In Proceedings of the 13th U.K. Workshop on Performance Engineering of Computer and Telecommunication Systems (UKPEW'97), July 1997, Ilkley, U.K.

Investigation of Cell Scale and Burst Scale Effects on the Cell Loss Probability using Large Deviations

Costas Courcoubetis*, Vasilios A. Siris,* and Richard Weber†

Abstract

Considerable research activity has been directed towards estimating the Cell Loss Probability (CLP) at the buffer of an ATM switch and understanding the ways in which it can occur. Much of this activity has been conducted by modeling the cells which enter the buffer as a continuous fluid. This model can capture the variability of a source at a burst level, but it ignores the fact that the workload actually arrives in discrete cells. Nonetheless the fluid model can give accurate estimates of the cell loss probability when the buffer size is not very small. If the switch has a very small amount of buffer per source then cell level effects can not be ignored. We consider both Constant Bit Rate (CBR) and periodic on/off Variable Bit Rate (VBR) sources, and apply large asymptotic techniques to a cell level model of an ATM output link. Our analysis simultaneously captures both the cell scale and burst scale effects, enabling us to study the boundary between regions in the parameter space where cell level effects are or are not significant. In addition to accurately computing the CLP, we are able to give an insightful qualitative description of how cell loss occurs in very small buffers.

Keywords: ATM, large deviations, burst scale, cell scale, random phases, fluid approximation

1 Introduction

The cell loss probability at the buffer of an ATM switch has an important impact on service quality. For this reason, it is interesting to know how it is affected by buffer size and switch bandwidth and to understand how buffer overflow can occur. The aim of this paper is to investigate how cell loss can arise from effects at various time scales. The investigation takes place in terms of the following model. We suppose there are N identical sources and that N is large. Each source is a periodic on/off source with deterministic lengths of "on" and "off" phases of duration T_{on} and T_{off}, respectively. During each "on" phase the source produces cells at rate h, i.e., when h is measured in cells per second, it produces one cell every $\tau = 1/h$ seconds. Cells

*ICS-FORTH and Department of Computer Science, University of Crete, P.O. Box 1385 GR 711 10 Heraklion, Greece. Email: {courcou,vsiris}@ics.forth.gr

†University of Cambridge, Statistical Laboratory, 16 Mill Lane, Cambridge CB2 1SB, U.K. Email: rrw1@statslab.cam.ac.uk

enter a shared buffer of size $B = Nb$ which is served by a deterministic server of rate $C = Nc$, i.e., when c is measured in cells per second, one cell is served every $1/(Nc)$ seconds. Parameters b and c are respectively the buffer size and service capacity per source.

We assume that $\tau = 1/h$ is very small compared to T_{on} and T_{off}. For each source, the start of the first "on" phase following time 0 is uniformly distributed on $(0, T_{\text{on}} + T_{\text{off}}]$, independently of other sources. Let $\bar{p}N$ be the average number of sources that are "on" and let m be the mean rate of each source, i.e., $\bar{p} = T_{\text{on}}/(T_{\text{on}} + T_{\text{off}})$ and $m = \bar{p}h$. Figure 1 shows the model for $N = 4$.

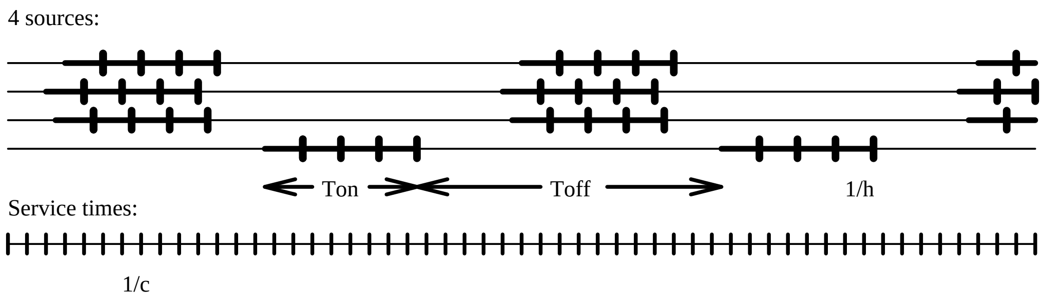

Figure 1: Cell level model. Cell scale effects occur due to the synchronization of cell arrivals from different streams.

Cell loss at an ATM output link can arise from both *cell scale* and *burst scale* effects [Rob91]. A burst scale effect takes place when the number of sources that are in their "on" phase, say aN, is such that their combined rate exceeds the server capacity, i.e., $aNh > Nc$; in this case, the buffer will start to fill. A cell scale effect takes place when the synchronization of the sources which are "on" is such that they deliver cells to the buffer at nearly the same time, e.g., as for the second and third sources in Figure 1. If a large number of sources are synchronized in this manner, then for an interval of time the aggregate arrival rate can exceed Nc, even when only the average number of sources $\bar{p}N$ are "on" and each has $h < c$. Finally, if $b < 1$, it is possible that the cell scale effect can be enough on its own to fill the buffer.

Much of the analysis of queueing and multiplexing in ATM switches has been conducted using a model in which the bursts of discrete cells are replaced by a *continuous fluid*. This fluid is either "on", at a constant rate h, or "off". The model captures the burst scale effect, but ignores any cell scale effect. Whether or not this approximation is good depends on the length of the typical time over which the buffer content increase from 0 to B just prior to cell loss. If this time, say t, is large in comparison to the cell scale $1/h$, then the fluid model is good. Typically, this is the case when b is large. If t is large compared to $1/h$) but small compared to T_{on} and T_{off}, then it can be a good approximation to assume that each source is fully "on" or fully "off" over the typical period during which the buffer fills; this gives the simplification that the cell loss probability (CLP) depends on the source statistics only through the peak and mean rates, and the bufferless on/off fluid model can be used [Hui88, Kel91, Kel96].

If the time for buffer overflow t is in the order of magnitude of $1/h$, which occurs when the buffer per source b is very small, cell scale effects can not be disregarded; in this case, use of the bufferless fluid model may underestimate the CLP by several orders of magnitude. As an

illustration, suppose we have 331 identical on/off sources with peak rate $h = 4.5$ Mbps and peak rate to mean rate ratio $h/m = 3$, which are multiplexed in a link with capacity $C = 622$ Mbps and total buffer $B = 30$ cells. The bufferless on/off fluid approximation gives a CLP less than 10^{-8}, when in fact the actual CLP is 10^{-6}.

To analyze the above model, we apply the continuous-time version of the large asymptotic techniques which were developed in [CW96]. Our approach simultaneously captures the effects at the cell scale and burst scale, and accurately computes the cell loss probability. This contrasts with other work which has addressed either the cell scale or the burst scale alone. Using a simple heuristic, we are able to investigate the qualitative nature of cell loss. We show that there is a critical buffer size above which cell scale effects are no longer active, in which case it is valid to adopt the fluid model. However, when the amount of buffer per source is very small, both burst and cell scale effects are present (see Figure 2). The burst scale effect occurs when the empirically observed proportion of sources which are in their "on" phase is above average; the cell scale effect occurs when the "on" sources of the above burst are "in phase", creating a large cell arrival rate for some short interval of time. In Figure 2, these are labeled *empirical mean deviation* and *random phases* effects, respectively.

Approximations for the cell scale component have been proposed [DRS91, RSKJ91, NRSV91, FLVO94], but at a time before the large deviation analysis of buffer overflow was well understood. These approximations did not correctly capture the effects at both times scales and can give erroneous estimates of the CLP. Fortuitously, they happen to be very accurate for the cases of practical interest, i.e., for CLPs less than 10^{-10}.

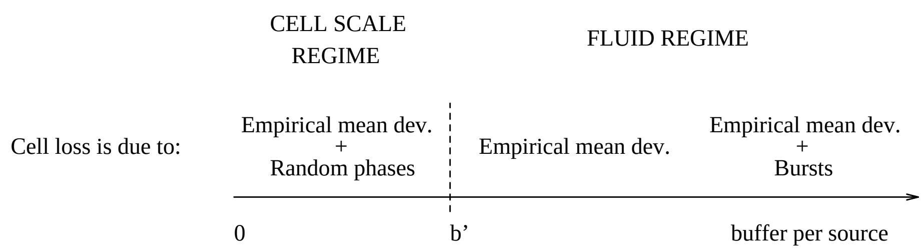

Figure 2: Cell loss regimes. The value of b' depends on the network and source parameters.

The rest of the paper is structured as follows. In Section 2, we review the discrete-time many sources asymptotic developed in [CW96], and show that it can be generalized to a continuous-time form which is able to capture the cell scale effects. In Section 3, we apply this asymptotic to constant bit rate (CBR) traffic (i.e., $T_{\text{off}} = 0$) and compare the CLP estimated using our approach with that estimated using other approximations. In Section 4, we apply this asymptotic to our model for periodic on/off sources. Specifically, in Section 4.1, we use a simple heuristic to investigate the qualitative nature of cell loss for very small buffers. The heuristic is based on standard notions of large deviation theory and is validated through numerical comparisons which show that it is very accurate. Using this heuristic, we describe the sharp boundary between regions where the cell scale effect is or is not relevant, and explain why previously proposed approximations for the cell scale component are accurate for practical situations. Section 6 concludes the paper.

2 The many sources asymptotic

In this section, we review the discrete-time many sources asymptotic developed in [CW96], and present a continuous-time version which can capture cell scale effects.

Our model for an output link is a single server of rate $C = Nc$ and a finite buffer of size $B = Nb$. The service discipline is assumed FIFO (First In First Out). Let $X[0,t]$ denote the number of cells generated by a single source in the interval $(0,t]$. We suppose sources are stationary, hence this represents a typical time interval of length t.

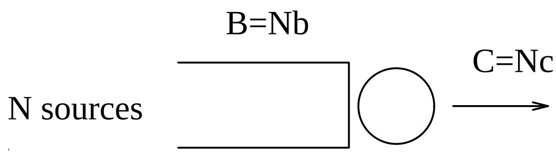

Figure 3: ATM output link model.

Let $\varphi(s,t)$ denote the log moment generating function

$$\varphi(s,t) = \log \mathrm{E}\left[e^{sX[0,t]}\right] . \tag{1}$$

Suppose we adopt a fluid model for the input traffic. In this model, we choose a small positive δ, fixed for all N, and imagine that $X[0,t]$ increases linearly over each interval $t \in (k\delta, k\delta+\delta]$, $k = 0, 1, \ldots$. Thus, the number of cells in the buffer is treated as a real-valued variable (although in actuality it is an integer). Service is also imagined to proceed at a constant rate, rather than cell by cell.

For this model of input traffic, the cell loss probability (CLP) has an exponential tail, as does the probability $\mathrm{P}\left(Q^N > Nb\right)$ that in a system with infinite buffer the occupancy is more than B. Furthermore, both can be approximated as, [CW96]

$$\mathrm{P}\left(Q^N > Nb\right) \approx \mathrm{CLP} \approx e^{-NI} , \tag{2}$$

where

$$I = \inf_{t=\delta,2\delta,\ldots} \sup_s \{s(b+ct) - \varphi(s,t)\} . \tag{3}$$

The approximation (2) is asymptotically exact for large systems in the sense that $\lim_{N\to\infty}(1/N)\log \mathrm{CLP} = -I$ and similarly for $\mathrm{P}\left(Q^N > Nb\right)$.

If we are to make a more refined analysis which can account for the fact that the data flow is not a fluid, but rather a stream of discrete cells, then one needs a continuous-time version of the (2)–(3) in which $\delta \to 0$. The extension is to replace (3) by

$$I = \inf_{t>0} \sup_s \{s(b+ct) - \varphi(s,t)\} , \tag{4}$$

where now t is a continuous variable. A sketch of the proof of the validity of (2) with (4) is given in [CSW96]. A similar result has been proved in [BD95].

3 Constant bit rate sources

In this section, we consider independent and identical constant bit rate (CBR) sources whose phases are randomly distributed. This is the model introduced in Section 1 with $T_{\rm off} = 0$. We suppose source i has a rate of h cells per second and generates a new cell at each of the times $\ldots, U_i - 1/h, U_i, U_i + 1/h, U_i + 2/h, \ldots$, where the U_i are independent and uniformly distributed on $[0, 1/h]$. For stability, we require $h < c$, equivalently $\rho = h/c < 1$.

3.1 The bufferless model

Suppose there is no buffer. The probability that cell loss occurs during an interval of length $\Delta = 1/C$, i.e., an interval during which the server can serve exactly one cell, is just the probability that the N sources together produce more than one cell during this interval. A single source produces either 0 or 1 cell during such an interval, with probabilities $1 - \rho/N$ and ρ/N respectively. Since the sources are independent the total number of cells that they produce in an interval of length $1/C$ has the binomial distribution $B(N, \rho/N)$. Thus

$$\begin{aligned}
\text{CLP} &= \text{P(more than 1 cell in } \Delta) \\
&= 1 - \text{P(0 cells in time } \Delta) - \text{P(1 cell in time } \Delta) \\
&= 1 - (1-\rho/N)^N - N(\rho/N)\,(1-\rho/N)^{N-1} \\
&\stackrel{N\to\infty}{\to} 1 - (1+\rho)e^{-\rho}
\end{aligned}$$

The limit is the probability that a Poisson random variable with mean ρ is greater than 1. Note that in this case CLP $\not\to 0$ as $N \to \infty$. The fluid model would give CLP $= 0$.

3.2 The buffered model

We now consider the case where the server has a buffer. During the interval $(0, t]$ the expected number of cells produced by a single source is ht. Let $k = \lceil ht \rceil$ be the greatest integer not exceeding ht. The actual number of cells which are produced by one source in this interval of length t is then either k or $k+1$, depending on the phase of the source relative to time 0. A little thought shows that for a random phase CBR source (1) has the evaluation

$$\begin{aligned}
\varphi(s,t) &= \log\left[(k+1-ht)e^{sk} + (ht-k)e^{s(k+1)}\right] \\
&= s(k+1) + \log\left[(k+1-ht)e^{-s} + (ht-k)\right] .
\end{aligned}$$

If we substitute this into (4) the minimizing t must be no greater than $1/c$. For suppose $t > 1/c$. Consider $t' = t - 1/c$ and note that for all s,

$$\begin{aligned}
s(b+t'c) - \log \mathrm{E}\left[e^{sX[0,t']}\right] &= s(b+tc) - s - \log \mathrm{E}\left[e^{sX[0,t]-sX[t',t]}\right] \\
&\leq s(b+tc) - s - \log \mathrm{E}\left[e^{sX[0,t]}\right] + s \\
&= s(b+tc) - \log \mathrm{E}\left[e^{sX[0,t]}\right] ,
\end{aligned}$$

since, by stability, a single source certainly produces no more than 1 cell in an interval of length $1/c$. Therefore $k = 0$ and

$$I_p(\rho, b) = \inf_{0<t\leq 1/c} \sup_s \{s(b+ct) + \log[(1-ht)e^{-s} + ht]\} \tag{5}$$

$$= \inf_{0<t\leq 1} \sup_s \{s(b+t) + \log[(1-\rho t)e^{-s} + \rho t]\} , \tag{6}$$

where now it is convenient to show I_p as a function of ρ and b. The suffix p is a reminder that this rate function is for CBR sources and measure a cell level effect due to their random phases.

3.3 Comparison with other approximations

Next, we compare the value of CLP obtained using the many sources asymptotic, (2) and (6), with the values obtained using some other approximations.

The CLP in a system with a finite buffer of size $B = Nb$ has the same asymptotic for large N as $\mathrm{P}\left(Q^N > Nb\right)$, the probability that in a system with an infinite buffer the occupancy is more than B. For the latter, there is an exact formula, [NRSV91]

$$\mathrm{P}\left(Q^N > Nb\right) = \sum_{Nb<i\leq N} \binom{N}{i} \left(\frac{i-Nb}{1/\rho}\right)^i \left(1 - \frac{i-Nb}{1/\rho}\right)^{N-i} \frac{1/\rho - N + Nb}{1/\rho - i + Nb} . \tag{7}$$

However, this formula gives little insight. We also have the following two approximations, which are not asymptotically exact.

$$\text{(i)} \qquad \mathrm{P}\left(Q^N > Nb\right) \approx e^{-\frac{2B^2}{N} - 2B(1-\rho)} = e^{-N(2b^2+2b(1-\rho))} \tag{8}$$

$$\text{(ii)} \qquad \mathrm{P}\left(Q^N > Nb\right) \approx -\frac{1-\rho}{\log(\rho)} e^{-\frac{2B^2}{N} - B(1-\rho-\log(\rho))}$$

$$\approx e^{-N(2b^2+b(1-\rho-\log(\rho))} . \tag{9}$$

The first of these is based on a Brownian bridge approximation and holds for large N and ρ close to 1 [DRS91, NRSV91]. The second appears in [FLVO94].

Figure 4 displays a comparison of (7), (8), (9), and the 'large N' approximation based on (2) and (6). These graphs are for an output link with $C = 155$ Mbps, sources with rate $h = 1$ Mbps, and utilization $\rho = 0.3$ (N=46 sources) and $\rho = 0.8$ (N=124 sources). All approximations are close when N is large and ρ is close to 1. One can show by power series expansions that (6), (8) and (9) differ only by $O((1-\rho)^2)$. Equation (8) is inaccurate when the utilization is small and (9) is accurate for the range of CLP that is of practical concern, i.e., about 10^{-10}. Even for much smaller CLP and larger b, (9) only slightly underestimates the CLP.

4 Periodic on/off sources

In this section, we estimate the CLP at a multiplexer serving a number of identical periodic on/off sources. Each source is either "on" (at rate h) or "off" (at rate 0). The probability

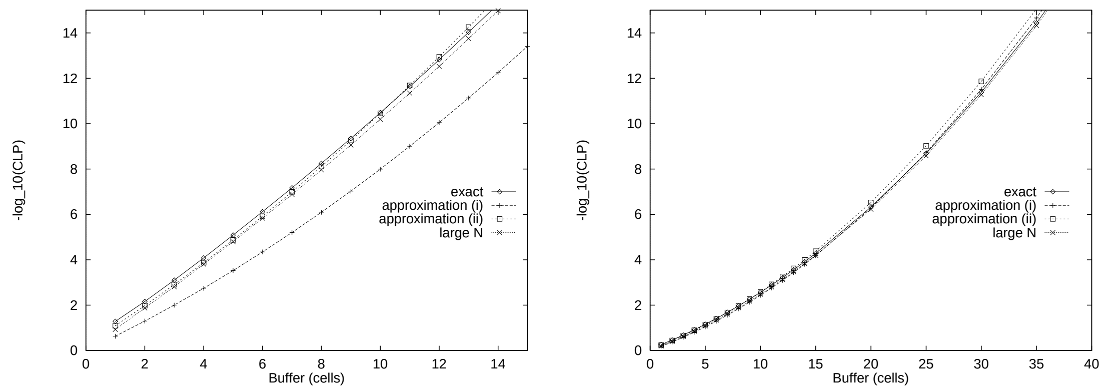

(a) $\rho = 0.3$

(b) $\rho = 0.8$

Figure 4: Comparison of the large N approximation with other approximations and exact calculation. (sources are constant bit rate with $h = 1$ Mbps)

a source is "on" is $\bar{p} = T_{\text{on}}/(T_{\text{on}} + T_{\text{off}})$ and $m = \bar{p}h$ is the source's mean rate. The system utilization is $\rho = m/c$ where, as previously, c is the capacity per source.

We consider two approaches. The first approach is heuristic: a simple but sound intuitive argument enables us qualitatively to explain the nature of cell loss and estimate the CLP when b is small (where cell loss is due to a combination of cell scale and burst scale effects). The second approach is brute force: we analytically compute $\varphi(s,t)$ and apply (2) and (6) (see [CSW96]). This method is algebraically messy because $\varphi(s,t)$ is complicated and the calculation can only be completed numerically. While the second method is accurate and valid for any buffer size, the heuristic method turns out to be very accurate also.

4.1 Heuristic for small buffers

For small buffers, overflows will occur when some number of sources are "on" and there is cell loss because cells from a number of different sources arrive very close together. We approximate the CLP with the probability of the most probable way this can happen. If aN sources are "on" during the time which the buffer fills, then the overflow occurs according to the model discussed in Section 3.2, where the the "effective utilization" is $\rho_a = ah/c = a\rho/\bar{p}$ and the "effective buffer" per "on" source is $b_a = b/a$. This is the case because, as we have already discussed in the introduction, we assume that T_{on} and T_{off} are large compared to the cell scale $1/(Nc)$ and, in the small buffer regime we are investigating, the time for buffer overflow is on the order of $1/(Nc)$. Hence, the number of sources that contribute to the overflow remain practically constant. Thus, an intuitive derivation is the following

$$\mathrm{P(cell\ loss)} \quad = \quad \int_{a=0}^{1} \mathrm{P(cell\ loss} \mid aN \text{ sources on}) \times d\mathrm{P}(aN \text{ sources on})$$

$$
\begin{aligned}
&\approx \int_{a=0}^{1} e^{-aNI_p(\rho_a,b_a)} \times e^{-NI_f(a)}\, da \\
&= \int_{a=0}^{1} e^{-N[aI_p(\rho_a,b_a)+I_f(a)]}\, da &(10)\\
&\approx e^{-N\min_a[aI_p(\rho_a,b_a)+I_f(a)]}\,, &(11)
\end{aligned}
$$

where $I_p(\rho_a, b_a) = 0$ for $\rho_a \geq 1$ and is given by (6) for $\rho_a < 1$; hence, $e^{-aNI_p(\rho_a,b_a)}$ estimates the cell loss probability in a system of aN CBR sources, with utilization ρ_a and buffer per source b_a. The term $e^{-NI_f(a)}$ is the large deviation approximation to the probability that at least aN out of N sources are on. $I_f(a)$ is given by

$$I_f(a) = a\log\frac{a}{\bar{p}} + (1-a)\log\left(\frac{1-a}{1-\bar{p}}\right)\,. \tag{12}$$

Finally, (11) follows from (10) using Laplace's argument that the large N asymptotic behavior of this integral is determined by the largest term in the integrand.

Accuracy of the heuristic

Figure 5 shows that the CLP given by the above heuristic equals the CLP estimated by the large N asymptotic (equations (2) and (6)), using the analytical expression of $\varphi(s,t)$ for periodic on/off sources.

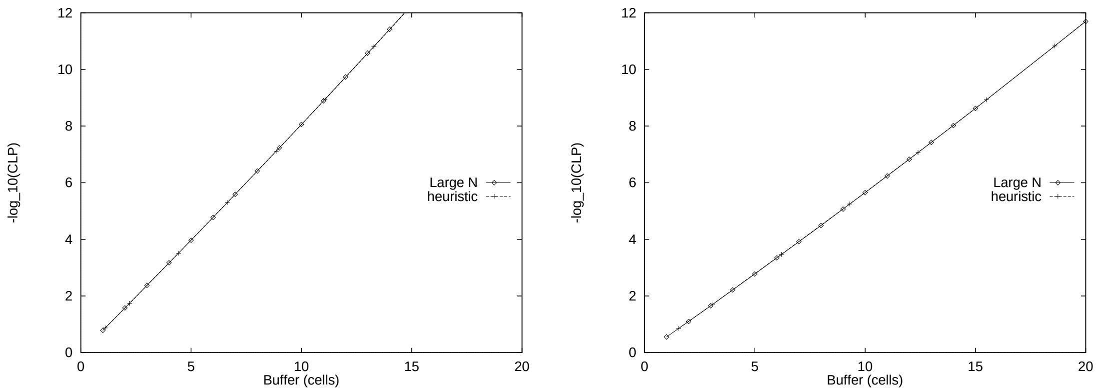

(a) $c/h = 0.7$ ($\rho = 0.36$) (b) $c/h = 0.5$ ($\rho = 0.5$)

Figure 5: Comparison of the heuristic for small buffers with numerical computation using direct application of the large N asymptotic. ($C = 155$ Mbps, $h = 1$ Mbps, $\bar{p} = m/h = 0.25$)

Dependence of a^* on the buffer size

Denote $a^* = \mathrm{argmin}_a[aI_p(\rho_a, b_b) + I_f(a)]$. Figures 6 shows that a^* increases with b. Furthermore, for small buffer sizes, $a^* \approx \bar{p}$, i.e., the number of "on" sources is close to the average and the cell loss is almost completely due to cell scale effects.

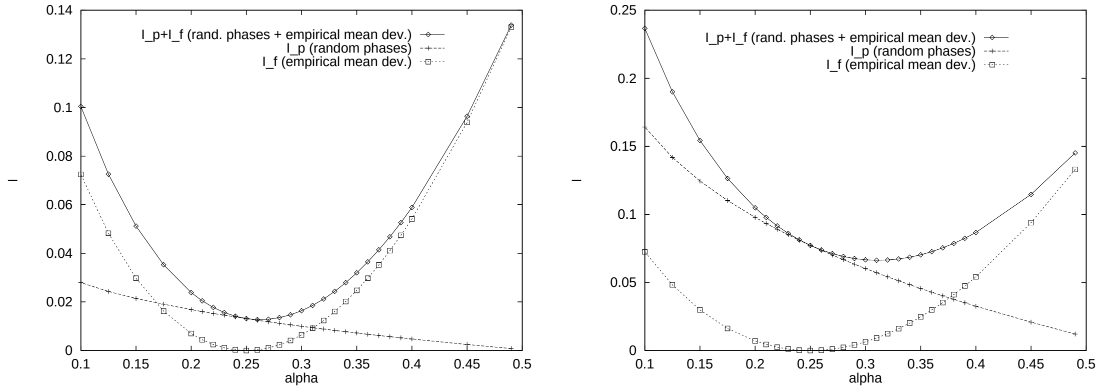

(a) $b = .01$ cells

(b) $b = .05$ cells

Figure 6: Heuristic for small buffers. As the buffer per source b increases, the contribution due to the burst scale component increases. For small b, the most probable way to have losses is to have the mean number of sources "on". ($\bar{p} = m/h = 0.25$, $c/h = 0.5$)

Boundary between cell scale regime and fluid regime

In the bufferless on/off model, cell loss occurs as soon as the aggregate rate of incoming cells exceeds the capacity of the link, i.e., if aN sources are "on" and a is such that $aNh > C$, or equivalently, $a > c/h$. Thus P(cell loss) $\approx e^{-NI_f(c/h)}$, where I_f is defined in (12). At the boundary of the cell scale and fluid scale regimes, $\min_a\{aI(ah/c, b/a) + I_f(a)\} = I_f(c/h)$. The boundary is shown in Figure 7. Observe that for CLP values of practical interest, the buffer per source is small. Specifically, for sources with $h = 1$ Mbps and $m/h = 0.25$, the buffer size above which cell scale effects are no longer important (Figure 2) is $b' \approx 0.018$ cells for $C = 622$ Mbps and $b' \approx 0.05$ cells for $C = 155$ Mbps. For such small buffers, $a^* \approx \bar{p} = m/h$, hence from (11) we have:

$$\text{P(cell loss)} \approx e^{-NI}, \text{ where } I = \bar{p}I_p(\bar{p}h/c, b/\bar{p}),$$

and using (9) (since it is the more accurate of (8) and (9)) we get, with $\rho = h\bar{p}/c$,

$$I = 2b^2/\bar{p} + b(1 - h\bar{p}/c - \log(h\bar{p}/c)). \tag{13}$$

In [RSKJ91, NRSV91, FLVO94], the cell scale component when N periodic on/off sources are multiplexed is approximated by the cell loss in a system in which there are N sources and each source is imagined to be CBR with rate $h' = m$. The utilization in this system takes the correct value, $\rho' = h\bar{p}/c$, but otherwise this approximation has no particular justification. However, (9) gives

$$I = 2b^2 + b(1 - h\bar{p}/c - \log(h\bar{p}/c)). \tag{14}$$

Equations (13) and (14) differ in one term: the first term on the righthand side of (14) is $2b^2$, where it appears that the correct asymptotic has $2b^2/\bar{p}$. However, recall from our previous

discussion that for CLP of practical importance the buffer per source is small: e.g., $b \approx 0.013$ cells for $C = 622$ Mbps, CLP $= 10^{-8.4}$, utilization $\rho = 0.6$, and $m = 0.25$ Mbps. For such values, the error in the first term of (14) is insignificant. In conclusion, for CLP values of interest, the rate function can be approximated by

$$I = b(1 - h\bar{p}/c - \log(h\bar{p}/c)) = b(1 - \rho - \log \rho)\,. \tag{15}$$

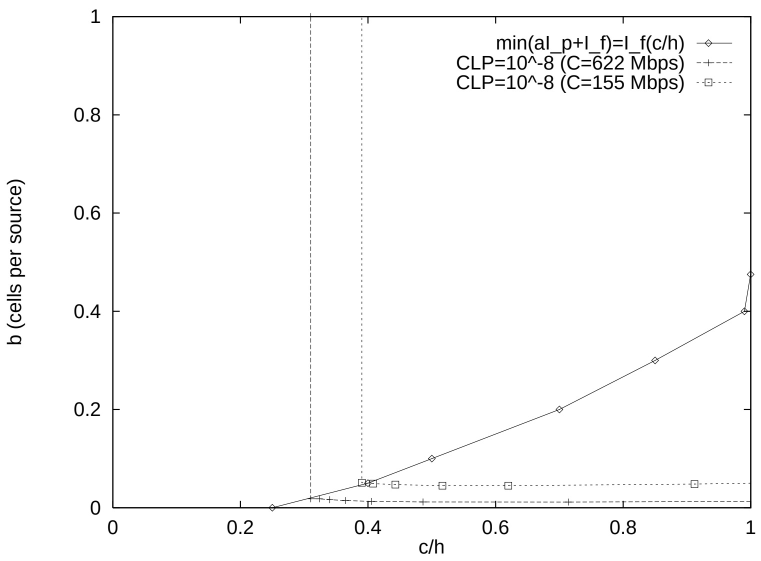

Figure 7: Area where random phases have an effect (lower right). On the boundary (bold line) $\min\{aI(ah/c, b/a) + I_f(a)\} = I_f(c/h)$. The area that corresponds to CLP range of interest is located near the x-axis.

5 Numerical results

When ATM switch designers architect their switches so that these can fit in a single VLSI chip, the amount of buffering available in the chip for real-time traffic is very much constrained and amounts to a few hundreds cells to be shared by all links. Hence an important question is if this amount of buffer suffices to compensate for the burstiness of the incoming traffic. Traditionally, since the amount of buffer is small compared to the burst sizes, one would neglect the cell level effects and predict cell loss from the fluid approximation for the bufferless case. In this section we investigate the accuracy of such an approach.

We consider periodic on/off sources with $h = 1$ Mbps and $m/h = 0.25$, and compare the CLP estimated using a bufferless on/off fluid model with the estimate using the many sources asymptotic. The results for different link capacities are shown in Tables 1, 2, and 3.

Observe that, in the cell scale regime, the CLP depends primarily on B and ρ, and is independent of C. This agrees with the approximation in (15). Also, for fixed ρ and CLP, a greater C allows a lesser buffer per source b. This is due to the more efficient statistical

multiplexing for larger link capacities. The total amount of buffer for which the CLP is less than 10^{-8} grows with the link capacity and the utilization, and amounts to 45 cells for a 1200 Mbps link with $\rho = 0.8$. Hence, if the total amount of buffer exceeds 100 cells, we can safely use the simple fluid approximation for the bufferless case. Of course, this assumes that the utilization due to real-time traffic will not approach 1, which will be the case if the links also carry second priority best-effort traffic.

6 Conclusions

We have applied a continuous-time version of the large asymptotic in [CW96] in order to understand when the cell scale effect is or is not significant. For a periodic on/off model of VBR traffic, we simultaneously capture the effects of both the cell scale and burst scale. By applying a simple heuristic, we have been able to give a qualitative description of the way cell loss can occur in very small buffers, namely, that cell loss occurs due to the combination of two events: a deviation that takes the number of "on" sources above its mean value, and a synchronization amongst the random phases of "on" sources. The heuristic is very accurate and is motivated by standard ideas of large deviation theory. An open issue is to discover if there are scenarios for which the heuristic is exact.

Some approximations for cell loss that have been proposed for this regime can be very inaccurate. However, they are accurate for the size of cell loss probability that is interesting in practice, i.e., CLP $\approx 10^{-10}$. This is due to the small values of buffer per source that is typical for such CLPs. When we leave the cell scale regime and enter the fluid regime, the exponential tail of the cell loss is due mainly to a deviation in which the number of "on" sources is above its mean value. The assumption that T_{on} and T_{off} are infinite is not a major restriction and the conclusions in the paper apply when "on" and "off" phases are of random lengths. We have assumed that all sources are identical, but it should be possible to consider mixtures of sources of different types. We have noted that the cell scale effect has a Brownian bridge approximation. It is easy to compute the log moment generating function for mixtures of Brownian bridges with different parameters.

Ongoing work also includes studying the statistical multiplexing gain, and how it is affected by the buffer size, when we multiplex real traffic sources, such as videotelephone and MPEG compressed video, which are policed by the leaky bucket mechanism.

References

[BD95] D. D. Botvich and N. Duffield. Large deviations, the shape of the loss curve, and economies of scale in large multiplexers. *Queueing Systems*, 20:293–320, 1995.

[CSW96] C. Courcoubetis, V. A. Siris, and R. Weber. Cell scale and burst scale effects: An asymptotic approach for calculating cell loss probability. Technical Report no. 181, ICS-FORTH, December 1996.

[CW96] C. Courcoubetis and R. Weber. Buffer overflow asymptotics for a switch handling many traffic sources. *Journal of Applied Probability*, 33, 1996.

[DRS91] L. G. Dron, G. Ramamurthy, and B. Sengupta. Delay analysis of continuous bit rate traffic over an ATM network. *IEEE J. Select. Areas Commun.*, 9(3):402–407, April 1991.

[FLVO94] G. Fiche, W. Lorcher, R. Veyland, and F. Oger. Study of multiplexing for ATM traffic sources. In *ITC - 14*, pages 441–452, North Holland, 1994. Elsevier Science B. V.

[Hui88] J. Y. Hui. Resource allocation for broadband networks. *IEEE J. Select. Areas Commun.*, 6:1598–1608, 1988.

[Kel91] F. P. Kelly. Effective bandwidths at multi-class queues. *Queueing Systems*, 9:5–16, 1991.

[Kel96] F. P. Kelly. Notes on effective bandwidths. In F. P. Kelly, S. Zachary, and I. Zeidins, editors, *Stochastic Networks: Theory and Applications*, pages 141–168. Oxford University Press, 1996.

[NRSV91] I. Norros, J. W. Roberts, A. Simonian, and J. T. Virtamo. The superposition of variable bit rate sources in an ATM multiplexer. *IEEE J. Select. Areas Commun.*, 9(3):378–387, April 1991.

[Rob91] J. W. Roberts. Variable-bit-rate traffic control in B-ISDN. *IEEE Commun. Mag.*, pages 50–56, April 1991.

[RSKJ91] C. Rasmussen, J. H. Sorensen, K. S. Kvols, and S. B. Jacobsen. Source-independent call acceptance procedures in ATM networks. *IEEE J. Select. Areas Commun.*, 9(3):351–358, April 1991.

	$\rho = 0.4, N = 248$		$\rho = 0.6, N = 372$		$\rho = 0.8, N = 496$	
Buffer (cells)	b	$-\log_{10}(CLP)$	b	$-\log_{10}(CLP)$	b	$-\log_{10}(CLP)$
5	0.020	3.57	0.013	2.10	0.010	0.97
10	0.040	7.25	0.027	4.29	0.020	2.02
15	0.060	11.05	0.040	6.57		
20	0.081	14.96	0.054	8.94		
on/off fluid		33.69		10.70		2.14

Table 1: C = 155 Mbps

	$\rho = 0.4, N = 995$		$\rho = 0.6, N = 1492$		$\rho = 0.8, N = 1990$	
Buffer (cells)	b	$-\log_{10}(CLP)$	b	$-\log_{10}(CLP)$	b	$-\log_{10}(CLP)$
5	0.005	3.53	0.0034	2.07	0.0025	0.95
10	0.0101	7.09	0.0067	4.16	0.0050	1.91
15	0.0151	10.67	0.0101	6.28	0.0075	2.90
20	0.0201	14.29	0.0134	8.41	0.0101	3.90
25	0.0251	17.93	0.0168	10.57	0.0126	4.92
30			0.0201	12.75	0.0151	5.96
35					0.0176	7.02
40					0.0201	8.10
on/off fluid		135.2		43.0		8.6

Table 2: C = 622 Mbps

	$\rho = 0.4, N = 1920$		$\rho = 0.6, N = 2880$		$\rho = 0.8, N = 3840$	
Buffer (cells)	b	$-\log_{10}(CLP)$	b	$-\log_{10}(CLP)$	b	$-\log_{10}(CLP)$
5	0.0026	3.52	0.0017	2.06	0.0013	0.94
10	0.0052	7.06	0.0035	4.14	0.0026	1.89
15	0.0078	10.61	0.0052	6.22	0.0039	2.85
20	0.0104	14.18	0.0069	8.32	0.0052	3.82
25			0.0087	10.43	0.0065	4.80
30			0.0104	12.55	0.0078	5.79
35					0.0091	6.79
40					0.0104	7.80
45					0.0117	8.82
50					0.0130	9.85
55					0.0143	10.89
60					0.0156	11.94
on/off fluid		260.79		82.86		16.53

Table 3: C = 1200 Mbps