Predicting bandwidth requirements of ATM and Ethernet traffic

Simon Crosby[1], Ian Leslie[1], Meriel Huggard[2,3], J.T. Lewis[2],
Brian McGurk[2], and Raymond Russell[2]

February 29, 1996

Abstract

In this paper, a technique based on the on-line estimation of the entropy of a traffic stream is proposed for the prediction of bandwidth requirements. We report on the results of experiments designed to test the method. These experiments show that conservative estimates, within 3% of the true value as determined experimentally, can be achieved for Ethernet traffic. Experiments using ATM traffic on the Fairisle ATM LAN show that estimates of the effective bandwidth requirements can be obtained within 5% of the true value.

1 Introduction

The concept of *effective bandwidth* was introduced by Hui [8]. Kelly [6] considered various simple models of a queue which serves traffic from a number of distinct sources and which is required to deliver a performance guarantee. He showed that an effective bandwidth can be associated with each source and that the queue can deliver its performance guarantee by limiting the sources served so that their effective bandwidths sum to less than the capacity of the queue. Kesidis *et al.* [9] showed the existence of effective bandwidths for multiclass Markov fluids and other types of source that are used to model asynchronous transfer mode (ATM) traffic. They discussed a call-acceptance algorithm for ATM networks based on calculated effective bandwidths. Courcoubetis *et al.* [2] addressed the issue of call-acceptance in ATM networks based on an algorithm which can be used by a switch to predict its spare capacity. In Duffield *et al.* [5], we exploited the analogy between the large deviation rate-function of an arrivals process and thermodynamic entropy to propose an algorithm for the estimation of cell-loss ratio at a buffered switch based on on-line measurements of cell arrival times. In this paper, we investigate the application of this idea to the on-line estimation of effective bandwidth.

2 The Effective-Bandwidth Approximation

Consider a single-server queue with finite waiting-space: if the arrival-rate can exceed the service-rate, then the waiting-space will sometimes be full and customers arriving during such times will be refused entry to the system. It is obvious that the bigger the waiting-space, the smaller the loss ratio (the ratio of the number of refused arrivals to the total number of arrivals).

Suppose for simplicity that the arrivals to the queue can be represented by a stochastic process and that the service-capacity is deterministic; again, it is obvious that the higher we set the service-rate, the smaller will be the loss-ratio. Let s be the constant service-rate, let b be the buffer-size (the size of the waiting-space) and let $\Gamma(b, s)$ be the loss-ratio. If the arrivals process is stationary and mixing, then the loss ratio decays exponentially for large buffer-sizes:

$$\lim_{b\to\infty} \frac{1}{b} \log \Gamma(b, s) = -\delta(s).$$

[1]University of Cambridge Computer Laboratory.
[2]Dublin Institute for Advanced Studies.
[3]Author presenting paper.

The decay-rate $\delta(\cdot)$ is determined by the service-rate s of the queue and the *Large Deviations* of the arrivals process:

$$\delta(s) = \inf\{\theta : \lambda(\theta) \leq s\theta\}.$$

Here λ is the *scaled Cumulant Generating Function* (sCGF) of the arrivals process, defined by

$$\lambda(\theta) = \lim_n \frac{1}{n} \log E \exp(\theta A_n),$$

where A_n denotes the number of cells arriving in an interval of length n and E denotes the expectation with respect to the probability measure which describes the process.
This suggests the *effective bandwidth approximation*:

$$\Gamma(b, s) \approx e^{-b\delta(s)}.$$

Making this approximation, we identify $\lambda(\theta)/\theta$ as the *effective bandwidth function* of the arrivals process. Suppose we have a fixed buffer-size b and we wish to ensure that the loss ratio $\Gamma(b, s)$ is less than an acceptable level Γ_0, then we must allocate a minimum service rate $s_{min}(b, \Gamma_0)$ such that

$$\Gamma(b, s) \leq \Gamma_0$$

whenever s exceeds $s_{min}(b, \Gamma_0)$. Using the effective bandwidth approximation for $\Gamma(b, s)$, we find that s_{min} can be approximated by

$$s_{eff}(b, \Gamma_0) = \frac{\lambda(\delta)}{\delta}$$

where $\delta = -\log \Gamma_0 / b$.
The function $\lambda(\theta)/\theta$ is a non-decreasing function of θ; for small values of θ, its value is the mean arrival-rate; for large values of $theta$, its value is the peak arrival-rate. Thus, for large buffer-size, $s_{eff}(b, \Gamma_0)$ is close to the mean bandwidth of the traffic, the bandwidth as measured over long periods of time; for small buffer-size, s_{eff} is close to the peak bandwidth. If the arrivals were not stochastic but constant, then the required service-rate would simply be the arrival-rate and there would be no need for a buffer. This is not the case: the arrival-rate fluctuates; if there were no buffer, then the service-rate would need to be close to the peak arrival-rate in order to cope with the fluctuations and ensure small losses. The buffer helps the system handle the short-term fluctuations in the arrival-rate; if the buffer is very large, then a service-rate close to the mean arrival-rate is sufficient since the buffer will absorb most of the fluctuations. Thus the effective bandwidth s_{eff} measures the demand for sevice-rate that the arrivals stream places on the server.

2.1 A Refined Approximation

The effective-bandwidth approximation may be refined by including a prefactor:

$$\Gamma(b, s) \approx p_0(s) e^{-b\delta(s)}.$$

It often happens that $\log \Gamma(b, s)$ is a convex function of b; in this case we can get an upper bound on $\Gamma(b, s)$ by taking $p_0(s)$ to be an upper bound on $\Gamma(0, s)$. Such an upper bound is obtained by noting that, in an unbuffered system, an arriving customer will be lost only if it arrives before the previous customer is finished being served, and setting

$$p_0(s) = \mathrm{P}\left(\text{interarrival time} < \text{service time}\right) = \mathrm{P}\left(\text{interarrival time} < 1/s\right).$$

This bound is not tight since, even if a customer arrives within $1/s$ of the previous customer, it may find the server free if the previous customer itself was lost. The resulting refined bandwidth approximation may now be defined by

$$s_{ref}(b, \Gamma_0) := \inf\left\{s : p_0(s) e^{-b\delta(s)} \leq \Gamma_0\right\}.$$

In a queue with a large buffer, a service-rate close to the mean arrival-rate will suffice to keep the loss ratio to acceptably small levels. At such a service-rate, the losses in an unbuffered system would be close to 100%, giving a prefactor close to 1. In which case, the dominant contribtution to the bound on the loss ratio is $e^{-b\delta(s)}$ and the refined approximation is very close to the effective bandwidth approximation. However, in a queue with a small buffer, a much higher service-rate is required to keep losses small. In the corresponding unbuffered system, losses are also comparitively rare and so the approximation $\Gamma(0, s) \approx \mathrm{P}\,(\text{interarrival time} < 1/s)$ is very accurate. In this case, we have a small and accurate prefactor $p_0(s)$ which, since b is small, is the dominant part of the refined approximation. Thus we expect that for small buffers the refined approximation should work better than the effective bandwidth approximation and that for large buffers they should both give similar results.

2.2 Estimating Bandwidth Requirements from Traces

If we are given a sample realisation $\{X_1, X_2, X_3, \ldots\}$ of the arrivals stream of a queuing system, we can estimate its bandwidth requirements on the basis of the two approximations developed above. We have developed methods of estimating the Large Deviation properties of the arrivals stream [5],[4]; a simple estimate of the sCGF is based on finding the asymptotic distribution of the block sums, $\tilde{X}_k$ defined for each block-size B by

$$\tilde{X}_1 := \sum_{k=1}^{B} X_k, \quad \tilde{X}_2 := \sum_{k=B+1}^{2B} X_k, \quad \ldots$$

Since the arrivals are assumed stationary and mixing, these block sums will be approximately I.I.D. for B large enough, and so

$$\lambda(\theta) \approx \lambda_B(\theta) := \frac{1}{B} \log \mathrm{E} e^{\theta \tilde{X}_1}$$

This suggests using the normalised CGF of the empirical distribution of the block sums as an estimator for λ; this in turn implies an estimator for δ:

$$\hat{\lambda}_B^n(\theta) := \frac{1}{B} \log \frac{B}{n} \sum_{i=1}^{n/B} e^{\theta \tilde{X}_i}$$

$$\hat{\delta}_B^n := \sup\,\{\theta : \lambda_B^n(\theta) \leq s\theta\}$$

We can use this to give us an estimate of the effective bandwidth function in the obvious manner:

$$\text{Effective bandwidth estimate} = \frac{\hat{\lambda}(\hat{\delta})}{\hat{\delta}}.$$

We can also use it to give us a refined estimate of the bandwidth requirements using the refined approximation:

$$\hat{s}_{ref}(b, \Gamma_0) := \inf\left\{s : \hat{p}_0(s) e^{-b\hat{\delta}(s)} \leq \Gamma_0\right\},$$

where $\hat{p}_0(s)$ is the empirical frequency with which interarrival times are less than the service-time, and $\hat{\delta}(s)$ is given by

$$\hat{\delta}(s) := \sup\left\{\theta : \hat{\lambda}(\theta) \leq s\theta\right\}.$$

The statistical properties of $\hat{\delta}$, the estimate of the asymptotic decay rate of Γ in b, have already been investigated in [5] using a two-state Markov chain for the arrivals stream.

3 Experimental Results

3.1 Networked "Doom"

Doom is a fast, popular PC game featuring a virtual environment in which computer- and user-controlled characters move and interact. A free implementation is included in the Slackware distribution of Linux, using UDP/IP rather than IPX as the networking protocol. In a network game, one machine is the master and the others are slaves. The state of the game is maintained by each machine broadcasting information to all the others. The master machine broadcasts packets at an almost constant rate, but each slave broadcasts at a variable rate which depends on the activity of its human user. A plot of the bitrate of a slave's broadcast is typically as shown in figure 1.

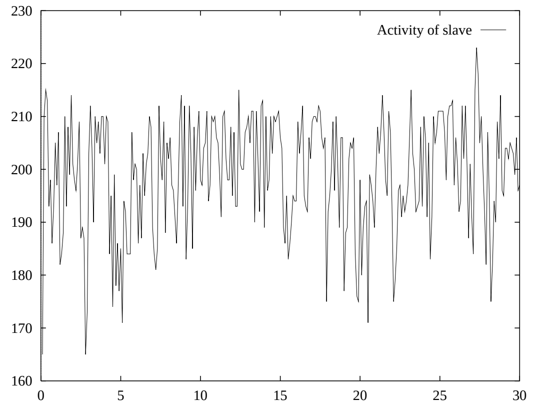

Figure 1: Activity level (cells per interval) vs time (min) for the slave broadcast in the two-player game

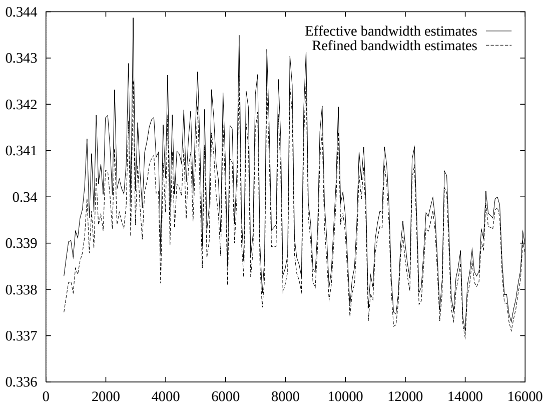

Figure 2: The two bandwidth estimates vs block-size for a two player game

We set up some multi-player network games on PC's running Linux over our Ethernet and monitored the network traffic thus generated using the Sun/OS command `etherfind` on a Sun workstation. This provided the source, destination and size of each udp packet, along with a time-stamp indicating to a resolution of one tenth of a millisecond when it was detected at the Sun workstation's ethernet interface.
The trace shown in figure 1 is that recorded during a half-hour long two-player game. We used our bandwidth estimates on this trace to test if we could predict accurately the bandwidth it would require in a queuing system.

Each of the bandwidth estimates is based on an estimate of the sCGF of the activity of the source. In order to form this latter estimate, we must choose an appropriate block-size, as explained in section 2.2. For a given statistical model of the activity, an optimal choice of block-size will, in general, exist and be calculable. It is not possible to determine such a block-size for a source on the basis of just one trace; however we find that the estimates are very similar over a broad range of block-sizes and, for each estimate required, we present the estimates formed with 200 different block-sizes.

First, we compared the effective-bandwidth estimate with the refined estimate and found that, for the parameters used, they agree almost exactly, with the the effective-bandwidth estimate slightly more conservative (higher) than the refined one. We estimated the service-rate required to achieve a packet-loss ratio of 0.0002 in a buffer capable of holding 100 packets and figure 2 shows the values of the two estimates against block-size.

The results are similar for other values of buffer-size and target loss ratio, so from now on we only present results for the refined estimate.

Figure 3 shows the refined estimate against block-size along with a solid line showing the true value of service-rate required. The latter is determined experimentally, repeatedly passing the trace through a software simulation of a queuing system at different service-rates until the minimum

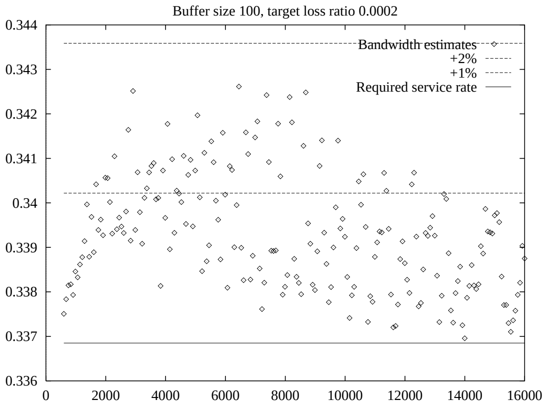

Figure 3: Refined bandwidth estimates vs blocksize showing the true minimum required service-rate

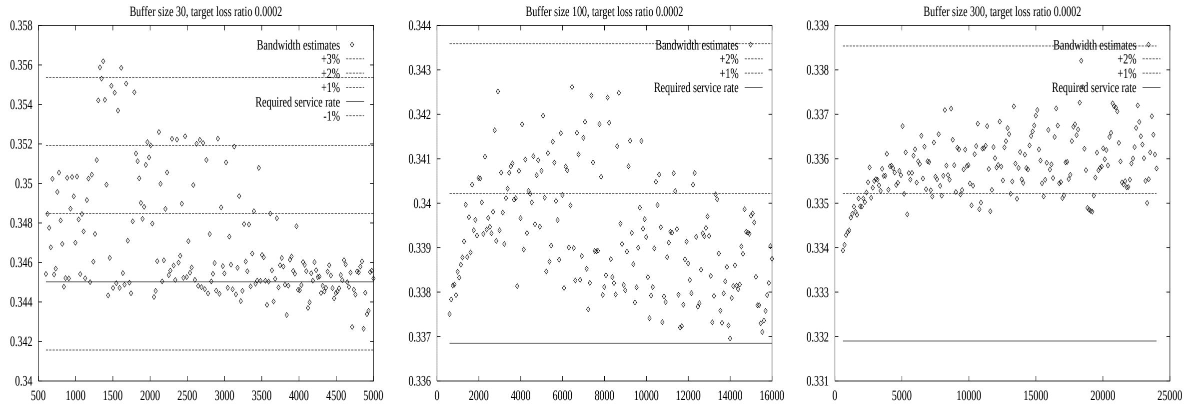

Figure 4: Refined estimates vs blocksize for different buffer sizes

acceptable rate is found. The fainter lines lie at 1 and 2 per cent of the minimum service-rate above it, showing that all the bandwidth estimates in this range of block-sizes lie within 2 per cent of the true required bandwidth and, furthermore, that they are all conservative. Thus, if we allocate service to this traffic on the basis of these estimates, we will achieve or exceed the desired quality of service and we will do so very efficiently, to within 2 per cent of the minimum resources necessary.

Figure 4 shows the refined bandwidth estimates against block-size for three different buffer-sizes: 30, 100 and 300. In each case the desired loss ratio is 0.0002. We see that for small buffer-sizes, the results are not as accurate; this is not surprising since the estimates are based on large-buffer approximations. They are still accurate to within a few per cent.

Figure 5 shows the refined bandwidth estimates against block-size for three different desired loss ratios: 0.001, 0.0002 and 0.00005. In each case the buffer-size is 100. Again we see that the estimates are very accurate and consistent over different block-sizes.

It is interesting to ask whether the bandwidth estimates can, not only measure the resources needed by a traffic source, but also predict its future needs. Another way of posing this question is to ask how long the network game must be running before we can get a good estimate of the resources it requires. Figure 6 shows the refined bandwidth estimates (for a buffer of size 100 and a target loss ratio of 0.0002) based on the first minute, 3 minutes, 6 minutes, 10 minutes and quarter of an hour of the duration of the game. We see that, as we might expect, accuracy increases as the estimates use more and more of the trace but that, even after just one minute, we can accurately predict the service-rate required with a conservative estimate.

These encouraging results are not a property of this one trace; this is not a just one particularly well behaved trace, it is characteristic of network games of Doom in general. To illustrate this, we also gathered a trace from a four-player game, and estimated the bandwidth requirements of the broadcast each of the slaves. A plot of the slaves' bitrate is shown in figure 7.

Figure 8 shows the results of estimating the bandwidth required by each of the slaves in buffer

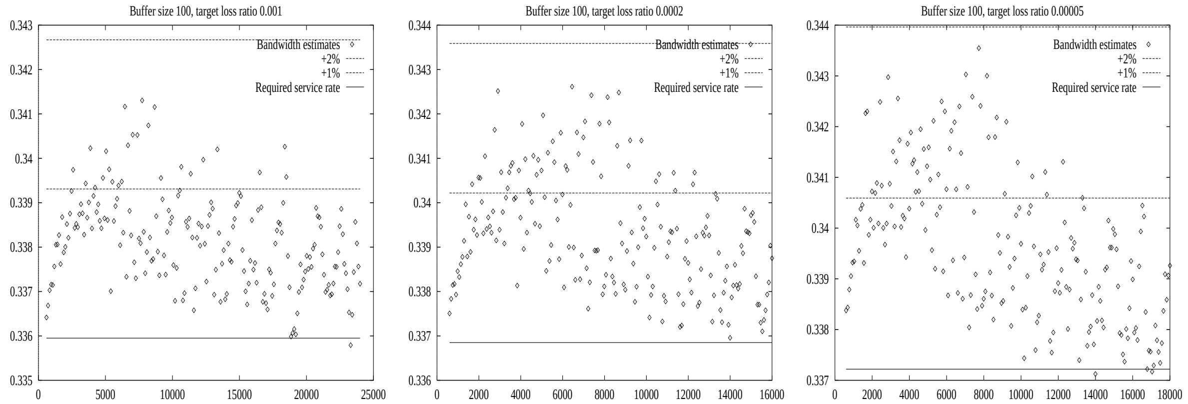

Figure 5: Refined estimates vs blocksize for different target loss ratios

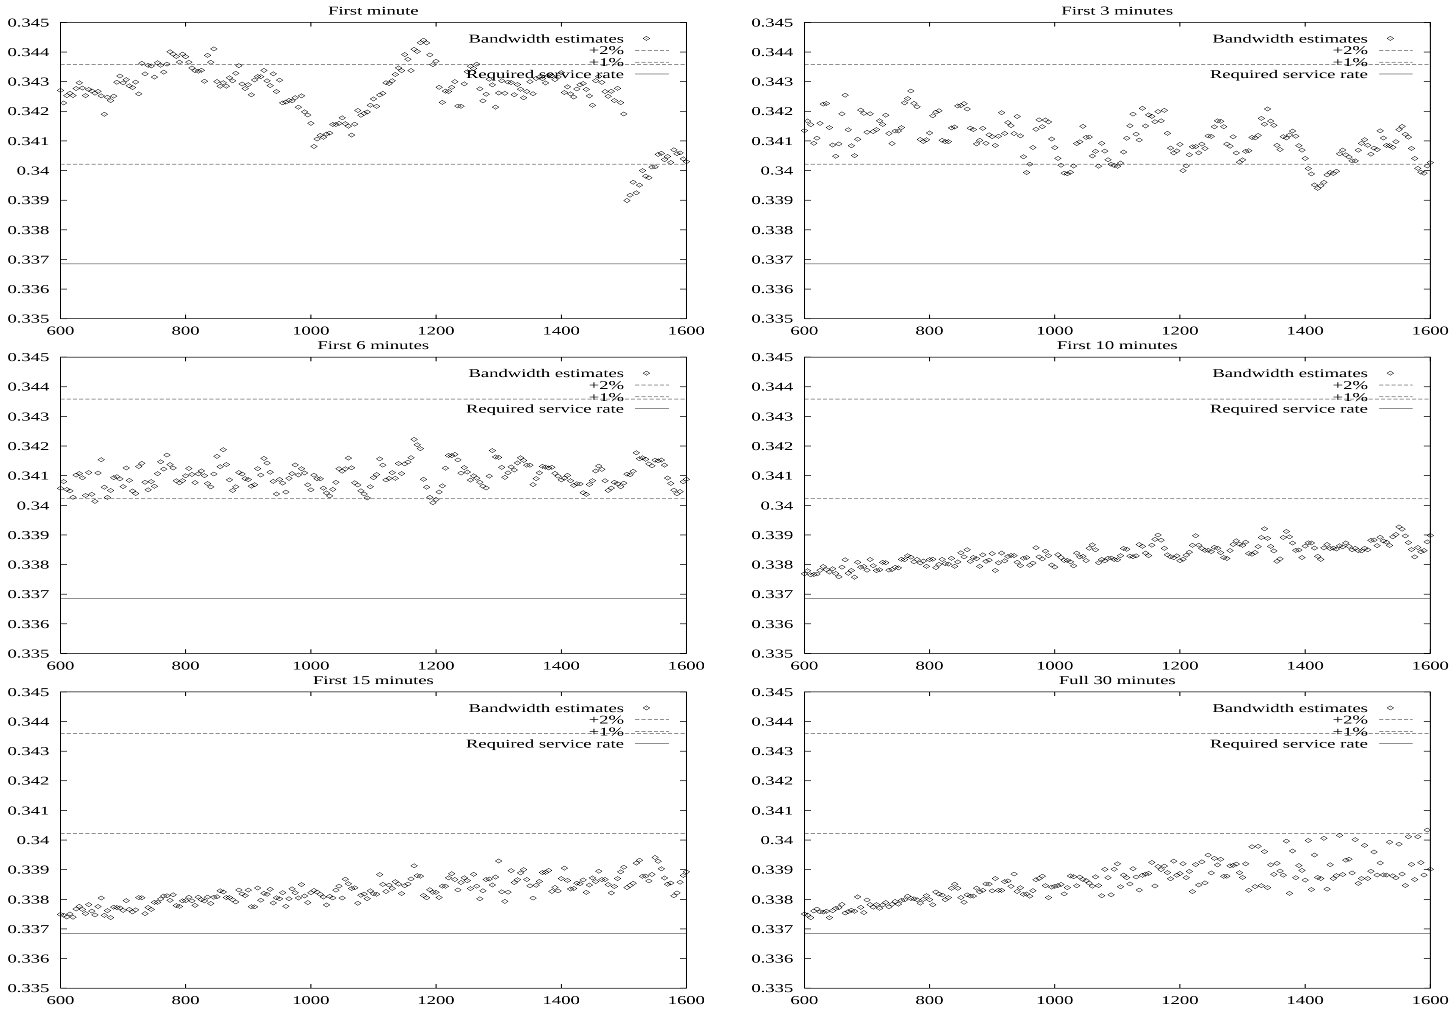

Figure 6: Refined estimates based on different lengths of observation

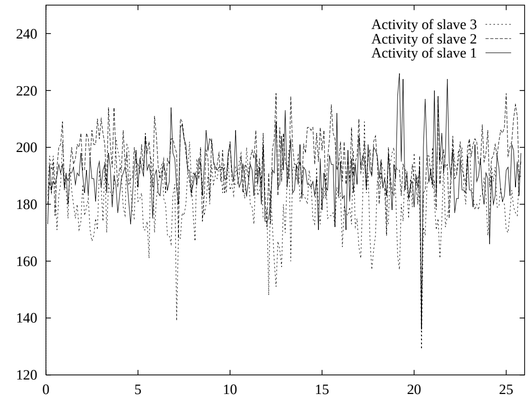

Figure 7: Activity level (cells per interval) vs time (minutes) for the slaves' broadcasts in the four-player game

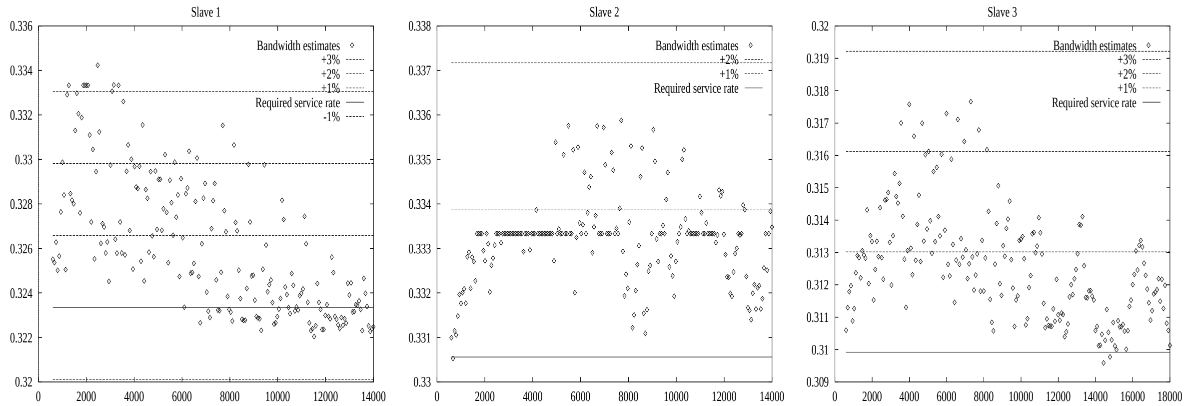

Figure 8: Refined estimates of the slaves' broadcast bandwidths

of size 100 with a target loss ratio of 0.0002.

3.2 Star Wars Video Traces

In a separate experiment, we investigated estimates of effective bandwidth on real ATM traffic produced on the Fairisle ATM network [1]. Fairisle is an experimental ATM LAN whose design permits a high degree of experimental flexibility, including high resolution clocks for timestamping and measurement of traffic, and a general purpose CPU and memory on every switch port.

Fairisle is equipped with a wide range of ATM traffic sources, including compressed video, audio and LAN data. In addition traffic can be generated from a pre-recorded trace of traffic activity. One such source is the Star Wars movie. The Star Wars data set comprises the information content, in bytes per frame, for 171,000 frames (approximately 2 hours) of the film, as transmitted by a Discrete Cosine Transform (DCT) based codec, with a compression algorithm similar to that of JPEG. This data set has been studied in some detail, as in [7]. More information on its exact content may be found in [3].

The data which we used was constructed from the Star Wars traces by encoding each slice of video as a single AAL5 PDU; each PDU was transmitted in such a fashion that its cells were evenly spaced throughout the slice time. This smoothing effect over the slice time serves to reduce the peak rate of the source from the line rate (100 Mb/s , if all cells in a slice are transmitted back-to-back) to about 24.2 Mb/s for the worst slice, and is typical of ATM video sources. The mean rate of the source is about 5.3 Mb/s.

We chose to work with the effective bandwidth estimate because it is much less computationally intensive than the refined estimate. Although it is, as in the case of the Doom data, slightly less accurate, it allowed us to perform our experiments with much greater speed and flexibility.

The Star Wars trace is more challenging than the Doom traces; it shows a much greater vari-

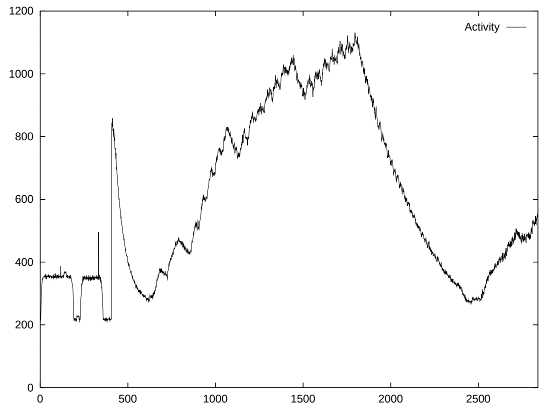

Figure 9: Activity level (cells per interval) vs time ($\frac{1}{24}$secs) for the opening section of the Star Wars film

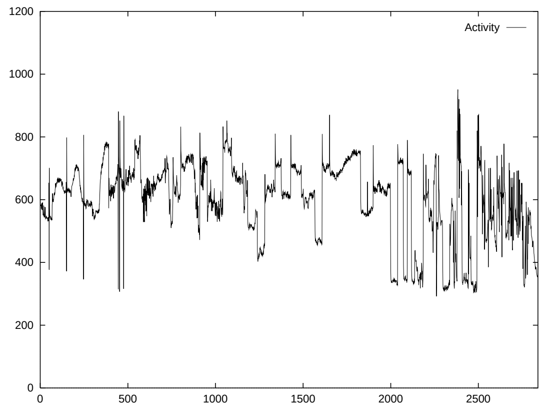

Figure 10: Activity level (cells per interval) vs time ($\frac{1}{24}$secs) for the second section of the Star Wars film

ability of characteristics than do the latter. Its bandwidth requirement changes greatly over short periods of time, making it difficult to measure effective bandwidths accurately; furthermore, the changes are by their very nature impossible to predict. Consider the plot of the bitrate of the first sixtieth of the trace, which comprises roughly the first minute and a half of the film, as shown in figure 9.

Those familiar with the film may recognise in the activity plot the two titles which are displayed right at the beginning, followed by some lines of text which scroll into the background until they disappear. The initial space-battle scene starts in the second sixtieth of the trace, whose bitrate is shown in figure 10. Unless one had seen the film before, it is impossible to know that the bitrate would show such a strong initial upward trend followed by an equally clear downward trend. This behaviour is completely unpredictable and, while it is not so strikingly evident in the activity plots, is mirrored throughout the trace. Such behaviour poses serious challenges to any statistical estimation or prediction scheme.

Nevertheless, it is possible to estimate effective bandwidths reasonably accurately for the Star Wars trace. Figure 11 shows the effective bandwidth estimates against block-size for a range of buffer-sizes/target loss ratios.

While the estimates are not as accurate as in the case of the Doom traffic (we cannot expect them to be), we see that they are still consistent over a wide range of block sizes. They are sometimes optimistic, underestimating the required service; if we allocated bandwidth to this video stream on the basis of these estimates, we would experience higher than acceptable loss ratios. However, in our experience with this trace, we find that we can consistently estimate to within five per cent accuracy on the optimistic side. If this is true of video streams in general, then it suggests the practical precaution of adding five per cent to all estimates as a safety margin.

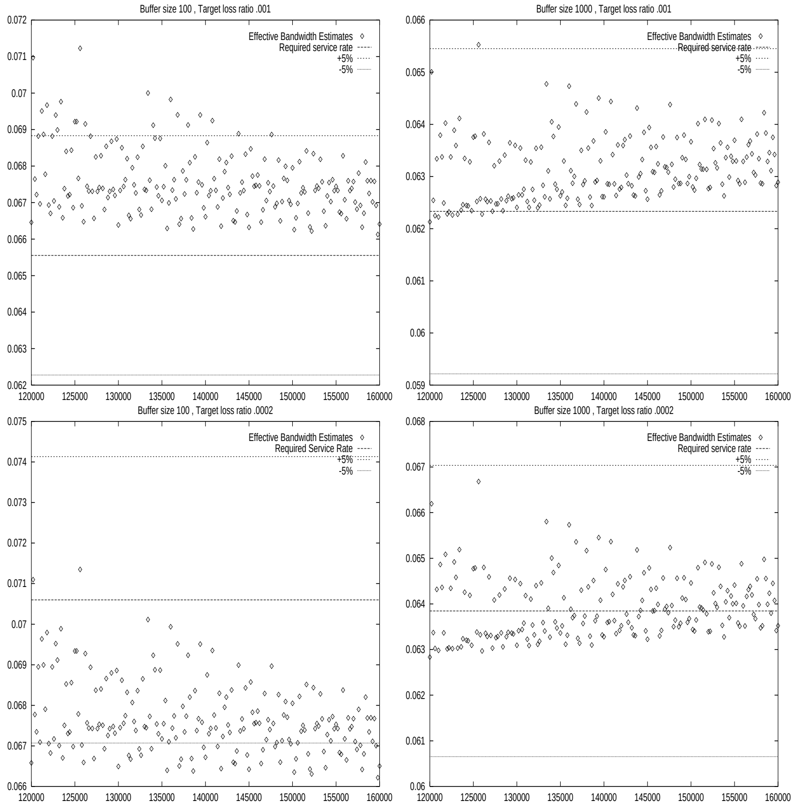

Figure 11: Effective Bandwidth estimates for the Star Wars Data

4 Conclusions

In this paper we have outlined a method for estimating the effective bandwidth required by traffic in a network. We described two experiments to estimate the effective bandwidth of traffic on a LAN: one on an ethernet and the other on the Fairisle ATM network. From these experiments we found that we can consistently predict the required effective bandwidth of the traffic to within 3% of the true value for the ethernet traffic and to within 5% for the traffic on the Fairisle ATM network. Moreover, in the case of the ethernet traffic, the estimates formed were found to be conservative.

This work demonstrates that the on-line estimation of entropy can lead to useful estimates of the effective bandwidth required by a given source, and that these estimates may be calculated using a very simple, on-line method. We have shown that performance guarantees could be delivered at the cost of marginally over allocating bandwidth.

5 Acknowledgement

The work presented in this paper has been carried out under the ESPRIT contract MEASURE: Resource allocation for Multimedia Communication and Processing Based on On-Line Measurement.

References

[1] R. Black, I. Leslie, and D. McAuley. Experiences of building an ATM switch for the Local Area. In *SIGCOMM*, volume 24(4), September 1994.

[2] C. Courcoubetis, G. Kesidis, A. Ridder, J. Walrand and R. Weber. Admission control and routing in ATM networks using inferences from measured buffer occupancy. In *IEEE Trans. Commun.*, **43**, pages 1778-1784, 1995.

[3] Simon Crosby, Ian Leslie, John T. Lewis, Neil O'Connell, Raymond Russell, and Fergal Toomey. Bypassing Modelling: an Investigation of Entropy as a Traffic Descriptor in the Fairisle ATM network. Proceedings of 12th UK Teletraffic Symposium.

[4] Simon Crosby, Meriel Huggard, Ian Leslie, John T. Lewis, Fergal Toomey, and Cormac Walsh. Bypassing Modelling: Further Investigations of Entropy as a Traffic Descriptor in the Fairisle ATM network. In *Proceedings of 1st Workshop on ATM Traffic Management*, pages 139-146, Paris 1995.

[5] N.G. Duffield, J.T. Lewis, N. O'Connell, R. Russell and F. Toomey. Entropy of ATM Traffic Streams: A tool for Estimating QoS Parameters. In *IEEE Journal on Selected Areas in Communications*, volume 13(6), August 1995.

[6] F.P. Kelly. Effective bandwidths at multi-class queues. In *Queueing Systems*, 9(1991), pages 9-16.

[7] M.W. Garrett and M. Willinger. Analysis, modelling, and generation of self-similar vbr video traffic. In *Proceedings, 1994 SIGCOMM Conference*, pages 269-280, London, 1994.

[8] J.Y. Hui. Resource allocation for broadband networks. In *IEEE Journal on Selected Areas in Communications*, **6**, pages 1598-1608, 1988.

[9] G. Kesidis, J. Walrand and C.-S. Chang. Effective bandwidths for multiclass Markov fluids and other ATM sources. In *IEEE/ACM Tran. Networking*,Volume 1(4), pages 424-428, 1993.