Rate based flow control with bandwidth information

O. Ait-Hellal, E. Altman
INRIA B.P. 93
2004 Route des Lucioles
06902 Sophia Antipolis Cedex
FRANCE

T. Başar*
Coordinated Science Lab.
University of Illinois
1308 West Main
Urbana, IL 61801
U.S.A.

Abstract

The ATM Forum has chosen the rate-based approach for flow control of ABR (Available Bit Rate) traffic in ATM. It is based on a reactive approach whereby the transmission rate of ABR sources can be adapted to the available bandwidth at a bottleneck link. The ATM forum has specified the behavior of the source and destination, as well as the manner in which feedback information (on the available bandwidth and on the congestion state of the network) should be conveyed back to the source. The decision on the precise control mechanism, however, has been left to the designer of the switches. We propose in this paper a reactive control scheme that is based only on information on the available bandwidth. We analyze its stability, and then test its performance by simulations in the presence of other higher priority CBR or VBR traffic.

1 Introduction

We focus in this paper on the stability and performance of rate-based flow control, where the controller determines the allowed transmission rate, based only on information on the available bandwidth. We consider a saturated source, i.e. a source that has always information to send. We consider, more specifically, the reactive control of ABR (Available Bit Rate) traffic in ATM. According to the ATM Forum *Traffic Management Specification*, Version 4.0 [1], RM (Resource Management) cells make periodically the round trip between the source, destination and back to the source, and inform the source about the allowed transmission rate. Each switch on the way along the circuit may change the control information, so that the allowable transmission rate is determined by the most congested switch. The behavior of the switches has not been standardized,

*Research supported by Grants ECS 92-20632 and ECS 93-12807 from the National Science Foundation. All correspondence to be sent to this author. Tel: +1 217 333 3607; Fax: +1 217 244 1653; Email:tbasar@isaacs.csl.uiuc.edu

however, and their design has been left to the manufacturer. Several schemes have been proposed in the literature; see, for example, [1, 2, 3, 7, 13, 14] and the references therein.

The purpose of this paper is to obtain some *qualitative understanding* of the performance of flow control mechanisms that are based only on information on the available rate. It has earlier been shown in [4] that such control schemes could be unstable, if they attempt to achieve 100% of the bandwidth utilization. In this paper we relax this restriction, and propose rate-based control schemes that use only a fraction of the available bandwidth. We prove stability of such schemes, and then evaluate their performance using simulations.

2 The model

We consider the discrete time model considered earlier in [4]. A time unit here corresponds to $(\theta+1)$th of a round trip delay. Let q_n denote the queue length at a bottleneck link, and μ_n denote the effective service rate available for traffic of the given source in that link at the beginning of the nth time slot. Let x_n denote the source rate during the nth time slot. The queue length evolves according to

$$q_{n+1} = (q_n + x_n - \mu_n)^+ \tag{1}$$

where $(a)^+$ denotes $\max(a, 0)$. Since several sources with varying transmission rates may share the same bottleneck link, and in particular, since higher priority traffic may be present (in particular, CBR or VBR traffic), the service rate μ_n available to the controlled source may change over time randomly. We denote the nominal value by μ, and the variations around μ by the process $\{\xi_n\}$, which represents the interference due to other sources. Hence, we have

$$\mu_n = \mu + \xi_n. \tag{2}$$

One possibility for the process $\{\xi_n\}$ is to take it as the output of an ARMA model driven by i.i.d. Gaussian random variables; see [2, 3, 4]. We will make this choice more precise in the next section.

We will assume that the controller receives a noisy delayed information on the available rate:

$$\hat{\mu}_n = \mu_{n-\theta} + e_{n+1}\,,$$

where e_n is the measurement noise, and θ is some nonnegative integer representing the delay. Thus, the input rate during the time slot $n+1$ is allowed to depend on $\hat{\mu}_j$ for $j \le n$. A further assumption is that the process $\{\xi_n, e_n\}$, defined on some joint probability space $\{\Omega, \mathcal{B}, \mathcal{P}\}$, is stationary ergodic.

In [4] the following pure rate matching algorithm of [11] was analyzed:

$$x_{n+1} = \alpha x_n + (1-\alpha)\hat{\mu}_n\,, \tag{3}$$

where $\alpha \in (0,1)$ denotes the forgetting factor in the standard framework of exponential forgetting mechanism. It was shown in [4] that even for the simplest case of $\xi_n = 0$ P-a.s. and e_n being i.i.d., this rate update mechanism leads to an unstable system.

In this paper, we analyze the following modified version of (3):

$$x_{n+1} = \alpha x_n + \beta \hat{\mu}_n - \epsilon \,, \tag{4}$$

where $\alpha \in [0,1)$, β is an arbitrary nonzero parameter (at this point), and ϵ is an appropriate drift term. We will prove in the next section that this update mechanism leads to a stable queue process under fairly general conditions on α, β, and ϵ, and the statistics of $\{\xi_n\}$ and $\{e_n\}$. These conditions will subsume two particular cases of interest, namely:

- Full matching with drift: $0 < \alpha < 1$, $\beta = 1 - \alpha$, and $\epsilon > 0$.
- Partial matching: $0 < \alpha < 1$, $\beta < 1 - \alpha$, and ϵ is arbitrary.

It will be shown that in the *partial matching* case a drift is not needed to stabilize the system, and hence one may choose without any loss of generality $\epsilon = 0$.

It is worth noting at this point that if we had worked with the linearized version of (1), i.e.

$$q_{n+1} = q_n + x_n - \mu_n \,,$$

it would not have been possible to stabilize it with the update mechanism (4) regardless of the values of the parameters α, β, and ϵ. To be able to stabilize (asymptotically) this marginally stable linear system, some feedback information on the queue length is essential. Hence, the linearized system cannot be used in this case to establish the stability of (1).

3 Main Results

We introduce the netput process (input minus output rate) as

$$\begin{aligned} \Delta_n &:= x_n - \mu_n = [\alpha x_{n-1} + \beta\mu + \beta\xi_{n-\theta-1} + \beta e_n - \epsilon] - [\mu + \xi_n] \\ &= \alpha\Delta_{n-1} + \alpha\xi_{n-1} + \beta\xi_{n-\theta-1} + \beta e_n - \xi_n + v \end{aligned} \tag{5}$$

where

$$v := (\beta + \alpha - 1)\mu - \epsilon \,.$$

Note that in terms of $\{\Delta_n\}$, the queue length process (1) can be rewritten as

$$q_{n+1} = (q_n + \Delta_n)^+ \tag{6}$$

which we will henceforth refer to in place of (1).

We first study the special case $\alpha = 0$, for which the stability proof is simpler.

Theorem 3.1 *Suppose that $\alpha = 0$ and $v < 0$. Then the queue length process couples to a stationary ergodic regime within finite time P-a.s.*

Proof. Since $\alpha = 0$ and the process $\{\xi_n, e_n\}$ is stationary ergodic, it follows that $\{\Delta_n\}$ is also stationary ergodic. The queue dynamics (6) is seen to be that of a G/G/1 queue [8, 9, 12] driven by the stationary ergodic sequence Δ_n. Since $E\Delta_n = v < 0$, the result now follows from [9] (Ex. 1 and Theorem 3 (p.21)); see also [8] (Section 4.6; p. 258) or [12]. ∎

To study the case $\alpha \in (0, 1)$, we now assume a Markovian setting. As in [2, 3], we let $\{\xi_n\}$ correspond to a stable ARMA process driven by an i.i.d. sequence $\{\phi_n\}$ with mean *zero* and variance *one*:

$$\xi_n = \sum_{i=1}^{d} \ell_i \, \xi_{n-i} + \sum_{i=1}^{\hat{d}} k_i \, \phi_{n-i}. \tag{7}$$

Here ℓ_i's are chosen such that $\{\xi_n\}$ is a stable process, and k_i's are some constants not all of which are zero. Without loss of generality we may assume that $\hat{d} = d$ (by choosing some of the ℓ_i's or k_i's to be zero).

We further assume that the distributions of both $\{\phi_n\}$ and $\{e_n\}$ are nonsingular with respect to the Lebesgue measure, and admit non-trivial densities.

Before stating and proving the main stability result, we rewrite the dynamics of the system in an equivalent matrix notation. Set $d' = \theta + d + 1$. Define the d'-dimensional vector η_n to be the vector $(\xi_{n-d'+1}, \ldots, \xi_n)'$ that would be obtained from (7) if $k_1 = 1$ and $k_i = 0$ for $i \neq 1$. Hence, η_n is generated by the first-order vector difference equation:

$$\eta_{n+1} = C\eta_n + \begin{pmatrix} 0 \\ \cdot \\ \cdot \\ \cdot \\ 0 \\ 1 \end{pmatrix} \phi_n \tag{8}$$

where

$$
C := \begin{pmatrix}
0 & 1 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\
. & . & . & \dots & . & . & . & . & . & \dots & . \\
0 & 0 & 0 & \dots & 0 & 1 & 0 & 0 & 0 & \dots & 0 \\
0 & 0 & 0 & \dots & 0 & 0 & 1 & 0 & 0 & \dots & 0 \\
0 & 0 & 0 & \dots & 0 & 0 & 0 & 1 & 0 & \dots & 0 \\
. & . & . & \dots & . & . & . & . & . & \dots & \\
0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & \dots & 1 \\
0 & 0 & 0 & \dots & 0 & \ell_d & \ell_{d-1} & \ell_{d-2} & & \dots & \ell_1
\end{pmatrix} \tag{9}
$$

Then the actual ξ_n generated by (7) is given by[1]

$$
\xi_n = (0, \dots, 0, k_d, k_{d-1}, ..., k_1)\, \eta_n. \tag{10}
$$

Let $\tilde{\Delta}_n := \Delta_n - v/(1-\alpha)$, and set

$$
Z_n := (\tilde{\Delta}_n,\ \eta'_{n+1})', \qquad \psi_n := (e_n,\ \phi_n)'. \tag{11}
$$

Then $\{Z_n\}$ is a Markov chain, generated by

$$
Z_{n+1} = AZ_n + D\psi_{n+1} \tag{12}
$$

where

$$
D := \begin{pmatrix} \beta & 0 \\ 0 & 0 \\ . & . \\ . & . \\ 0 & 1 \end{pmatrix}, \qquad A := \left(\begin{array}{c|c} \alpha & A_{12} \\ \hline & \\ 0 & \mathrm{C} \\ & \end{array}\right),
$$

A_{12} is a row vector of size $d + \theta + 1$, given for $\theta > 0$ by

$$
A_{12} = (0, \dots, 0, \beta, 0, \dots, 0, \alpha, -1) \begin{pmatrix}
0 & ... & . & . & . & ... & 0 & 0 & 0 \\
. & ... & . & . & . & ... & 0 & 0 & 0 \\
. & ... & . & . & . & ... & 0 & 0 & 0 \\
0 & ... & . & . & . & ... & 0 & 0 & 0 \\
k_d & ... & . & . & . & ... & 0 & 0 & 0 \\
. & & & & & & & & \\
. & & & & & & & & \\
0 & ... & k_d & k_{d-1} & k_{d-2} & ... & k_1 & 0 & 0 \\
0 & ... & 0 & k_d & k_{d-1} & ... & k_2 & k_1 & 0 \\
0 & ... & 0 & 0 & k_d & ... & k_3 & k_2 & k_1
\end{pmatrix},
$$

[1]The reason for this is that, due to the linearity of the system (7), its response to the input $\sum_{i=1}^{d} k_i\, \phi_{n-i}$ equals to the sum (over i) of individual responses to $k_i\, \phi_{n-i}$. Furthermore, due to the linearity and time homogeneity of the system, a response to $k_i\, \phi_{n-i}$ equals to $k_i\, (0, \dots, 0, 1)\, \eta_{n+1-i}$.

where β is at the d'th location; and for $\theta = 0$ by

$$A_{12} = (0, \ldots, 0, 0, 0, \ldots, 0, \alpha+\beta, -1) \begin{pmatrix} 0 & \ldots & . & . & . & \ldots & 0 & 0 & 0 \\ . & \ldots & . & . & . & \ldots & 0 & 0 & 0 \\ . & \ldots & . & . & . & \ldots & 0 & 0 & 0 \\ 0 & \ldots & . & . & . & \ldots & 0 & 0 & 0 \\ k_d & \ldots & . & . & . & \ldots & 0 & 0 & 0 \\ . & & & & & & & & \\ . & & & & & & & & \\ 0 & \ldots & k_d & k_{d-1} & k_{d-2} & \ldots & k_1 & 0 & 0 \\ 0 & \ldots & 0 & k_d & k_{d-1} & \ldots & k_2 & k_1 & 0 \\ 0 & \ldots & 0 & 0 & k_d & \ldots & k_3 & k_2 & k_1 \end{pmatrix}.$$

To lay the ground for the main result, we first state a key result from the text by Meyn and Tweedie [10].

Proposition 3.1 *([10] p. 385) Consider a Markov chain $\{X_n\}$ on a state space $\mathcal{X}$ with transition probabilities $P : \mathcal{X} \times \mathcal{B}(\mathcal{X}) \to [0,1]$. Assume that there exists a set $\mathcal{K} \in \mathcal{B}(\mathcal{X})$ and a function $g : \mathcal{X} \to \mathbb{R}$, $g(\cdot) \geq 1$, such that*
(i) for some constants $\hat{\epsilon} > 0$ and $\overline{b} < \infty$,

$$E[g(X_{n+1}) - g(X_n)|X_n] \leq -\hat{\epsilon}\, g(X_n) + \overline{b}\, 1\{X_n \in \mathcal{K}\}, \tag{13}$$

for $X_n \in \mathcal{X}$;
(ii) $\mathcal{K}$ is a petite set, and $\{X_n\}$ is an aperiodic Ψ-irreducible Markov chain, for some maximal measure Ψ.
Then, $\{X_n\}$ is g-uniformly ergodic, and

$$\sum_{r=1}^{\infty} r^n \, \|P^n - \pi\|_g < \infty \tag{14}$$

for some $r > 1$, where π is the invariant measure corresponding to P.

Remark 3.1. A definition for the g-norm used in (14) above can be found in [10] p. 382. A g-uniformly ergodic Markov chain is a Markov chain which is Harris ergodic; moreover, it is geometric ergodic (in the sense that the probabilities converge to the steady state (invariant) measure at a geometric rate); finally, the rate of convergence is "uniform" in the initial state (up to some multiplicative factor given by the function g). In our main stability result, Theorem 3.2, we shall only need the fact that *a g-uniformly ergodic Markov chain is Harris ergodic.* We therefore do not cite here the precise definitions of g-uniform ergodicity and g-norm. The definitions of "petite set", Ψ-irreducibility, and aperiodicity can also be found in [10] (pp. 121, 87, 118, respectively), which we shall not explicitly make reference to either. ∎

We next state two useful lemmas.

Lemma 3.1 *Suppose that $\alpha \in [0,1)$, $\beta > 0$, and that $\{\eta_n\}$ is generated by (8), where C is a stable matrix (all eigenvalues inside the unit disk). Then, the Markov chain Z_n is Ψ-irreducible, aperiodic, and any compact subset of the state space is a petite set.*

Proof. Here we only sketch the proof; details are provided in the Appendix. Ψ-irreducibility of the Markov chain and smallness of compact sets follow from Proposition 6.3.5 of [10], since A is stable, $(A,\, D)$ is controllable, and the noise process $\{\psi_n\}$ has a spreadout distribution. Aperiodicity follows from Proposition 7.3.4 and Theorem 7.3.5 of [10], and from the observation made on p. 158 of the same text. ∎

Before stating the second lemma, let us introduce an $m \times m$ positive-definite matrix Q, where $m :=$ dim (Z), with the property that for any vector $z \in \mathrm{I\!R}^{\mathrm{m}}$,

$$|Az|_Q^2 \leq \gamma \, |z|_Q^2 \,, \tag{15}$$

for some $\gamma < 1$, where $|z|_Q^2 := z'Qz$. Since A is stable,[2] there exists such a Q (i.e. such a norm under which A is contracting); see [6]; p. 237.

Lemma 3.2 *Suppose that $\alpha \in (0,1)$, $\beta > 0$ and that $\{\eta_n\}$ is generated by (8), where C is a stable matrix. Then, the netput process is g-uniform ergodic with $g(z) = |z|_Q^2$, and*

$$\lim_{n\to\infty} E\Delta_n = \frac{v}{1-\alpha} \,. \tag{16}$$

Proof. The limit (16) follows directly from (5). For the first part of the Lemma, we simply show that the conditions of Proposition 3.1 hold. Fix some $\rho \in (1-\alpha,\, 1)$ and define the set

$$\mathcal{K} = \left\{ z : |z|_Q^2 \in \left[0\,,\ \frac{E|D\psi_{n+1}|_Q^2}{\rho} \right] \right\} , \tag{17}$$

which is petite because it is a compact subset of $\mathrm{I\!R}^{\mathrm{m}}$. With $\{\psi_n\}$ defined as in (11), we have from (12), and using (15),

$$\begin{aligned} E\left[|Z_{n+1}|_Q^2 - |Z_n|_Q^2 \,\middle|\, Z_n \right] &\leq (\gamma - 1)|Z_n|_Q^2 + E|D\psi_{n+1}|_Q^2 \\ &= -\rho\,|Z_n|_Q^2 + (\gamma + \rho - 1)\,|Z_n|_Q^2 + E|D\psi_{n+1}|_Q^2 \\ &\leq (\gamma + \rho - 1)\,|Z_n|_Q^2 + E|D\psi_{n+1}|_Q^2\, 1\{Z_n \in \mathcal{K}\}. \end{aligned}$$

[2]This is because $|\alpha| < 1$ and the ARMA process in (7) is stable (i.e. all eigenvalues of the matrix in (9) – and hence of C – are in the interior of the unit disk).

This implies condition (i) of Proposition 3.1, with $\hat{\epsilon} = 1 - \rho - \gamma > 0$ and $\overline{b} := E|D\psi_{n+1}|_Q^2$ which is clearly finite. Condition (ii) of the proposition holds by Lemma 3.1. ∎

This now brings us to the main theorem:

Theorem 3.2 *Assume that $\alpha \in (0,1)$, and that $\{\eta_n\}$ is generated by (8), where C is a stable matrix. If $v < 0$, then the queue length process generated by (6) couples within a time that is finite P-a.s., to a process which is stationary ergodic.*

Proof. Since the Markov chain $\{Z_n\}$ is Harris recurrent (by Lemma 3.1, in view of Remark 3.1), it couples within a time that is finite P-a.s. to a stationary ergodic regime ([5] p. 157). In view of the fact that the queue dynamics (6) is a G/G/1 queue, and that $E\Delta_n < 0$, it again follows from standard results as in the proof of Theorem 3.1, that the queue length process couples in a time that is finite P-a.s. to a stationary ergodic regime. ∎

4 Simulation Results

We have tested the family of control schemes analyzed here in the framework of the following scenario: There is one (or more) controlled ATM source(s) in the presence of another exogenous higher priority video source. The video source has a maximum rate of 2Mbps, and the rate at which the controlled source can transmit is upper bounded by 1Mbps; this bound may correspond to the link capacity or to a negotiated peak cell rate (PCR). The information for the ABR source is updated once every 30 msec (which is done by sending an RM cell every 30 msec), and we have taken a minimum negotiated rate of 13.8 Kbps corresponding to one cell every 30 msec (with RM cells not taken in account). The bandwidth that is further available for the ABR traffic (in addition to the min rate) is the leftover by the video traffic, i.e. 2Mbps minus the video transmission rate.

We used the trace of the first 101 seconds of JPEG video of the "star war" (one frame every 30 msec). We considered first a single controlled source with a round trip delay of 61.2 msec.

We are interested in the effect of the control parameters (such as α, β) on the following features: (i) The rate at which the input source adapts to the available transmission rate; (ii) the queue size; (iii) the variability in the transmission rate.

In general, it is desirable to have a high tracking rate between the input and the available rate, low queue sizes (which then implies lower cell loss rates), and low variabilities in the transmission rates.

The average available bandwidth over the duration of the simulation was taken as 196.996 Kbps with a deviation (i.e. the square root of the variance) of 311.605 Kbps.

We considered several cases of α, β and ϵ; in particular, we considered both stable ($v < 0$) and (potentially) unstable ($v = 0$) conditions . We simulated 9 different scenarios, which are depicted in Figures 1-9. The throughputs, queue lengths and netputs (both averages as well as the standard deviations) corresponding to the first 8 figures, are given in the 8 rows of Table 1. In all these cases, a buffer of size 1000 was used, and no losses occured (for the controlled ABR traffic).

The pure rate-matching case (3) for which $v = 0$, is considered in Figs. 1 and 3. In these two cases we see that the throughput is indeed the highest; the available throughput is fully used. However, the queue length for these cases are seen to be very large: 247.1 and 117.8 respectively. The maximum queue sizes are 415 and 214 respectively. These results are compatible with the fact that the pure rate matching has unstable behavior, as shown in [4]. In particular, it was shown there that the queue size converges in distribution to infinity if the variance of the measurement noise is nonzero. In our case, we have no measurement noise, but still, we can conclude from the simulation results that the pure rate matching should not be used, as it results in very large queues.

In contrast, cases 2, 4, 5, 6, 7, 8, all pertain to stable situations, where either full rate matching with drift (Figs. 7 and 8) or partial rate matching has been used.

In cases 2, 4, 5, 7, 8, not more than 90% of the available rate has been used. The queue length varies in these cases between 10.22 and 36.03. The maximum queue sizes vary between 73 (case 7) and 152 (case 5). In case 6, 96% of the available bandwidth is used, yielding an average queue length of 48.45, and a maximum queue size of 163.

The importance of avoiding the pure rate matching is especially important for smaller buffer sizes. Fig. 9 illustrates the pure rate matching for $\alpha = 0.8$, and $\beta = 0.2$; here a buffer of size 400 has been used, yielding a loss rate of 0.031%.

The influence of α on the tracking between the input and output rate is clearly seen from the last column of Table 1, which shows the standard deviation in the netput. The larger the standard deviation is, the worse is the quality of tracking. Thus, high standard deviation of the netput indicates that the input does not track well the output, and thus, the resources are not well used. We see from Table 1 that as α increases from 0 to 0.8, the standard deviation in the netput decreases from 405 to 276. The tracking quality is best seen visually in the figures themselves. Thus, the lower α is, the better is the tracking.

On the other hand, the choice of α has an effect over the variability of the throughput. In cases where it is desirable to have a smooth transmission rate, it is preferable to choose a larger value of α. This is seen in Table 1 in the column describing the standard deviation of the throughput. For α close to 0, the variability (deviation) of the throughput is almost the same as the variability of

N^O	parameters			simulation results					
	α	β	$\epsilon Kbps$	throughput		buffer occupancy		netput	
sim				mean	deviation	mean	deviation	mean	deviation
1	0.8	0.2	0.0	196.419	174.456	247.1	111.9	-0.627	276.44
2	0.8	0.1	0.0	98.9984	86.4767	10.13	14.32	-98.04	281.56
3	0.4	0.6	0.0	196.363	221.074	117.8	49.99	-0.683	320.68
4	0.4	0.5	0.0	164.326	185.030	21.83	18.99	-32.72	308.04
5	0.0	0.9	0.0	175.754	272.415	36.03	20.79	-21.29	395.43
6	0.01	0.97	0.0	188.111	287.157	48.45	28.05	-8.935	405.71
7	0.4	0.6	204.8	69.2178	135.433	10.51	13.65	-127.8	313.72
8	0.4	0.6	40.96	121.882	189.263	15.98	15.24	-75.16	315.25

Table 1: Simulation results for a single controlled source: Throughput, buffer occupancy, and netput

the available bandwidth. The standard deviation in cases 5 and 6 are about 90% of the deviation in the available bandwidth. On the other hand, in case 2, in which $\alpha = 0.8$, the standard deviation of the throughput is less than 30% of the available bandwidth.

The mean netput agrees well with the calculated mean netput (see Lemma 3.2) in all cases except for cases 5, 6 and 7. In the latter, the calculated mean netputs are -19.69, -3.97, and -341 respectively, whereas the simulated ones are -21.29, -8.93 and -127.8, respectively. Accordingly, the mean throughput agrees well with the calculated mean throughput in all cases except for cases 5, 6 and 7; the calculated mean throughputs are 177.29 and 193.02, respectively, for cases 5 and 6; due to the choice of ϵ (it is chosen larger than the mean available bandwidth), the calculated throughput is negative in case 7, which explains the differences in throughputs and netputs in that case. In cases 5 and 6, these differences between the calculated and the actual mean rates are due to the restriction of 1Mbps maximum rate for the source.

Finally, we observe that there is a flat horizontal line from 20 to 24 secs, as can be seen in all queue length plots. This corresponds to the time when the send rate equals to the available bandwidth which equals to the minimum negotiated rate (one cell every 30 msec); see Fig. 1.

We also considered the case of two controlled sources. We adopted a fair control scheme whereby each source adapts to half the estimated available bit-rate. In other words, the rate of source i, $i = 1, 2$, is

$$x^i_{n+1} = \alpha_i \, x^i_n + \beta_i \, \hat{\mu}^i_n / 2 - \epsilon_i, \tag{18}$$

where x^i_n is the rate of source i, and $\hat{\mu}^i_n$ is the estimated available bit-rate by source i. We chose $\mu^1_n = \mu^2_n$, since in the control of ABR traffic it is the switch that has to estimate the available bandwidth (since it is the switch that informs the sources about the control actions). We chose

$\alpha_1 = \alpha_2 = \alpha, \beta_1 = \beta_2 = \beta$, and $\epsilon_1 = \epsilon_2 = \epsilon/2$. Moreover, we took again the buffer size to be 1000, and each source is limited by a maximum cell rate of 0.5 Mbps. The minimum negotiated rate guaranteed to each source is 6.9 Kbps.

This special choice of parameters implies that, in the case when the two sources are *equidistant*, the sum of the two controlled sources behave as the previous single source. More precisely, the total instantaneous netput (i.e. the sum of input rates from both sources minus the available bandwidth), the (instantaneous) buffer occupancy, the maximum buffer size, and the cell loss rates are practically the same as those obtained for a single controlled source (with the parameters as in the beginning of this section) that follows (4). Slight differences in the queue size of the order of a single packet do occur due to the synchronization between the sources (indeed, even if the sum of the throughputs of the two sources is the same as that of the single source we had before, two packets might be sent by the two sources at the same time, whereas in the case of a single source, two consecutive packets are always spaced by at least 1/PCR, where PCR is the peak cell rate).

We thus concentrated on cases when two sources are not equidistant. The round trip delay for source 1 was taken as 0.1212 sec, and for source 2 as 0.0412 sec. Table 2 lists the maximum queue sizes and compares them with the corresponding figures for a single controlled queue. The exogenous traffic and the duration of the simulation have been taken to be the same as in the previous simulations. Figs. 10-17 depict the trajectories of input rates for both sources. As can be seen from Table 2, the maximum queue length is larger in this case as compared with the case of a single controlled source.

In cases 1-8 we did not obtain any losses. Next, we chose the same parameters as in case 1, but decreased the buffer size from 1000 to 400. The loss rate was then 0.133 % for source 1 and 0.117 % for source 2. We conclude that larger delays result in larger loss probabilities.

5 Conclusions

We have analyzed in this paper reactive rate-based flow control mechanisms that rely only on information on the available bit rate. The rate of the controlled source is to be adapted to the available bandwidth, which may vary in time due to higher priority traffic. We have shown that control schemes that use only rate information may still lead to small queue sizes, and thus, to small loss rates, provided that we attempt to use only a fraction of the available bandwidth.

In contrast, we have tested the pure rate-matching (3), which was shown in [4] to be unstable: the queue sizes were shown in [4] to converge in distribution to infinity if the variance of the measurement noise is nonzero. We observed that, indeed, very large queues build up in these cases, even though we do not have any measurement noise.

	parameters			*buffer size*				*thrughput*	
				single source		*two sources*			
N^O	α	β	ϵ_i	*max*	*mean*	*max*	*mean*	*source1*	*source2*
1	0.8	0.2	0.0	415	247.1	460	273.9	98.207	98.245
2	0.8	0.1	0.0	98	10.13	107	10.53	49.497	49.517
3	0.4	0.6	0.0	214	117.8	249	142.7	98.167	98.236
4	0.4	0.5	0.0	137	21.83	144	26.541	82.15	82.208
5	0.0	0.9	0.0	152	36.03	159	40.774	87.842	87.934
6	0.01	0.97	0.0	163	48.45	194	61.858	94.017	94.118
7	0.4	0.6	102.4	73	10.51	85	12.012	34.625	34.659
8	0.4	0.6	20.48	125	15.98	135	19.537	60.924	60.986

Table 2: Buffer size and throughput for the case of two sources (the buffer size is compared to that of a single source)

Our simulation analysis illustrates the playoff between the choice of the parameters: small values of α were shown to yield good tracking, while large values of α yield smoother throughput.

6 Appendix: Proof of Lemma 3.1

Controllability of the pair (A, D) can be seen in various ways. The most standard approach is to define the *controllability* matrix:

$$R := [\, D \,|\, AD \,|\, A^2 D \,|\, ... \,|\, A^{d'} D \,].$$

and to show that it is full rank, that is that

$$\text{rank } R = d' + 1. \tag{19}$$

A routine computation shows that this is indeed the case. However, there is a more direct method here, because of the special structures of A and D. Clearly the AR process η_n (see (8) is controllable from its input $\{\phi_n\}$ (see e.g. [10] p. 96). Furthermore, since it feeds into the first-order system (5) which is trivially controllable from its input $\{e_n\}$ (since $\beta \neq 0$), it readily follows that the combined system (11), viewed as a deterministic system, is controllable from the two-dimensional input $\{\psi_n\}$.

Now, (19), together with the fact that all eigenvalues of A are in the interior of the unit-disk, and the assumption that the noise process $\{\psi_n\}$ has i.i.d. components with spreadout distribution, imply that the Markov chain $\{Z_n\}$ is Ψ-irreducible, and that all compact sets are “small sets”, which are further (by their definition in [10]) petite (see Proposition 6.3.5 of [10]).

It remains to establish aperiodicity. For any fixed initial state z_0, the set M of reachable points is the whole state space $\mathbb{R}^{d'+1}$. This follows from (19), and from the observations made on p. 16 of [10]. This set is the unique minimal set, as defined in [10] p. 158, and is obviously connected. Thus, by Proposition 7.3.4 of [10], it is aperiodic. The aperiodicity of the Markov chain then follows from Theorem 7.3.5 of [10]. ∎

Acknowledgment The authors wish to thank Daniel Kofman and Sean Meyn for helpful discussions, and Jean Bolot, for providing them with the video traces of "Star-war".

References

[1] The ATM Forum Technical Committee, *Traffic Management Specification*, Version 4.0, October 1995.

[2] E. Altman and T. Başar, "Optimal Rate Control for High Speed Telecommunication Networks", *Proceedings of the 34th IEEE Conference on Decision and Control*, pp. 1389-1394, New Orleans, LA, December 13-15, 1995.

[3] E. Altman and T. Başar, "Optimal Rate Control for High Speed Telecommunication Networks: the case of delayed information", First Workshop on ATM Traffic Management, WATM'95, IFIP, WG.6.2 Broadband Communication, Paris, pp. 115-122, December 1995.

[4] E. Altman, F. Baccelli and J. C. Bolot, "Discrete-time analysis of adaptive rate control mechanisms", *High Speed Networks and their performance*, Perros and Viniotis Eds., North Holland, 121-140, 1994.

[5] S. Asmussem, *Applied Probability and Queues.* J. Wiley and Sons, 1987.

[6] T. Başar, "Decentralized multicriteria optimization of linear stochastic systems", *IEEE Transactions on Automatic Control*, AC-23(2):233–243, April 1978.

[7] F. Bonomi and K. W. Fendick, "The rate-base flow control framework for the Available Bit Rate ATM service", *IEEE Network*, pp. 25-39, March/April 1995.

[8] A. A. Borovkov, *Asymptotic Methods in Queuing Theory*, John Wiley & Sons, 1984 (translated from Russian).

[9] A. A. Borovkov and S. G. Foss, "Stochastically recursive sequences and their generalizations", *Siberian Advances in Mathematics*, **2**, No. 1, pp. 16-81, 1992.

[10] S. P. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*, Springer Verlag, 2nd printing, 1994.

[11] K-T. Ko, P. Mishra, "Interactions among virtual circuits using predictive congestion control", to appear in *Computer Networks and ISDN Systems*, vol. 25, no. 6, Jan. 1993.

[12] R. Loynes, "The stability of a queue with non-independent inter-arrival and service times", *Proc. Cambr. Phil. Soc.* **58**, No. 3, pp. 497-520, 1962.

[13] H. Ohsaki, M. Murata, H. Suzuki, C. Ikeda and H. Miyahara, "Rate-based congestion control for ATM networks", *Computer Communication Review, acm-sigcomm*, special issue in ATM, ed. R. Jain and K. Y. Siu, Vol 25, No. 2, pp. 60-71, 1995.

[14] K. Y. Siu and H. Y. Tzeng, "Intelligent congestion control for ABR service in ATM networks", *Computer Communication Review, acm-sigcomm*, Vol. 24, No. 5, pp. 81-106, 1995.

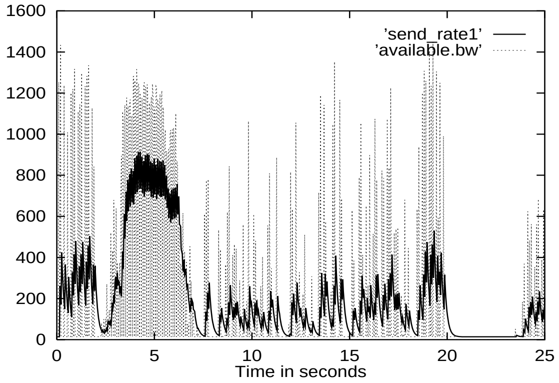

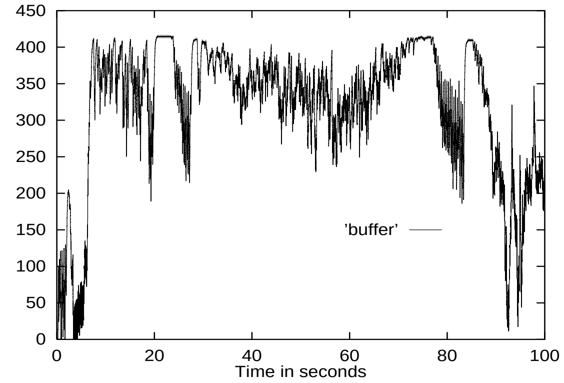

Figure 1: Send rate, available bandwidth (Kbps) and number of cells in the buffer as function of time for $\alpha = 0.8$, $\beta = 0.2$, $\epsilon = 0$ and buffer size = 1000 cells.

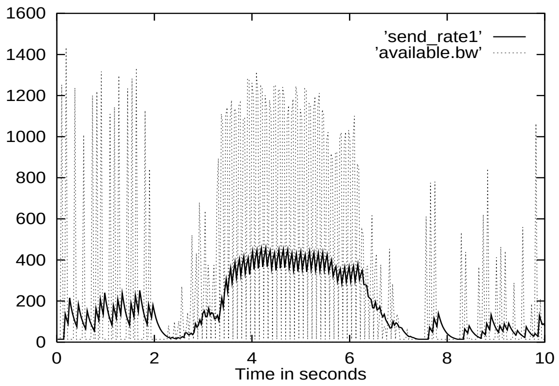

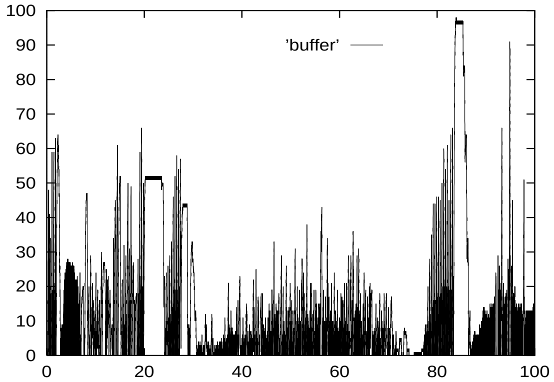

Figure 2: Send rate, available bandwidth (Kbps) and number of cells in the buffer as function of time for $\alpha = 0.8$, $\beta = 0.1$, $\epsilon = 0$ and buffer size = 1000 cells.

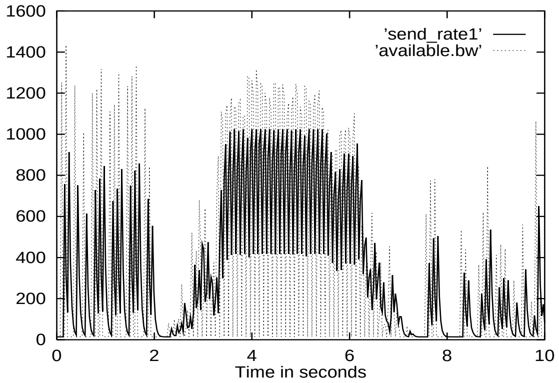

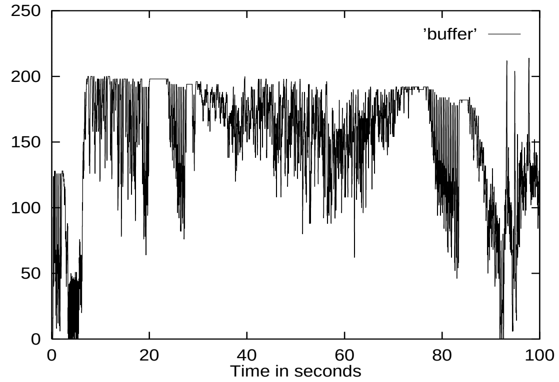

Figure 3: Send rate, available bandwidth (Kbps) and number of cells in the buffer as function of time for $\alpha = 0.4$, $\beta = 0.6$, $\epsilon = 0$ and buffer size = 1000 cells.

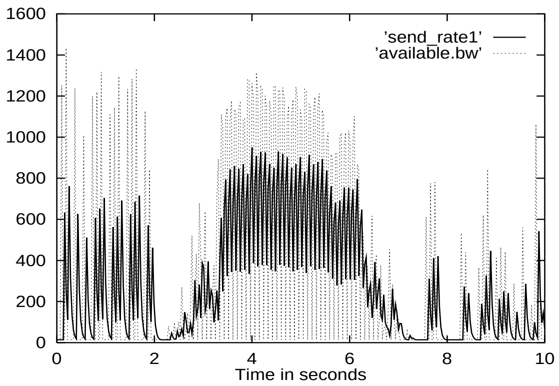

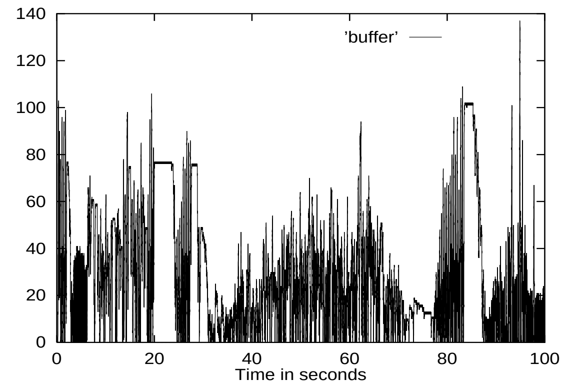

Figure 4: Send rate, available bandwidth (Kbps) and number of cells in the buffer as function of time for $\alpha = 0.4$, $\beta = 0.5$, $\epsilon = 0$ and buffer size = 1000 cells.

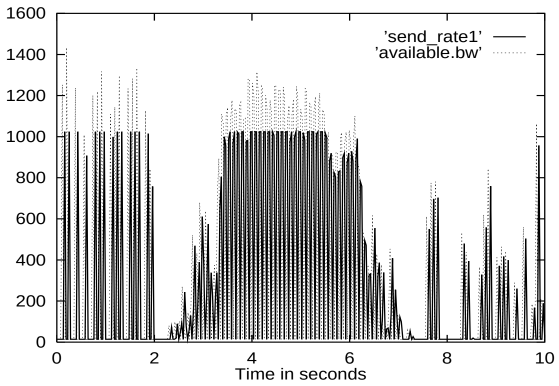

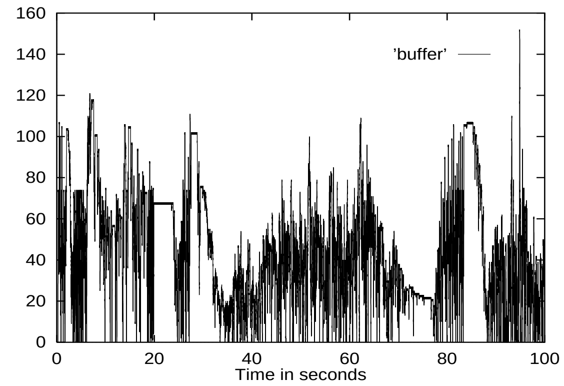

Figure 5: Send rate, available bandwidth (Kbps) and number of cells in the buffer as function of time for $\alpha = 0$, $\beta = 0.9$, $\epsilon = 0$ and buffer size = 1000 cells.

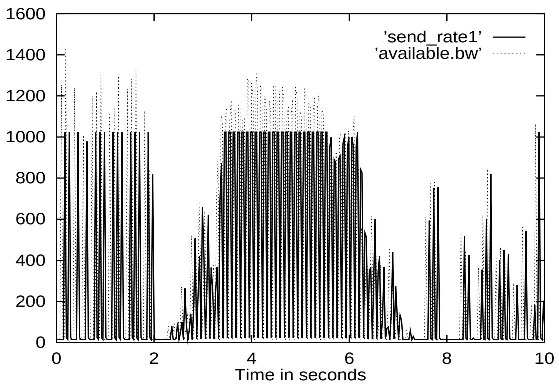

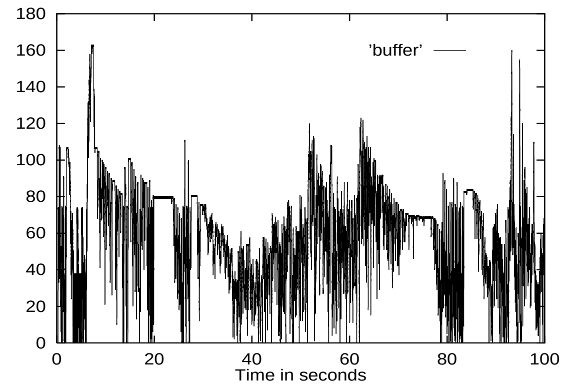

Figure 6: Send rate, available bandwidth (Kbps) and number of cells in the buffer as function of time for $\alpha = 0.01$, $\beta = 0.97$, $\epsilon = 0$ and buffer size = 1000 cells.

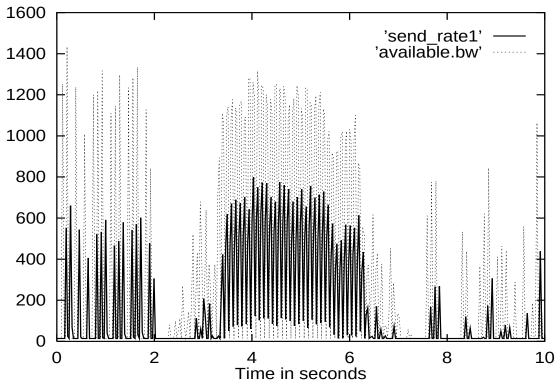

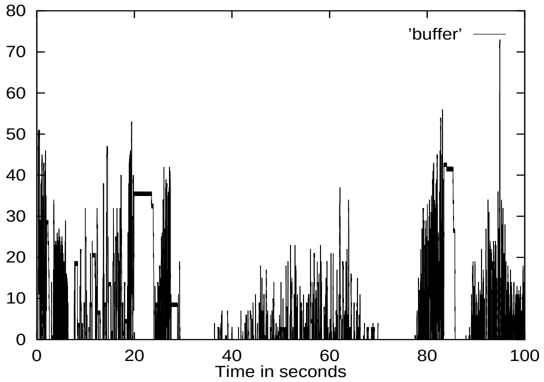

Figure 7: Send rate, available bandwidth (Kbps) and number of cells in the buffer as function of time for $\alpha = 0.4$, $\beta = 0.6$, $\epsilon = 10\%\mu = 204.8Kbps$ and buffer size = 1000 cells.

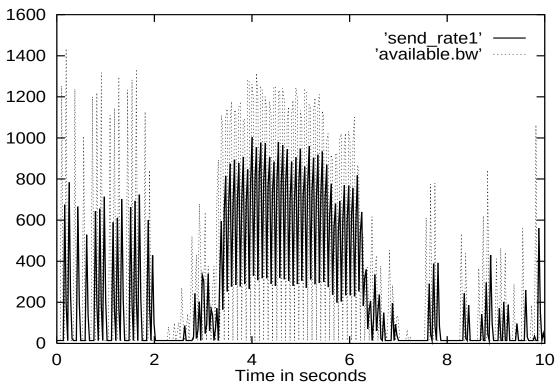

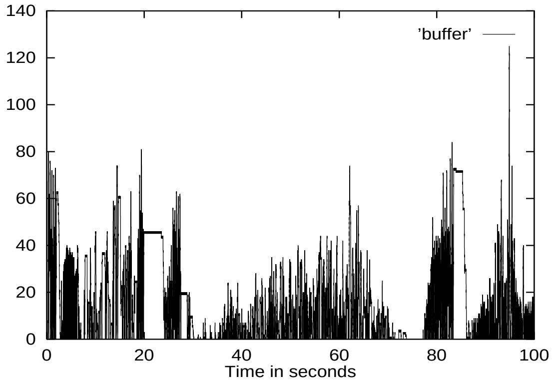

Figure 8: Send rate, available bandwidth (Kbps) and number of cells in the buffer as function of time for $\alpha = 0.4$, $\beta = 0.6$, $\epsilon = 4\%\mu = 40.96Kbps$ and buffer size = 1000 cells.

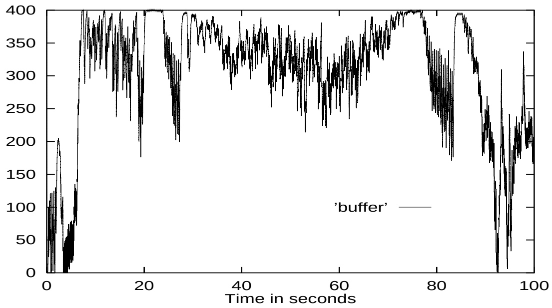

Figure 9: Number of cells in the buffer as function of time. Buffer size = 400 cells, $\alpha = 0.8$ $\beta = 0.2$ and $\epsilon = 0$. Loss rate = 0.031% for 101 sec of simulation.

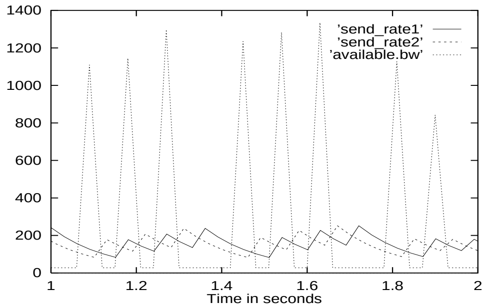

Figure 10: Send rate, available bandwidth (Kbps) as function of time for $\alpha = 0.8$, $\beta = 0.2$, $\epsilon = 0$ and buffer size = 1000 cells.

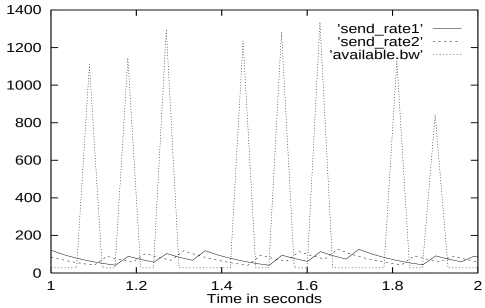

Figure 11: Send rate, available bandwidth (Kbps) as function of time for $\alpha = 0.8$, $\beta = 0.1$, $\epsilon = 0$ and buffer size = 1000 cells.

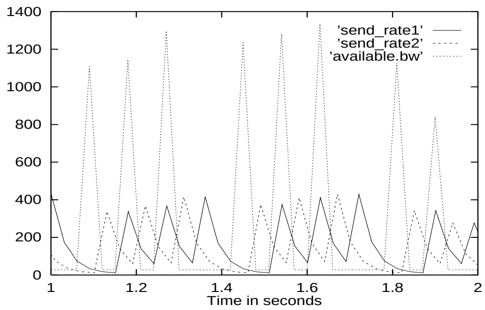

Figure 12: Send rate, available bandwidth (Kbps) as function of time for $\alpha = 0.4$, $\beta = 0.6$, $\epsilon = 0$ and buffer size = 1000 cells.

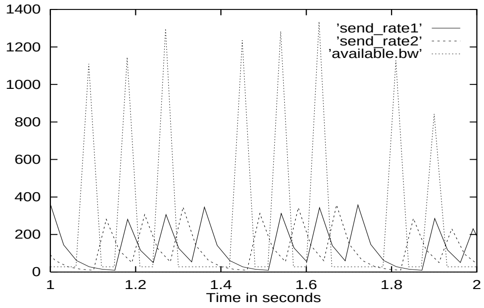

Figure 13: Send rate, available bandwidth (Kbps) as function of time for $\alpha = 0.4$, $\beta = 0.5$, $\epsilon = 0$ and buffer size = 1000 cells.

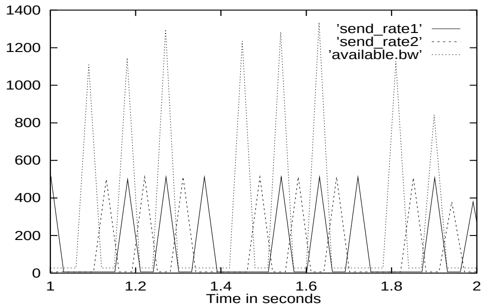

Figure 14: Send rate, available bandwidth (Kbps) as function of time for $\alpha = 0$, $\beta = 0.9$, $\epsilon = 0$ and buffer size = 1000 cells.

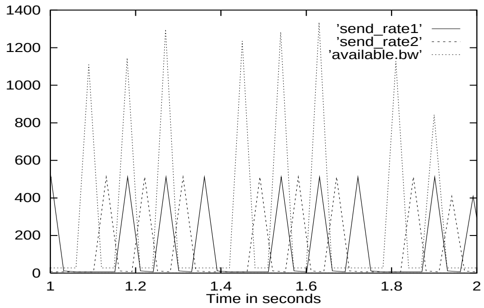

Figure 15: Send rate, available bandwidth (Kbps) as function of time for $\alpha = 0.01$, $\beta = 0.97$, $\epsilon = 0$ and buffer size = 1000 cells.

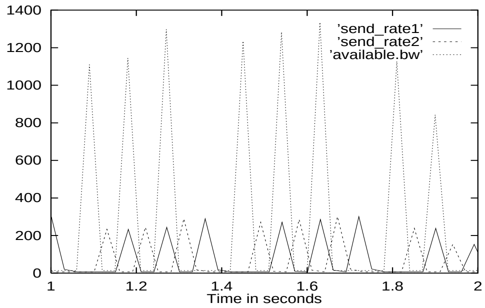

Figure 16: Send rate, available bandwidth (Kbps) as function of time for $\alpha = 0.4$, $\beta = 0.6$, $\epsilon = 10\%\mu = 102.4Kbps$ and buffer size = 1000 cells.

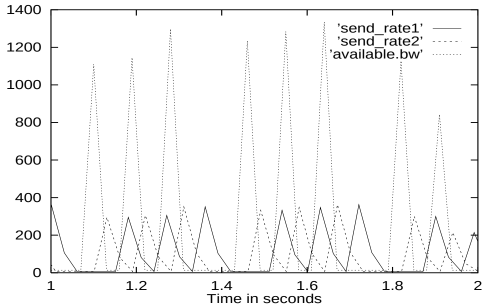

Figure 17: Send rate, available bandwidth (Kbps) as function of time for $\alpha = 0.4$, $\beta = 0.6$, $\epsilon = 4\%\mu = 20.48Kbps$ and buffer size = 1000 cells.