header. The fixed size of the cells reduces the variance of delay making the networks suitable for integrated traffic consisting of voice, video and data. Proper traffic management helps ensure efficient and fair operation of networks in spite of constantly varying demand. The problem is especially difficult during periods of heavy load particularly if the traffic demands can not be predicted in advance as in the case of data traffic. Hence congestion control is critical in ATM networks and although only a part of the traffic management, is its most essential aspect. When two or more bursts arrive simultaneously at a node, the queue lengths may become large very fast resulting in buffer overflow. At the points where the total input rate is larger than the output link capacity, congestion becomes a problem. Several congestion control schemes for ATM networks have been developed [1] - [10].

The remainder of the paper is organized as follows: in Section 2 the assumed network structure and how it operates is described; in Section 3 the design of AS switch algorithm is outlined; in Section 4 the simulation models and obtained results are reported and in Section 5 conclusion remarks are made.

2. Network configuration

It is assumed that two types of traffic share the ATM network resources (Fig. 1). The first type consists of Variable Bit Rate (VBR) traffic which is delay sensitive i.e. which has stringent delay requirements and is not subjected to any control and hence is uncontrollable. The second type of traffic requires good loss performance, but agrees to have its bit rate controlled when necessary, thus accepting the possibility of longer mean delays during periods of high demand, i.e. cell loss is avoided at the expense of delay and hence is controllable. All delay insensitive, or Best Effort, i.e. Available Bit Rate (ABR) traffic can join this category which is designed for normal data traffic such as file transfer and email. In these experiments delay sensitive traffic was given a higher priority. The ACRs of the ABR traffic sources take the values being conveyed by the feedback messages which are carried from the switch to the sources' terminals through backward Resource Management (RM) cells. Each source could also be a switch acting as a virtual source and generating RM cells periodically as if it was a source. The switch acts as a virtual destination and returns all RM cells received from the source as if the switch was the destination. In this study RM cells are "in-rate" since they are counted toward the ACR, what

means that the total rate of data and RM cells should not exceed the ACR of the source. During idle periods a source does not generate RM cells. A RM cell is always sent as the first cell after an idle period for the source. In the compared schemes each individual lower priority (controllable) traffic virtual connection (VC) is guaranteed by the network only a minimum bandwidth of its minimum cell rate (MCR). However, it is allowed to dynamically share the available resources left over from the higher priority (uncontrollable) traffic, i.e. the lower priority traffic grabs whatever remains of the network resources (bandwidth and buffer space) after the higher priority traffic is served. Thus the ABR capacity is not fixed, but varies according to the VBR load. Switch complexity does not depend on the number of VCs since there is no per VC queuing. In the switch, output buffering is used, which means that the queuing happens at the output port. The cells of both traffic classes are treated here, in the switch output buffer, by equal priority and are served with first-in-first-out (FIFO) discipline. The same service discipline applies also to all sources' queues. In order to support the feedback control schemes a simple Connection Admission Control (CAC) policy ensures that the peak cell rate of the total higher priority (VBR) traffic is within the ATM switch server capacity in every moment and that the long term mean cell rate of all the traffic (higher and lower priority together) is within the ATM switch server capacity. AS and ERICA schemes operate at the ATM switch output queuing point.

3. Description of Adaptive Stochastic (AS) switch algorithm

The traffic flowing from the sources to the ATM switch is assumed to be a process which takes place in random environment. The design of the control algorithm has been based upon input-output model of the process which allows analysis over the transfer function and realization of the control algorithm with minimal a priori knowledge about the dynamics and parameters of the process. The transfer function is defined as

$y(k)u^{-1}(k) = B(z^{-1})A^{-1}(z^{-1})$

where z^{-1} is delay operator meaning

$z^{-1}y(k) = y(k - 1), \quad z^{-1}u(k) = u(k - 1)$

and integer k is the control interval's time index. Since the process is modeled here as the first order linear dynamical system:

$A(z^{-1}) = 1 + a(k)z^{-1}, \quad B(z^{-1}) = b(k)z^{-1},$

it is described by the stochastic difference equation

$y(k+1) = - a(k)y(k) + b(k)u(k).$

y(k) is the output equal to instantaneous queue length minus the reference level (threshold). The scalar u(k) is the control input generated by the algorithm. It is a function of the queue size measurement and is used for identification of parameters and for attaining the desired system response. u(k) determines the ACRs of the ABR sources with the goal to maintain the ATM switch output buffer occupancy as close as possible to a certain reference level. u(k) is partitioned among active ABR sources accordingly. An ABR source seen by the switch as inactive is allocated its MCR only. An active ABR source's ACR is then formed as the sum of the source's long term mean cell rate and its share of u(k). In proportion to its absolute value, a negative u(k) reduces and a positive u(k) increases the ABR sources' ACRs. Hence, the ACRs fluctuations are dictated by u(k) but can not go below the MCR (agreed during the connection setup) or above the link rate. The unknown parameters are a(k) and b(k). Their actual values are random and time varying. The uncertainty in b(k) corresponds to uncertainty in the system gain while the uncertainty in a(k) corresponds to uncertainty in the system time constant. Not knowing a(k) exactly is equivalent to not knowing the pole location of the system function. Since the system is operating in a stochastic environment, both y(k) and u(k) are random sequences. The objective of the problem is to find in each control interval the optimal control input that minimizes the scalar real-valued cost functional i.e. the performance criterion

$J(u) = E\{qy^2(k+1) + ru^2(k) + qy^2(k)\} \qquad k = 0, 1, ...$

where E denotes expected value. q and r are positive scalars (weighting factors). As the only information about the process comes from the observation of the queue size, since the queue size is the only directly measurable quantity at the switch, the optimal control policy can be interpreted as follows: at the end of the k-th control interval an optimal control input $u^o(k|k)$ (superscript o stands for optimal) is generated, based on the queue length measurement. k on the left side of | denotes the control input's time index and k on the right side of | means that the measurement of the queue length was taken during the control interval k. The next control input $u^o(k+1|k+1)$ is again computed, but based on the information about the queue size available at the control interval k + 1. Thus, the optimal deterministic control shall be recomputed after new information becomes available at each control interval. The switch gives only one feedback value in a single control interval which is equal to the round-trip delay because there is little point in sending further control feedbacks to the sources until previous feedback has had time to take effect. The overall adaptive controller can be viewed as combined of an identifier and a zero-memory controller in the feedback loop. The former is the learning device where the estimates of the parameters are generated in real time

and can be designed independently of the controller objectives. The later computes the on-line optimal control for a deterministic system and is parameter adaptive. Accordingly, as this adaptive control problem involves nonlinear parameter estimation, the extended Kalman filter is used to generate the approximate conditional means and the associated error covariances of the unknown parameters. This approach led to the design of a simple proportional controller that has the desired adaptive properties. The closed loop feedback optimal control is given by

$u^o(k|k) = -g(k)y(k)$ $\quad k = 0, 1, ...$

The adaptive feedback gain g(k) depends not only on parameters a(k) and b(k) estimates ($a^{est}(k|k)$ and $b^{est}(k|k)$ resp.), but also on conditional error covariances

$e_a(k|k) = a^{est}(k|k) - a(k)$ and $e_b(k|k) = b^{est}(k|k) - b(k)$,

i.e. it is modulated by the current parameters uncertainty (vectors and matrices are denoted in bold letters, superscript T stands for 'transposed')

$$g(k) = -\{[R(k|k) + \mathbf{b}^T(k|k)\mathbf{Q}\mathbf{b}(k|k)]^{-1}\mathbf{b}^T(k|k)\mathbf{Q}\Phi(k|k) + R^{-1}(k|k)\mathbf{d}^T(k+1|k)\}[1, 0, 0]^T$$

where $\mathbf{b}^T(k|k) = [b^{est}(k|k), S_{bb}(k|k), S_{ab}(k|k)]$,

$\mathbf{d}^T(k+1|k) = [qS_{ab}(k|k), qa^{est}(k|k), 0]$,

$\Phi(k|k) = \mathbf{F}(k|k) - \mathbf{b}(k|k)R^{-1}(k|k)\mathbf{d}^T(k+1|k)$,

$$\mathbf{F}(k|k) = \begin{bmatrix} a^{est}(k|k) & 0 & 0 \\ S_{ab}(k|k) & a^{est}(k|k) & 0 \\ S_{aa}(k|k) & 0 & a^{est}(k|k) \end{bmatrix},$$

$S_{aa}(k|k) = e_a^2(k|k)$, $S_{bb}(k|k) = e_b^2(k|k)$,

$S_{ab}(k|k) = S_{ba}(k|k) = e_a(k|k)e_b(k|k)$,

$$R(k|k) = r + qS_{bb}(k|k), \qquad \mathbf{Q} = \begin{bmatrix} q & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

The parameters estimates are generated via the equations

$\mathbf{p}^{est}(i+1|i+1) = \mathbf{p}^{est}(i|i) + \mathbf{G}(i+1)[y(i+1) - \mathbf{c}^T(i)\mathbf{p}^{est}(i|i)]$,

$\mathbf{p}^{est,T}(i|i) = [a^{est}(i|i), b^{est}(i|i)]$, $\mathbf{c}^T(i) = [-y(i), u(i)]$.

The filter gain matrix $\mathbf{G}(i+1)$ satisfies the relation

$\mathbf{G}(i+1) = \mathbf{S}(i|i)\mathbf{c}(i)[\mathbf{c}^T(i)\mathbf{S}(i|i)\mathbf{c}(i)]^{-1}$ $\quad i = 0, 1, ... , k -1$

$\mathbf{S}(i+1|i+1) = [\mathbf{S}(i|i) - \mathbf{G}(i+1)\mathbf{c}^T(i)\mathbf{S}(i|i)] / ff$,

$$\mathbf{S}(i|i) = \begin{bmatrix} S_{aa}(i|i) & S_{ab}(i|i) \\ S_{ba}(i|i) & S_{bb}(i|i) \end{bmatrix},$$

with the initial estimates given by

$\mathbf{p}^{est,T}(0|0) = [a_0 , b_0]$,

$$\mathbf{S}(0|0) = \begin{bmatrix} S_{a0} & 0 \\ 0 & S_{b0} \end{bmatrix}.$$

The AS control algorithm expresses its dual nature by simultaneously performing two functions: perpetual parameters estimation i.e. "learning" and process control. The forgetting factor 'ff' allows the controller's capacity for adaptation on changes in the system parameters to be preserved, by preventing error covariance matrices to become approximately equal to zero. Estimating of parameters could continue only if the matrices are different from zero, i.e. if some level of parameters uncertainty remained. Owing to the forgetting factor, the identification of parameters even if they for some time remain constant, will never be completed. Fortunately, the exact identification of the unknown parameters is not necessary for good control. Through the parameters uncertainty the controller expresses its "cautiousness".

ERICA switch algorithm - see [10].

4. Simulation model and results

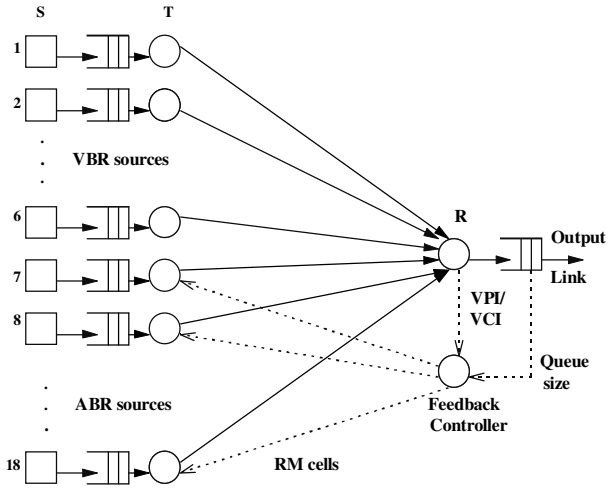

Figure 1: The simulation model

Both congestion control algorithms were tested and compared on the same network simulation model (Fig. 1), with the same MCRs. Each source S generates cells that are queued in its buffer waiting there for transmission by the transmitter T. Each of the transmitters T independently removes cells from its source queue and transmits them to the receiver R. R is a point of congestion somewhere in the network, e.g. an output buffer on an ATM switch. There is a transmission delay between each transmitter transmitting a cell and the cell being received at the receiver. The transmission delay represents the combination of the propagation delay and the switching delay from the source to the point of congestion. For simplicity, all the sources are equidistant to the ATM switch and hence the propagation delay is the same between the ATM switch and all the sources.

AS scheme's control interval as well as ERICA scheme's averaging interval were equal to the round trip delay, i.e. the control rate was matched to the feedback rate. This illustrates a fundamental principle in control theory, which says that the system is unstable when the control is faster than feedback. But the system is unresponsive if the control is slower than feedback. The performance of the algorithms was examined over source - switch distances (link lengths) 1 km, 10 km, 100 km and 1000 km, with corresponding control/averaging intervals 10 μsec, 100 μsec, 1 msec and 10 msec, respectively (1 km implies 5 μsec propagation delay). The schemes were not tested and compared on distances longer than 1000 km where the round trip delays can be very large and performances of the schemes consequently would deteriorate. This problem could be fixed by segmenting the feedback loop in smaller pieces and letting the switches act as "virtual source" and/or "virtual destination". To prevent too many RM cells being generated, one exception from the above stated principle was made: on 1 km link instead of 10 μs, 100 μs control/averaging interval was used, i.e. the same as on 10 km link. Large buffers are assumed and the effect of cell losses is not modeled. The terminals are connected to the switch through links running at 155 Mbps (155 Mbps = 365566 cells/sec). The switch receives traffic from terminals that are attached to its input ports. All input ports have traffic flowing to a common output port operating at 155 Mbps. The input to the ATM switch is a superposition of streams from 18 (12 ABR and 6 VBR) different ON/OFF sources, also referred to as Interrupted Bernoulli Process (IBP) which alternate active periods, transmitting at their PCR and idle periods, transmitting at zero bit rate. Both periods of burst and silence are exponentially distributed with constant average lengths during a simulation. Two different burst sizes were used in the simulations: small bursts in average contained 128 cells and large bursts in average contained 6144 cells. The switch server rate is assumed to have a fixed ratio with respect to the long term mean arrival rate. As AS scheme has been primarily designed for ATM LANs where ideal value of the switch target utilization parameter was determined to be 0.95 [10], ABR traffic sources were given 60% and VBR traffic sources 35% bandwidth. In order to have more than a million cells generated, a simulation run time was 3 seconds. The higher priority (VBR) traffic was generated by the sources 1, 2, 3, 4, 5, 6 with PCR 50 Mbps, 40 Mbps, 30 Mbps, 20 Mbps, 10 Mbps, 5 Mbps, respectively, totally 155 Mbps, therefore equal to the switch server rate and with source utilizations 0.4, 0.5, 0.2, 0.2, 0.3, 0.25, respectively. The lower priority (ABR) traffic was generated by the sources 7, 8, ..., 18 with PCR 155 Mbps each, what means that each ABR traffic source was capable of filling the entire server capacity of the switch. The ABR traffic sources' utilization was 0.10 for sources 7 and 13, 0.08 for sources 8 and 14, 0.05 for sources 9 and 15, 0.04 for sources 10 and 16, 0.02 for sources 11 and 17, 0.01 for sources 12 and 18. The aggregate peak rate of the total traffic thus was 2015 Mbps. The purpose of the experiment was to examine how successfully the compared schemes can cope with such stressful conditions when arrival rate to the switch output buffer can exceed 13 times the server capacity/output link rate, while preserving the design goals. The output buffer threshold value used in AS scheme experiments was 100 cells. Total MCR value used was 1/10 of the portion of the bandwidth allocated to the ABR traffic class, that is 6% bandwidth. This amount was divided in proportion to ABR traffic source utilizations: MCR = PCR/100 for sources 7 and 13, PCR/125 for sources 8 and 14, PCR/200 for sources 9 and 15, PCR/250 for sources 10 and 16, PCR/500 for sources 11 and 17 and PCR/1000 for sources 12 and 18. In practice, it is desirable to allow some minimum capacity for processing aggregate ABR traffic when there is contention. As all ABR traffic sources had the same PCR = 155 Mbps, in AS scheme experiments each of the active ABR traffic sources received periodically the same amount of the feedback equal to u(k)/(number of currently active ABR traffic sources). In the case of different PCRs, the formula above could be easily adjusted. In AS scheme as well as in ERICA scheme experiments, each inactive ABR traffic source was allocated its MCR only, for two reasons mainly: in order to drain out their queues rapidly and because of the transmission delay the sources which may seem as inactive, could be in fact, active. In examining AS scheme the forgetting factor equal to 0.96 was used; initial estimates were $S_{a0} = S_{b0} = 1$ and $a_0 = b_0 = 0.1$.

Performance evaluation of the schemes showed that ERICA scheme, compared to AS scheme, achieved average utilization (Fig. 16) when small bursts were used: on 1 km link 0.2% higher, on 10 km link 0.15% higher, on 100 km link 0.8 % higher and on 1000 km link 2.5% higher; when large bursts were used: on 1 km link 0.7% lower, on 10 km link the same, on 100 km link 1.1 % higher and on 1000 km link 5.6% higher. As any MCR > 0 contradicts the initial stipulation that VBR traffic can take the entire server capacity whenever it needs it, the MCRs employed in this experiment brought about conflicting demands by VBR and ABR traffic classes, as VBR traffic class, albeit able to fill the entire server capacity (if all VBR traffic sources are simultaneously transmitting) was guaranteed only 94% of it. Obviously, although the MCR was guaranteed to all 12 ABR traffic sources, due to their bursty nature they did

not use it all the time and not all of them at the same time, what allowed the VBR traffic to have enough bandwidth to flow smoothly. The number of VBR traffic cells received (Fig. 14) in the case of small bursts for both schemes was practically identical on all four link lengths. In the case of large bursts both schemes performed equally on 1 km and 100 km links, AS scheme performed better on 10 km link but much worse on 1000 km link. With regard to the number of ABR traffic cells received (Fig. 15) in the case of small bursts both schemes performed on shorter links (1 km and 10 km) almost equally, but on longer links (100 km and 1000 km) ERICA scheme performed better; in the case of large burst the results were mixed: AS scheme performed better on 1 km link and 1000 km link but worse on 10 km link and 100 km link. However, the utilization metric gives only a partial picture of efficiency. Ideally, an efficient scheme should also control queuing delays within acceptable limits. The delay varies according to load, but there is always a non-zero queuing delay. ERICA scheme (Fig. 2) had the mean queuing delay in the switch output buffer when small bursts were used: on 1 km link 5% shorter, on 10 km link 5% longer, on 100 km link 2.1 times longer, on 1000 km link 3.6 times longer; when large bursts were used: on 1 km link 20 times longer, on 10 km link 52.5 times longer, on 100 km link 11.7 times shorter, on 1000 km link 42 times shorter. ERICA scheme (Fig. 3) had the queuing delay variance in the switch output buffer when small bursts were used: on 1 km link 4.3 times bigger, on 10 km link 5.1 times bigger, on 100 km link 3 times bigger, on 1000 km link 2.6 times bigger; when large bursts were used: on 1 km link 922 times bigger, on 10 km link 671 times bigger, on 100 km link 103 times shorter, on 1000 km link 1927 times shorter. ERICA scheme (Fig. 11) had the mean queue length in the switch output buffer when small bursts were used: on 1 km link 4% shorter, on 10 km link 5% times longer, on 100 km link 2.1 times longer, on 1000 km link 3.7 times longer; when large bursts were used: on 1 km link 20 times longer, on 10 km link 58 times longer, on 100 km link 11.6 times shorter, on 1000 km link 60.2 times shorter. ERICA scheme (Fig. 8) had the maximum queue length in the switch output buffer when small bursts were used: on 1 km link 2.4 times longer, on 10 km link 2.9 times longer, on 100 km link 2% shorter and on 1000 km link 1.7 times shorter; when large bursts were used: on 1 km link 9.1 times longer, on 10 km link 12.4 times longer, on 100 km link 8.5 times shorter and on 1000 km link 37 times shorter. The corresponding results for the buffers of sources 7 and 12 are presented because the former is one of the two with the highest and the later one of the two with the lowest source utilization. ERICA scheme's queue lengths, mean queuing delays and queuing delay variances in the ABR traffic source 12 buffer (Figures 6, 7, 10, 13) were smaller than AS scheme's (during the simulation run the ABR traffic source 12 did not generate any large bursts); in the ABR traffic source 7 buffer (Figures 4, 5, 9, 12) in the case of large bursts on all four link lengths as well as for small bursts on shorter links (1 km and 10 km) both schemes performed approximately equally, but on longer links (100 km and 1000 km) ERICA scheme performed better. However, due to ABR traffic's delay insensitivity, this ERICA scheme's advantage is insignificant. With the sort and configuration of traffic generators employed throughout these experiments the ABR service class whose capacity and demand thus varied dynamically, was always in a transient state. For such a case there is currently no known model for quantifying fairness. Hence the issue of fairness could not be addressed in this study and for the issues that were addressed, the simple network model (Fig. 1) was appropriate. As the system was in a state of high variation and had no steady state, the aim was to keep the average utilization high, the average queue size small and prevent queuing delays from becoming unbounded. Under such stressful conditions both schemes were found to be robust, not allowing essential metrics to degrade to unacceptable levels. It should be pointed out that the numbers of ABR and VBR cells received, as well as the average utilization, could have been on 1 km link, compared to those attained on 10 km link, rather higher than lower if the correct control/averaging interval equal to the round-trip delay (10 μs) was implemented. Indeed, in that case ERICA scheme's and AS scheme's average utilizations were found to exceed 96% . However, RM cells' share in the total number of generated cells then was 27.3% what means that data cells utilization fell to 70%. When the 10 times round trip delay (100 μs) control/averaging interval was implemented, RM cells constituted only 2.7% of all cells generated. In that case data cells utilization of both schemes reached for small bursts 91% and for large bursts 93.5% (ERICA) and even above 94% (AS), whereby the violation of the above cited control theory fundamental principle in order to limit the number of RM cells was justified.

5. Summary

Basic ERICA and proposed AS congestion control schemes for ABR service in ATM networks are compared in the presence of two traffic classes - ABR and VBR, under heavy load conditions exceeding 0.95. The discussion relates to the traffic management at the segment of an ATM network between terminals

generating traffic and the switch. The schemes performed differently. When the system was fed with large bursts, AS scheme, in relation to all tested requirements, performed better than ERICA scheme on shorter links (1 km and 10 km), but worse on longer links (100 km and 1000 km). When the system was fed with small bursts, in relation to the mean queue length and mean queuing delay both schemes performed mainly equally on shorter links (1 km and 10 km), but on longer links (100 km and 1000 km) AS scheme performed better. In relation to the maximum queue length AS scheme performed better on shorter links (1 km and 10 km) but worse on longer links (100 km and 1000 km). In relation to the queuing delay variance AS scheme performed better on all four link lengths. In relation to the output link utilization AS scheme performed mainly equally on shorter links (1 km and 10 km) but worse on longer links (100 km and 1000 km). Hence, it could be concluded that AS scheme suits the purpose of congestion control on shorter links better than ERICA scheme, thus recommending itself for ATM LANs, but not so much for ATM WANs.

References

[1] The ATM Forum Technical Committee, "Traffic Management Specification Version 4.0", April 1996.

[2] K. Kawahara, Y. Oie, M Murata, H. Miyahara, "Performance Analysis of Reactive Congestion Control for ATM Networks", IEEE JSAC, vol. 13, May 1995.

[3] W. Matragi, K. Sohraby, "Combined Reactive/ Preventive Approach for Congestion Control in ATM Networks", IEEE ICC, May 1993.

[4] C. Ikeda, H Suzuki, "Adaptive Congestion Control Schemes for ATM LANs", IEEE INFOCOM, June 1994.

[5] G. Pecelli, B. Kim, "Dynamic Behavior of Feedback Congestion Control Schemes", IEEE INFOCOM, vol. 1, 1995.

[6] A. Kolarov, G. Ramamurthy, "End-to-End Adaptive Rate Based Congestion Control Scheme for ABR Service in Wide Area ATM Networks", IEEE ICC, June 1995.

[7] J. Ren, J. Mark, "Design and Analysis of a Credit Based Controller for Congestion Control in BISDN/ ATM Networks", IEEE INFOCOM, vol 1, 1995.

[8] A. Arulambalam, X. Chen, N. Ansari, "Allocating Fair Rates for Available Bit Rate Service in ATM Networks", IEEE Communication Magazine, November 1996.

[9] D. Gaiti, N. Boukhatem, "Cooperative Congestion Control Schemes in ATM Networks", IEEE Communications Magazine, November 1996.

[10] S. Kalyanaraman, R. Jain, S.Fahmy, R Goyal, B.Vandalore, "The ERICA Switch Algorithm for ABR Traffic Management in ATM Networks", IEEE/ACM Transactions on Networking, November 1997.

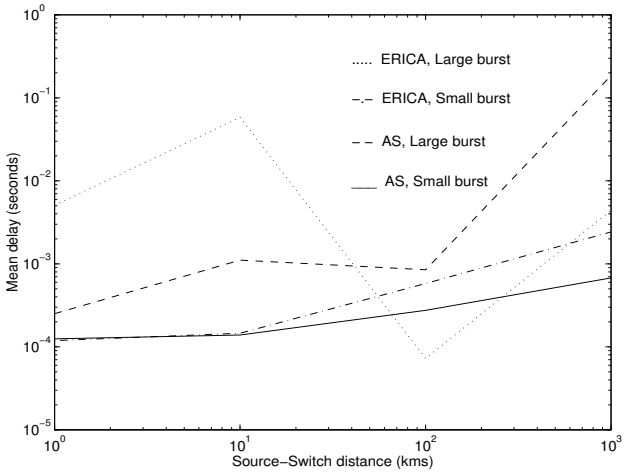

Figure 2: Mean queuing delay in the ATM switch output buffer

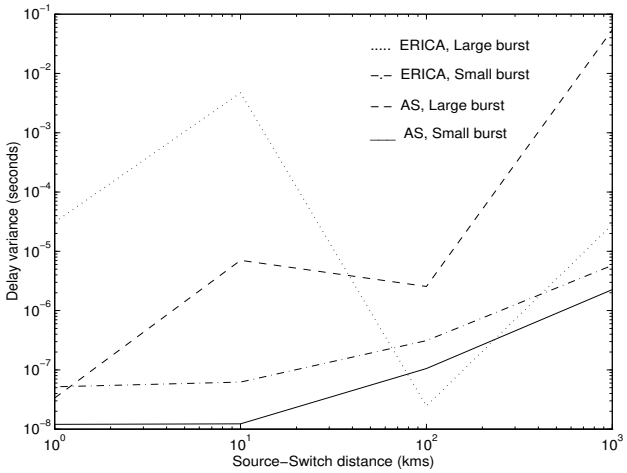

Figure 3: Queuing delay variance in the ATM switch output buffer

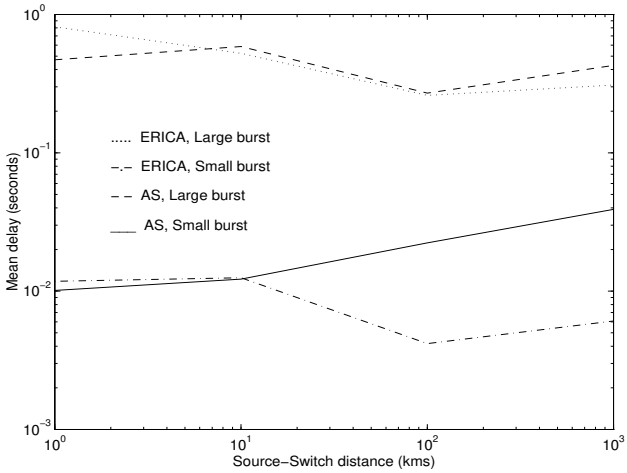

Figure 4: Mean queuing delay in the ABR traffic source 7 buffer

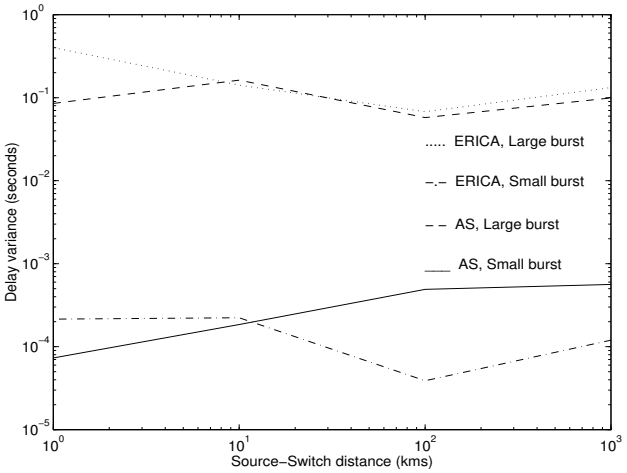

Figure 5: Queuing delay variance in the ABR traffic source 7 buffer

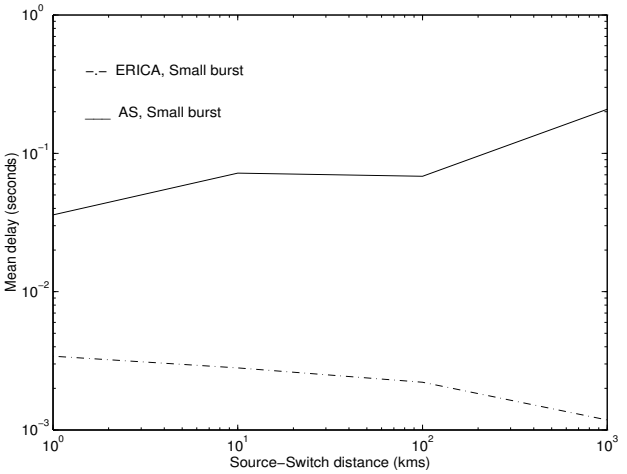

Figure 6: Mean queuing delay in the ABR traffic source 12 buffer

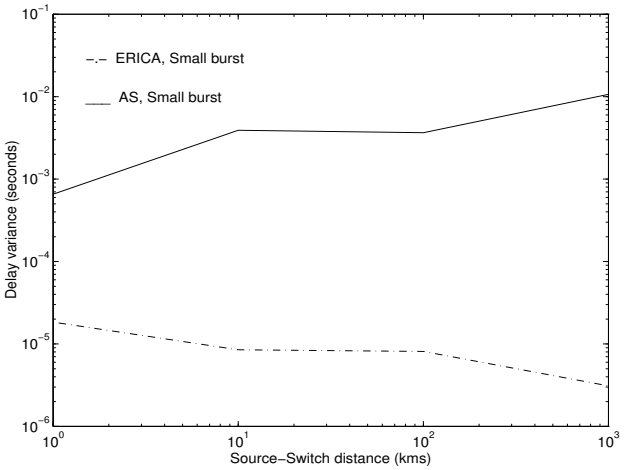

Figure 7: Queuing delay variance in the ABR traffic source 12 buffer

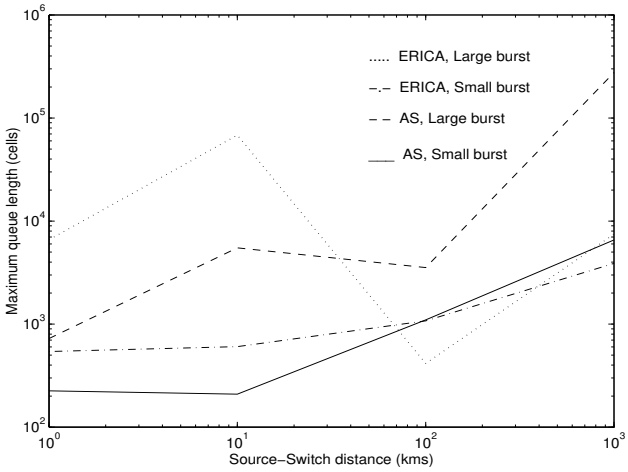

Figure 8: Maximum queue length in the ATM switch output buffer

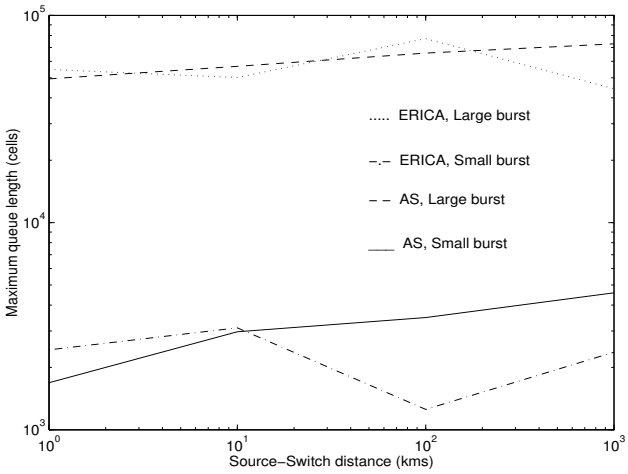

Figure 9: Maximum queue length in the ABR traffic source 7 buffer

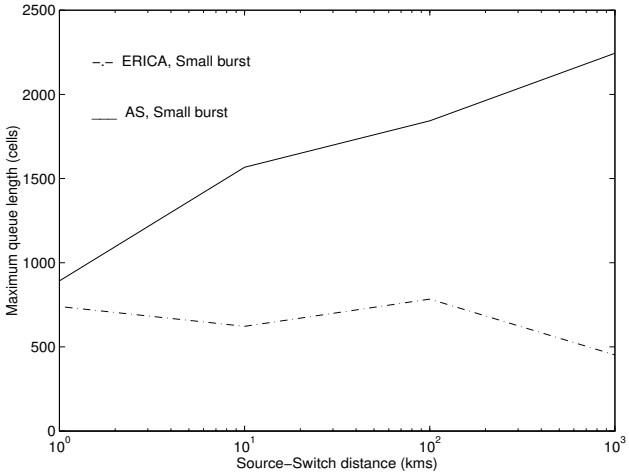

Figure 10: Maximum queue length in the ABR source 12 buffer

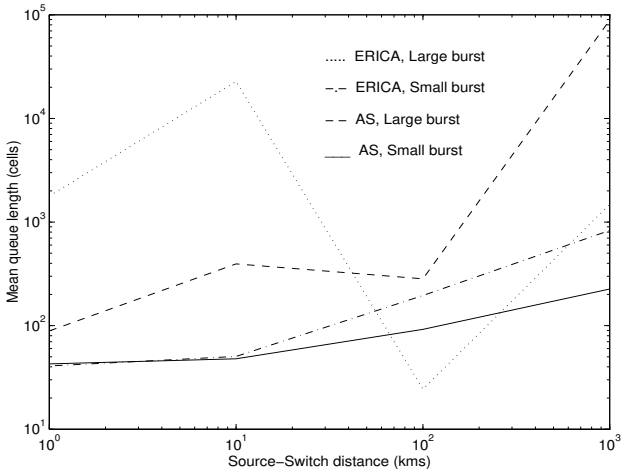

Figure 11: Mean queue length in the ATM switch output buffer

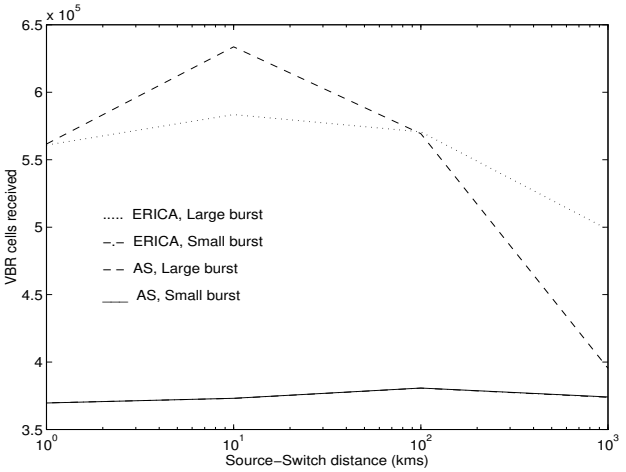

Figure 14: Number of VBR traffic cells received

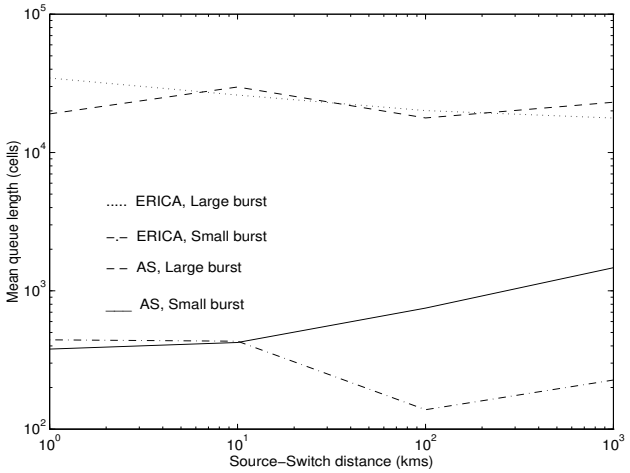

Figure 12: Mean queue length in the ABR traffic source 7 buffer

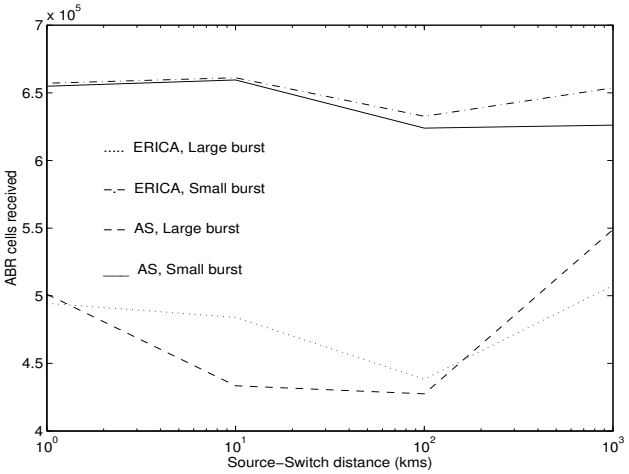

Figure 15: Number of ABR traffic cells received

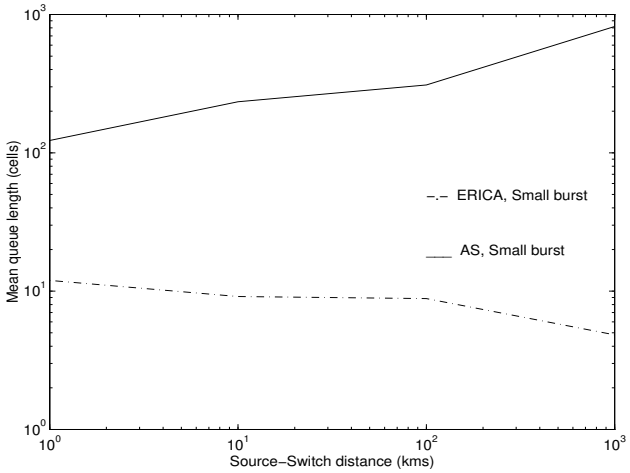

Figure 13: Mean queue length in the ABR traffic source 12 buffer

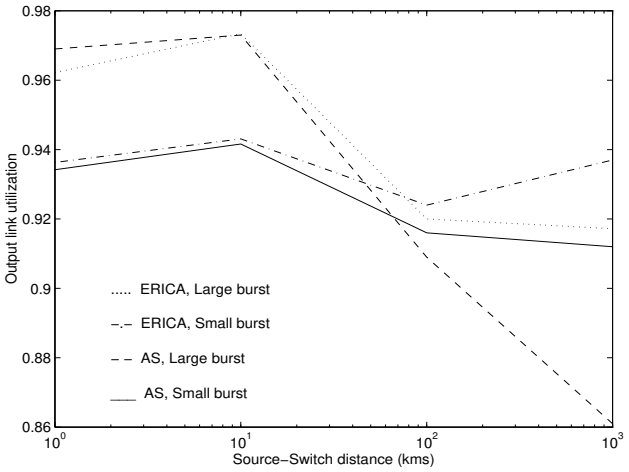

Figure 16: Server/output link utilization