In-call Renegotiation of Traffic Parameters

Simon Crosby
University of Cambridge Computer Laboratory
Pembroke Street, Cambridge, CB2 3QG, U.K.

Abstract

We construct an approximate model of an ATM multiplexer to permit comparison of the performance of static and dynamic resource allocation schemes. Static allocation parameterises a source at connection setup only, whereas dynamic allocation requires a source to inform the multiplexer whenever its bandwidth requirements change. We analyse a simple dynamic "renegotiation" protocol and derive bounds on the performance of an optimal dynamic protocol. Markovian source models parameterised from working applications are used to contrast the behaviour of the three control schemes.

1 Introduction

In an ATM network sources are required to "reserve" the resources they require before they commence transmission. The reservation may be expressed in terms of the peak and mean bandwidth requirements and some measure of the average or maximum burst length of the source [3]. Additional parameters describing the Quality Of Service (QOS) required by the source may specify the maximum tolerable cell loss rate, and bounds on jitter and delay [1, 11, 17]. The Call Acceptance Control (CAC) uses these traffic descriptors to decide whether or not to accept the call. In accepting a call the network is obliged to ensure that its specified QOS is met.

It has been suggested that allocation of resources (bandwidth and buffer space) at connection setup may result in inefficient use of the network [3, 8] and/or reduce the QOS experienced by sources [6, 5]. For Constant Bit Rate (CBR) sources static allocation of resources is clearly adequate, but for bursty sources it may result in wasted capacity because the sources only rarely transmit at their peak rates. Dynamic resource allocation attempts to address this problem by requiring a source to inform the network whenever its resource requirements change. The network may reject a request for additional resources if it has insufficient capacity or if meeting the request would reduce the QOS experienced by other sources [9, 8, 3, 7]. Fast reservation protocols which can quickly allocate the peak rate to bursts of traffic within a connection have been proposed [4, 16]. Following [8] we refer to in-call parameter negotiation as "renegotiation".

The benefits of renegotiation are twofold: it prevents congestion and consequently data loss, by exerting direct control over the sources; it also permits the network to optimise its resource usage whilst meeting QOS guarantees. Linking charges to resource usage provides an incentive for sources to release capacity they do not require. The drawback of renegotiation is its cost: its implementation requires substantially more protocol overhead in terms of signalling between sources and the network. In addition the renegotiation delay per burst increases the end-to-end delay for the transmission of a fixed amount of data. The utility of renegotiation thus depends on the dynamic behaviour of the sources and the efficiency of the protocol used.

Doshi and Heffes [9] analyse a scheme which permits a source to renegotiate for an increase in window size and buffer space at network nodes when it has a burst of data to transmit. The buffer allocation depends on the amount of surplus buffer space, and is allocated for the duration of the burst. They conclude that renegotiation performs better than static allocation when buffer space is limited. A similar scheme is discussed in [8] which proposes a Fast Reservation Protocol (FRP) for use in broadband networks. An analysis using Erlang's formula yields the probability that a burst request will be blocked. Assuming Poisson burst arrivals, a queueing model is used to calculate the utilisation of nodal buffer space and end-to-end transmission delay. The results indicate that whilst renegotiation permits the network to operate with fewer buffers, the bandwidth of the network tends to be underutilised.

When the ratio of the source peak rate to the link rate is high, static allocation performs poorly for bursty traffic. In [7] renegotiation is proposed as an alternative to static allocation for high speed file transfer. Sources alternate between periods of activity and inactivity; when active a source transmits small files with a short, random interval between each. Sources renegotiate for buffer space either on transition from the inactive to active state or between file transfers. Results obtained from a queueing model indicate that whilst renegotiation before each file transfer yields the best blocking performance, under heavy load it leads excessive renegotiation and hence delay. Renegotiation per activity period reduces delay, but increases the blocking probability.

In this paper we contrast the performance of static allocation and renegotiation in an ATM link. Each source is modelled by a multi-state Markov process. Under static allocation a source is parameterised at connection setup only and thereafter varies its resource requirements randomly. Transmitted traffic in excess of the link capacity is lost. The performance measure of interest is the loss probability of a burst of traffic. Under renegotiation a source informs the network every time its requirements change. Traffic loss is prevented and following [8] we measure network performance by the probability that a renegotiation request will be rejected. We compare static allocation to a simple "on the fly" renegotiation scheme, before constructing bounds on the performance of a renegotiation scheme which is optimal in the sense that it minimises the cost of renegotiation failure.

Our results confirm the observation of [9] that when the source peak to link rate ratio is high, renegotiation is beneficial. We contend, however that in high capacity networks where a significant degree of statistical multiplexing occurs, a single parameter describing the *effective bandwidth* of a source is sufficient to permit the multiplexer to decide statically whether or not to accept an additional call: if the sum of the effective bandwidths of all connected sources is less than the capacity of the link, then the required performance guarantees can be met. Consequently the benefits of renegotiation are small: the gain in link utilisation is not sufficient to warrant the implementation of a renegotiation protocol. In the local area, where network capacities are smaller, the renegotiation gain can be significant, particularly for the optimal scheme. Such networks are frequently under-utilised, however, and therefore the additional complexity of a renegotiation protocol may not be warranted.

We note that the effective bandwidth is not a suitable parameter for use in charging algorithms because the effective bandwidth of a class of traffic is a function of the total load on the multiplexer, of all traffic

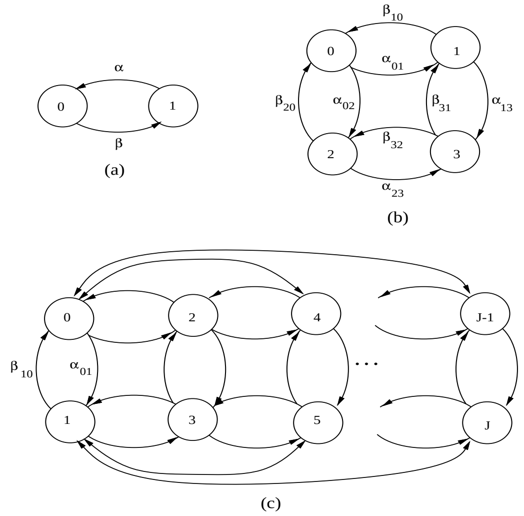

Figure 1: Type 1 Markov sources

types, and varies according to the degree of statistical multiplexing attainable at each link.

2 Source Models

We distinguish between sources according to the time scales of their burstiness. Denote the arrival rate of new connection requests by λ and the arrival rate of bursts from a single source by ν. Consider a link carrying data from n independent sources. The following interesting cases arise:

1. $\lambda \geq n\nu$: New calls arrive at a higher rate than bursts. An example is a data source involved in large file transfers with long silence periods [9, 8]. If lightweight virtual circuits are used [14, 15], renegotiation may be more expensive than setting up a new connection for each burst.

2. $\lambda < n\nu$: New calls arrive less frequently than bursts, but at a rate which justifies the implementation of an efficient renegotiation protocol such as XON/XOFF [12], FRP [3, 9], or FRP/DT [16]. It is this case which is analysed in this paper.

3. $\lambda \ll n\nu$: Bursts arrive too fast to permit the use of renegotiation, for example in silence suppressed voice traffic, where a "talkspurt" arrives approximately every 220 ms.

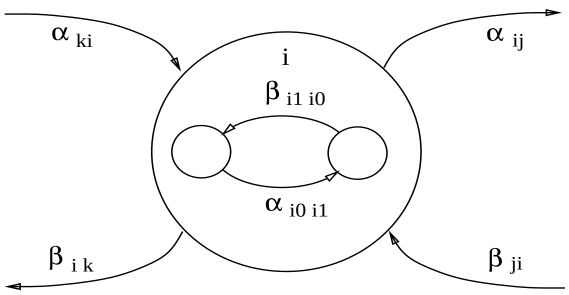

Figure 2: State i in a type 2 Markov source

Each source is represented as a Markov process with J states, in which transitions between states of the Markov process represent changes in the transmission rate of the source. Previous models of renegotiation have used only the "on-off" source [1], a two state birth-death process which transmits data at a (constant) rate r in state 1 and is idle in state 0.

We consider two source types: a type 1 source transmits at a constant rate in each state, whereas a type 2 source transmits at a variable rate. Figure 1 shows three type 1 sources used in the numerical investigation of section 6: (a) the "on-off" source, (b) a simple 4 state source representing a video conferencing application and (c) a more complex, rate-switchable video source. For a type 1 source the equilibrium state probability distribution is denoted by $\vec{p} = (p_0, \ldots, p_J)$. The birth rate from state i to j is denoted by α_{ij}, and the death rate from state j to i by β_{ji}. The mean bandwidth requirement is given by $\bar{r} = \sum_{j=0}^{J} r_j p_j$, where r_j is the bandwidth requirement in state j, and the peak bandwidth requirement is denoted by $\max_j(r_j)$.

Type 2 sources are bursty on two time scales – a slow scale similar to that of the type 1 sources and a faster scale corresponding to fluctuations in transmission rate within a state. Each state of the type 1 model is replaced by a "miniprocess" [1], and we refer to bursts on the fast time scale as "minibursts". Figure 2 shows a single state i of a type 2 source with a miniprocess representing the fluctuations in transmission rate within the state. The source description includes $\vec{q_j} = (q_{j0}, \ldots, q_{jm_j})$ for $(1 \le j \le J)$, the equilibrium probability distribution of the ministates, where m_j is the number of ministates in state j, $\vec{p}$, and $\vec{r_j} = (r_{j0}, \ldots, r_{jm_j})$, the bandwidth requirements in each state. The mean bandwidth is given by $\bar{r} = \sum_{j=0}^{J} p_j \sum_{i=0}^{m_j} q_{ji} r_{ji}$ and the peak by $\max_{ji}(r_{ji})$.

In the following sections we analyse the multiplexer under static allocation and renegotiation, for both source types. Section 3 introduces the Large Deviations Approximation and derives a model for the system under static allocation. Section 4 derives approximate bounds for the performance of the multiplexer under a simple renegotiation protocol. Section 5 constructs bounds on the performance of all optimal renegotiation schemes, and Section 6 contrasts the performance of the three schemes by means of a numerical investigation.

3 System with Static Allocation

3.1 Type 1 Sources

Consider a link of capacity C, subjected to a traffic load from n identical type 1 sources. The Large Deviations Approximation (LDA) [13] for a single, unbuffered resource is constructed as follows: denote by $\boldsymbol{P}()$ and $\boldsymbol{E}[]$ the probability measure and expectation respectively. Let $X_1, X_2, \ldots X_n$ be iid random variables describing the bandwidth requirements of each of the n sources. For each X_i define the logarithmic moment generating function

$$M(t) = \log \boldsymbol{E}\left[e^{tX_i}\right],$$

where $t \ge 0$, and let

$$\inf_t M(t) = \rho.$$

The LDA for the sum of n random variables X_i states that

$$\lim_{n\to\infty} \frac{1}{n} \log \boldsymbol{P}\left(\sum_{i=1}^{n} X_i \ge 0\right) \le \inf_t M(t).$$

The approximation is asymptotically exact in the limit $n \to \infty$ provided that $\boldsymbol{E}[X_i] < 0$ and $\boldsymbol{P}(X_i > 0) > 0$. Let $S_n = \sum_{i=1}^{n} X_i$, where $X_i = r_j$ with probability p_j, $1 \le j \le J$. Then

$$\boldsymbol{P}(S_n \ge na) = \boldsymbol{P}\left(\sum_{i=1}^{n} (X_i - a) \ge 0\right),$$

and Chernoff's theorem applies for $\bar{r} < a < \max_j(r_j)$. Here

$$\begin{aligned} M(t) &= \boldsymbol{E}\left[e^{t(X_i - a)}\right] \\ &= e^{-ta} \sum_{j=1}^{J} p_j e^{t r_j}. \end{aligned}$$

Minimising $M(t)$ entails finding the roots of the polynomial

$$e^{-ta}\sum_{j=1}^{J} p_j(r_j - a)e^{tr_j} = 0.$$

Bean [2] has shown that this polynomial in e^t always has exactly one nonnegative root, τ, so we are assured that the infimum exists. In this case an upper bound on the probability of overload is given by

$$\boldsymbol{P}(S_n \geq na) \approx e^{n\rho}, \qquad (1)$$

where $\rho = e^{-\tau a}\sum_{j=1}^{J} p_j e^{\tau r_j}$.

The LDA suggests the existence of an *effective bandwidth* $\hat{r}$ for a bursty source, which is dependent on the source characteristics and the degree of statistical multiplexing achieved at the link, such that if $n\hat{r} < C$ then the required performance guarantee (1) can be met [11].

3.2 Type 2 Sources

Here the source model includes additional burstiness due to the ministates within each state of the source. to obtain the probability of overload, P_o. Let $S_n = \sum_{i=0}^{n} X_n$, as before. Assume that X_i is distributed as Y_j with probability p_j, such that $\boldsymbol{P}(Y_j = r_{ji}) = q_{ji}$, for $0 < i \leq m_j$. Here i is the index of the ministates within each state j of the source model. Now

$$\boldsymbol{P}(S_n \geq na) \approx M(a, \tau)^n \qquad (2)$$

for $\bar{r} < a < \max_{ji}(r_{ji})$. Here

$$M(a,t) = e^{-ta}\sum_j p_j \sum_{i=0}^{m_j} q_{ji}e^{tr_{ji}}.$$

Thus τ satisfies

$$e^{-ta}\sum_j p_j \sum_{i=0}^{m_j} q_{ji}(r_{ji} - a)e^{tr_{ji}} = 0.$$

Equations (1) and (2) describe the performance of the system with static capacity allocation. For given capacity C and maximum probability of burst loss P_o we can determine the maximum number of sources n which the link can accommodate.

4 System with Renegotiation

In this section we analyse a simple renegotiation protocol which functions as follows: whenever a source moves to state requiring more bandwidth it sends a renegotiation request to the network. If the link has sufficient capacity to grant the request then it is accepted, otherwise it is refused. A source whose request is refused is assumed to wait for a random length of time before again changing state and requesting more bandwidth.

4.1 Type 1 Sources

The system with renegotiation avoids burst loss because sources are prevented from transmitting if their traffic would congest the network. Consider the Markov process S with state space $\mathcal{S}$, formed from the superposition of n independent Markov source processes of the types described in section 2. A state $s \in \mathcal{S}$ represents a tuple $(n_0, \ldots, n_J)$, where n_j is the number of sources in state j. The bandwidth requirement of state s is $R_s = \sum_{j=0}^{J} n_j r_j$. Using the renegotiation scheme described above, the network will refuse a request for more bandwidth only when it is in a state such that $R_s > C - \Delta$ and the request is for an additional Δ units of bandwidth. Thus under renegotiation the allocated capacity of the link can be represented by a Markov process T formed by truncating the state space of S so that it contains only states s with $R_s < C$. Denote the state space of T by $\mathcal{T}$.

Define $\sigma_{ij} = \{s \in \mathcal{S} | R_s < C, R_s + \Delta_{ij} \geq C\}$, and $\Delta_{ij} = (r_j - r_i)^+$, for $0 \leq i, j \leq J$. Here Δ_{ij} is the increase in bandwidth requested by a source when it moves from state i to state j. The set σ_{ij} contains all states in $\mathcal{S}$ from which a single renegotiation request of size Δ_{ij} will be rejected. There is one such set for every different request size. Define $\theta = \{s \in \mathcal{S} | R_s \geq C\}$ as the set of states in which the link is overloaded, and let $\pi_\theta = \sum_{s \in \theta} \pi_s$, where π_s is the equilibrium probability of state $s \in \mathcal{S}$. The exact long term probability of renegotiation failure, P_b, could be computed using the process S, but its size and complexity make this infeasible. Instead we derive approximate bounds for P_b using the LDA and the truncated superposed source process T.

Let $\pi_{\sigma_{ij}} = \sum_{s \in \sigma_{ij}} \pi_s$ denote the equilibrium probability that the superposed process is in the set σ_{ij}, and let $\Delta_{\max} = \max_{ij}(\Delta_{ij})$, corresponding to the set $\sigma_{\max}$. Then for all i, j the relation $\sigma_{ij} \subseteq \sigma_{\max}$ holds, and so $\pi_{\sigma_{\max}} \geq \pi_{\sigma_{ij}}$. Denote the summation $\sum_{i=0}^{J}\sum_{j=0, j\neq i}^{J}$ by $\sum_{i,j}$. Clearly

$$\begin{aligned} P_b &\leq \sum_{s \in \sigma_{\max}} \pi_s \sum_{i,j} \frac{\alpha_{ij} n_i}{\sum_{k,l} \alpha_{kl} n_k} \\ &= \pi_{\sigma_{\max}} \end{aligned} \qquad (3)$$

where n_i is the number of sources in state i. Thus $\pi_{\sigma_{\max}}$ is an upper bound on P_b. Similarly let $\Delta_{\min} = \min_{ij}(\Delta_{ij})$, corresponding to the set $\sigma_{\min}$. Then for all i, j the relation $\sigma_{\min} \subseteq \sigma_{ij}$ holds, and so $\pi_{\sigma_{\min}} \leq \pi_{\sigma_{ij}}$. Now

$$\begin{aligned} P_b &\geq \sum_{s \in \sigma_{\min}} \pi_s \sum_{i,j} \frac{\alpha_{ij} n_i}{\sum_{k,l} \alpha_{kl} n_k} \\ &= \pi_{\sigma_{\max}} \end{aligned} \tag{4}$$

and it is clear that $\pi_{\sigma_{\min}}$ is a lower bound on P_b.

Let $\Theta = \sum_{s \in \mathcal{S} \setminus \mathcal{T}} \pi_s$. Approximate values of the bounds (3) and (4) on P_b can be computed using Equation 1 as follows:

$$\pi_{\sigma_{\max}} \approx [L(C - \Delta_{\max}, n) - L(C, n)]/\hat{\Theta} \tag{5}$$

$$\pi_{\sigma_{\min}} \approx [L(C - \Delta_{\min}, n) - L(C, n)]/\hat{\Theta} \tag{6}$$

where $L(x, n) = \boldsymbol{P}(S_n \geq x)$ is the LDA (1) of section 3.1 and $\hat{\Theta} = 1 - L(C, n)$. Note that using the lower bound $\pi_{\sigma_{\min}}$ for P_b will mean that *more* sources will be admitted to the link in the system with renegotiation than would be admitted if the exact value of P_b were used. Thus the gain due to renegotiation will be overestimated. Similarly using the upper bound $\pi_{\sigma_{\max}}$ will underestimate the gain. In the case of on-off sources the upper bound Equation (3) is exact.

4.2 Type 2 Sources

The bounds derived in section 4.1 are easily extended for type 2 sources. Here we use Equation (2) for $L(x, n)$ in Equations (5) and (6). Two cases arise: in the first, the values of $\Delta_{\max}$ and $\Delta_{\min}$ are chosen taking the bandwidth changes between ministates of the source model into account. This models a scheme with renegotiation at the miniburst level. As stated in the introduction, renegotiation is impractical on this time scale. Instead we model a second scheme, in which renegotiation a the miniburst level is ignored, and a single renegotiation request is made for the duration of transmission from a particular state. In this case the bandwidth requirement of the state is the peak bandwidth required in that state.

5 Bounds on Optimal Performance

The simple renegotiation scheme described in the previous sections is not optimal. Requests are processed "on the fly" and accepted or rejected depending only on their size and the remaining capacity. If

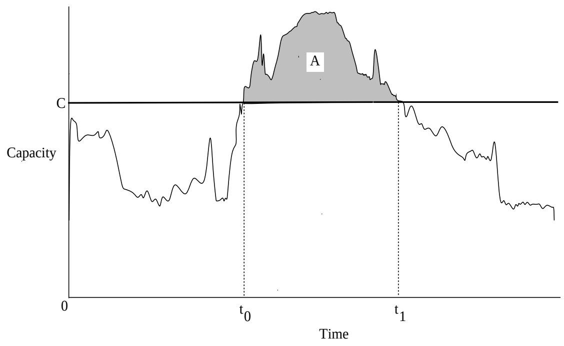

Figure 3: Bandwidth requirements of process S

single request of size Δ is accepted at the expense of m requests of size Δ/m then the rejection probability experienced by the sources will be higher than if the m smaller requests were accepted. An optimal renegotiation scheme might be able to avoid such situations. It might also permit sources to prioritise their requests according to their size or urgency. In this section we describe a model which places bounds on the performance of all such optimal schemes. The details of particular optimal schemes do not concern us, but bounding their performance permits us to calculate the maximum benefits of renegotiation.

Consider the model of an ATM multiplexer and the superposed source process S of section 4, in a network with static allocation. Figure 3 pictures a typical plot of the bandwidth requirement of S over a period of time. Occasionally the system enters overload, as in the interval (t_0, t_1). Denote the portion of traffic lost in (t_0, t_1) by A. Under renegotiation, the same source activity would cause the the traffic A to be rejected. An optimal renegotiation scheme can only minimise the cost of renegotiation (e.g. the number of requests rejected), given that the traffic A must be rejected anyway.

Let $f(c, n) = \boldsymbol{P}(S_n = c)$, and $F(c, n) = \boldsymbol{P}(S_n \leq c)$. The mean capacity in excess of C required to accommodate n sources without loss is given by

$$\begin{aligned} O &= \int_C^\infty (c - C) f(c, n) dc \\ &= \int_C^\infty (1 - F(c, n)) dc \end{aligned} \tag{7}$$

where $(1 - F(c, n)) = P(S_n \geq c)$. An approximate value of O can be computed using the LDA by substituting $L(c, n)$ for $1 - F(c, n)$ in Equation (7), where $L(c, n)$ is Equation (1) or (2) for type 1 and 2 sources respectively. In the long term a proportion $O/(O+C)$ of the offered traffic will be lost due to overload.

On average, the total number of renegotiation requests of size Δ_{ij} made in a period of time δ is $R_{\Delta_{ij}} = \delta\alpha_{ij}m_i$, where m_i is the number of sources in state i when the superposed process S is in its equilibrium state $\vec{m} = (p_0N, \ldots, p_JN)$, and p_j is the equilibrium probability that the source is in state j. A source spends on average $1/\mu_j$ units of time in state j, where μ_j is exponentially distributed with mean $\sum_{i\neq j}\alpha_{ji}$. With each request size Δ_{ij} associate a value $v_{ij} > 0$ which represents the cost incurred by the network when it rejects a request of that size. The most obvious assignment sets $v_{ij} = \Delta_{ij}$, but other assignments are possible.

The optimal renegotiation scheme can do no better than to deny a portion O of the traffic in such a way that the least cost is incurred through rejecting excess traffic. By rejecting a request for more bandwidth the multiplexer prevents transmission of an additional Δ_{ij} units of bandwidth for a period $1/\mu_j$. Let $\delta \to \infty$. The objective of the optimal renegotiation scheme is to find a set of values l_{ij} which

$$\begin{array}{ll} \text{Minimise} & V = \sum_{i,j} l_{ij}v_{ij}, \\ \text{Subject to} & l_{ij} \leq R_{\Delta_{ij}}, \\ & \sum_{ij} l_{ij}\Delta_{ij}/\mu_j = O. \end{array}$$

The probability of renegotiation failure under an optimal scheme is given by

$$P_{\text{opt}} = \frac{\sum_{i,j} l_{ij}}{\sum_{i,j} R_{\Delta_{ij}}}. \qquad (8)$$

We note in passing that in the solution of the above model for the sources described in the following section, the rate of arrival of renegotiation requests with the lowest cost $v_{\min}$ is sufficient to account for the total proportion of rejected traffic. This suggests that an optimal scheme, should one exist, could ensure that only requests of least cost are denied.

6 Numerical Results

In this section we present results of a numerical investigation using the models described above. We use three source models, pictured in Figure 1, whose parameters are summarised in Table 1. Source bandwidth requirements are in Mbit/s. The burstiness is defined to be the ratio of the peak to mean bit rates [3], and for type 2 sources the number of states is the total number of ministates. We assume a link capacity of 100 Mbit/s except where stated.

The "on-off" source is used here to represent a data source, such as a file transfer application. "Xvconfer" and "xvlook" are parameterised from the characteristics of working video phone applications from the Pandora system [10]. "Xvconfer" permits two-way video conversations with a choice of 2 picture sizes and 2 frame rates. "Xvlook" permits a user to "look" through the camera attached to a remote Pandora's Box with a choice of 2 picture sizes at frame rates between 0 and 25 fps. In the type 1 case sources transmit at a constant rate in each state. The type 2 parameters specify the additional burstiness due to the "on-off" behaviour between frames in each state.

To permit a direct comparison of static allocation (SA), simple renegotiation (SR), and optimal renegotiation (OR), we use two measures of performance: the gain due to renegotiation and the link utilisation. Results for SA use a burst loss probability of $P_o = 10^{-8}$, which for a link multiplexing 100 connections results in a much lower cell loss probability per source [11]. Values of P_b (for SR and OR) are chosen between 10^{-4} and 10^{-1}, permitting comparison of our results with those of [9, 8] for "on-off" sources.

Figure 4 plots the gain due to SR for type 1 parameters for "on-off" and "xvlook as a function of the link capacity. Each graph plots the gain for values of P_b, from 10^{-1} to 10^{-4}. The gain is the ratio n_r/n_s where n_r is the number of sources which can be accommodated under renegotiation (for given P_b) and n_s the number for given P_o under SA. Here n_r is obtained from the *lower* bound on P_b (except for "on-off"); thus the gain is optimistic – favouring the system with renegotiation. The gain decreases rapidly with increasing capacity. Although "on-off" and "xvlook" have similar means and variances, "xvlook" has a peak rate double that of "on-off" and consequently it is less easy to multiplex at low capacities and therefore gains far more from SR. At higher capacities, the inflated gain due to the higher peak rate of "xvlook" is lost, and both types have similar performance. With increasing values of P_b, the gain due to SR is far greater at low capacities, but comparatively small for high capacities. In all cases the gain declines rapidly to less than 10% with links in excess of 1 Gbit/s, confirming that in high speed trunks renegotiation is not worthwhile.

With type 2 sources the gain due to renegotiation is greater due to the increased burstiness and variance of the sources. Figure 5 compares SR and OR for the type 2 source parameters, for "xvconfer" and "xvlook". Here $P_b = 10^{-3}$. It is clear that SR performs poorly in comparison to OR at low bandwidths. With increasing link bandwidth, however, the perfor-

Type	Name	Peak	Mean	Burstiness	Variance	States
type 1	on-off	1.0	0.5	2	0.25	2
	xvconfer	2.2	1.09	2.018	0.277	4
	xvlook	2.0	0.405	4.935	0.234	14
type 2	xvconfer	2.2	0.724	3.038	0.547	8
	xvlook	2.0	0.277	7.177	0.243	28

Table 1: Source Parameters

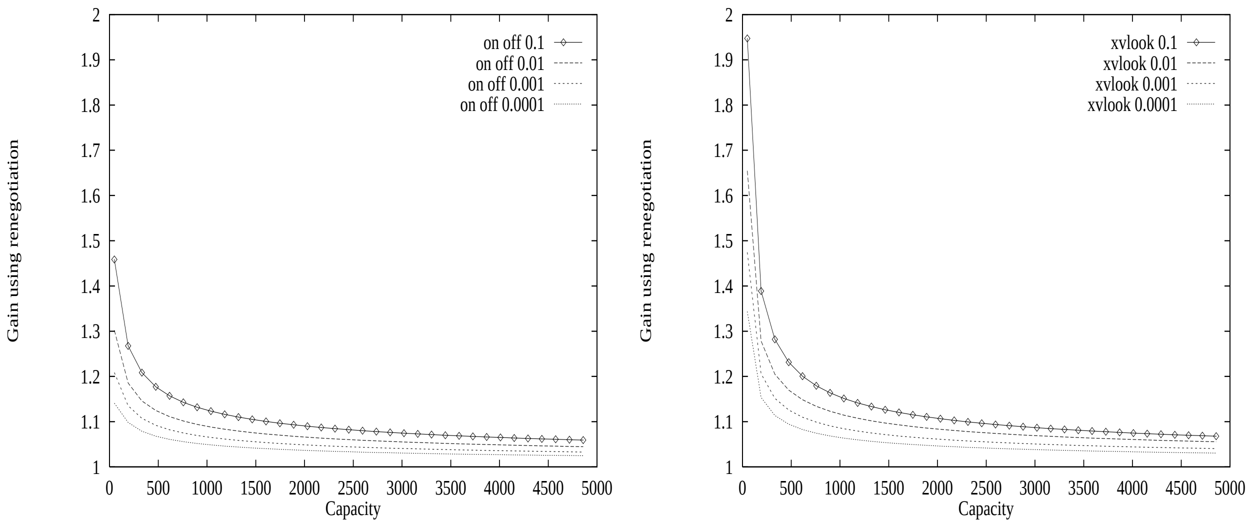

Figure 4: Gain due to renegotiation, type 1 sources

mance for both sources under OR converges, and the difference between SR and OR becomes minimal. The results suggest that if renegotiation is to be used in a low bandwidth link, simple renegotiation may not perform very well.

Of importance to the network operator is the link utilisation which can be attained under different control schemes. Figure 6 shows the link utilisation resulting from SA and SR for "xvlook". Four different values of P_b, indicated in the key to the graph, are used for the simple renegotiation scheme. Most obvious is the distinctly lower utilisation using SA than with SR. This improves somewhat with an increase in link capacity. Under SR a tolerance for high values of P_b permits very high utilisation of the link, even at low capacities, but for more reasonable values of P_b, say 10^{-3}, the utilisation for high bandwidths is greater than 90% for all link capacities of interest, suggesting that implementation of renegotiation will not result in excessive waste of capacity. For very high speed links in excess of 1 Gbit/s the utilisation under SA exceeds 90%, and we conclude again that at these rates renegotiation is not attractive.

Figures 7 and 8 compare, for type 1 and 2 source parameters, the link utilisation under SA, SR and OR. These two graphs summarise very effectively the conclusions of this study: the type 1 schemes perform better in all cases than their type 2 counterparts. This is because the additional variability of the type 2 sources requires the multiplexer to be more conservative when accepting calls, so fewer calls can be in progress. In Figure 7 there is a distinct difference in performance between the different source types for a given control scheme. Thus "xvlook" and "xvconfer" differ significantly under SA, under SR, and at low bandwidths under OR. In Figure 8 it is clear that the differences within each scheme are minimal, and that there are three limiting regimes: Optimal renegotiation performs best, followed by simple renegotiation and then static allocation. For type 2 sources static allocation

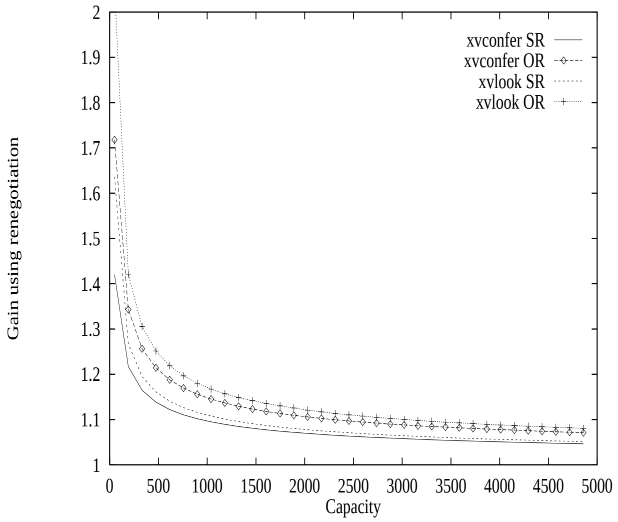

Figure 5: Gain due to SR and OR, type 2 sources

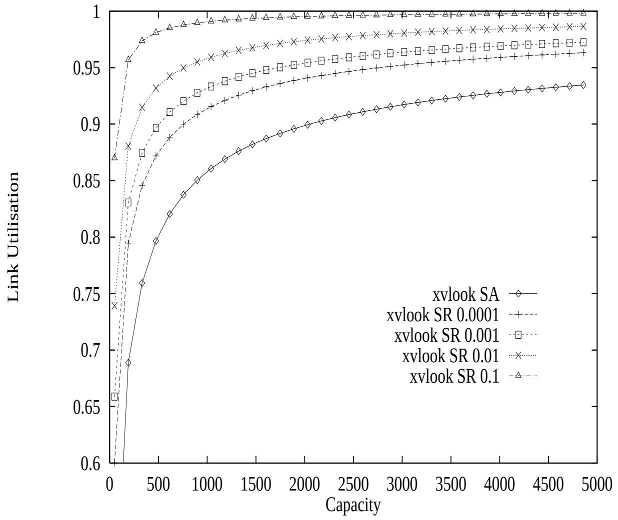

Figure 6: Link utilisation, SA and SR, “xvlook”, type 1

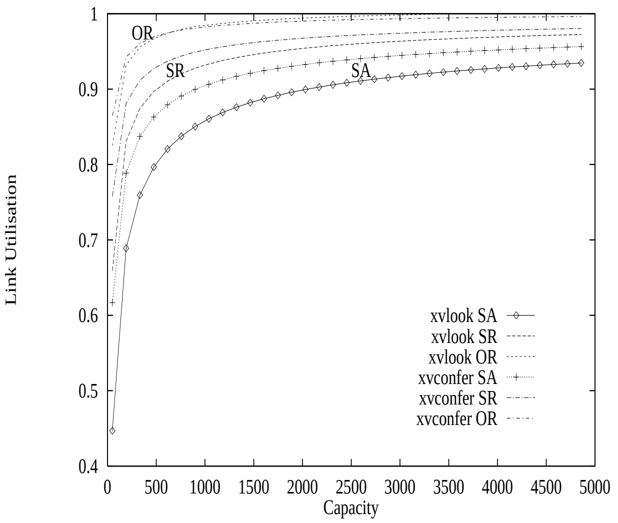

Figure 7: Link utilisation under SA, SR and OR, type 1 sources

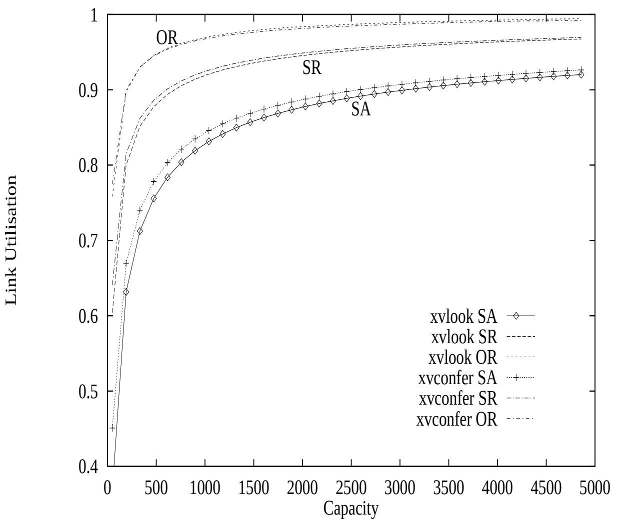

Figure 8: Link utilisation under SA, SR and OR, type 2 sources

yields a utilisation greater than 80% for links in excess of 750 Mbit/s, whereas both renegotiation schemes achieve 90% utilisation in this range. Static allocation has a utilisation limit of just over 90% whereas optimal renegotiation achieves almost 100% and simple renegotiation about 96%.

7 Conclusion

We have presented approximate models of an ATM multiplexer subjected to bursty traffic from general Markovian sources, under static allocation and simple and optimal renegotiation. Models of the three control schemes have been constructed using the Large Deviations Approximation for an unbuffered resource.

A comparison of the three schemes using sources parameterised from working applications shows that whilst renegotiation permits more sources to be accommodated by a link of fixed capacity, yielding consequently higher link utilisation, the gain to the network operator is small if the source peak to link rate ratio is small. For higher values of this ratio, renegotiation becomes more attractive. We note however that local area networks are frequently under-utilised and therefore suggest that implementation of renegotiation on such networks would not be useful. Simple renegotiation is far from optimal in performance and should renegotiation be implemented, better algorithms should be sought.

The author wishes to gratefully acknowledge the enthusiastic help of Nigel Bean, Frank Kelly and Ian Leslie.

References

[1] J. Bae and T. Suda. Survey of traffic control schemes and protocols in ATM networks. *Proc. IEEE*, 79(2):170–189, Feb. 1991.

[2] N. Bean. Mathematical modelling of packet switched communication networks. Technical report, University of Cambridge Statistical Laboratory, 16 Mill Lane, Cambridge CB2 1SB, 1992.

[3] P. Boyer. A congestion control for the ATM. In *Broadband Technologies*, New Jersey, Oct. 1990. Proc. Seventh ITC Specialist Seminar.

[4] P. Boyer, L. Louvion, and D. Tranchier. Intelligent multiplexing in an ATM network. In *Multimedia '90*, Bordeaux, France, 1990. Third IEEE COMSOC.

[5] A. Campbell, G. Coulson, F. Garcia, and D. Hutchison. A continuous media transport and orchestration service. In *Proc. SIGCOMM*, Maryland, USA, 1992.

[6] A. Campbell, D. Hutchison, H. Leopold, and N. Singer. Distributed multimedia communication system requirements. Technical Report OSI95 Deliverable ELIN-1/C/V1, Esprit Project 5341, 1991.

[7] B. Doshi. Performance of in-call buffer/window allocation schemes for short intermittent file transfers over broadband packet networks. In *Proc. INFOCOM*, Florence, May 1992.

[8] B. Doshi and S. Dravida. Congestion control in high speed wide area packet networks: In-call parameter negotiations. In *Broadband Technologies*. Proc. Seventh ITC Specialist Seminar, Oct. 1990.

[9] B. Doshi and H. Heffes. Performance of an in-call buffer–window reservation scheme for long file transfers. *IEEE JSAC*, 9(7):1013–1023, Sept. 1991.

[10] A. Hopper. Pandora – an experimental system for multimedia applications. *Operating Systems Review*, 24(2):19–34, Apr. 1990.

[11] J. Y. Hui. Resource allocation for broadband networks. *IEEE JSAC*, 6(9):1598–1608, Dec. 1988.

[12] M. G. Katevenis. Fast switching and fair control of congested flow in broadband networks. *IEEE JSAC*, SAC-5(8), Oct. 1987.

[13] F. Kelly. Effective bandwidths at multi-class queues. *Queueing Systems*, 9:5–16, 1991.

[14] I. Leslie. *Extending the Local Area Network*. PhD thesis, University of Cambridge Computer Laboratory, Pembroke Street, Cambridge, CB2 3QG, Feb. 1983.

[15] D. McAuley. *Protocol Design for High Speed Networks*. PhD thesis, University of Cambridge Computer Laboratory, Pembroke Street, Cambridge, CB2 3QG, Jan. 1990.

[16] D. Tranchier, P. Boyer, Y. Rouaud, and J.-Y. Mazeas. Fast bandwidth allocation in ATM networks. In *Proc. XIV Int. Switch. Symp.*, volume 2, pages 7–11, Leige, Belgium, Oct. 1992.

[17] K. Tutufor. On admission control and policing in an ATM based network. In *Broadband Technologies*, New Jersey, Oct. 1990. Proc. Seventh ITC Specialist Seminar.